

LINUX

THE UK'S BEST-SELLING LINUX MAGAZINE!

FORMAT

TURN OVER
FOR YOUR
FREE
LINUX PRO
MAGAZINE

SPEED UP LINUX!

Maximise your performance
with our expert guide to
memory, drives, hardware &
software optimisation **p51**

WHAT ON EARTH IS BSD?

Get the lowdown on some of Linux's closest cousins – and try
FreeBSD now, on this issue's CD or DVD **p58**

ARMARI OPTERON SERVERS ON TEST

64-bit powerhouse
reviewed in-depth **p20**

WHICH WEBSERVER?

What's the best option for
serving http? The Linux
contenders go head-to-head
in our Roundup:

- Apache ➤ Boa ➤ Caudium
- Roven Webserver ➤ thttpd
- Tux ➤ Zeus Webserver **p38**

WHAT NOW FOR MOZILLA?

Mitchell Baker on post-AOL
lizard development **p12**

ZEND STUDIO 3

Latest iteration of the leading
PHP IDE revealed **p32**

SERVER SCHOOL

Practical guide to running a
server – new series **p88**

"Initiatives like GNOME and
Mono will only improve with
Novell behind them"
– Miguel de Icaza **p6**

PLUS Yellow Dog Linux • MySQL and PHP optimization
• Scribus DTP • User's guide to GNOME • Advanced Blender

DVD issue also available Printed in the UK
LXF45 OCTOBER 2003 \$5.99
 9 771470 423019 
1.0
(13.99 outside UK and Republic of Ireland)

www.linuxformat.co.uk

Redline your Linux box

Whenever I consider my long history of using computers (like many, I started on a Sinclair ZX80), there always seem to be a few constants to the whole experience. One is obviously, that by the time you have set up a new computer and got it doing all the things you want, it is hopelessly out-of-date. The second is that nothing ever runs fast enough. I was considering this the other day as I waited for a kernel recompile. I started off with a 3.25MHz processor in the ZX80 (though as each instruction took four clock cycles, actually it was really less than 1MHz), but I can't say I waited around for it any longer than I do the 1.8GHz machine currently on my desk. Sure, I couldn't do as much, but there still seems to be just as much wasted time.

We could throw in the towel and suggest that it's just one of those fundamental laws – no matter how fast your computer gets, you spend the same amount of time waiting for it to do things. Or, scorning such apathy, we could tweak and tune our boxes to try and turn every beat of

the crystal into a productive one. That's the aim of our feature this month – whatever you use Linux for, you should be doing it faster!

The news that AOL was pulling the plug on the Mozilla project caused some consternation – is it really the end of the road for this flagship Open Source project? Mozilla has already given the world a lot of things, including a world-class browser, bugtracking system and mail client. But the story shouldn't end here – some even say having rid themselves of the shackles of corporate oversight, Mozilla should go from strength-to-strength. Make up your own mind with our exclusive story, including inside info from chief lizard-wrangler, Mitchell Baker.

There's definitely a server feel to this issue too – as well as our roundup of webserver software, this issue our new *Server School* tutorial series starts – it's what many of our readers have been asking for, so don't miss it.

As ever, enjoy this issue, and if you have any ideas, comments or suggestions, mail me!



Nick Veitch EDITOR

Squeeze every last computational drop out of your Linux box – tips and tricks!
p51

What next for Mozilla – we get the inside story, straight from the lizard's mouth
p12

Webserver software rounded up and rated
p38



12



38



51



AIMS OF THE MAGAZINE

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

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

MEET LINUX FORMAT'S TEAM OF WRITERS...



Andrew Channelle
Andy is the LXF chap with a nose for news, and is the novices' best friend with his great Beginners' tutorials.



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



Richard Drummond
A born tinkerer, Rich has been trying to perfect his Mac setup – by installing Linux.



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



Paul Hudson
SkYLang develops apace, now it can count and run loops. Taking over the world is surely not far behind.

David Cartwright
Veteran journalist and Linux consultant, he eats and sleeps real-world Linux usage and Open Source advocacy.

Hoyt Duff
Author of several *Red Hat Unleashed* books and long-time LXFer, if you can't find him, he's probably gone fishing.

Mike Saunders
Now considered an LXF veteran, he's the only writer to have caused the editor to laugh out loud. Intentionally...

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

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

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

LXF45 October 2003

Contents



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



COVER HIGHLIGHTS

Speed up Linux

Tweak and optimise your performance to use every last clock cycle to the full!

51

12 What now for Mozilla?

Mozilla.org is dead, long live the Mozilla Foundation!

32 EXCLUSIVE! Zend Studio 3

The leading PHP IDE ships loads of new features

38 Which webserver?

Old contenders and new pretenders battle it out

88 Server School

NEW SERIES has all you need for your own Linux server



'I may have a fight figuring out how to make the thing even start up in the first place. But hey, the client's paying me by the hour...'

11 David Cartwright becomes the Alpha geek

'The tools seem just above the sophistication of QuarkXPress 3, while the interface owes more to Adobe Pagemaker.'

29 Hold the front page – it's Linux DTP!

LINUXPRO

Storage Cost advantages of each option **Stop spam** IP Barrister David Harris on the law, junk email and you **Hosting focus** New in-depth series – this issue: 1&1 **LinuxExpo2003** LXF & LXPro will see you there!



20 Amari Opteron server – power at a paltry price!



6 News

Analysis of the latest Linux stories

14 Mailserver

Your points of view and tips

20 Amari Opteron Server review

64-bit cheap as pomme frites!

24 Yellow Dog 3.0 review

Linux on your Apple Mac

26 SNAP Graphics 2.0 review

If X is too heavy, this is much lighter

28 Revolution 2.0 review

Cross-platform RAD

29 Scribus 1.0 review

Open Source desktop publishing

30 Infomap Navigator review

Travel Europe with Linux

31 OpenGroupware.org review

Collaborative working with Linux

35 Book reviews

FreeBSD, Cryptography, Java & more

44 HotPicks

The best new Open Source software

58 What on Earth is BSD?

Open Source Linux alternative

64 Beginners' tutorial

GNOME desktop for novices

68 Red Hat Package management

Alternatives to the official tools

72 Compiler writing – flex & bison

SKYLang takes over the main mag!

76 Blender tutorial

Advanced modelling – Bezier curves

80 The GIMP tutorial

Create your own functional skins

84 PHP tutorial

Simple rules to accelerate queries

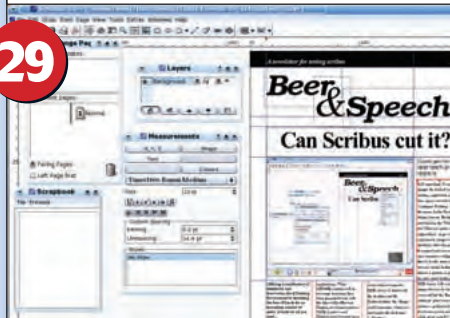
94 Answers

Your queries answered

110 User Groups

Local support and fun

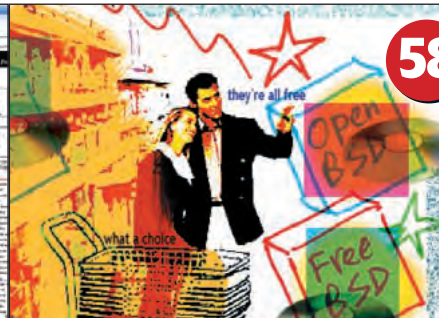
29



35



58



80



COVERDISCS

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



CDS A AND B

FreeBSD 4.8 It's not Linux, but another UNIX-alike distro for you to explore – BSD extras are on CDB; **K3b** Small name for a big CD burning program for audio and data; **Kognition** Optical character recognition; **Beneath A Steel Sky** Revolution's classic game is now playable on Linux with **ScummVM** virtual machine; **Java Runtime Environment** run code & apps; **Thunderbird** Mozilla's email client



DVD

Topologilinux Linux from Windows without repartitioning; **JRE Software Development Kit** create Java progs; **WifiScanner** find those hotspots for totally free Web surfing!

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



SAVE MONEY!

SUBSCRIBE TO LINUX FORMAT

See page 102 or phone 0870 4448645



Newsdesk

● Kroupware released ● Protect your system from Linux viruses ● SCO News ● Help GnuCash
● Native Savage and Puzzle Pirates games for Linux ● Embedded Linux News ● PDA voice recognition

MONKEYING WITH THE LINUX BUSINESS MODEL

Novell buys Ximian

Novell is building on its recent commitment to Linux with the surprise purchase of Ximian, the monkey-obsessed desktop development team, for an estimated \$30 million. Attracted by Ximian's foothold in the business desktop sector and the potential of its Mono project, Novell claims the acquisition would help drive wider adoption of Linux in the enterprise.

CEO Jack Messman said that while Linux was the fastest growing platform in the market, there were areas that needed attention. "Customers still face two key business issues: how to provide cost-effective management and maintenance of Linux systems, and how to deploy and support low-cost Linux desktops within the organisation," he said. The acquisition of Ximian would enable Novell to offer market-leading solutions for both.

"Just as important, Ximian brings Novell unparalleled Linux expertise – helping us not only deliver more value to customers, but also strengthening our ability to work with and leverage Open Source initiatives more constructively."

More resources

Novell aims to plough resources into Ximian's GNOME-based desktop and, perhaps more importantly, Mono, the program to bring Microsoft's .NET framework to Linux and other Unices. Ximian's Red Carpet application management software would also be tightly integrated into future versions of Novell's ZENworks product line, and Evolution should, in time, become a valid Groupwise client.

"Mono is great. I like the idea very much of being able to run .NET applications on Linux and recompiling things written in C# to run on Linux. We will become a big advocate for Mono. I hope Microsoft views it as a good thing."

CHRIS STONE, NOVELL VICE-CHAIRMAN



Ximian will now go under the less elegant name Novell Ximian Services business unit at Novell. David Patrick, former CEO of Ximian and newly crowned general manager of Novell Ximian Services, said the marriage of the two companies offered "tremendous benefits" to existing and potential customers. "Our breakthrough management and desktop products strongly complement Novell applications on Linux and Novell Enterprise Linux Services."

Novell's profile as an enterprise company, Patrick said, made them the ideal choice to "drive growth for the Linux market as a whole." Indeed some analysts have suggested that a Novell/Ximian product has the potential to grab significant market share in the enterprise sector, though

Novell's chief technology officer Alan Nugent said there is no specific company or product in their sites. "We're not going after Microsoft," he told newsfactor.com. "And we're not going after Sun." But some competition was inevitable: "It doesn't matter what technology you pick, you compete against Microsoft."

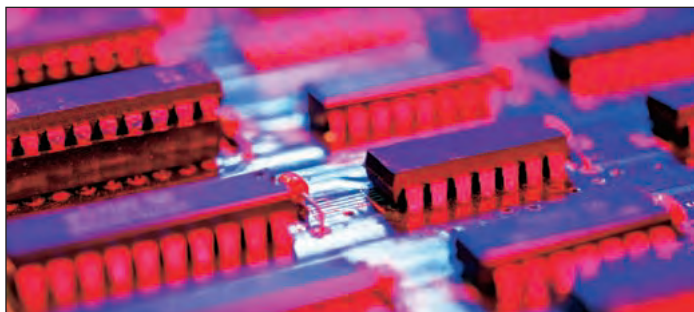
Extra funding

Ximian's Miguel de Icaza, when he attended the recent Linux World Expo in San Francisco, said the Novell buyout would bring much-needed extra funding to the company's main projects, and that it was a big step forward for the whole Linux community. "Initiatives like GNOME and Mono will only improve with Novell's resources behind them," he said.

Novell will absorb almost all of the 75-strong Ximian workforce, though they will continue to work from their Boston base for the foreseeable future.

De Icaza's co-founder Nat Friedman, who becomes Vice-President of Research and Development at Novell Ximian Services, said the two companies shared a common goal: reducing the barriers to Linux adoption in the enterprise. "This combination can do just that, and make Novell the number one Linux solutions company in the process."

"We're extremely excited to join with Novell to pursue a shared vision of leadership in a full range of integrated Linux desktop and server software for the enterprise, backed by world-class management tools, support and services."



OPTERONS AND ITANIUM

Supercomputer race

China's dawning IT industry is to build one of the first supercomputer systems around AMD's new Opteron CPU. It will contain 2000 processors and, with ten Teraflops of power will be the third-most powerful computer system on Earth (for a few weeks at least). AMD also claims to be working on a 10,000 Opteron machine called Red Storm with the US Dept of Energy. Red Storm's 40 Teraflop rating will easily steal the crown of the current most powerful computer from NEC's Earth Simulator (36 Teraflops).

Meanwhile in Australia's University of Queensland, Earthquake researches have a new SGI Altix supercomputer

powered by over 200 Intel Itanium 2 processors, called the Australian Computational Earth Systems Simulator (ACCeSS). It uses SGI's NUMAflex memory system to manage its 208GB of main memory and also utilises the company's new flagship graphics system, *Onyx4 UltimateVision*. It will carry out large scale simulations of faultlines in the Earth's crust with the intention of predicting major earthquakes.

Not to be outdone, the National Centre for Supercomputing Applications will soon have 1,450 Dell PowerEdge blade servers, each packing a pair of 3.06GHz Xeon processors. The cluster has a theoretical rating of 17.7 Teraflops.

PROCESSING SPEEDS UP

Fastest Linux Supercomputer... Ever!

Japan's National Institute of Advanced Industrial Service (AIST) has ordered a new supercomputer designed to handle 11 trillion calculations per second. The supercomputer, built by IBM, is made up of 1058 1U IBM eServer 325 machines, containing 2116 Opteron CPUs, will run SuSE Linux Enterprise Server 8, making it the world's most powerful Linux-based supercomputer.

David Turek, the VP of IBM Deep Computing, said "The eServer 325 powered by the Opteron processor offers strong performance and extended memory addressability while ensuring backward compatibility that preserves customers existing 32-bit software investments. Many high performance computing customers have expressed interest in acquiring this technology."

AIST's plans for the new supercomputer are quite wide - they

hope to accelerate research using grid technology in a wide variety of projects including the search for new materials to be used for superconductors and fuel cell batteries, and also the search for new compounds that could be the basis for a cure for various malignant diseases.

Marty Seyer, the VP of AMD's Microprocessor Business Unit, had this to say: "The announcement of the IBM eServer 325 paired with the immediate demand from a renowned institution such as AIST shows that the market is hungry for the performance the AMD Opteron processor provides for both 32-bit and 64-bit computing."

IBM also announced the intention to deliver an Opteron-based workstation next year. With recent announcement that its DB2 Universal Database for Linux is now available for the Opteron, there is clearly increasing commitment to AMD's new Itanium-beating CPU at IBM.

NEWSBYTES

■ *Linux Format* and *Linux Pro* will both be present at LinuxExpo 2003 to be held at London's Olympia on 8-9 October. Registration is free at www.linuxexpouk.co.uk/

■ In the absence of any independent research, German consulting firm Relevantive, set about determining the usability of KDE against that of Windows XP. The result was very close, with most tasks accomplished in KDE just 'a little' behind Windows. "A couple of tasks were even easier and faster to solve on Linux." KDE was criticised for the wording and, occasional lack of clarity, in programs, menus and interfaces. Get the entire report at www.relevantive.de/Linux_e.html.

■ The latest HD-based media player from Rio (the \$399 Rio Karma) includes support for the Open Source ogg vorbis format.



■ Mainconcept is in the process of updating its highly regarded film editing software and has a preview edition of MainActor 5.1 available for Linux users. The preview has most of the features expected to make the grade for the final version which should be available towards the end of 2003, including transitions, fades, advanced title effects and a range of codecs. www.mainconcept.com. Also, testers can visit the forums to relate their (good and bad) experience with preview software.

■ A recent survey suggests that ten per cent of all PCs sold in India are shipping with Linux. The figure is up from virtually zero in January. Germany's SuSE Linux AG has signed a worldwide agreement with SAP that means users of SAP's range of enterprise software on SuSE's OS products can contact either company for support.

■ In Australia, the job of surveying Victoria's 22,000 miles of highway has been given a Linux upgrade, meaning the two man survey team can now cover the ground at 80-100kph, instead of the 20kph that the old Mac system was capable of.

■ The overcrowded distro market may soon get a new player in the form of InceptionOS. It will be built from scratch and be able to switch between RPM and DEB package management. The project is in need of coders. <http://sourceforge.net/projects/maxlinux/>

Jono Bacon

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



COMMENT

Contribute!

“ As A Linux user, I am very grateful to my fellow members of the community. It is the tireless work of these free software contributors that keeps us in good free software. One thing I have been pondering recently is how we should contribute.

Free software projects often have a PayPal donation button. It seems clear that many hackers want help with development fees and suchlike. I think this is a great idea, but may be misguided; after all, free software often separates financial dependency from the act of creating software. Obviously we all need readies to live, but maybe more emphasis should be placed on non-financial donation?

A financial contribution is a very easy way to feel like you are giving to the community. While I respect this notion, a much more kindred contribution is something such as a bug report or writing docs. A contribution should directly help the project in question, and though £5 may pay for a bit of Internet access for a developer, it merely helps that single developer contribute to the project.

This brings us onto the issue of contributing to distros. Mandrake, for example, set up MandrakeClub so people can help MandrakeSoft stay afloat. This I feel is a good idea, but once again focuses too much on money. Maybe MDK could reduce costs with contributions such as using a Wiki for application help files and my oft-cited idea of setting up a P2P network for package distribution - I would be happy to contribute bandwidth to a distro - a legitimate use for P2P networks in an age where they are condemned for supporting piracy.

I think we need to stress the cost-free contributions that require no programming knowledge, and a real difference could be made.

OPEN SOURCE REPLACEMENT FOR MS ACCESS

Kroupware bears fruit in less than a year

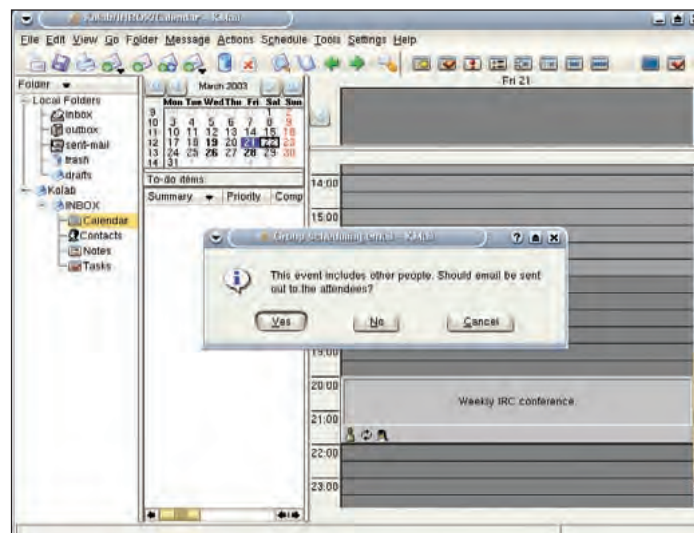
Less than a year after its inception, the *Kroupware* project has announced the completion of a stable version of the software. Based completely on Open Source solutions and marketed initially under the *Kolab* brand, the application suite is available from www.kroupware.org. Commercial support is already available.

Kroupware deals with email, group scheduling/calendar and tasks on both the client and server ends, and with extensive adherence to open standards should integrate into many existing groupware solutions. The core components of the client suite include *KMail*, *KOrganizer*, *KAddressbook* and *Pilot*, while the server is designed to

be very scalable benefiting from 'already matured Free Software components like *Cyrus IMAPd*, *Postfix*, *Apache* and *OpenLDAP*'.

Bo Thorsen, of Klaralvdalens Datakonsult AB, said the key to creating such a robust and usable product lay in having great foundations and the talents of many KDE developers. "We closely cooperated with the KDE community and the high quality of GUI components maintained by the community made it possible for us to assemble a feature-rich groupware client in such a short timeframe."

LXF will be featuring *Kroupware* in our *HotPicks* next issue, and there'll be a version on the coverdiscs too!

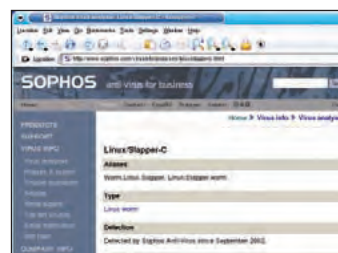


Kroupware is based on common KDE components.

Linux Web Watch/



linuxsecurity.com – patch apps.



Sophos – don't get infected.



Run as root? No thanks!



Guarddog keeps crackers out.

Aftermath of Blaster/LoveSan Windows worm

Gloat by all means, but don't get complacent...

The latest worm - Blaster/LoveSan - to target Windows computers (see www.colorado.edu/its/news/blaster.html) may make Linux users feel not a little smug, but there are other threats out there. Fortunately, there are also excellent sources to keep your Linux box safe from harm.

One of the best places for news on the latest vulnerabilities, exploits and problems is www.linuxsecurity.com/.

The site is sponsored by Guardian Digital and is constantly updated. As well as advisories, the site also has a pretty good selection of security-based features and tutorials.

While Windows may be subject to virus threats everyday, Linux is not immune. In fact, there is thought to be as many as 14 viruses which target Linux according to www.sophos.com. Of course, unless you run as root in

your day-to-day computing, the potential for destruction of viruses is reduced, though this risk not to be dismissed. If you do need to run as root (or in *su* mode) for any significant time, have a look at www.helpdesk.umd.edu/topics/security/136/ for some advice. This site is also useful if you need to configure secure remote access to a root account.

Last but not least, you'll need to sort

out your firewall. Most distros ship with a decent firewall configuration utility, but they're not always the most usable things. *Guarddog* by Simon Edwards is a small application that takes the pain out of starting and defining the rules to your firewall. www.simonzone.com/software/guarddog/. The firewall is complemented by *Guidedog* (IP forwarding) and *Watchdog* (real-time monitoring).

SCO News

■ IBM finally got round to responding to the SCO allegations, claiming that SCO has released (and continues to give away) its own version of Linux, therefore its allegations are bogus. Furthermore, Big Blue's patent lawyers have been scouring their archives and uncovered four alleged violations in almost the entire SCO product line – raising the prospect of an injunction which could effectively shut down the whole of SCO's software business.

■ Exhibiting huge generosity, Red Hat has launched a new fund to support Open Source developers who are facing copyright or patent infringement cases. The distro vendor has kicked \$1 million in to start things off. See www.redhat.com/opensourcenow/



■ More serious allegations are being thrown at SCO executives including claims that the whole exercise is merely a 'pump-and-dump' scheme, whereby the share price is pumped up by the constant stream of press releases and allegations, while inside shareholders quietly dispose of their shares at the inflated prices.

■ It has been claimed that an unnamed 'Fortune 500' company has purchased a SCO licence.

■ SCO fleshed out its allegations of IBM's theft (while 'finally' terminating its license to sell AIX). Oddly, the code in question is said to be 148 files (some 168,000 lines) of Sequent UNIX code, which would suggest that IBM is being sued for contributing its own code into Linux, as it purchased Sequent in 1999.

■ Finally, (honest!) some bright young thing has noticed that all of the disputed code relates to NUMA, RCA and other multi-processor systems; and as this is not usually compiled into the kernel for a single CPU system, SCO efforts to extort – sorry, secure – licence payments from the average user are irrelevant.

Account Name	Description	Total
Bank		\$13,278.61
ABC Checking		\$3,967.00
ABC Savings		\$2,020.34
Mega Brokerage Settlement		\$1,291.47
Savings		\$0.00
Securities		\$1,556.00
Mega Brokerage		\$1,556.00
HAL Stock		\$0.00
SHRT Stock	SHRT Corp is a good one for	\$0.00
THJ Stock		\$1,556.00
Income		\$10,297.61
ABC Bank Interest	Interest from ABC Bank	\$32.34
Test	This is a test account	CAD 0.00
PGR Consulting	paychecks	\$3,968.00
Stock Dividends & Distributions	Stock Dividends & Distribution	\$277.47
Expenses		\$3,213.00
Gas		\$0.00
Groceries		\$30.00
Rent		\$0.00
Other Expenses		\$2,883.00

GnuCash is regarded as a critical application for wider acceptance of Linux.

MONEY APP DEVELOPMENT

GnuCash needs you!

The *GnuCash* project is facing a crisis which, developers say, could lead to the application into a "dead end". In a call to arms, project leader Benoit Grégoire said *GnuCash* was suffering for a lack of exposure and, crucially, talent. "I think most everyone agrees that *GnuCash* is a critical piece of software for the Linux desktop. It's also one the largest free software projects.

"Although *GnuCash* comes up in every discussion of needed software to get Linux on the desktop, but the project currently has only about seven active developers." However, he said there was a danger of the application stagnating without more manpower – and not just developers.

Grégoire said that in the team may have overstretched themselves in the quest to implement the latest features. "Choices had to be made," he said, "and a lot of important things are currently being neglected. If the *GnuCash* project can't manage to attract more contributors and refocus the efforts of those it already has, it's going to become unmanageable. We often say that Linux would survive even if Linus got hit by a bus. Well, right now I am not too certain that *GnuCash* would currently survive if Derek Atkins got hit by a bus." To get involved, visit http://gnucash.org/en/state_of_the_gnucash_project.p.html to see what needs to be done.

HOLLYWOOD

Disney's Photoshop efforts

The position of Linux as desktop of choice for filmmakers was confirmed last year when Walt Disney Pictures migrated their systems away from proprietary software. The one sour note was the lack of *Photoshop*. Speculation grew that Disney and others would exert enough pressure on Adobe to port their image editing application. But when this didn't happen a trio of studios, Disney included, committed resources to CodeWeavers WINE efforts specifically to make *Photoshop 7* usable.

In all, Disney's Jack Brooks speculates, the endeavour cost the company \$15,000 compared to a projected \$50,000 spend on Windows licenses and a further \$40,000 per year on support. Brooks said it was a win-win situation for all involved. "I didn't



Disney's Photoshop to Linux efforts benefit the whole community.

have the resources to chase that project internally. This way, the open-source community got the product, and we got what we needed cheaper than we could have done it ourselves."

Like many others in Hollywood, Disney is now a multi-platform shop, though there is an obvious bias towards Linux.

Hoyt Duff

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



COMMENT

Survival?

“ The first sign of impending doom appeared at Linux trade shows – fewer ponytails and T-shirts, more suits. Floor displays became ultra-professional and the .org group booths were shuffled to one side. As established IT players got interested in Linux, they also inculcated their culture into the community. Loose confederations of developers became indulgers in ego-fuelled disagreements. This wasn't supposed to happen.

I'm reminded of the SCO IP debacle, allegation of subversion in the XFree86 development process, the fights over kernel developers using *BitKeeper*, the ongoing vitriol with closed-source drivers; the KDE versus GNOME nastiness.

Linux, the GNU generation and the Open Source philosophy were meant to change the world. Information was to be intellectually free and shared; people would be cooperative; unconventional ideas would find a safe place to grow; we would all be happy.

As is true with any endeavour, instead of changing the system, Linux changed to the system. There's no one individual to blame, we all chose to play in the arena thinking we couldn't be corrupted, but found ourselves co-opted. To the world, the Linux culture, the GNU generation and Open Source are becoming like Microsoft, like IBM, like Novelle, like Sun, like SCO. Is it over? Should we move on to the "next" Linux?

No! Linux is ready for the desktop, ready for the server. It's ready for Mum and mad Uncle Ned and obnoxious 12-year old flammers. It's getting corporate support, mainstream publicity, hardware manufacturer awareness, and it's still free, still open and anybody with the initiative can build their own special OS from it. Forget the corporate nonsense; we'll always have Linux.

GAMES ON LINUX

Native Savage released!

Savage: The Battle for Neweth is to become the latest game to ship in a dual-format package, meaning the Linux edition will be available right out of the box. The game, an attractive looking real-time strategy (RTS)/first person shooter (FPS) affair, has been

developed by S2 and designed from the ground-up as a multiplayer, team-based experience. Teams of 32 players (one in RTS mode, 31 'on the ground') compete across the primitive, though beautifully rendered, landscapes for supremacy. The Linux edition of the software

will be 'functionally identical' to the Windows version and will ship on the same CD. S2's Jesse Hayse said there were a couple of factors that influenced their decision to develop on both Windows and Linux.

"We have a bunch of Linux lovers on board here at S2, so we

understand the importance of having great games available for that community," he said. "We've also found that a great deal of gamers would like the option to play or set up servers on the Linux OS, so we went that route with *Savage*." The game is scheduled for release on August 25th.

Shiver me timbers!



Puzzle Pirates is a new massively multiplayer puzzle game built around the world of, you guessed it, pirates. Developed in Java, the game is Linux and Windows compatible, with OS-X support scheduled for later in the year.

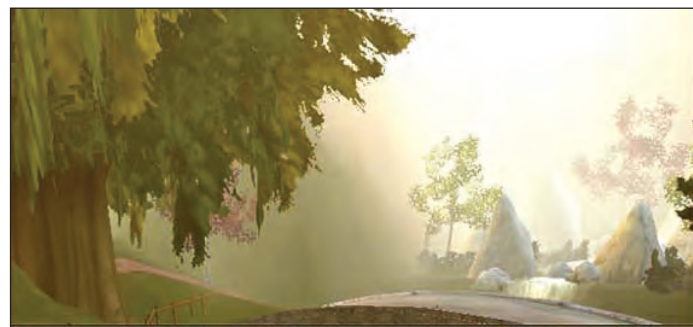
The game is currently in beta phase, and you should be able to sign up at www.puzzlepirates.com. Subscription prices should be announced very soon; designer Daniel James said the sub would be a "modest monthly charge."



It's a blend of RTS and FPS, but the mythical characters also give this strategy game some of the atmosphere of a Role-Playing Game too.



LXF's Art Editor Julian has a body like this, but where he keeps it is a secret.



Many multiplayer server owners favour Linux, so this release makes sense.

Embedded Linux News

- A new reference design for embedded applications based on **Motorola's** PowerQUICC II microprocessor, has been created by **Arabella Software**. The package, which is free, supports both the new PQ2FADS-ZU evaluation board and the earlier 8260FADS boards and includes a CD of software from Arabella. The CD is available separately.

www.arabellasw.com/Motorola.htm



- IBM, Toshiba and Sony are reportedly on the brink of completing designs for the new 'Cell' processor that will form the heart of the **Playstation 3**. Dubbed 'supercomputer on a chip', it could go into production as early as January 2004. Cell's designers say that, while the hardware is almost ready, the task of creating software (and Linux will probably be the first ported OS) is far more difficult, thanks to the fundamental differences between this and 'traditional' CPUs.

- **SCO's** licensing scheme includes a demand from embedded Linux developers for \$32 per device. Astonishingly they are attempting to levy this charge on the makers of TiVo (which runs a pre-2.4 kernel).

- **Metrowerks**, which was recently acquired by Motorola, has added kernel level debugging tools to its CodeWarrior embedded development environment.

www.metrowerks.com/

NEWSBYTES

■ Aprelimum has updated its Abyss X1 web server to fix a number of bugs and improve support for PHP 4.3.2. In common with previous version, the server is available free for both Linux and Windows from www.aprelimum.com.

■ The latest innovation from Lindows.com (after last month's hard-diskless machine) is the KooBox, a complete LindowsOS-based PC with 1.2 Duron processor, 256MB RAM and 20GB HDD. The machine also features a TFT monitor and a \$449 price tag. No European distribution has been sourced yet.



■ IBM has joined forces with Threshold Digital Research to create a series of CGI movies to rival the likes of *Toy Story* and *Shrek*.

■ Red Hat has expanded on its recent decision to pull out of the boxed retail sector, claiming that the distro that was RedHat 9 will become a '24/7 development' release, constantly updated. To replace it, the company says it will soon announce a new desktop orientated and small server distro based on the Enterprise Server line.

■ Trolltech has announced the latest release of its cross-platform application development platform, Qt3.2. Qt forms the core of the KDE desktop environment and releases of KDE usually shadow Qt. Important improvements include full support for script languages such as Hindi and Bengali, and a completely rewritten font rendering engine.

■ KDE's multi-protocol instant messaging system *Kopete* has received an update. *Kopete* allows users to communicate with users of various systems including *Windows Messenger*, AOL's *Instant Messenger* and the universally popular ICQ chat software, as well as *Jabber*.

■ Sun Microsystems has signed a support deal with SuSE which echoes its tie up with Red Hat. Sun's attitude to Linux has been described by some unkind commentators as being 'schizophrenic' in recent months, and some have claimed that the overtures to SuSE may have been as a result of market demands rather than any real desire to form a partnership. Sun will sell and provide support for SuSE Enterprise Server on both Sun x86 and Intel hardware.

LINUX MULTIMEDIA

Music to your ears

RealNetworks is expanding its

Open Source Helix project to include a brand new media player. Real says the *Helix Player* project will give Linux and other Unix users the same access to web-based media as that currently enjoyed by Mac and Windows users. The player will have support for a number of codecs, including *RealAudio* and *RealVideo*.

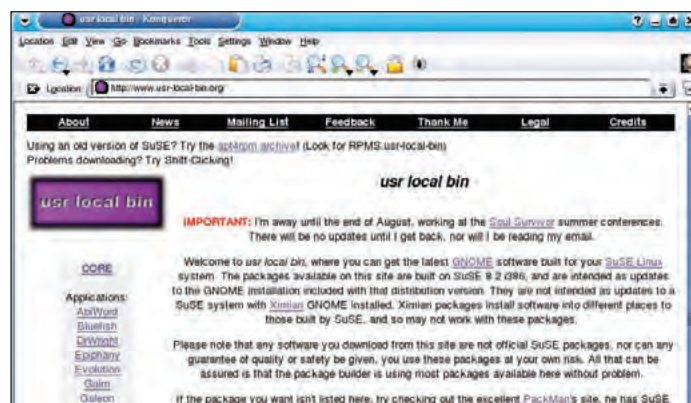
Nagesh Pabbisetty said that Real

wanted to take its experience on the server side and apply it to client software. "We are thrilled to now build, in collaboration with members of the Helix Open Source community, the best media player for Linux, as yet more evidence of our long-standing commitment to Linux."

The Helix community has signed up its 20,000th registered developer, see – www.helixcommunity.org

DESKTOP DEVELOPMENTS

Better GNOME for SuSE



ULB has extended its remit from applications to a drop-in GNOME replacement for SuSE 8.2.

In recent years *usr local bin* has been creating fully formed GNOME application packages for SuSE, including up to date editions of *Abiword*, *Bluefish* and *Galeon*. Now the project has taken the next logical step and launched an effort to create a full 'drop-in replacement' for SuSE's GNOME implementation. Maintainer James Ogley said ULB GNOME goes beyond packaging applications, but shouldn't be regarded as a full

GNOME distribution. "I've been gradually building more and more core GNOME components, and have now decided to take the next step and plug some of the gaps." Changes to the standard SuSE desktop include removing much of the branding, menu themeing, improved panel themeing and changes to default applications, including setting *Galeon* and *Evolution* to handle web and email comms.

www.usr-local-bin.org/

PDA VOICE RECOGNITION

Talk to me Tux!

IBM has released a new toolkit for its *Websphere* development studio which brings advanced speech recognition to Linux. The new Multi-modal Toolkit makes it possible to work with various forms of communication on the same device, such as issuing voice commands to a PDA and having the response returned as a text message. Analysts are predicting that

voice recognition – not a new technology by any stretch of the imagination – is on the cusp of acceptance thanks to the prevalence of devices available without a traditional keyboard. One of IBM's planned applications, built on a combination of XHTML and Voice XML, is a smart helpdesk system which can respond to natural-language questions.

David Cartwright

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



COMMENT

Alpha Linux

“A client came to me the other day, saying: “What can I use my Alpha server for?” He's had this socking great Alpha-based machine running Debian Linux for a while, running the first release of his company's software offering, but now that release two has been developed (and runs on a Windows machine) he's said: “Can we run Windows on it?”

The answer, of course, is: “No”. (NT 4 for the Alpha architecture is an end-of-life product and Win2K+ won't support Alpha architecture.)

I was, therefore, reassured to discover that although most the Linux distros seem to be Intel-only, Debian is at the forefront of cross-platform support. I was surprised, but pleasantly so, to find that Debian 3.0 supports not just x86 machines but also Alpha, ARM, PA-RISC, IA-64, 680x0, MIPS, PPC, S/390 and SPARC kit. And – it's £17 to buy the seven-CD set!

Praise is due to the Linux development community for continuing Alpha support. After all, unlike SPARC, PA-RISC etc, it's not forward-looking (Compaq canned the processor) and yet we have a load of developers willing to continue producing Alpha-compliant code for those stuck with old hardware. And well done to Debian for including it in its list of distros, it takes significant effort to bundle all this stuff together and support it – significant commercial resources have been employed for what must be negligible financial gain.

There's just one downside. I've fiddle with this machine before, I know that the EEPROM-level (equivalent to the PC BIOS level, I guess) command set is less than trivial to comprehend, so I may have a fight figuring how to make the thing even start up in the first place. But hey, the client's paying me by the hour...



EMAIL CLIENT/BROWSER SPECIAL UPDATE!

All change at Mozilla

There are significant upheavals afoot at the Mozilla project. But what effect will AOL/Time Warner's departure have on one of the most high-profile Open Source projects? **Andy Channelle** investigates...

The long history of the Mozilla project has not been without its problems. And yet, in the background during these tumults, the support of AOL/Time Warner, which purchased Netscape at the height of the browser wars, was constant. Despite a seemingly overwhelming competitor, AOL/Time Warner continued to pump time and money into Mozilla, betting on the Open Source development method long before it began edging into the mainstream. But no more.

The media giant has pulled the plug on its support for Mozilla and also announced there will be no more Netscape browser releases. The Netscape brand will continue as a web portal, but the application which played such a big part in helping kick off the

web revolution we take for granted today has effectively fallen off its perch.

For a commercial application this would mean almost certain death, but the early decision to turn Netscape over to an Open Source group called Mozilla.org means that this is one technology that won't wither and die. In fact, its developers claim, this could be the opportunity Mozilla needs, to move out of the shadow of its monolithic 'parent' to stand or fall on its own considerable merits.

As a parting gift, AOL/Time Warner dropped two million dollars into the coffers of the newly formed Mozilla Foundation, and also donated Mozilla related intellectual property, trademarks and, crucially, hardware to the project. What Mozilla loses, though, is the efforts of the 50 full-time coders that AOL has

either laid off or reassigned. A significant drop in development personnel.

Announcing the creation of the Mozilla Foundation, Mitchell Baker, who has headed the Mozilla project for the last two years and has been appointed President of the new foundation, said it had been an ambition to forge a genuinely independent organisation to move the *Mozilla* technology forward. "*Mozilla* was developed in an Open Source environment and built by harnessing the creative power of thousands of programmers on the Internet. Going forward, we will continue to partner with developers and industry leaders to keep content on the Web open." She said project leaders were grateful for the past and ongoing support of America Online – "and we look forward to continuing to

work with AOL over coming years" – but that time was right to move on.

Positive Thinking

Baker told *Linux Format* that the core group were very excited about the creation of an independent group to oversee the future of *Mozilla*. "We won't have AOL to employ people, but we also have an enormous set of new opportunities." She said the project has needed this type of organization for a long time, not so much for individual contributors but to attract the attention of commercial entities. "Not having an easily understandable organization has made commercial adoption and participation much harder," she says. Mozilla.org was a 'loose coalition', the Mozilla Foundation gives the operation 'not-for-profit' status, which will make

A lizard divided

Separate Firebird browser and Thunderbird email client

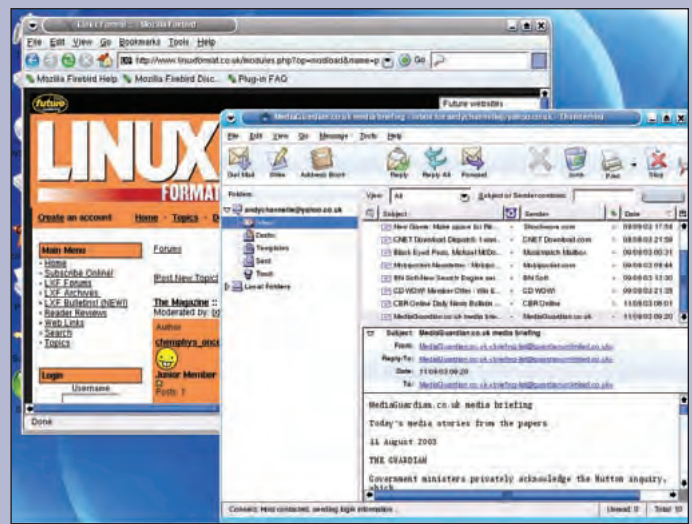
One of the results of the structural change at Mozilla will be a new focus on end users rather than relying on third-parties for promotion and advocacy. The *Mozilla* browser suite, always regarded as a development tool, will be faded into the background, with promotional efforts concentrated on the *Firebird* browser and *Thunderbird* mail client.

The first demonstration of this change is a redesign of the Mozilla.org front page, which is also reachable on the www.mozilla.org address, to make finding and downloading the core apps easier. And, for the first time, *Mozilla* will be available on a branded CD from the organisation itself. Spokesman Bart Decram said a deal had been reached with E-FLO (The company that previously managed Netscape's CD distribution program) which would allow the Foundation to offer CDs at "about \$10."

Taking into account the time it takes to ramp up production, Decram says, "the most likely result is that we'll start offering CDs around the time of the *Mozilla 1.5* release – initially we'll probably just ship the latest stable release of *Firebird* and *Thunderbird*." The Foundation aims to have a more user-orientated product available by the end of the year.

However, in a recent developers' meeting, Brendan Eich suggested that 'a shippable product', based on *Firebird* and *Thunderbird* would take "about a year."

In addition to the two main components, the project has also thrown up a world class bug reporting tool (*Bugzilla*), an adequate HTML authoring tool and – the latest addition – *Sunbird*, a standalone calendar application that is in its testing stage.



Firebird and *Thunderbird* will replace the traditional *Mozilla* browser suite for both Linux and Windows end users.

soliciting funds from the likes of Red Hat and IBM more effective.

The immediate result, she says, is less people paid to work full time on the project. But then: "We've reached a point where we need fewer people. The underlying components are stable, and *Firebird* and *Thunderbird* require fewer people than building the application suite did." The money from AOL, pledged over two years, would also be used to retain a core set of engineers.

"The Mozilla Foundation is hiring a key set of people," Baker says. "I'm sure we will wish we had more people, but we certainly won't be the only development organization wishing that." Three members of the current Mozilla team – Baker, Brendan Eich and Christopher Blizzard – will stay on the Board of Directors, and they will be joined by OSFA (see boxout) founder Mitch Kapor – Baker has also joined OSFA in its Chandler development

effort. "We'll talk about the specific people before too long. It's a great group, and very excited."

AOL vs Open Source?

It is ironic that AOL has made this announcement at a time when the rest of the world is investigating or wholeheartedly embracing the Open Source development model, but most of the people we spoke to at Mozilla were, understandably, unwilling to speculate

on the cause of their parent company's *volte-face*. ("As to AOL's reasoning, that's a topic for AOL; I have no special information on this," Baker maintains.)

Tight-lipped official sources haven't stopped the conspiracy theories emerging on the Net, chief among them the idea that AOL used its support of Netscape to leverage favourable terms in its recent negotiations with Microsoft on the integration of *Internet Explorer* with AOL's client software. These negotiations led to the two companies abandoning their long-running legal battles. Microsoft agreed to pay AOL \$750 million and provide royalty-free access to *Internet Explorer* technology for the next seven years. The payout may have grabbed the headlines, but access to *IE* may be the most significant part, especially as *IE* Project Manager Brian Countryman said in a recent webchat that *IE6* (SP1) would be the last standalone version of the browser.

Whatever the reason, recent announcements by Apple, which is shipping its own KHTML-based browser, and Microsoft, will accelerate the evolution of the browser sector. The options for 'non-embedded' cross-platform browsers are effectively narrowed down to two: *Opera* and *Firebird*. And only one of these is free (in terms of both beer and speech). The question that faces the new Mozilla Foundation is: will users be happy to stick with what they're given, even if what they're given is inferior? **LXF**



The frontpage's makeover should appeal to Blaster worm victims looking for an *MS Internet Explorer* alternative!

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 copy of the O'Reilly book of their choice!

On Safari

Having bought many, many books on technical subjects over the years, I was astonished to find that O'Reilly has started an online service called Safari (<http://safari.oreilly.com>) which puts the entire contents of over 1,300 books online for your reading pleasure.

Not only can you search the entire 1,300+ book base for any terms that you're looking for, but you can also read/browse an entire book (including all of its images) if you're so inclined. It really is a proper electronic library. The downside is that it's not free, but this isn't a bad thing. For \$14.99 a month you can check out 10 books, which I think is quite a reasonable sum. You borrow each book for 30 days before it's returned to the system and you can then replace it with another book. The other downside is that if you find something juicy and you want to look at the full article, you're going to have to check the whole book out to do so. You



You don't have to rely on reader reviews like some other book buying sites – you can browse entire books for a fairly nominal charge!

also don't get any files that appear on the CD that may accompany the printed book.

The good thing about it is that no Acrobat or other proprietary PDF reader is required to browse the books – just a web browser is needed. I'm recommending to our employers to get a subscription as it would save us much time, money and generally trying to find whose borrowing the office copy of

Mastering Regular Expressions!

Martyn Drake, www.drake.org.uk

Thanks for your suggestion. The Safari site has been running for some time, and does offer a good service if you need to check up on lots of Linux- or coding-related topics. As a special prize this month, email us with the name of the O'Reilly book you most want, and we'll get it over to you (whether you share it with your colleagues is up to you)!

and many more. As for your comments on WinCVS being difficult to use – this was a claim within the context of a GUI application and not within any kind of application – sure it is easier to work with for many users than the command-line client, but is certainly not what I would consider easy-to-use. *Linux Format* will be featuring a *What On Earth* feature on CVS as some point over the next few issues – if you want to know what it is and why you should be using CVS, don't miss it!

Linux hardware

I've just read Roger Gibson's email to you (LXF42) and note your reply concerning hardware suitability. I don't want to be too critical but frankly this answer, which comes down to "find out for yourself", is less than useful.

What I would like to see is a list of those companies that actually do support Open Source, which would exclude nVIDIA as it, as far as I know, does not. All their drivers are apparently strictly closed source.

I know it's a forlorn hope but it would be really nice if Mandrake, Red Hat, SuSE or Debian could put a list of those hardware configurations which they have actually tried in a file on the first CDROM of their distros so that people like me who are planning to buy a new base unit will at least know what will certainly work even if we can't get hold of it.

TE Groves, via email

One of the problems facing developers, distro vendors and magazines when it comes to hardware is that there is just so much of it. As a least common denominator, most of it has support for Windows and little else. The companies that actively support their own drivers AND Open Source them would fill a very, very short list – at least nVIDIA actively

CVS concerns

I'd like to highlight a few points from Jono Bacon's Review of the subject application. Jono writes, "There is no doubt that the nearest competitor to *Seapine SCM* is *CVS*". I'm sure Jono has heard of *Subversion* and *Bitkeeper*. IMHO *Subversion* is a far better product than CVS even although it is still considered beta quality. *Bitkeeper* is what Linus and the gang use to keep track of

revisions to the Linux kernel. Also it is used by the folks at MySQL. The *Bitkeeper* GUI is really awful, but *Bitkeeper* itself is really easy to use. As for *WinCVS* being difficult to use, I couldn't agree less. It's easier than using CVS from the command line. Even though *Subversion* does not have a GUI at the moment I believe that there are several projects in progress which aim to fill that void. John Rygannon, via email

Jono Bacon writes: Hi John. Thanks for your comments – allow me to clarify a couple of points. I compared *Seapine SCM* primarily to CVS, as the latter is currently the most popular form of source control for Linux systems. One of the main reasons I chose CVS as a comparison is that most Open and closed-sourced projects use CVS; yes, the kernel uses *Bitkeeper*, but CVS is used by KDE, GNOME, XFree86, Blender, Mozilla



Most premier Linux distros like Mandrake, Red Hat and SuSE will recognise your hardware without too much trouble these days.

develop drivers for Linux.

Mandrake do have a hardware compatibility database on their website (www.mandrakelinux.com/en/hardware.php3), which in my opinion is a better and more accurate way of dealing with queries. If you were to list every available piece of hardware and which versions of Linux they worked on, you would fill a years-worth of *Linux Format* magazines!

Dual head function

Just bought the July issue of *LXF* and it's great, especially the articles about MySQL, X and steganography – good job!

Thanks for mentioning something so brilliant as *Twin* in the *HotPicks* section. I've been using *Twin* as long as I can remember and I'm on a 1GHz, 256MB, 20GB, TNT2 system running it in a 1024x768x16 console (tty1 to tty6). The last few weeks I've been using X and *Konqueror* to test some PHP apps, but I still run *Emacs* in a *Twin* session and when I hack non-web stuff then I switch back to the console. (There is an extensive amount of software that makes the console usable, eg *Emacs*, *MPlayer* with *fbdev* support, *cdp*, *aumix*, etc)

I think you should cover more console applications in the future, perhaps an article about Dual Head functionality in text mode – is that possible? There actually are a few console lovers out there, at least I'm one of them...

Per Edin, *via email*

Well, we aren't ignoring console apps. They do appear quite regularly in *Hot*

Picks and other parts of the magazine.

The console is eminently useable and I don't think anyone at *LXF Towers* runs a Linux box without at least a terminal window open. We'll consider your tutorial idea along with all the others that any readers care to send in!

Reliable MDK 9.1

Having used Linux at home for the last 18 months, originally with Red Hat 7.1, but now up to SuSE 8.1 Professional. I thought I'd push my luck and install the free Mandrake 9.1 you provided on the cover of the magazine on my computer at work (a self-built affair with an MSI K7N2G nForce mobo, complete with onboard GeForce 4 graphics). Having installed *Tux* from scratch before I was prepared for some issues (mainly graphics, chipset and LAN driver problems).

During a particularly dull (for dull, read "when the boss wasn't in") afternoon I decided to partition my hard disk and go for it. Imagine

"We aren't ignoring console apps; the console is eminently useable, and I don't think that anyone at LXF runs a Linux box without at least a terminal window."

my surprise when Linux installed, detected everything but my onboard LAN, configured my 19" Liyama monitor and ran perfectly first attempt (It's interesting to note that my *bought* (60 quid!) version of SuSE 8.1 took FOUR install attempts as it continuously crashed whilst setting up my graphics in X – another, older MSI mobo with separate GeForce 2 card).

Expecting trouble I downloaded the mobo and 3D drivers from nVIDIA. And guess what? Nothing! No trouble! Installed first time, let me play *Tux Racer* straight away. Whilst sending you this email, I'm in the process of installing Ximian Evolution 1.4 via download (*wget*) – again, attempting to do this on my SuSE distro merely resulted in numerous error messages (as did trying to upgrade to KDE3.1!).

In short, I have to manage than SuSE 8.1 Pro – if only you'd given it us in January I could've saved £60. I think I'll go home and install MDK on a separate partition – maybe it can teach SuSE a thing or two!

Tony, *via email*

A lot of the experience comes down to particular hardware. I'm sure we could find readers who have had the opposite experience, but fair play to Mandrake, who make great efforts on their installer. It certainly has to be one of the most reliable and easy-to-use.

Munich's choice

In David Cartwright's column in *LXF43*, he seems to imply that the city of Munich chose Linux to save money. From what I've read about the issue this is far from clear, especially when you note the fact that the Microsoft bid was about 30% lower than the IBM/SuSE bid. It would seem more that it was a case of *trust*; I get the impression that Microsoft fudged itself with its voluminous terms and conditions. "OK, we'll only put software X on so many computers, we'll have a 'special' deal and not force you to upgrade your license in X years, etc, and in the end we'll only charge you *this* much." In my opinion that left the Munich council full of fear, uncertainty, and doubt about what Microsoft will deliver and how much it will cost in the long term.

Whether or not it *will* cost more in the long term is not at all clear.

Dr. Cirilo Bernardo, *via email*

From the reports in various places, money was the initial consideration, but even when Microsoft lowered their price, the decision remained firm. It's easy to speculate what the reasons might be, we just don't know, but it doesn't seem far-fetched to suggest that the multiplicity of Linux providers and Open Source nature of the OS were factors.



Cost or trust? In Munich and elsewhere, it's difficult to see the division.

SUBMISSION ADVICE

WHAT WE WANT:

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

WHAT WE DON'T WANT:

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

WRITE TO US AT:

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

READER TIPS

GRUB

In response to "Time Tinkering" in *Mailserver*, LXF42, I also had similar problems with my Gigabyte Motherboard with KT400 chipset, onboard sound and Ethernet. When I installed Mandrake 9.1, the Sound would not work. When I installed SUSE 8.2, I could not get a network connection. After much Googling, I found a glimmer of hope in the SUSE Newsgroups. I disabled ACPI and my problems were solved. I left SUSE on that PC. Not a hiccup since. I use Grub as my Bootloader and this is pretty easy to turn off. Use the option `ACPI=off`.

Chuck Coxhead, *via email*

CARD RECOGNITION

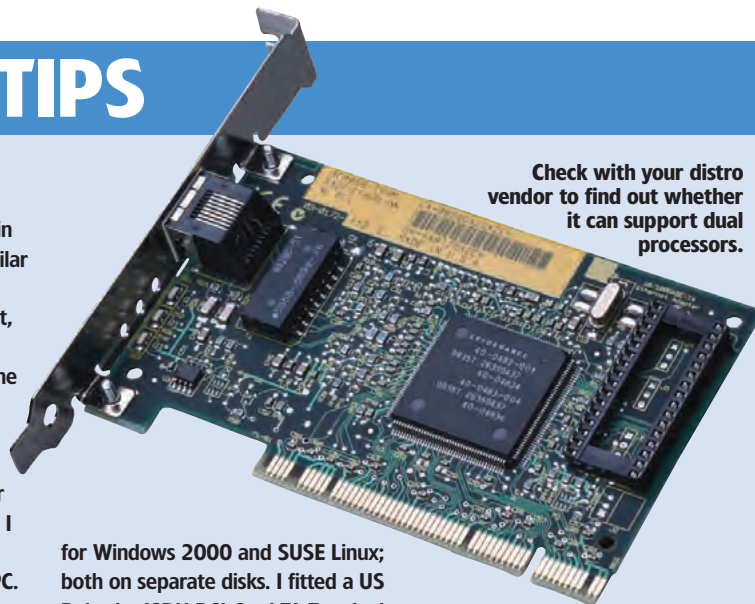
I read with great interest the letter in LXF43 from Trevor Collier-Moore concerning problems with card recognition in SuSE 8.2. My story has been running for a lot longer than two days – nearer two months in fact! Here's an extremely abbreviated version:

My PC is a twin processor Dell Workstation 530. This is dual boot

for Windows 2000 and SUSE Linux; both on separate disks. I fitted a US Robotics ISDN PCI Card TA Terminal Adapter which SuSE claims to support. This card was working perfectly with Windows 2000. During installation of SuSE Linux 8.2, I found that `YaST2` was hanging when I tried to configure the card.

After installation had completed, there was an error message with each reboot to the effect that the ISDN module would not load. I could not load it manually from a command prompt either. I could not connect to the Internet using the ISDN card at all. The free SuSE installation helpline told me that this was "too advanced" for them

Check with your distro vendor to find out whether it can support dual processors.



and was not really an installation issue (though I think that it is) and gave me an '0900' number to call for advanced help. I've met with this advice before – quite simply: this number does not exist, nor has it ever existed (even back in the days of SuSE v7.3).

So I called SUSE's UK HQ and was told that one problem was caused by the fact that SuSE could not cope with a twin-processor machine. This is because it uses virtual IRQs. A normal PC has IRQs from 0-15, my machine also has IRQs 16-23. This can theoretically

be solved by typing in a kernel parameter at install and boot time to inform the SuSE system that there are 2 processors (argument 'APIC'). It did not cure the problem.

There was obviously still some conflict between SuSE and the Dell ACPI BIOS. Yes this BIOS does a lot more than just power management. It routes IRQs so as to facilitate sharing rather than conflicts. As I understand it (in a nutshell) what to Windows is sharing is seen by SuSE Linux as a conflict.

It was apparent to me that the PCI card had 2 problems:

1 Installed on a virtual IRQ - 18.

2 Sharing with another device.

SuSE support then (belatedly) told me that I should be using the kernel parameter: `'pci=biosirq'` to get IRQ assignments direct from the BIOS. This suggests that, on a dual boot machine, the IRQ assignment was being taken by default from the Windows IRQ table – including virtual IRQs. This is bizarre! Using this parameter should have solved the first problem.

Entering setup after reboot, I could not change the BIOS IRQ assignments so that the USB card

Helpdex

shane_collinge@yahoo.com



would not share with anything else; whichever free PCI slot it was installed in. So the second problem still remains. The ISDN card still does not work. So now I have had to give up and go back to using an external 56K modem connected to the serial port; SuSE can configure this OK!

I cannot believe that SuSE Linux is unable to configure a commonly used TA card on what must be one of the most widely used Workstations on the market. Given how very critical Internet connection speed is these days for large downloads, this is a major failing in this product. My commiseration to Trevor Collier-Moore. If you publish this (and I hope very much that you can) – please, not beside a SuSE advert! Richard Myers, *via email*

DUAL BOOT

Firstly; great magazine ! Too bad here in Denmark, it's roughly £15 !

In LXF43, the issue of Linux users occasionally needing to run that one, essential application that's still only available for the Windows platform is touched upon in *Letter Of The Month*, and your reviews of *Win4Lin* and *VMware Workstation 4.0*. It's a problem I suspect, many of us are familiar with. It thought I'd share my own, rather elegant dual-boot solution – well I think it is anyway.

A year ago, I bought a Gigabyte GA8-IHXP

motherboard, which comes with an onboard RAID chip from Promise, which worked a treat with Red Hat 7.;;, unfortunately Red Hat 9.0 doesn't support the Promise chip (or is it that Promise doesn't support Red Hat 9.0 ?!), which is too bad, because the RAID function makes for a substantial increase in performance. Anyway, that left me with two identical IBM hard drives, each 60 GB, which could only be run as two separate, normal IDE-connected drives.

At the same time, I do very much enjoy the occasional bout of *CounterStrike*, but can't be bothered to install it on my Linux OS and there's no way, I can fill 120 GB with Linux stuff! So here comes the clever bit: I've installed both hard drives on the same IDE cable, set both of them as the Master drive and inserted a 2-poled 'flick'-switch to the 5v and 12v power wires. One hard drive has a 'clean-single-boot' installation of Red Hat 9.0, the other Windows. Before bootup, I just flick the switch to choose which hard drive gets the power connection, thus I choose which OS to run... It has worked without fault for several months. Obviously, there is no hard drive partition shared by the two OSes, so exchanging files the easy way is not possible, but I do enjoy the non-polluted feel of 'pure Linux'. Hope this hint comes in useful to others, but don't complain to me about any hardware failures!

Soren O'Neill, *Denmark*

MOUSE MADNESS

In response to Marcel Bosch's question on page 95 of LXF43, I've experienced a similar problem, and I offer some tips that might help.

I run SuSE 8.1 Pro, and use the other virtual consoles a lot (as in Alt+Ctrl+F1) and switch between them quite regularly.

International readers – stop paying silly prices for LXF – turn to page 102 and subscribe to **SAVE MONEY!**

When I switch back to X (I typically run KDE) the mouse sometimes takes on a life of its own, tending towards the bottom left corner, and it jumps around, sometimes rearranging, or even deleting icons on my *Kicker* bar; but I've come to accept this problem as part of my Linux experience, as I also experienced this with SuSE 8.0 Pro on another machine.

However, with this experience I've learnt the following:

■ When this happens, don't try and overpower it – it only makes things worse. The way to get back to normal mouse functions is to switch back to a text console and then back to X.

■ It appears to be a problem with X, not KDE, as it happened one time when only *kdm* was running (though, as it's *kdm* rather than *xdm*, it could still maybe be a KDE-related problem).

To avoid these problems I've found it's best not to give any input to the mouse until about half a second after you switch back to X. More often than not, if I'm moving the mouse at the time, the cursor will behave oddly. But if I do nothing for that half second, there tends to be no problem.

I hope this might shed some light on the problem, or at least provide a workaround to make life a bit simpler for Marcel.

Aoife Nic Aodh, *Dublin, Ireland*

BEARPAW SCANNER

In LXF44 Dom Dorris asked for help with a Mustek BearPaw 1200 CU scanner.

I have a Mustek 1200 UB Plus – which is very similar to this one – particularly since it uses the gt68xx chipset in the scanner mechanism. Have a look at the following for more information on this: www.meier-geinitz.de/sane/gt68xx-backend/

To get a gt68xx-based scanner working you essentially have to download the firmware from the link above and copy it to something like `/usr/share/sane/gt68xx` so that the SANE backend can pick it up. I managed to get things working

with Red Hat 9 but I had to hack `/etc/sane.d/gt68xx.conf` to force it to use my particular scanner and not a more generic one.

I'm now using SuSE 8.2 and that, I think, managed to spot the exact type of scanner and did everything apart from download/install the firmware file which is when you get the bizarre errors messages about 'invalid argument' which really means "where have you put my firmware file?"!

OK, the above requires quite a lot of non-newbie-type hacking around and understanding of things like vendor and device IDs for USB scanners (you need to know what these are to pick the correct firmware file), but ultimately you can get there and have a perfectly working and usable scanner.

Oh, and I think you might be able to skip the searching for the firmware if you already have a Windows driver disk and can find the relevant file on that.

Andy Bianchi, *via email*

LINUX MODEM

For some time I ran a Lucent-based winmodem with *lt-modem*. There were few problems except that the Phoenix BIOS on my Asus CUSL2 does not always detect PCI boards (random failure). The problem with the winmodem-based solution is that not all Lucent modems can be driven by *lt-modem* and it is nigh-on impossible to get any information from suppliers about the products on offer. See LXF14.

In an attempt to isolate the pci detection problem I installed a MultiTech MT5634ZPX-PCI hardware modem which performs faultlessly. Installation is straightforward and instructions for Linux (simpler than those for Windows!) are included.

I am a Linux only user NEVER having run Windows. My OS was DOS 3.1 until I put together a new machine and installed Mandrake 7.2. Since then I have taken to compiling the kernel and use 2.4.19, which gives me the ability to run an Epson Scanner with USB2.

Peter Antonelli, *via email*



Reviews

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

LXF VERDICT EXPLAINED

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

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

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

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

Value for money/Documentation: Whichever is most appropriate!

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



10 The close-to-perfect product.



8-9 Good, but has a few niggles.



6-7 Does the job, but needs work.



5-4 Average.



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

THE TOP STUFF AWARD

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



WHAT'S NEW...

Armari RM-064-1AE Opteron Server >>

The old adage "you get what you pay for" just isn't true! **p20**

Yellow Dog 3.0

OS X isn't the only UNIX-like operating system that you can run on an Apple Mac **p24**

SNAP Graphics 2.0

If you need a lightweight graphics solution for something like an embedded app, this could be what you need **p26**

Revolution 2.0

The cross-platform rapid application development environment formerly known as *Runtime Revolution* **p28**

Scribus 1.0

Until now, Linux lacked a full-featured desktop publishing program – how does it compare with *QuarkXPress*? **p29**



Infomap Navigator

Another first for Linux – a route planner of Europe with handy features for both commercial and leisure travellers **p30**

OpenGroupware.org

Collaborative working in an MS Access style, only Open Source? A laudable aim indeed, but... **p31**

Zend Studio 3

EXCLUSIVE! The leading PHP development tool gets shiploads of new functionality **p32**

Books

Practical Cryptography, *Complete FreeBSD*, *Building Embedded Linux Systems*, *Java Extreme Programming Cookbook* **p35**

LINUX FORMAT BENCHMARKS EXPLAINED

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

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

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

BENCHMARKS

hd:	0.83
apache:	1.22
mysql:	1.11
compile:	0.96
oggenc:	1.71
Overall:	1.17

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

NETWORK HARDWARE

Armari RM-064-1AE 1U Opteron Server



Most would agree that Opterons lead the way in high performance computing, but what do you get when you take the fastest Opterons currently available, stir in 8GB DDR RAM, and bring to the boil? **Paul Hudson** finds out...

BUYER INFO

HPC 1U server designed to crunch numbers; can be customised for the more price-conscious. For smaller uses, consider the cheaper SR113 (reviewed in *LXF40*), or the NEC Express 5800 (in *LXF41*)

■ **SUPPLIER** Armari
 ■ **WEB** www.armari.com
 ■ **PRICE** £4694 (inc. VAT)

The specification for 1U servers defines a maximum width, height, and depth that a 1U server can be – any larger than that, and machines just won't fit into the rack slot provided for them. As a result, there's a limit to the size of 1U machines – if you want to add more power, you often need to juggle around other components to make it all fit. Many suppliers resort to 2U systems to make room for all the hardware inside, which is why we were very surprised to see a dual Opteron 244 machine (that's 1.8GHz apiece) with 8GB of RAM and a 40GB hard drive all wrapped up inside a small and relatively light 1U case. That much power generates a lot of heat, which made us wonder how they managed to dissipate it all in such a small case. Of course, when we turned it on we saw how it was all done...

Power-packed

In all our tests, Opterons have hammered (sorry!) the competition, and that's without having tested the

best Opteron CPUs available. When Armari called us up asking whether we'd like a dual-244 machine backed with 8GB of 333MHz DDR RAM, we thought we had died and gone to heaven, and yet the machine arrived a week later as promised, fully loaded with SuSE Linux Enterprise Server 8.

Armari has a long history of producing high-quality workstations and rack-mounted servers, with cutting-edge hardware being its hallmark. Its machines have won a number of performance accolades in their time, so it's really no surprise to see Armari taking such a confident lead with Opteron – the question is, has the company applied their magic to the new CPU from AMD?

While the box isn't attractive by any means, it's certainly functional – a floppy drive, CDROM, two USB ports, and an array of multi-coloured lights all combine together to make for an ugly, but helpful front face.

Around the back, there's the usual keyboard, mouse, and VGA ports, as well as two more USB ports and 2 Ethernet ports. The machine comes with two Broadcom gigabit LAN cards built in, which puts it at the top end of the networking scale. SLES 8, as always, worked like a charm – all the hardware came preconfigured, even booting into KDE 3 on start up.

Also inside the box were two manuals – one for the Rioworks HDAMA motherboard, and another for the ethernet controllers – as well as a

contents list. It also came with a driver installation CD for the motherboard (for the onboard graphics card and network cards), but, sadly, no installation CDs for SuSE Linux Enterprise Server 8 (SLES) – we contacted Armari about this apparent mishap, and were pleased to hear that full machines (that is, non-review boxes) are all supplied with a full copy of SLES 8 in its original packaging – all the CDs and the excellent SuSE manual too.

Sound and software

The power button is a tiny little dot at the front of the box, which is fair enough – however, having pressed it, everyone within twenty metres must have wondered where the hovercraft-like sound was coming from. Loud? You don't know the half of it – this machine has three 9.5cm chassis fans in, which combine together to make a colossal noise. However, these powerful fans do mean that the server is cool to the touch – I'd even go so far as to say it's cold, which is remarkable considering how hot these machines usually run. The price you pay is that the machine is twice as loud as anything else *LXF* has reviewed in for a long time, if ever.

Bundled with the machine was SuSE Linux Enterprise Server 8, an excellent server distribution that makes it easy for even novice admins to get to grips with their hardware. It was well-configured on the machine, although we had to resort to our own

SLES CDs because none were provided. Armari are happy to supply Linux with any of their servers, and support it, too – they can even help you set up SuSE technical support to handle any server issues you run into. Of course, if any part of the machine fails, you're covered by a basic return to base warranty, so whatever you do with this thing you are rarely far from a helping hand.

Armari had preconfigured the machine that we reviewed with *Apache* and *MySQL* servers, so it can simply be dropped into a corporate environment and start working immediately. The server was configured to boot into X-Windows, which is perhaps an unusual choice for a system like this, however it wasn't difficult to re-arrange things to get it to stay at the shell as most people would expect.

Performance

There's no doubt that this machine was designed for high-performance computing – putting this machine to use as a file and print server is quite frankly sacrilege. The hard drive bundled wasn't anything special – it's all about RAM and CPU speed, and that's where the machine tears away. These machines are designed to work





with many others running through huge calculations – Armari deploy these machines in large racks, usually also agreeing a bespoke maintenance contract to make sure that each customer gets full support.

Because the machine is geared so clearly towards high-speed calculations and computing, it wasn't surprising that the hard drive performance was less than stellar. While a score of 3.63 is very good considering that the machine comes with a plain IDE hard drive (40GB Western Digital Caviar, for the purists out there), it's not really in the league of the Xinit Sharq machine we reviewed in *LXF42* – of course, it got its score of 8.55 by having multiple SCSI hard drives, a feature this machine didn't come with. However, if hard disk performance is something you need, Armari can provide you with whatever you need – you can configure your machine to be IDE or SCSI, depending on your preference.

In the real-world tests, oggenc and kernel compilation both stand out as something for Armari to be proud of – these are the two best scores in this category we've ever seen, again proving

the quality of AMD's flagship CPU, and also saying volumes about the trouble Armari have gone to to make this machine leave the world in its dust. Perhaps 8GB of RAM is a little over the top, but it's all down to what kind of data you intend on working with, and at this price it's hard to complain!

The biggest win for this machine is its Apache result, which is a world-beating 7.33 – that's what you get when you take a machine designed for high-performance computing and set it to work on a task as "mundane" as web serving. The MySQL performance would almost certainly benefit from a better hard drive – the switch to SCSI alone would see the performance spike a great deal. However, given that the goal of this server is to run everything from RAM – and, really, with 8GB to spare it's not so hard! – the hard drive performance can be excused.

Overall, performance for this machine is very high – yes, it's dragged down by its hard drive, but even *with* that poor area, the sheer speed of the CPU/RAM combination means that the machine turned out a very respectable score of 3.72. We'd be very interested

in seeing a machine that has a better hard drive (preferably with RAID), as the chances are that's the only thing holding this thing back from a score like we've never seen before.

Conclusion

The Opteron continues to astound and delight here at *LXF Towers*, and it seems like its performance knows no bounds. The more servers we see powered by AMD64, the more we realise that it wasn't really a risk at all for AMD – they must have realised from the start that they were onto a sure-fire winner that will guarantee them the number one spot in benchmarks for years to come.

Armari has a history of producing machines that work above and beyond competitors at the same price point because the company always takes the time to make sure each component is tuned to work perfectly with the others around it. This machine is no different – while the hard drive is a bit suspect, the rest of the machine is head, shoulders, chest, and hips above its competition, leaving Armari as deserved leaders in the field.

At just under £4000 before VAT, this is high-performance computing on a shoestring budget, with the difference that you're getting no-compromise hardware – if you want to spend a little

All the connectivity you'd expect – keyboard, mouse, VGA, a pair of USB sockets and two Ethernet ports, plus there's two Broadcom gigabit LAN cards nestling in there.

more and invest in a more powerful hard drive, this server's only a few clicks away on the Armari site. In fact, the worst thing about the entire system is its name – "RM-064-1AE" doesn't exactly roll off on the tongue!

So, to conclude: more laurels for AMD and their all-singing, all-dancing Opteron; a hearty "thank you" to Armari for showing all the other manufacturers why it's called *high-performance* computing; and for the rest of us, this machine is one more good reason why, with Opteron, the adage "you get what you pay for" just isn't true any more – you get a whole lot more. **LXF**

BENCHMARKS

hd	3.63
ab	7.33
mysql	0.70
compile	2.14
oggenc	4.81
Overall	3.72

VERDICT

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

HPC has a new yardstick, and it gives you change from £4000

LINUX FORMAT RATING
9/10



Reviews YellowDogLinux

POWERPC LINUX DISTRO

Yellow Dog 3.0

Richard Drummond rejoices that Mac OS X isn't the only Unix-like OS that you can run on a Mac.

BUYER INFO

Other vendors producing PPC distros include Debian, Mandrake and Gentoo.

- **DEVELOPER** Terra Soft Solutions
- **PRICE** See *How much...* box
- **WEB** www.yellowdoglinux.com/

Common sense should tell you that the Mac and Linux would go together like, well, apples and oranges. After all, the Apple Macintosh is the epitome of the computer as an appliance, while Linux is diametrically opposed to this idea as the hacker's favourite OS. Common sense, though, as is often the case, is wrong. Linux actually runs rather well on the PowerMac, and here we have a Linux distribution to prove it: Yellow Dog Linux.

Red Hat don't support the PowerPC platform (the hardware architecture that the Mac is built upon) but Yellow Dog Linux (YDL) is the next best thing: YDL 3.0 is in essence a PowerPC port of Red Hat 8.0. The previous release of YDL was derived from Red Hat 7.3, and many of the improvements in this version are directly derived from improvements to its progenitor, Red Hat 8.0. For example, YDL 3.0 offers a slick,

graphical installer, based on Red Hat's *Anaconda*, and unified GNOME and KDE desktops with a common graphical theme (called Wonderland, the earlier name for Red Hat's Bluecurve theme). Other benefits are wholly its own, however. YDL 3.0 includes the latest PPC kernel, to better support newer Mac hardware, and the latest version of XFree86, which can drive the latest Mac graphics cards such as the ATI Radeon 9000 and those based on nVIDIA's GeForce4 chipset. What's more, YDL has grown in size from two binary CDs to three, thus offering a much greater selection of software.

Dog on it

Thanks to the new installer, installing YDL 3.0 is simple. Actually, there are a few potential problems to be faced, the biggest of which may be getting your Mac to boot Linux from the installation media (see box, *Old vs New*). Users with New World machines should have few difficulties, but Old World Macs can be problematic. One of the difficulties to be faced is support for graphics cards. Earlier YDL releases could fall back on a text-mode installer and thus required only that cards be supported by Open Firmware (the Mac's superior equivalent of the



Yellow Dog 3.0's graphical installer is a tweaked version of Red Hat 8.0's *Anaconda* and is comfortable and straightforward to use.

PC's BIOS). YDL 3.0 just has the graphical installer, though, which needs a display capable of 800x600 in 15-bit depth. Depending on your graphics hardware, getting YDL to boot into a suitable screen mode may require some trial and error, and the documentation supplied is rather thin in this area.

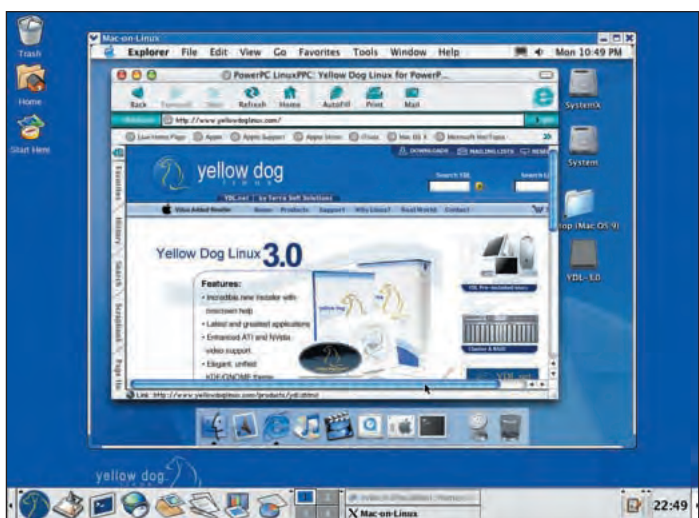
The other problem to be faced during installation is finding free space for YDL to install into. YDL has no tools for non-destructively resizing partitions, so you need to do this before installing. The YDL installer can automatically use any free-space or you can use its *Disk Druid* to manually set-up your Linux filesystem; only *ext2* and *ext3* partitions are supported. A point worth noting is that YDL provides no kernel support for HFS+ filesystems (as used in Mac OS 8 and later), so it's a good idea to create a small HFS partition if you want to swap data between Mac OS and Linux. (Alternatively, you can use the cumbersome *hfsplus* shell commands.) This is particularly an issue with Old World machines since, for instance, when you upgrade a kernel, you need to copy the new kernel image across to where Mac OS and BootX can find it. (The YDL installer no longer supports the *Quik* boot-loader which can boot Linux on an Old World Mac directly from Open Firmware.)

Hardware detection is generally good in YDL. I successfully installed it on a range of machines from an old all-in-one 5400 to an early G4 with no serious difficulties. The fact is that the Mac platform doesn't enjoy the

overwhelming abundance and variety of hardware that the PC platform does, but, conversely, fewer developers work on supporting hardware in Linux PPC. Consequently, most common Mac hardware works well in YDL 3.0, but problems can arise. As I mentioned above, support for graphics cards may be an issue, especially on older Macs. Thanks to the XFree86 4.3.0, support for the latest Radeon and nVIDIA chipsets is now included, but, generally, ATI Rage 128 and older Radeon chips work the best. (Only these two chipset have 3D drivers in YDL 3.0). The Mac's built-in sound, SCSI and ATA devices should present no problems. When compared to the PC, a surprise bonus of running Linux on a Mac is that power management works well. Without the complexities of APM or ACPI to deal with, YDL 3.0 on an iBook, for instance, will go to sleep like a baby and makes for a great mobile Linux platform.

Hat and dog

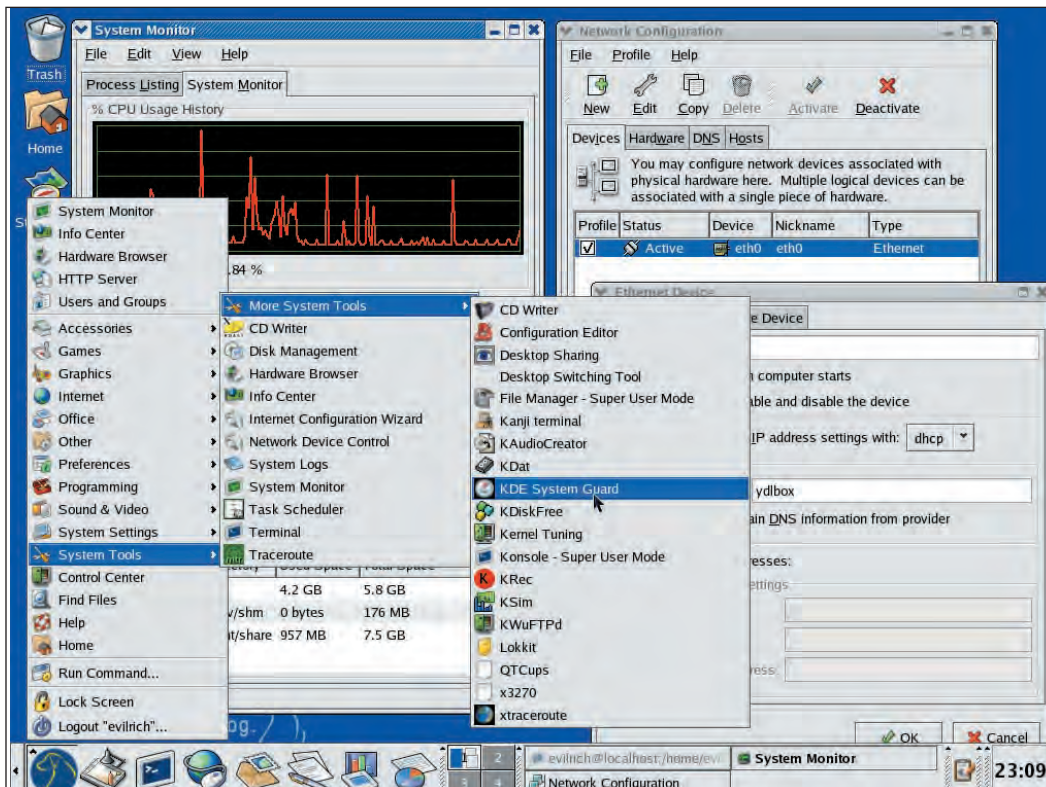
Once installed, YDL creates a comfortable and fairly complete desktop environment. It borrows Red Hat's controversial unification of the GNOME and KDE desktops, along with



Mac-on-Linux – the killer app of PowerPC Linux – lets you install and run Mac OS X, even on machines not normally supported by Mac OS X.

Details

Kernel	2.4.20-ben2
Glibc	2.3.1
GCC	3.2.2
XFree86	4.3.0
KDE	3.1
GNOME	2.2
Mozilla	1.2.1
Apache	2.0.40



YDL includes a wealth of software, but oddly there aren't any graphical tools for managing and installing packages.

the unified Bluecurve, sorry, Wonderland theme. This rather begs the question – why bother with both GNOME and KDE if you are going to make them look and act the same? But it gives a polished look and feel to the desktop and makes navigation simpler. However, Terra Soft has skimped in a few areas. YDL includes a weak selection of fonts – only the standard X and Ghostscript fonts – and the CD auto-mounter doesn't work.

YDL comes stocked with a generous selection of software, such as Mozilla, Evolution, OpenOffice.org, The GIMP and so on. All the usual server tools and development tools are included too. One notable omission, however, is any kind of video player – there's no Xine, MPlayer, VLC or anything. And while a Java run-time is included, it's the out-dated version 1.3.1. A selection of Red Hat's configuration tools are supplied for things like sound, network, printer and server configuration, but otherwise system management tools are weak – especially when compared with Mandrake PPC. The life-saving remote admin suite, Webmin, isn't featured.

Naturally, YDL does include PowerPC Linux's killer application, Mac-on-Linux (MOL). This virtualizes the Mac hardware – in a manner similar to VMware on x86 hardware –

and let's you install and run multiple Mac operating systems simultaneously. The latest version of MOL lets you install and run Mac OS X – even on Mac hardware not supported by Mac OS X (of course, you still need enough horsepower and memory). MOL is particularly useful for providing a fast and stable foundation on which to run the classic Mac OS, and one that is not fraught with networking niggles which beset real Macs.

For package management, YDL features a port of Connectiva's RPM version of Debian's famous *apt* (Advanced Package Tool). This can be used to keep your system up-to-date with bug and security fixes and to fetch new software from the web – and unlike Red Hat's *up2date* service it requires no registration. However, there are number of problems with *apt* on YDL. Firstly, the *apt* sources configuration file is broken, and you'll need to modify this (the main YDL *apt* repository is called 'main' not 'base') to

get it working. Secondly, Terra Soft include no front-end of any sort for *apt*. While *apt* is a powerful system, it can be a pain to use manually. (I recommend you pop along to <http://freshrpms.net> and grab a copy of *Synaptic For YDL*, a handy GUI front-end for *apt*. *freshrpms.net* also maintains a repository of extra software for YDL 3.0 obtainable via *apt*.) Thirdly, the YDL install discs themselves cannot be used as an *apt* source, so you have to install any additional software from CD the hard way. There's not even a graphical RPM installer such as *KPackage* to aid with dependency grappling.

Old dog, new tricks?

Oversights such as these mar what would otherwise be an excellent product. Despite its flaws, though, Yellow Dog Linux is still very much the leading Linux distro for the Mac. It is easier to install than its rivals, and runs well on a wide variety of Mac

Old vs New

The PowerMac generation gap

Apple calls all Macs since the iMac “New World” machines; earlier Macs have been retrospectively termed “Old World”. Why should you care? Well, the principal difference between Old and New is the way they boot an operating system, and this is a big deal when it comes to installing Linux. Old World Macs have a substantial portion of Mac OS in ROM and are designed to boot from that; when the New World dawned, Apple did away with the physical ROM, and instead New World Machines boot a ROM image (or boot-loader) from disk or CD-ROM.

Since they can boot directly from CD-ROM, New World Machines are much easier to install Linux on. If you have an older beige-box PowerMac, then typically the simplest way to install Linux is via the Linux boot-loader BootX (not the Mac OS X boot-loader of the same name). This is a plug-in or ‘extension’ for the classic Mac OS which can run early on during the Mac OS boot sequence and is a graphical equivalent of Grub or LILO on the Intel platform. It lets you pick a kernel to boot and supply boot parameters – or continue loading Mac OS. The disadvantage of BootX is that you need Mac OS installed to use it.

hardware. And YDL 3.0 is a huge improvement on its predecessor, YDL 2.3, with benefits including better hardware support, a greater selection of software and newer software versions. YDL compares less favourably with current x86 distros and with its competition on the Mac, Mac OS X. If YDL is to stand out in that crowd, Terra Soft needs to pay a lot more attention to detail. Solve these issues and nobody will be able to justifiably ask that question again, “But why run Linux on a Mac?” [LXF](#)

VERDICT

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

Yellow Dog 3.0 is the slickest and friendliest way to get Linux installed on a Mac, but there's still a lot of work to be done for it to compete with either Mac OS X or x86 Linux distros.

LINUX FORMAT RATING

7/10

How much is that doggie?

- **Geek Edition \$24.95**
6 CDs (3 binary, 3 source) plus sticker.
- **Box Set \$54.95**
As above plus “Getting Started with Yellow Dog Linux” book.
- **Box Set with Installation Support \$84.95**
Same as Box Set but includes T-Shirt and 60 days of installation support.

GRAPHICS DRIVERS

SNAP Graphics 2.0 beta

Whether X configuration is getting you down or you need a lightweight graphics solution for that embedded app, **Richard Drummond** thinks SciTech has the answer.

BUYER INFO

SNAP Graphics works with XFree86 and Qt Embedded. A competing Open Source product is *DirectFB* (see www.directfb.org/).

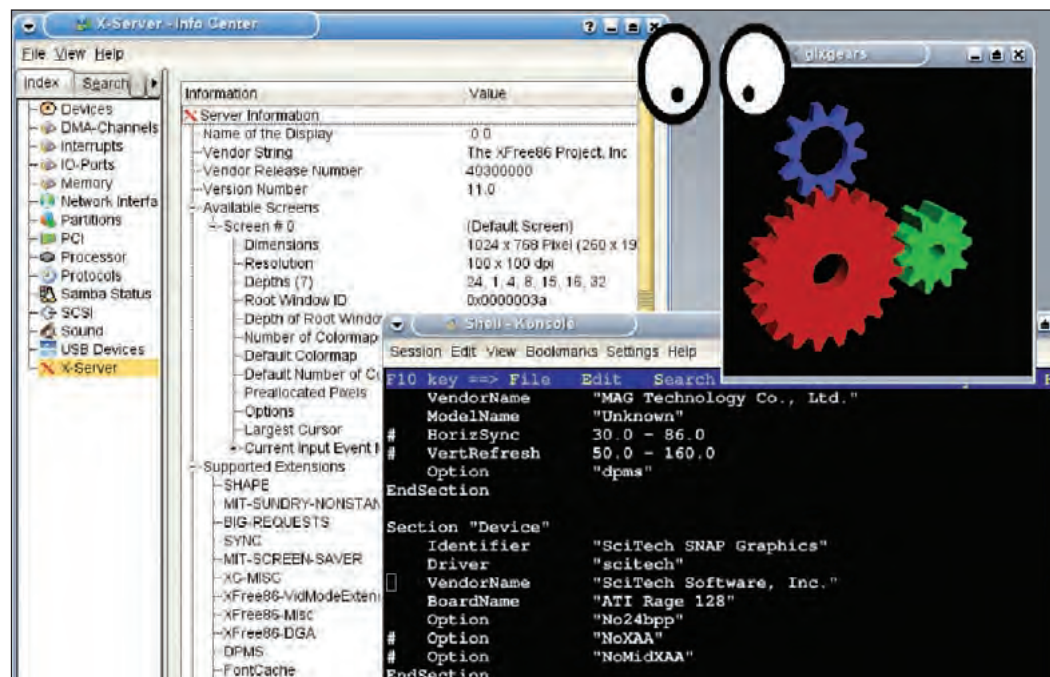
■ **PUBLISHER** SciTech Software
 ■ **PRICE** Price \$19.95 (preorder)
 ■ **WEB** <http://www.scitechsoft.com/>

Since the release of the version 4.x series, XFree86, the most commonly deployed X server on Linux, is easier to configure than ever. Nevertheless, it can still be a thorn in the side of sysadmins, distro vendors and end users. Wouldn't it be great if X could automatically configure itself for your graphics card and monitor and use an appropriate screen mode without all that tedious mucking about with configuration files? Well, with SciTech's *SNAP Graphics* drivers it can.

It's a snap

SNAP (System Neutral Access Protocol) is the name for SciTech's cross-platform driver architecture which allows the same binary driver to be used on a range of operating systems (running on the same CPU architecture, of course) and with multiple APIs. *SNAP Graphics* uses the *SNAP* architecture to provide a robust, high-performance 2D graphics framework for embedded and desktop environments. Linux users may not have heard of *SNAP* before, but it has a long pedigree on OS/2 and DOS; rumour even has it that AmigaOS 4.0 will ship with *SNAP Graphics* drivers.

The Linux version of *SNAP* is still in the late beta-testing phase, but a time-limited version can be downloaded from SciTech's site for evaluation. Currently, this package features 2D-accelerated drivers for a wide range of modern and not so modern graphics chipsets including the latest Radeon and nVIDIA chips (see SciTech's site for the full list of compatible hardware). Also included in the package is series of XFree86



The SNAP Graphics X driver works transparently with your X server and requires minimal configuration.

driver modules – for XFree86 4.0.2 to 4.3.0 – which allows XFree86 to make use of the *SNAP* drivers to drive your X display – rather than the standard XFree86 drivers.

SNAP supports multiple APIs not just X, though. *SNAP Graphics* drivers will also work with *Qt Embedded* (the small-footprint version of the *Qt* toolkit which doesn't require X), and SciTech itself offers an advanced 2D drawing toolkit, called *MGL*. In addition, the *SNAP SDK* is available if you want to develop directly to the *SNAP Graphics* architecture. And, while the *SNAP* drivers are a proprietary, non-free product, both *MGL* and the *SNAP SDK* are available under LGPL and commercial licenses. A port of the cross-platform GUI toolkit *wxWindows* has been made targeting *MGL*.

Low maintenance

Installing *SNAP* is easy. Simply unpack the archive, launch the installer and everything will be magically done for you. Restart X, and your X server will now be using *SNAP*.

The installer modifies your X config (making a backup first) creating new 'Device', 'Monitor' and 'Screen' sections. If you take a look at the changes the installer makes, then you will see that very little configuration is actually required. Typically, if you have a DDC-capable (plug-and-play) monitor and chipset (and everybody does these days) then all that is really required is to tell X to use the 'scitech' module in the 'Device' section of the config rather than the XFree86 chipset driver. No monitor or screen-mode configuration is strictly necessary, since *SNAP* will probe for your monitor capabilities and choose an appropriate screen mode automatically. Of course, you can add 'Mode' definitions if you want *SNAP* to use a particular screen size, or add a 'DefaultDepth' option, but any additional configuration required can be performed via the command-line and menu-based tools that ship with *SNAP*, such as *gactrl* and *gamode*.

The *gamode* tool, for instance, allows you to fine-tune the display modes that your graphics adapter

uses, to tweak screen sizes and positions on screen and refresh rates. It's text menu-driven, but it's much easier and quicker to use than its counterpart in XFree86, *xvidtune*. Also useful is *gaoption*, which lets you manipulate various aspects of *SNAP*'s and your card's behaviour. This can be used, for example, to set your card's AGP rate, to disable hardware acceleration, or to rotate the screen.

A useful corollary of *SNAP*'s plug-and-play nature is that when used with DDC-capable devices, your X config doesn't need to explicitly define your graphics chipset or monitor capabilities. For example, in the 'Device' section, it will just associate the Scitech driver with the card via its bus ID. Thus you can swap your graphics card or monitor, reboot your system, and everything will magically just work with the new hardware – without any modification. Very few Linux distros can do this kind of thing so fluidly or reliably.

So *SNAP* is easy to use, but how well does it perform? Firstly, let's mention some limitations. The *SNAP*

Graphics X plug-in currently does not support the XVideo extension or GL direct-rendering, thus it doesn't provide hardware-acceleration for video playback or 3D graphics. The GLX extension itself will still work (albeit via software rendering only), as will most other common X extensions – such as *Render*, *DGA*, *Shape*, etc. The new *XRANDR* extension also seems to function with *SNAP*, but rather curiously, it doesn't report refresh rates for any screen modes that it lists; the new *ARGB* cursors in XFree86 are supported.

Multi-headed setups

Another limitation is lack of support for multi-headed set-ups. Multi-headed Matrox cards are supported, but multiple, discrete graphics cards are not, it seems. The *snaptx86config* tool won't create multi-headed config files for you, and when I tried various multi-head configurations a whole slew of problems occurred, from conflicting I/O spaces between the cards to complete system lock up. *SNAP* doesn't currently seem able to soft-boot the BIOS on secondary cards like XFree86 can – but that it could just due to my other problems.

As we said, *SNAP Graphics* provides driver coverage for a wide range of hardware – but support for older chipsets is notably lacking. For example, I tested an S3 Virge-based card and an old Cirrus Logic Alpine card (the latter is supposed to be supported), but *SNAP* fell back to using its unaccelerated VESA driver for both. Support for newer chipsets is better than XFree86, however, and this is one area in which *SNAP* may score. The slow release schedule of XFree86

How snappy is SNAP?

A comparison of SNAP with other D drivers

CHIPSET	DRIVER	X11PERF TEST (GXcopy x1000 ops/sec)				XMARK
		RECT100	POLYTEXT	SCROLL100	PUTIMAGE100	
MGA G200 (AGP/8MB)	XF86 VESA	10.8	1040	8.8	3.80	29.73
	XF86 MGA	22.1	3140	14.3	5.39	57.54
	SNAP	27.0	2860	14.2	5.36	53.68
	XDIRECTFB	21.7	790	11.9	3.11	15.93
RAGE128 (AGP/16MB)	XF86 VESA	10.4	1030	8.7	3.91	29.55
	XF86 R128	50.6	3190	19.4	6.13	60.82
	SNAP	51.1	3100	19.3	6.03	51.06
	XDIRECTFB	44.9	804	15.8	4.00	16.36

All tests were run on a 1200MHz/256MB Athlon system running Debian unstable and XFree86 4.3.0. Version 1.0RC5 of XDIRECTFB was used, built against XFree86 4.3.0. Each X11perf run was executed on a newly booted, completely unloaded system, running only the relevant X server (and no window manager). See *man x11perf* for details on the various operations that x11perf tests and *man Xmark* to find what

the Xmark means. Note that the XFree86 VESA driver has no hardware acceleration, while the XDIRECTFB server currently features minimal acceleration (and is based on a very different rendering model).

See the main text of this review for a discussion of the significance of these results. Remember that the Xmark is not on the whole a reliable gauge of real-world performance.

means support for new cards often lags the state of the art significantly. Not only that, *SNAP* effects a separation of the driver from the core X infrastructure, and thus allows the user to simply update his or her driver – to fix bugs or to support new chipsets, perhaps – without having to update the whole of X. It has been suggested that XFree86 project should release their drivers separately from the X server for precisely these benefits. Well, Scitech is already there.

Scitech claims improved performance for the *SNAP Graphics* architecture over plain old XFree86 and has published some Xmark benchmarks on its website to prove it. I similarly ran some Xmark tests on a selection of X drivers under XFree86 4.3.0 and on XDIRECTFB (see *How Snappy is SNAP?*

above) to try gauge what performance benefits are possible. Rather than the 10% gains claimed by Scitech, I found that – on the full Xmark suite at least – that the *SNAP* drivers were marginally slower. Admittedly, I tested with different chipsets and with XFree86 4.3.0 rather than 4.2.0.

What do those numbers mean in the real world? As is usual with benchmarks, very little. In fact, the Xmark test is rather more artificial than most benchmarks, since it tests the complete range of X primitives – with *GXcopy* and *GXxor* raster operations – not just the ones that are commonly used. Most of these, especially when used with the *GXxor* operation, will only every be used by software designed to run on a monochrome display, so it is not worthwhile for a developer to optimise a driver for them. That's why I have listed the results from a selection of individual *x11perf* tests in addition to the Xmark score, which hopefully gives a better picture of how more modern GUI toolkits will perform with the various drivers. As you can see, there's really not much difference between the *SNAP* drivers and the standard XFree86 drivers.

Is *SNAP Graphics* worth paying for, then, simply to replace the free XFree86 drivers? It depends on who you are and your requirements. If you are using Linux at home and X is working perfectly well on your hardware, then I can currently see little reason to buy and install *SNAP* – not

until *SNAP* shows some more marked performance gain over XFree86 and at least supports the XVideo extension (and ideally GL direct-rendering, too). However, in the corporate world, the minimal maintenance required by *SNAP* could be a real winner. I could also see distro vendors who are already not shy of bundling proprietary software, such as Xandros or Lindows, adopting *SNAP* simply to reduce user support issues. Finally, *SNAP* makes a great deal of sense in the embedded space, where its small footprint, good hardware support (much better than *DirectFB* and *Qt Embedded*, for instance) and support for multiple APIs make it an attractive solution.

The final decision, however, must be left until a final release is made and a price announced. For comparison, the OS/2 version currently costs \$36.95 including user support, and Scitech say the Linux edition will be priced similarly or lower. That sounds like a good deal to me. **LXF**



You can preorder SNAP from <https://store.scitechsoft.com>

VERDICT

Features	7/10
Performance	8/10
Ease of use	9/10
Value for money	9/10

SNAP can solve your configuration and user support issues, but a lack of video and 3D support may deter some users.

LINUX FORMAT RATING
8/10

RAD ENVIRONMENT

Revolution 2.0

Richard Drummond finds out that developing database-driven and XML-parsing applications needn't be hard work, as this user-friendly development platform proves.

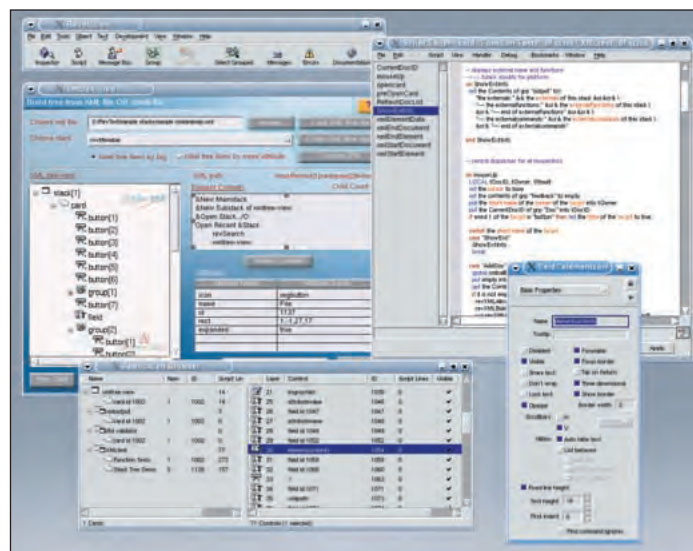
BUYER INFO

Cross-platform object-based rapid app development environment; try also *Kylix* and Java.

- **DEVELOPER** Runtime
- **PRICE** See Pricing box below
- **WEB** www.runrev.com

Revolution is a cross-platform, rapid development suite based on the MetaCard scripting engine, itself a derivative of Apple's pioneering HyperCard system, which supports developing and deploying on Linux, Windows, Mac OS (Classic and X) and various Unices. *Revolution* includes an IDE and toolkit for object-based development and allows the user to create a program or "stack" by assembling pre-built components with a familiar point-and-click interface and bringing them to life with scripts written in the English-like Transcript language. The emphasis for MetaCard and HyperCard was to enable the user with little programming experience to easily create multimedia presentations. Given this heritage, it is unsurprising that *Revolution* comes stocked with a powerful set of multimedia tools. But, with *Revolution*, Runtime is trying to re-target the HyperCard technology for real-world, business applications. Thus, for example, *Revolution* adds support for networking and for database access.

We reviewed *Revolution 1.0* back in LXF40, and found it to have much potential. Unfortunately, that release was fraught with bugs and the database tools were rudimentary. *Revolution 2.0* solves many such



Getting data into *Revolution* projects is easier than ever due to improved database access functions and support for parsing and validating XML.

problems and adds a whole host of new features, including improved database access, support for Unicode characters, and an XML parser.

Up the Revolution!

Revolution 2.0 feels a lot more solid than its predecessor. The scripting engine in the first version was liable to throw exceptions frequently for no good reason, which was a real killer for productivity. This version clears up those bugs, but it still seems flaky when run with some window managers, particularly KDE's *Kwin*. Another foible is that the mouse wheel doesn't scroll windows in the usual direction.

The new interface layout in *Revolution 2.0* provides a more fluid workspace. The essentials are still the same: the Application Browser provides an overview of the

components in your project, and to add components, you simply pick can them off the component palette, place them on your project window and modify their properties with the property editor, now called the Property Inspector. Navigation is easier, thanks to improvements in the Application Browser – it now actually provides a true hierarchical list of stacks and the cards (or panels) that they contain. Also, the selected card's controls, groups, animations etc are all shown in a flat, sortable list on the right – rather than separately as before. The property editor has been revised, too, and the script editor is now a separate window in its own right. The script editor now offers debugging controls for setting breakpoints and watching variables and, with the ability to list a component's handlers (methods) down the left of the window, better navigation.

Database support has been greatly improved with a new and full-featured toolkit for database access. This provides a library of functions which you can use to establish a database connection, perform queries, navigate result sets and handle transactions. Alternatively, a graphical query builder

is provided which lets you create a database connection and build queries by point-and-click. These appear in your project as non-visible components: they don't show up in your project's GUI, but can be treated just like any other *Revolution* object. Connections to MySQL, PostgreSQL, Oracle, Valentina or any database with an ODBC interface are supported.

Revolution 2.0 adds a lot more features than we have space to cover here. One other aspect of note, though, is the new multimedia tools: the ability to record video and audio streams and support for text-to-speech voice synthesis aren't included on Linux! The fact is that *Revolution* on Linux and Unix platforms enjoys a poorer range of multimedia features than on Windows or Mac OS, where much of this functionality relies on QuickTime. Moreover, that's the only area I can fault with *Revolution's* otherwise excellent and copious documentation. Just what works and what doesn't on Linux? It also highlights the fact that *Revolution* does a less than perfect job of abstracting away the underlying platform – especially if you compare *Revolution* to Java, for instance.

Nevertheless, when taken in context, *Revolution 2.0* is an attractive product. It's not in the league of *Kylix* or Java, and isn't suitable for back-end work. But with support for XML (and hence SOAP) and improved database access, *Revolution* makes an ideal tool for creating graphical front-ends to your business applications. **LXF**

Revolution Pricing

Scaled according to needs

Developers can choose between three versions of *Revolution* depending on their needs and budgets. Prices start at \$149 for the basic version, *Revolution Express*, which includes a license for developing on one of the supported platforms. *Revolution Studio* adds cross-

platform support for \$199. The top of the range is *Revolution Enterprise*, which also adds tech support, updates, printed documentation and more, and is \$1199. Discounts are currently available, and an evaluation edition of the program can be downloaded for free.

VERDICT

Features	8/10
Performance	7/10
Ease of use	9/10
Value for money	8/10

Revolution 2.0 corrects and improves upon its predecessor and establishes this as a serious development platform.

LINUX FORMAT RATING
8/10

DESKTOP PUBLISHING

Scribus 1.0

Until the advent of this app, Linux lacked a fully featured DTP package. But can this Open Source project cut it in the real world? **Andy Channelle** finds out.



BUYER INFO

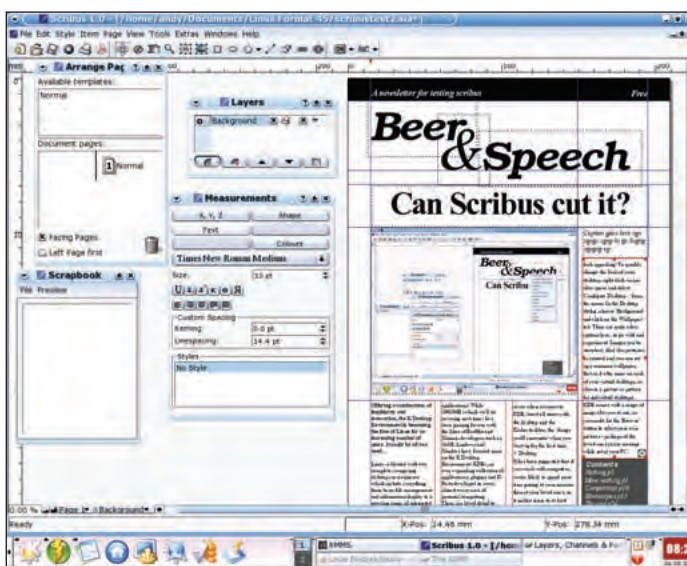
Native Desktop Publishing, with little competition beyond Open/StarOffice.

- **DEVELOPER** Franz Schmid
- **LICENSE** GPL
- **WEB**
<http://web2.altmuehlnet.de/fschmid>

Scribus is an Open Source attempt to create an alternative to the likes of *QuarkXPress* and *Adobe InDesign*, both of which offer an enormous range of design tools but with price tags way out of reach of the average community newsletter editor. But can a small project compete with the army of coders dedicated to making their app king of the presses? Well, *Scribus* is a valiant attempt. It covers the basics of design and text control with finesse, and while lacking some of the more recent additions to the mainstream, such as automatic drop shadows and transparency, there is no shortage of usable tools. For the sake of comparison, the tools seem to be just beyond the level of sophistication of *QuarkXPress 3*, while the interface owes more to Adobe's neglected *Pagemaker* application. Getting this far is in itself an achievement, but factor in the extensive PDF support, and you have a potential killer app capable of creating output suitable for professional production.

Binaries for the latest version are not yet available, but compiling from source didn't pose any serious problems, though you will need libraries and headers for a recent release of Qt. In my case, this involved installing the *Qt-devel* package from my distro's CD.

Users of recent development versions will notice a welcome tweaking of the UI. For instance, text formatting tools are now much more prominent in the 'measurements' pallet, instead of being hidden away, though there are still areas of confusion where I can't work out if a feature has been omitted, or I'm just not able to find it – the two most important being resetting the datum point for the rulers and automatic page numbering. The documentation, as is



Scribus has a wide range of design tools, but some you'll have to hunt for.

the case with many Open Source projects, is less than comprehensive.

Prints charming

There is much to admire here in both the tools on offer and the UI. *Scribus* is frames-based – you make space for page elements before importing them. A variety of image formats including .tif, .png and .jpg are supported, as well as .eps and .pdf, while text import is limited to just plain text. Text import is one area it is lacking, especially if you expect to receive documents from a variety of sources. Something like *OpenOffice* can handle the main formats (and copy/pasting from there worked well), but a few extra import filters would be very welcome.

There appears to be a bug that prevents text selection across frames, meaning long passages need formatting in a single frame before moving them onto a real page. This could effectively rule the app out for anything beyond a few pages, but doesn't impede the creation of short pieces too much, but can get frustrating if you intend to do a lot of post-import editing.

There is a good selection of both character and paragraph formatting, and the facility to post-define reusable styles will save a lot of effort on

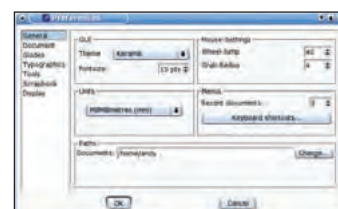
periodical projects which rely on a consistent look and feel. Reusable page elements – mastheads, flashes etc – can also be dragged onto the *Scribus* Scrapbook and then dropped on pages in the future. For more robust automation, the Python extension, can be used for quite complex scripting.

Graphical tools on offer include rectangles, ellipses, lines, polygons (including stars) and Bezier curves and these shapes, once created can be converted to image or text frames, or filled with solid colours or gradients. Adding colours and text styles to the list is a little cumbersome, recalling the 'Edit x' dialog of *Pagemaker*, but it's not too long-winded and, in the end, makes for more streamlined design.

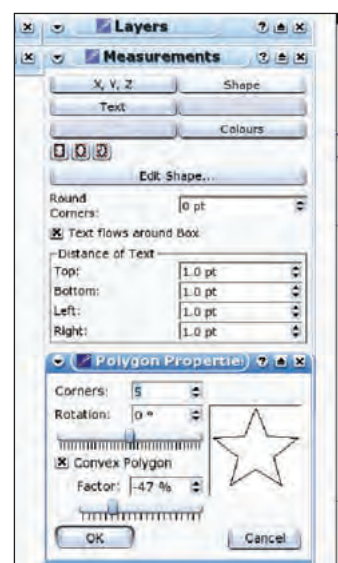
Conclusion

As a work in progress, *Scribus* is very usable – it handled a real-world production job and output a .pdf suitable for printing by my local bureau (printer's marks and bleeds will be great additions in future) with results comparable to mainstream packages – thanks mainly to the excellent pdf implementation.

My experience was blighted by a few random crashes and a hunt for features that are either missing or very well hidden, but overall this is very usable. If



The Preferences are extensive, easily understood and accessible.



The UI has been refined recently.

your needs are modest – a student mag, community newsletter, leaflet etc – fine as alternative to lower-end products from the likes of Microsoft and Serif, but not quite up to *QuarkXPress* standard though. But when you consider *XPress* will set you back at least £1000, *Scribus* is an exciting eye-opener. **LXF**

VERDICT

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

A good application in a field of one. At this rate of development, *Scribus* is definitely one to watch. It has the potential to go onto great things!

LINUX FORMAT RATING
7/10

FINDING YOUR WAY

InfoMap Navigator Europe

Do you need to know where you're going, or what to do when you get there?
Andy Channelle tests the first commercial route planner for Linux.

BUYER INFO

Route planning software with competition from websites such as www.mappoint.msn.com and www.streetmap.co.uk.

- **DEVELOPER** Directions Ltd.
- **PRICE** Price: £34.99 (UK/Ireland). £49.99 (Europe)
- **WEB** www.directions.ltd.uk

There are two reasons, arguably, why it has taken so long for Linux to get a route planner. Firstly pulling together so much information is labour-intensive; and secondly, sourcing the data is extremely expensive. *InfoMap Navigator* gets around both these issues by being a dual-format product developed using Qt, so the comparatively small Linux audience is bolstered, potentially, by Windows users. But being Qt-based it fits in perfectly with any KDE themes you may have installed.

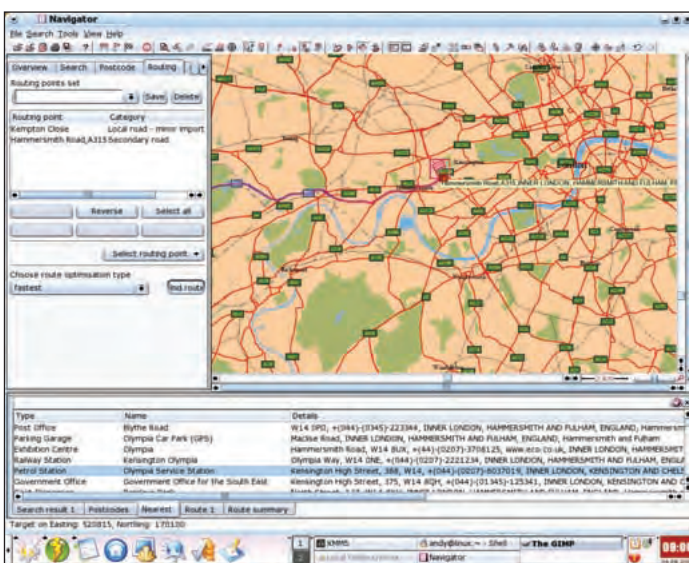
The product itself features over four million kilometres of road, 2.5 million streets, 400,000 sites of interest (including petrol station, cash points, supermarkets etc) and access to 1.7 million postcodes. Furthermore, users can add additional points of interest using pushpins similar to those found in Microsoft's *Autoroute* product.

Strap in and go

In use, the interface can look a little overwhelming with almost too much info available, and no part of the screen is

GPS

For about £100 you can bag yourself a GPS device which, when plugged into your laptop, provides verbal directions as you navigate your pre-prepared route. Unfortunately we were unable to test this facility, but the combination of decent software (which this is) and a real GPS system can make even complex navigation a breeze. The advantages of knowing not just *where* you should be going, but also *where you are* cannot be overstated.



The interface looks a bit busy to start with but you quickly get used to it.

spared: even the title bar is used to display the longitude and latitude of the main map's centre! Docked search and routing tabs on the bottom of the screen add yet more layers of data, but once you grasp the basic working method you can grab whatever you need fairly quickly. The online manual and tooltips help a great deal here, and if you do get overwhelmed by visuals, the UI is very configurable so you could, for instance, remove the 'Routing' tab, but still keep access just a right-click menu away.

The routing system itself is usable, with an unlimited number of way points – ideal for freight controllers – definable preferences for motorway, major roads etc and a cost calculator which takes the various consumption figures for your vehicle and the current cost of fuel to give you a pretty good estimate of how much you can expect to spend. There are a couple of things missing; the ability to choose specific areas to avoid, such as town centres, and automatically sort a list of waypoints into the most efficient route. It is possible to fudge the former using waypoints, but it would be nice, for instance, to just be able to highlight an area, right-click and select Avoid.

The routes provided by *Navigator* have been satisfactory, displaying none of the idiosyncrasies – yet – that seem

to bug most systems. The big problem with the European version of the software is that each national map is discrete – you can't program in a route from Bath to Bruges as this would span maps. So the trick is to know, in this instance, UK/France and France/Belgium border crossing points. For casual (or national) use, this isn't a problem, but for more serious routers, it sort of defeats the object of the software unless you're familiar with all of Europe's borders.

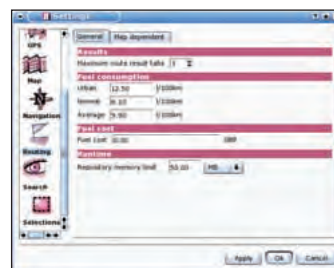
One of the selling points of the software is the sheer amount of information contained within the maps, and being able to highlight all the petrol stations, hotels or cashpoints within a certain area makes planning journeys and overnight stays a lot easier.

Conclusion

If you need route-planning software for your Linux box, there is no competition. *Navigator* is sufficiently better than most of the route planning websites to justify its cost: it may struggle to compete in the topicality: postcodes change, businesses come and go and new roads are built, so a fairly rapid upgrade cycle would be desirable. Perhaps this is one area where an annual subscription makes sense? The town plans and local info are all very useful and the cost calculator is



The maps contain a lot of useful extra information besides roads.



The fuel calculator is a great feature that'll be appreciated by both home and business users.

a nice touch, as is the dedication to covering the whole of Europe (and your hard disk will give thanks for the load-on-demand maps). Less satisfactory is the island nature of the European maps, and I'd guess an integrated version is in the pipeline. Still, if you're just interested in the UK iteration, this is a great buy.

Performance on my average specced machine was tolerable, but given the amount of data being shifted around a reasonable machine (1GHz plus) would be sensible, and if you want to install the whole of Europe, 5GB of disk space will just about manage it. **LXF**

VERDICT

Features	8/10
Performance	6/10
Ease of use	7/10
Value for money	8/10

The UK version is well worth the money, but extensive Euro-travellers would be well advised to wait for a more integrated product.

LINUX FORMAT RATING
 /////////////// 7/10

COLLABORATION SUITE

OpenGroupware.org

With *OOo* taking off like never before, many were interested to see OpenGroupware.org launched, an Open Source collaboration system. Paul Hudson looks it over, and discovers that all is not as it seems...

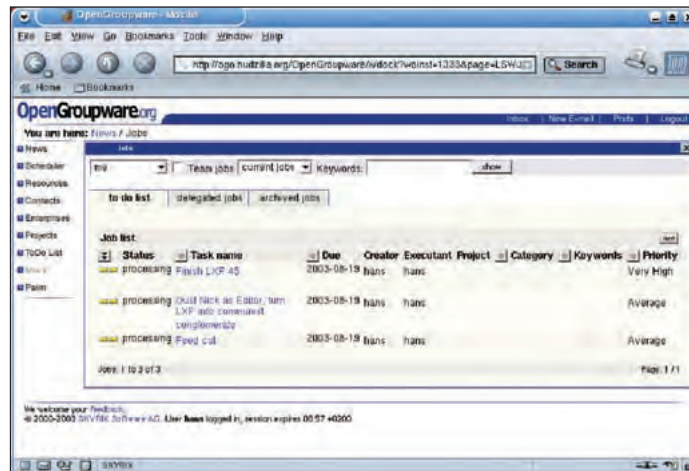
BUYER INFO

First release of a collaboration suite attempting to rival *MS Exchange's* dominance. See also *Kroupware*.

- **DEVELOPER** OpenGroupware.org
- **LICENCE** GPL/LGPL
- **WEB** <http://opengroupware.org>

We all know that *OpenOffice.org* is a great product – a little slow, sure, but, it has become the industry-standard word processor for Linux, and with good reason. But, there has long been a lack of an Open Source groupware system that would take the place of Microsoft's *Exchange* product. The problem always lies in the fact that, to be frank, *Exchange* is very, very good, and it puts the barrier for entry for new products very high.

Two products have come into the limelight: *Kroupware*, a German-government sponsored KDE suite, and *OpenGroupware.org*, a product seeking to capitalise on the success of *OpenOffice.org* by reusing the name style but otherwise unrelated. "What," I hear you say, "*OpenGroupware.org* and *OpenOffice.org* aren't related?" That's correct – they have nothing to do with each other, but the names have made people automatically give the *OGO* project more respect because of *OOo*. Groupware is the term used for a



To say "the interface needs work" is an elephantine understatement...

product that enables people to work together. The name usually refers to a package that includes email handling, calendaring, and To-do lists, although various other pieces of functionality are often added on top.

Grokking groupware

OpenGroupware.org used to be known as *SKYRIX 4.1*, a proprietary groupware system. However large parts of it has now been open-sourced (although it's written in Objective C), excluding load balancing, support for *Oracle* and *Sybase*, and *Outlook* support – that all comes with the new iteration of *SKYRIX*, a proprietary groupware system based on *OGO*. If you were thinking "Great, I

can get all the *Outlook* users in my company to connect to an open-source server", I'm sorry to disappoint.

Like *Kroupware*, *OGO* makes use of several off-the-shelf components, such as the *Cyrus IMAP* server, but unlike *Kroupware*, *OGO* has not made its own Linux client code available yet. While the *Kroupware* project is integrating its code directly into *KMail* and *KOrganizer*, which should mean full compatibility for users on KDE 3.2, *OGO* is yet to proceed along those lines, although it is likely in the future.

The web interface

While there is some limited support for existing clients, the only two that are able to make use of all of *OGO's* functionality are the *ZideLook* plugin for *Microsoft Outlook*, not free, and the *Ximian Connector* for *Exchange*, also not free. So, in order to be able to use all the functionality offered by *OGO*, you're forced right now to use their web interface. Using its web interface, you could surmise that lots of work has gone into making *OGO* look as malformed and unusable as possible – or at least so we hope, as we'd like to believe this level of user unfriendliness could not be achieved by accident!

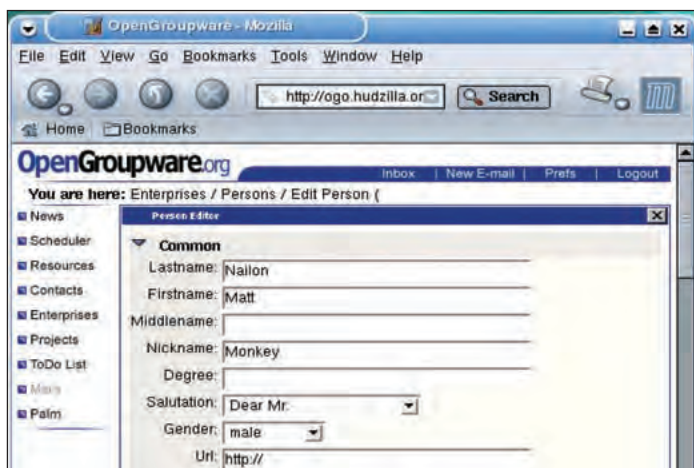
Many of the screens, for example, are littered with tiny little icons that apparently represent something, meaning you either have to hover over

each of them for a few seconds to try to figure out what they do, or you just click them and hope for the best. The most important functionality, such as 'Create new item' (whether that be an email, a task item, etc), is usually a very small 'New' button in one corner. The screenshots you see on this page are indeed the two most simple pages in the system, and even then it took us a while to make them fairly uncluttered.

Conclusion

What amazes us most is that this is apparently equivalent to version 4 of *SKYRIX* – one wonders how painful the early releases must have been! However, this is only the first release of *OGO*, and on that basis it works fairly well. The design is awful, yes, the documentation is non-existent, true, and the only real claim to being good that this software has is that its name sounds like the name of a good product. But, it's only the first release, and to be frank, *OOo* wasn't all that hot in its first release either.

Overall, this is a piece of software that was released prematurely; while it might be worth considering in six months, you'd do best to ignore its existence for the time being. Once it reaches some level of maturity, this may indeed become popular; at the very least, it will provide a running mate for *Kroupware*. Here at *LXF Towers*, we naturally want all Open Source projects to succeed, however this one doesn't really feel *entirely* Open Source as yet, and should be approached with caution. [LXF](#)



The Add New Contact dialog is sadly the least cluttered screen we found.

VERDICT

Features	6/10
Performance	6/10
Ease of use	1/10
Documentation	2/10

OpenGroupware.org is a sheep in sheep's clothing. Use *Kroupware* instead.

LINUX FORMAT RATING

3/10

Reviews**ZendStudio3**

PHP EDITOR

Zend Studio 3

EXCLUSIVE! It's the leading PHP development environment, so what has Zend changed to justify the version hike? **Paul Hudson** is bowled over as he puts the new release through its paces!

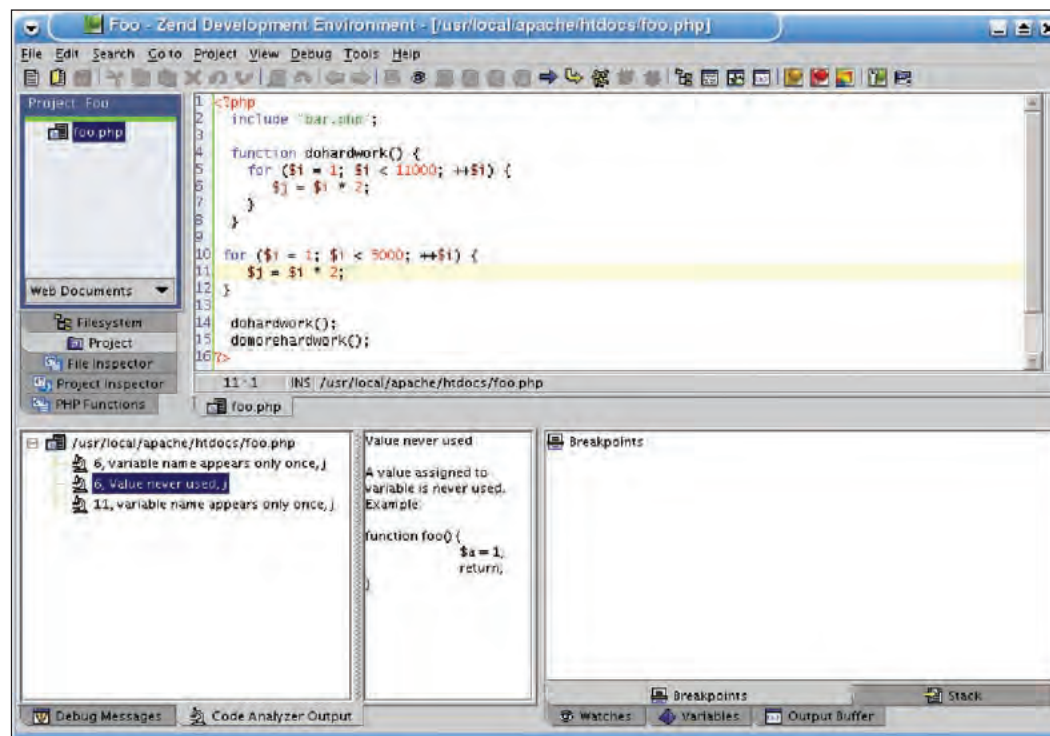
BUYER INFO

Top of the line PHP IDE with integrated debugging and profiling. Also consider, *PHPed*, *phpmole*, and *Quanta*.

- **DEVELOPER** Zend
- **PRICE** Starting at £195
- **WEB** www.zend.com

Zend Studio has been the most advanced PHP editor ever since its release – to say that it's ahead of its competition is an understatement! As such, we're always eager to see what new features have been added to make it worth upgrading. Version 3, we think, is the biggest upgrade yet – we're certain this will revolutionise the way people write PHP.

But we're getting ahead of ourselves – let's first examine the pedigree of *Zend Studio*. From relatively basic beginnings shortly after the release of PHP 4, the *Zend Studio* rose to prominence in its 2.x series – 2.5 scored 8/10 in *LXF31*, and 2.6 scored 9/10 in our review in *LXF40*. The developers have continually tweaked their existing software to enhance existing features, and occasionally adding new ingredients to the mix. For example, 2.6 saw the debut of complete CVS integration, cross-file searching, and also conditional breakpoints – excellent



The top half of the interface hasn't changed much since the previous version, but the real magic is all in the lower half – check out the Code Analyzer in the bottom-left corner.

features in themselves, but more *evolutionary* than *revolutionary*.

Zend Studio 3 has new ideas that will truly make a difference to developers – there are tweaks and fixes too, yes, but the features that will really knock you off your seat are fresh, new ideas that aren't

likely to see the light of day in other PHP editors for a long time yet.

The best addition in our opinion is built-in profiling for scripts. If you've only ever programmed using PHP, this will be new to you, but most other languages have had it for a long time.

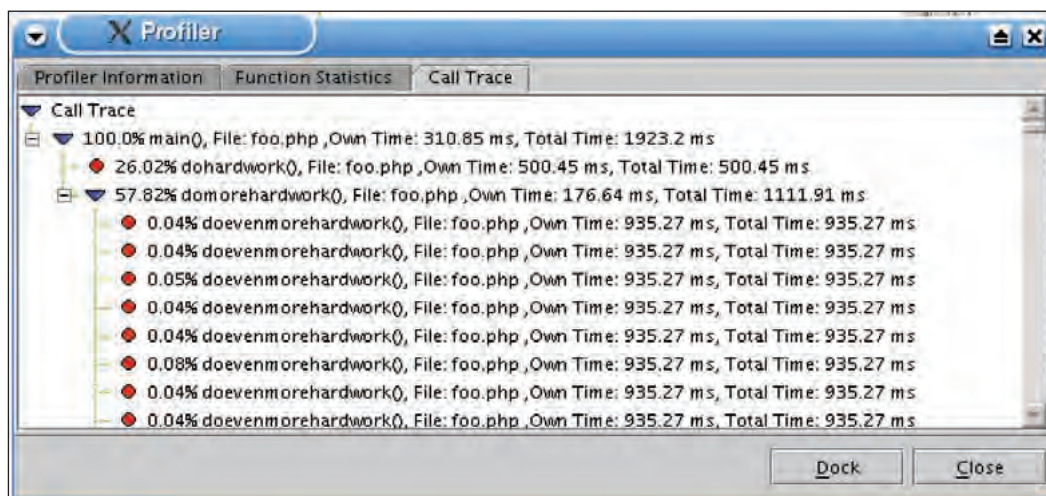
Code profiling is where *Zend Studio 3* will run your script for you, then report back how long each part of the script took to execute.

Also new to *Studio 3* is the *Zend Code Analyzer*, which is a design-time code scanner that attempts to spot errors in your code by examining it directly. If you've used *lint* for C, you'll understand how useful this can be.

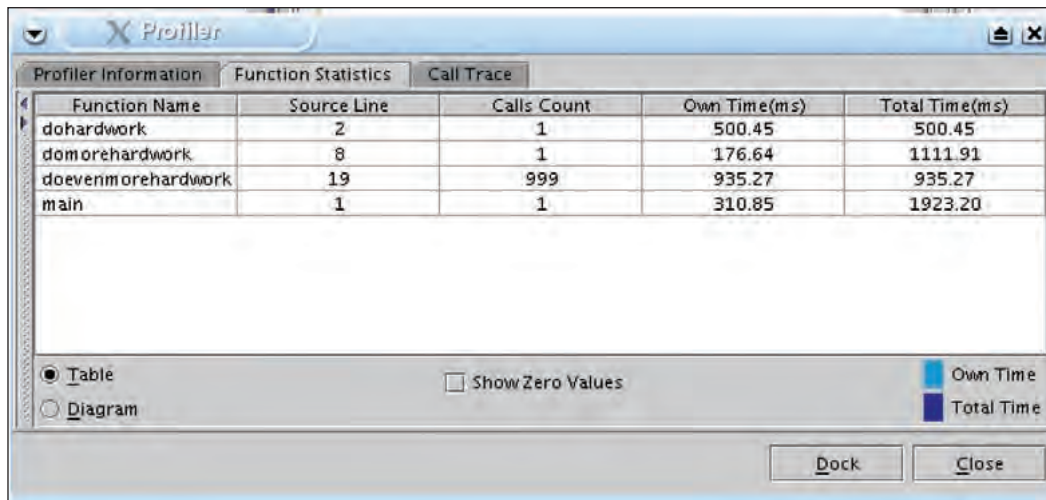
Finally, along with the usual complement of bug fixes, the *Zend Studio 3* team has plugged in full support for PHP 5 syntax, which means early PHP 5 adopters can get this extra insight simply by upgrading.

Profiling your code

If you've ever had to write very complicated scripts (over 5,000 lines or so), you'll know how hard it is to achieve maximum performance. The old rule of thumb that 80% of the work is done in 20% of the code is still true, which means that attempts to optimise code based upon statements



Real hardcore programmers can make use of the new Call Trace feature to track individual function calls.

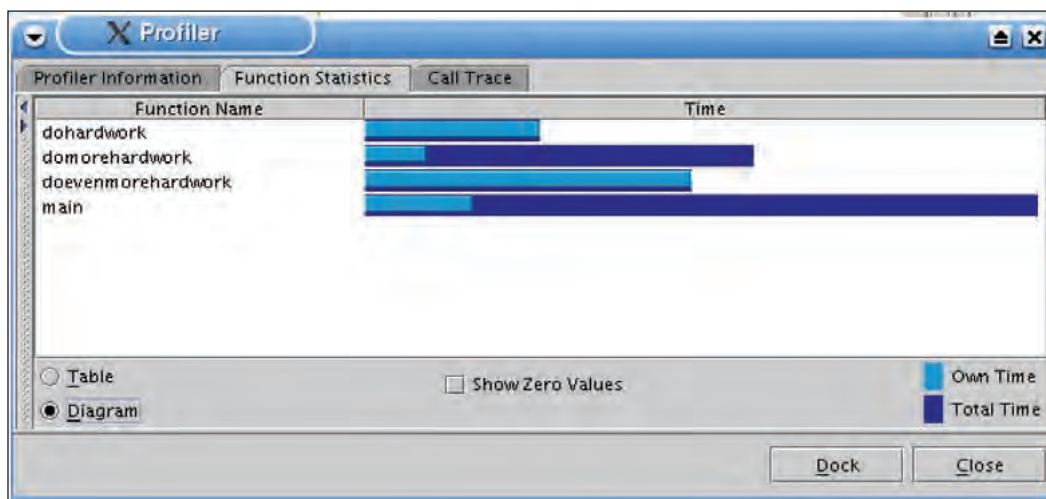


Function Name	Source Line	Calls Count	Own Time(ms)	Total Time(ms)
dohardwork	2	1	500.45	500.45
domorehardwork	8	1	176.64	1111.91
doevenmorehardwork	19	999	935.27	935.27
main	1	1	310.85	1923.20

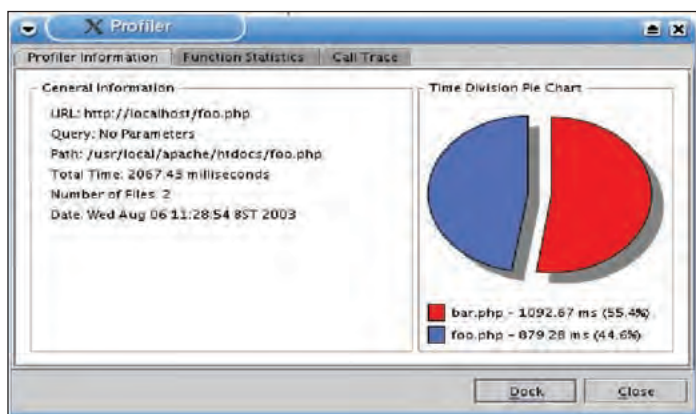
☒ Table
 ☐ Show Zero Values
 ☐ Own Time
 ☒ Total Time

Dock Close

The Function Statistics tab shows calltime information down to the millisecond.



Alternatively, you can view the same information in an easily comparable graph form.



For more complicated scripts, the Time Division Pie Chart breaks down how much time was spent in each file that makes up your script.

like "I bet that code must run slow..." are almost certainly a waste of time.

To counter this, more mature languages such as C and Java have profiling tools that literally measure how long your program spent executing various procedures so that you can spot where the slow points are. With *Studio 3*, Zend has finally brought this

technique to PHP, and what's more is that they've made it as easy as just selecting an option from the menu.

After running the script, the ZDE returns with a full breakdown of how the script ran – how much time was spent in each file that made up the finished script, how much time each function took (including child

functions), as well as an excellent call trace tracking that tells you where each function was called from.

Code analysis

The static code analyser tool runs a series of checks on your code to attempt to detect possible flaws. While the analysis isn't perfect, it does work remarkably well – it reports back variables that are unused, variables that are only used once, and a lot more.

Being a *static* code analysis, the code isn't actually run and so isn't foolproof. For example, using variable variables won't be detected properly, so you may receive spurious warnings from time to time. However, for the most part the system is eerily accurate, and should make things a great deal easier for everyone.

PHP 5 support

With PHP 5 just over the horizon, Zend has taken the lead and integrated full PHP 5 support into

Studio 3. While PHP 5 is not yet released, and won't be released for a few months just yet, the new syntax has already been finalised, meaning that the support included with *Studio 3* should be 100% correct when the final PHP 5 is released. Of course, if the language is tweaked slightly between now and release, you can bet that Zend will follow up with a patch to *Zend Studio*, keeping it in line.

As Zeev Suraski, CTO of Zend, told us exclusively last month, "While you can edit PHP 5 code with *Studio 2.6*, it's not going to give you any insight into PHP 5-specific code". We tested out various parts of the new syntax and were very pleased with the results – it all works perfectly.

While the need for PHP 5 isn't really going to be in big demand until at least February 2004, it's a smart move from Zend to support it now so that early adopters have yet another reason to use *Zend Studio*.

Conclusion

This is a "knock me down with a feather" release – the code profiling is a feature that we here at LXF Towers *already* cannot live without, and the code analysis tool, while fairly immature just now, is set to really push the development envelope forward. The new features in this release are so substantial that we haven't had space to briefly go over some of the features that *Zend Studio* has long done well – output buffer previewing, code step through, multi-file search and replace, and more.

If you use and love *Studio 2.6*, but don't think you need the code profiling technology, the code analysis, and the PHP 5 support, as well as the bug fixes and much-improved speed, then clearly you have no need for the upgrade. However, the other 99.99% of you should upgrade without delay – it will keep you smiling for weeks. **LXF**

VERDICT

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

Leading the pack, powering professional PHP development, and slicing through even the hardest tasks – *Zend Studio* has hit warp speed with v3!

LINUX FORMAT RATING
 9/10

Practical Cryptography

Without a secure computer system, your business won't make money, says **Paul Hudson**.

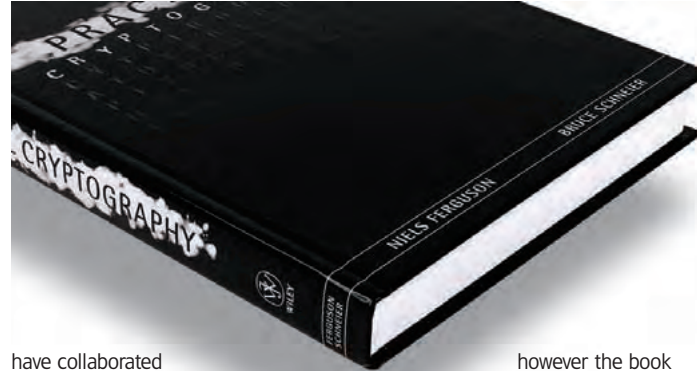


BUYER INFO

■ **AUTHOR** Niels Ferguson and Bruce Schneier
 ■ **PUBLISHER** Wiley
 ■ **ISBN** 0-471-22894-X
 ■ **PRICE** £34.95 **PAGES** 410

Dozens of books have been released in past years detailing cryptography concerns and solutions, but precious few if any have actually discussed practical applications of the knowledge presented – it's generally lots and lots of theory, with readers being left to their own devices for everything else. This is a shame, as it's such an important topic. Without your computer system being secure, your business or personal information could become known to malicious parties, causing you embarrassment and even complete loss of livelihood.

The sheer fact that Schneier and Ferguson, two names that make any security expert sit up and take notice,



have collaborated on one book is enough alone to make this book shoot up into bestseller lists the world over, and that's hardly surprising – Schneier designed the infamous Blowfish encryption algorithm, and both Schneier and Ferguson worked on the Twofish algorithm that was one of the leading contenders for AES.

The book starts with a fairly chatty approach, with titles such as *Cryptography is not the solution*, *Cryptography is Very Difficult*, and *Cryptography is the Easy Part*,

however the book then takes a dive into the technical parts of encryption for several chapters, and readers should come out from that a fully indoctrinated cryptographer – it really does go into a great deal of depth about how various algorithms work.

Just shy of half-way through the book, the content surfaces back at the less technical level, and that's where it stays for the remainder – high-level topics are covered here, such as deploying PKI servers, generating randomness, and how to store secret

information. The authors have also sprinkled various comments and hints throughout the book regarding problems they encounter frequently, and this makes it easy to directly apply knowledge you learn here to your own systems or communications policy.

Overall, the book is quite heavy on the mathematics and technical front for the very casual reader, but if you have any interest in cryptography that you'd like to develop further, this is likely to be just what you were looking for. The authors are masters of their art, and it shows in their writing – even relative novices can relax and enjoy being taught by leaders of the cryptography field.

VERDICT

An unbeatable team of well-respected authors, skillfully and effectively covering an important topic without glossing over any areas.

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

The Complete FreeBSD

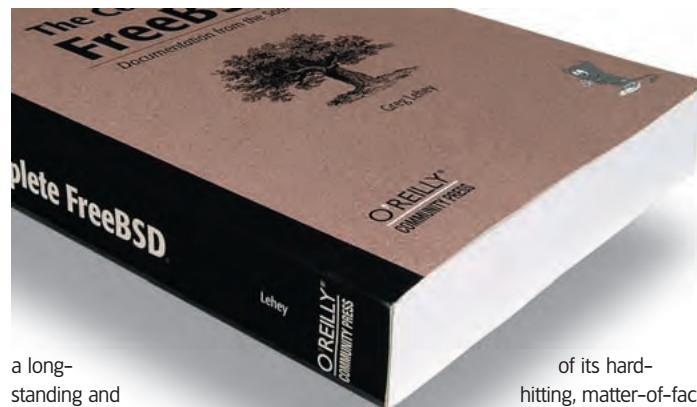
How do you go from newbie to master? **Paul Hudson** reviews 4th Ed of FreeBSD's classic text.

BUYER INFO

■ **AUTHOR** Greg Lehey
 ■ **PUBLISHER** O'Reilly
 ■ **ISBN** 0-596-00516-4
 ■ **PRICE** £31.95 **PAGES** 677

BSD is one of the foundations of the modern Unix world, having heavily influenced Windows, Mac OS X, and even Linux itself. Of all the BSDs, FreeBSD has long been the most popular, mainly due its focus on Intel CPUs and consumer software. Since the release of FreeBSD 5.x, FreeBSD usage has surged forward thanks to the large collection of new features that are available.

Linux users often are often pleasantly surprised to find excellent Linux binary compatibility in FreeBSD, and also popular tools and programs such as KDE available to them just as if they were using Linux. Naturally, such



a long-standing and mature OS requires an equally long-standing and mature book to cover its workings, and FreeBSD fans are blessed to have *The Complete FreeBSD*, by Greg Lehey. Now in its fourth edition, this tome is fully updated for FreeBSD 5, however there hasn't been an increase in size, thankfully – in fact, Lehey managed to cut the overall size of the book by around 100 pages without losing any

of its hard-hitting, matter-of-fact content.

The best word to describe this book is probably "jumpstart" – while it does start off at a fairly low level, it's a no-holds-barred tour of the world of FreeBSD, with every appropriate topic explained along the way in enough concentrated detail to satisfy even the most demanding of readers. For the most part, the book reads like annotated man pages, which is unsurprising

considering its heritage. On the whole, the format works well – man pages have always been terse, so Lehey's additional text explaining *why* things work as they do, and also the considerable lengths he has gone to to make sure the topics are ordered logically for ease of learning, make the book a fairly entertaining read.

Admittedly, you're as likely to read this book cover-to-cover as you are to read a dictionary in the same way, but that's no bad thing – as a reference, this book leads the way for both dyed-in-the-wool BSD zealots and the curious alike, and as an introduction and tutorial you certainly won't be disappointed.

VERDICT

At its fourth edition and still improving, this is documentation straight from the source.

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

Java Extreme Programming Cookbook

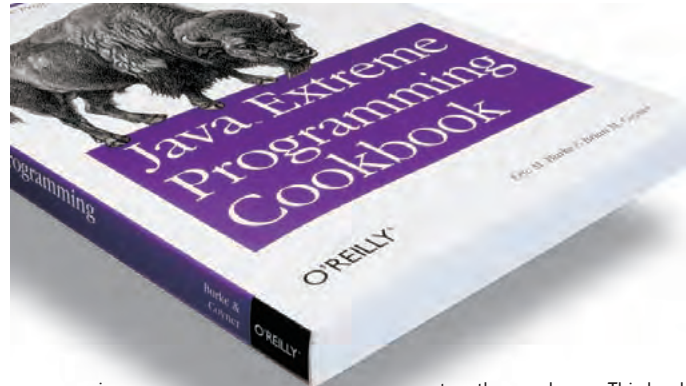
Richard Cobbett wants to code on the wild side, but finds the programming less extreme than promised.

BUYER INFO

■ **AUTHOR** Eric M. Burke and Brian M. Coyner
 ■ **PUBLISHER** O'Reilly
 ■ **ISBN** 0-596-00387-0
 ■ **PRICE** £24.95 **PAGES** 288

This is a book about XP. Not, we hasten to add, a certain operating system that you can't easily buy a new laptop without it encumbering its hard drive these days, but *Extreme Programming* – a buzzword that the majority of us have so far been fortunate enough to remain ignorant about. Ultimately, it's not even a Java book – instead, this cookbook focuses on the more esoteric sides of the Open Source programming community, such as *Ant*, *JUnit* and *Cactus*, to add spice to your creations, at least as far as they relate to XP, anyway.

Which begs a fairly obvious question: what is XP? In short, it's a programming process rather than a



programming method, involving rather less surfing down mountains on keyboards than you might originally expect. It's all about testing – but testing before you write the code rather than afterwards – and working collectively. Instead of having your own pet piece of code, every programmer on the project dips a hand into everything, constantly working with a partner (and preferably a dual screen, although that could easily become dual screen) to put

together packages. This book isn't however a guide to how to do it – only the first chapter – just how to do it using Java and basic packages.

Once the book gets properly underway, it's helpful to have a decent understanding of Java and its available addons, if not necessarily the ones under discussion. Each receives a very quick summary (sometimes maybe too brief), and the program listings are short and easy to read, but the book isn't intended to focus on vanilla usage – HTTPUnit for instance, includes sections

on simulating hyperlinks and writing testable HTML over its more immediate parsing and connectivity functions, and this isn't one of the dedicated testing tools on offer. A significant number of internal cross-references successfully link corresponding recipes and tools improves the usefulness of the book beyond a simple primer.

Should your idea of 'recipes' be more general help and advice on the covered programs, look elsewhere, but the wide range of programs on offer more than makes up for this when focusing on testing. Ultimately, the decision is going to be made based on your view of XP rather than the individual programs within.

VERDICT

An excellent way of boosting your testing methodology, but only if you've already adopted the practices at its heart.

LINUX FORMAT RATING

7/10

Building Embedded Linux Systems

Nick Veitch examines a text on what is arguably Linux's fastest growing area.

BUYER INFO

■ **AUTHOR** Karim Yaghmour
 ■ **PUBLISHER** O'Reilly
 ■ **ISBN** 0-596-0222-X
 ■ **PRICE** £31.95 **PAGES** 394

While pretty much everything in this book is described from the ground up, it will probably help if you have some experience in embedded system design, or at least a sound academic understanding of the processes. The book really describes the set of circumstances in which Linux is ideal for embedded applications, and the considerations you will have to make.

The first section kicks off with a good description of what the book covers and how it is organised. The architectures referenced throughout are mentioned here, which include x86 (unsurprisingly), PPC (used in a



fair few embedded hardware boards, ARM (as made popular by many PDAs including the iPAQ and SuperH (used in, amongst other devices, the Sega Dreamcast).

The following chapters cover all the major decision making areas of designing an embedded system – which development tools to use, how to build a kernel image (and what to put in it), filesystems, networking and

even useful methods for debugging your embedded app. Although a lot is being covered, some areas of the book are left feeling a little rushed. The chapter on setting up a root filesystem for example, does touch on the pros and cons of RAM disks, JFFS, CRAMFS and more, but never seems to go into enough thorough detail on them. JFFS2 gets a meagre page for example. As this is often one of the more critical design

decisions, it probably could have done with a bit more space.

The book rounds off with some sample worksheets, which may seem like a bit of a space-filler, but which are actually pretty useful reminders if you are planning on implementing such a system for real.

While much of the information that is contained here can be found for free elsewhere (and indeed, you may have to consult some of these sources to get enough detail on some aspects of building an embedded system), the book is at least easy-to-read and explains most things unambiguously.

VERDICT

Not a must-buy, but a nice to have.

LINUX FORMAT RATING

7/10

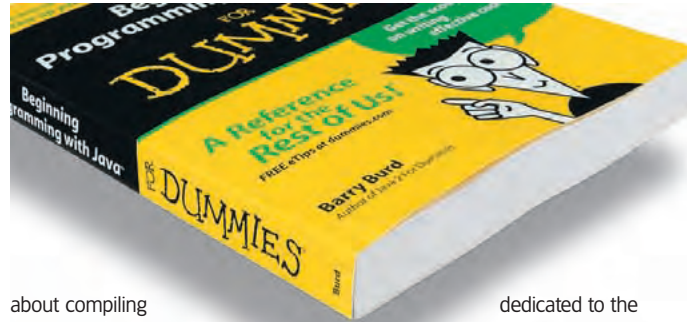
Beginning Programming with Java for Dummies

Richard Cobbett brews up a fresh cup and sits down with the latest idiot-proof guide...

BUYER INFO

■ **AUTHOR** Barry Burd
 ■ **PUBLISHER** Dummies
 ■ **ISBN** 0-7645-2646-4
 ■ **PRICE** £17.50 **PAGES** 398

Dummies books usually make for good reads, but are hard to review fairly. After all, if you've read one, you more-or-less know what you're getting as soon as you take the next off the shelf – a comprehensive, easy to read guide to a new subject. This is absolutely no different. In the course of its pages, this book doesn't simply teach elementary Java, but programming from scratch – from first 'Hello World' to how to use the labyrinthine API documentation which you need to go further. There is a slight leaning towards Windows, but not to any serious degree, and we were pleased to see a short section on how you go



about compiling and running your applications on both Linux and Mac OS. This isn't a sign of a radical sea-change towards Linux for this range, more that one of Java's strengths is its cross-compatibility.

The lessons themselves follow the standard pattern, kicking off with input and output, then running through variables, numbers, decision making and file input/output, although GUI applications are noticeable by their absence. There are only seven pages

dedicated to the subject, and one of those is a discouraging warning about how hard it is. While it's hard to argue, and a new user is certainly better off focusing on the command line, a few more pages on this subject would not have gone amiss – and preferably without muddying the water by venturing some way off-topic about the war between *Swing* and *AWT*.

The shorter applications follow the normal Dummies pattern of telling a short story, and then explaining the code

to do it – making them much easier to follow than the straight codedumps favoured by those other Reference guides with a capital R. For/Next loops are explained with the old advice about chewing your food x number of times without swallowing, before moving onto more conventional loops with the assistance of a ten-roomed hotel metaphor. Both the code itself and its functionality is very easy to read, doing a splendid job of teaching both Java itself and the basic programming principles that you need to take things further. **LXF**

VERDICT

4/10 for the long-winded title, but an excellent book that will really get users started if you've any interest in learning the Java programming language.

LINUX FORMAT RATING

8/10

Who will keep you on the Information Highway?

To date **Hostway** services the hosting needs of over 100,000 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.



30 day money-back guarantee...

	KEY FEATURES	DISK SPACE	MONTHLY TRAFFIC	SHELL ACCESS	PHP4.3.2 SCRIPTING	MYSQL3.23 DATA AREA	ADDITIONAL FEATURES	ACTIVE CGI LIBRARY	CRONJOBS	XML SUPPORT	PYTHON	FRONTPAGE EXTENSIONS	POP3/SMTP	EMAIL ACCOUNTS	UNLIMITED ALIASES	EMAIL FORWARDING	MEDIA FEATURES	HTTP: STREAMING	REAL MEDIA PROFESSIONAL	MONTHLY FEE*	YEARLY FEE*	SET UP FEE*
LINUX SILVER	100	1	Yes	Yes	0	Yes	Yes	Yes	Yes	Yes	Yes	5	Yes	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	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	Yes	No	£19.95	£179.95	Free		
LINUX PLATINUM	500	10	Yes	Yes	2	Yes	Yes	Yes	Yes	Yes	Yes	50	Yes	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	Yes	£89.95	£809.95	Free		

* excluding VAT & subject to change.

Roundup

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

OUR SELECTION AT A GLANCE

- Apache
- Boa
- Caudium
- Roxen Webserver
- thttpd
- Tux/kHTTPd
- Zeus Web Server

Web servers

David Coulson takes a look at the various Web servers available for use with Linux.

Having a website for friends or customer to view to is almost as commonplace as an email address or even a traditional storefront nowadays. It is the first impression people see of you as an individual or as a company, not to mention the fact that it is a vital means of communication between yourself and those who travel through cyberspace to visit. Most of us don't give any thought to how the site works, how the content gets to our system, and simply put the URL into our web browser and hope for the best.

Websites, as indeed anything that is handled via HTTP, is provided by a service known as a 'Web server.' Without getting too much into the

details of HTTP, a client will connect to the Web server, ask for a specific document, and the Web server will send it to the client. For a single page, without images, there is only one request performed by the client, most frequently a web browser, but for each image located on a site, a separate request must be performed.

Within each individual request, the server may simply pull a file from disk and pass it through to the browser. However, on more complex sites, such as those that involve databases or dynamic content, there is generally a significant amount of processing being done by the web server to execute scripts or binaries for each request to a page containing such content.

With a Linux system, it is possible to run a Web server on the most basic of hardware, depending upon traffic. Depending upon service provided by the ISP, it is perfectly acceptable for a low-end system on a cable or DSL line to provide HTTP services for a number of sites. Of course, this system can be scaled all the way up to massive services located at backbone providers, although the concept is almost identical.

We're going to look at a selection of popular Web servers – and some not so well-known – which we can use to run our websites, from simple pages about our dog's litter of puppies to real profit-generating enterprise like web stores and news services.

'With Linux systems, it is possible to run a Web server on the most basic of hardware... to provide HTTP services for a number of sites.'

Apache

The burgeoning Linux domination of the server market continues apace...

■ URL <http://httpd.apache.org/> ■ COST Free (Apache License)

Apache is easily the most well-known Web server available which runs on Unix, although it will run on almost any computing platform out there, including Microsoft Windows. As of June 2003, *Apache* was used by 63% of the 42 million websites tested by netcraft.com, more than double those using *Microsoft IIS*. With such a huge following and commercial interest, *Apache* is a tried-and-tested platform which can scale to handle many hundreds of requests per second, but also is very efficient when called on to provide service for smaller sites which may only get a few dozen hits per day.

There are currently two stable releases of *Apache* available for use. The 1.3.x release of *Apache* is based on earlier code, although it is still used by many websites and hosting providers due to its stability. The 2.0.x release has a complete core rewrite, allowing it to scale much further and adds in many capabilities which are unavailable in 1.3.x due to the architectural limitations of the core httpd code. Both trees continue to be maintained by the Apache Foundation, so those of us still running 1.3.x get all the security fixes and updates, although there are no planned feature additions to this tree.

Apache is a totally modular Web server, so it is very easy for a third-party to provide an additional capability to the server without having to go in and patch code. Unfortunately, modules must be written for a specific version of *Apache*, which is one reason why the uptake of 2.0.x has been slow while *Apache* modules are ported from the earlier version. Popular *Apache* modules include *mod_php4* and *mod_perl*, both of which allow code to be compiled and executed by the Web server process without forking another interpreter to run the code.

On the configuration front, everything is handled through the text based *httpd.conf* file. No fancy web interfaces here. Writing an *Apache* configuration from scratch isn't much fun, but adapting the default is

```

</VirtualHost>

<IfDefine SSL>
AddType application/x-x509-ca-cert .crt
AddType application/x-pkcs7-crl .crl

<IfModule mod_ssl.c>
SSLPassPhraseDialog builtin
SSLSessionCache dbm:/usr/local/apache/logs/ssl_scache
SSLSessionCacheTimeout 300
SSLMutex file:/usr/local/apache/logs/ssl_mutex
SSLRandomSeed startup builtin
SSLRandomSeed connect builtin
SSLLog /usr/local/apache/logs/ssl_engine_log
SSLLogLevel info
</IfModule>

NameVirtualHost 10.1.7.5:443
<VirtualHost 10.1.7.5:443>

<Directory "/home/*/personal_web">
    Day On
  
```

Apache is the most popular Web server available, and supports a wide range of third-party modules.

particularly straightforward. There is a huge array of documentation available at <http://apache.org>, plus as *Apache* has been around more-or-less forever, there is a cornucopia of example configuration file segments available on Usenet and various websites. Usually, it's as easy as cutting and pasting the pieces of code that you want to use into the *httpd.conf* file and modifying the paths accordingly, so it's not overly complex to hack up a configuration manually. One can configure *Apache* through *Webmin*, although it does not include support for all of the capabilities which are offered by *Apache*, so if you're wanting to make your system do anything a little out of the ordinary, then the text editor is the only way to go.

As with all modern Web servers, *Apache* supports both IP and name-based virtual hosting 'out of the box', as well as SSL support in 2.0. For SSL support with *Apache 1.3.x*, one needs to install the *mod_ssl* module, which is available for free download from www.modssl.org. There are a number of other SSL modules for *Apache*, including *Stronghold* and

Apache-SSL, although *mod_ssl* was merged into the 2.0 tree. *Apache-SSL* is freely distributed, but *Stronghold* is an enterprise *Apache* variant, so costs money. Generally, someone wanting *Stronghold* will want some of the other enterprise features, so it's not a fair comparison to simply look at the SSL capabilities.

Certainly one of the most powerful features of *Apache* is *mod_rewrite*, which we looked at in *LinuxPro* issues 38 and 39. This allows URLs to be rewritten using regular expressions, so that complex locations can be accessed with simple URLs. Many servers have the ability to alias directories or files to a location, but *mod_rewrite* gives us complete control over how the browser can access our content. We're sure almost everyone who uses *mod_rewrite* couldn't live without it if they were forced to switch to another server.

On the scalability front, *Apache* can be tuned to handle large number of processes, allowing massive web sites to be served from a single box. For stability and to avoid memory leaks, processes can be killed after

serving a specific number of requests, so the server can run without any attention for quite a while.

If you ask anyone for their Web server of choice, most will say "Apache" – with its flexible feature set, range of third-party modules, along with the fact that it's Open Source and therefore easily modified, *Apache* is generally the top of everyone's list. With third-party commercial modules, even ASP and *Coldfusion* can be executed through *Apache*, making it a near-seamless replacement for sites running *Microsoft IIS*.

VERDICT

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

Apache is a rock-solid server, with a wide range of capabilities, but if you're not happy modifying text files by hand, it might be too much.

LINUX FORMAT RATING
 /////////////// 8/10

RoundupWebServers

Boa

For those instances where size *is* important...

■ **URL** www.boa.org/ ■ **COST** Free (GPL)

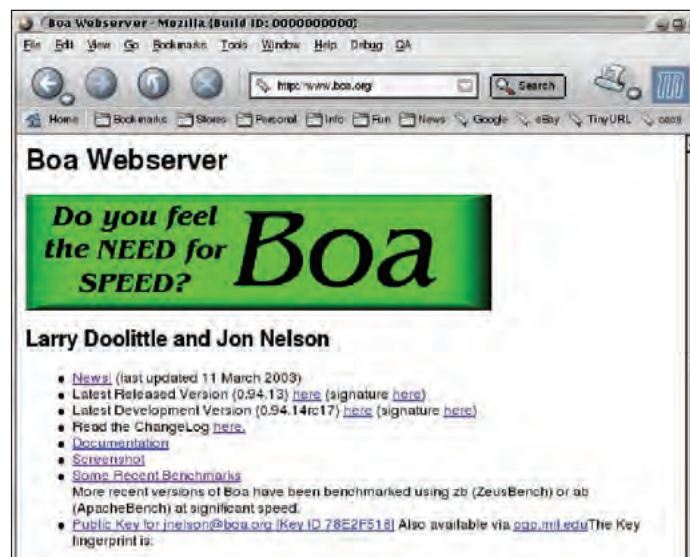
While many Web servers are required to serve huge sites to millions of users, there is also a sizeable market for very small and efficient servers for embedded systems or for raw web serving power. *Tux* is certainly a great example, however it does not handle CGI's and its capabilities are somewhat limited. If you simply want to operate a site with static content, then *Tux* would work fine, but *Boa* represents a great alternative to *Tux* for many systems, as it support CGI's, while avoiding the overhead of running a large httpd service, such as *Apache* or *Zeus*.

Compared to many Web servers, *Boa* can handle a massive number of requests per second. Even with a competitor such as *Apache*, if you're needing to handle large amounts of static content, *Boa* can certainly handle almost any traffic you throw at it. It runs as a single binary, so there is no overhead involved with running hundreds of processes. *Boa* also has

memory usage under 1700KB, making it a perfect web service for use with embedded systems, including routers, firewalls and portable devices.

The *Boa* configuration is very simple and looks particularly similar to that of *Apache*, although it is by no means identical. With support for aliases and user public_html directories, *Boa* makes for a featureful alternative to a big Web server. With the more current releases of *Boa*, virtual host support is also included, although there is no method for dynamic virtual host management.

We can handle dynamic content with *Boa* through the use of CGI's, although embedded capabilities, such as *mod_perl* and *mod_php4* are currently not supported. A number of individuals have produced patches for *Boa* to provide this functionality, but this certainly isn't in the mainstream *Boa* tree and probably will never be. Anything that is executable can be run



Many embedded systems, including routers and portable devices, use *Boa* for their web interface due to the small memory footprint.

as a CGI, from PHP and Perl, to C and a shell script.

Installing *Boa* is a very simple task and requires very little modification to get an 'out-of-the-box' *Boa* system up and running. For anyone just serving a simple site, or making use of a smaller system, *Boa* should be among their first choices of Web server.

VERDICT

Features	4/10
Ease of use	5/10
Documentation	9/10
Performance	9/10

Boa is fantastic for both static content and CGI's. However, its main rival httpd has a few more features.

LINUX FORMAT RATING

7/10

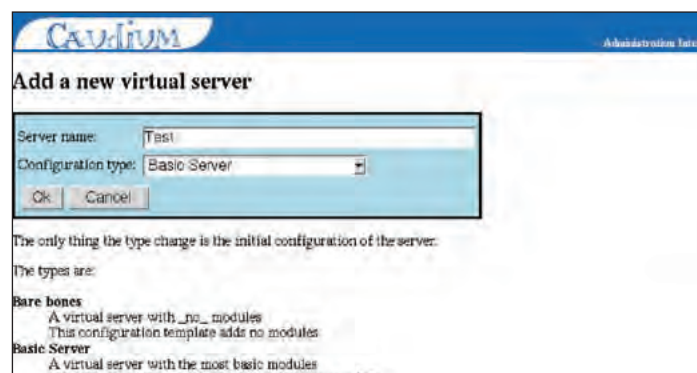
Caudium

Open fork of what used to be a closed program.

■ **URL** <http://caudium.net/> ■ **COST** Free (GPL)

Forking Open Source projects is generally not a positive occurrence, as we have seen with *Samba* in the past. Fortunately, the developers of *Caudium* have worked on their fork of the *Roxen*

Webserver 1.3, and succeeded in refining the code produced by the *Roxen* team. The original reason for the fork was due to incompatibilities between *Roxen 1.3* and *Roxen 2.0* back in 2000 caused by



Originally a branch of *Roxen Webserver 1.3*, *Caudium* has become a very powerful and capable Web server.

Roxen Webserver's inclusion in a commercial product. Since all the code was GPL, it was easier just to start a new development tree based on *Roxen 1.3*, rather than start again from scratch.

As far as the feature set is concerned, *Caudium* is very similar to *Roxen Webserver*, in that it has Pike and RXML support. Code from *Roxen* can be moved straight onto a *Caudium* system without any modification. This is a great choice for everyone running *Roxen 1.3*, which no doubt are few and far between now, although if you're already running *Roxen 2.0* then it will involve a rewrite of a decent chunk of RXML to get it running under *Caudium*.

Much like the Linux kernel, *Caudium* has both stable and development trees using odd and even numbers. One of the great features in their 1.3 development tree is the SNMP module, which allows the Web server to be queried by a SNMP client. If you've ever tried to parse out access logs with Perl, or mess around with pipes, then this is a great way to monitor a server. Coupled with a graph tool, such as *MRTG* or *Cacti*, and you can easily have funky looking

graphs of number of requests and bytes sent or received versus time.

Compared with *Apache* and *Zeus*, *Caudium* has similar advantages and disadvantages to *Roxen*, although *Caudium* is generally more accepted by the *Roxen* and *Pike* communities. Unfortunately, *Roxen* as a program has not always come across as being the most supportive of the community, particularly when moving from 1.3 to 2.0. *Caudium* attempts to solve many of these problems and at the same time works on many technical aspects where *Roxen* has been lacking.

VERDICT

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

Caudium makes life easy if web interfaces are your thing, just don't go mad with graphical text and graphs.

LINUX FORMAT RATING

7/10

Roxen Webserver

Free version of an originally proprietary platform.

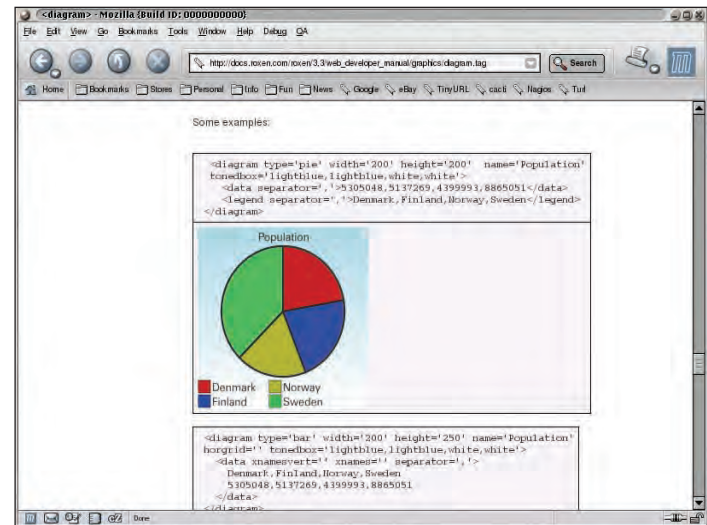
■ **URL** www.roxen.com/products/webserver/ ■ **COST** Free to download

Many will know *Roxen* as the highly graphical Web server, and while this is not exactly inaccurate, *Roxen* has a wide range of features beyond those provided by *Apache* and other comparable platforms. Originally built by Roxen Software as a commercial web platform, there has been an Open Source distribution of *Roxen Webserver* for over five years now, and many of those who started using *Roxen* 'back in the day' find it difficult to return to *Apache* afterwards.

Built around the object-oriented Pike language, *Roxen* has its own markup language known as RXML. Much like PHP with *Apache*, the server parses out the non-HTML from the pages before sending them to the server, although the format of RXML is based on SGML, rather than having a whole separate language embedded within the HTML. Of course, one can use Pike if needed and it is possible to use both RXML and Pike within the

same pages, if for no other difference than that the RXML is processed into Pike before execution. With a whole array of built-in functions, including automatic generation of images from anti-aliased text, on-the-fly image processing and comprehensive filesystem and database functions, one can build complex websites without a great deal of effort and knowledge of anything more than HTML and some high-level language understanding, using *Roxen*.

However, it is the administrative front-end where *Roxen* really shines. No ugly textfiles to edit – unless you really want to, of course! Everything is point-and-click, with a very clean and neat web-based tool. All of the virtual hosts are built from pre-built templates, although one can easily create customised templates for a specific purpose. Everything from a basic server, SSL server and CGI server are handled within these templates, so



With RXML and Pike, *Roxen* gives extensive capabilities to site developers.

as soon as the virtual host is added, it has instant capabilities.

With the image processing capabilities of the program, *Roxen* is certainly not suitable for those with low-end systems, but with the easy to use RXML tools, along with the flexibility of Pike for lower-level code and custom RXML functions, it is a great option for anyone needing something very easy to maintain.

VERDICT

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

Like *Caudium*, *Roxen* has a great web interface to configure the server. Just don't expect to migrate to another server without a lot of code rewriting!

LINUX FORMAT RATING
 7/10

thttpd

Run separate processes from the same config file.

■ **URL** www.acme.com/software/thttpd/ ■ **COST** Free

Some Web servers are designed to be flexible, others try to occupy a small memory footprint. *thttpd* manages to do both of these, while adding in a few features of its own which set it apart

from other Web servers. As with *Boa*, *thttpd* is a very basic server without any of the high-end capabilities found in *Apache*, *Zeus* and other similar services. A single process that listens on a single



The only HTTP server with throttling capabilities, *thttpd* is a fantastic solution for an image or download server.

IP address, with the ability to handle significantly more requests on very basic hardware than most other Web servers.

As with most servers, *thttpd* is configured via a very simple text file, which is the basic 'option=value' style. Everything can also be specified on the command line, so it's perfectly possible to run three or four *thttpd* processes from shell scripts without a separate configuration file for each one. Within each server, virtual hosts can be configured, although they occupy an intricate directory hierarchy, making it difficult to actually setup customised virtual hosting services. Of course, if one is building a *thttpd*-based hosting system from scratch, it's not a major limitation.

One of the major features in *thttpd*, and certainly something that sets it aside from many other servers, including *Apache*, is its throttling capabilities. For individual URLs, including wild cards, we can set a specific bytes/sec throughput ceiling, so if we are serving lots of images, or some users are pushing video or music files through the server, we can carefully manage how bandwidth is used. *Apache* may have

mod_throttle, but it does not work on a bytes/sec basis and will simply refuse more connections if the limit for the preset time period is hit.

thttpd compares to *Boa* very well. While it does not have many of the fancy features found in *Boa*, including the neat virtual hosting, it may prove a great alternative to both *Tux* and *Boa* if bandwidth is limited or otherwise restricted. Along with *chroot* capabilities that come as standard, as well as simple, yet effective, security features, *thttpd* is a great option for a static or CGI hosting server.

VERDICT

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

If throttling capabilities are what you need,

LINUX FORMAT RATING
 8/10

RoundupWebServers

Tux/kHTTPd

A Web server as a loadable kernel module.

■ URL <http://kernel.org/> ■ COST Free (GPL)

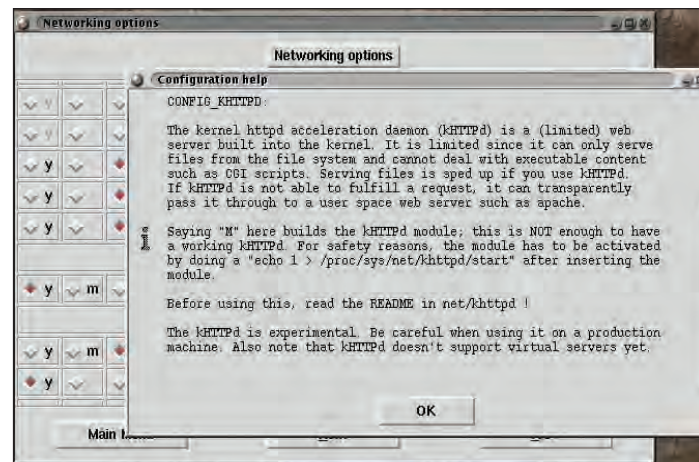
If an award existed for the most uniquely implemented Web server, *Tux* would certainly be close to the top of the list, or maybe even the outright winner. Unlike *Apache*, *Boa*, *Roxen* and all other Web servers that are available which run as a user-space process on the server, *Tux* is a loadable kernel module which provides HTTP services at the kernel level. Not only does this make it very quick, but it also significantly reduces the resource overhead required to run it, including memory consumption and scheduler time wasted through context switching.

Tux is only appropriate for static content serving, although it has a pass-through mode so that it can send the request over to another server, such as *Apache*. *Tux* itself won't handle any embedded code, including SSI, PHP or Perl, and it won't do anything exciting like CGI's and complex access

privileges. However, if you're serving lots of static content, such as images or a mirror of a FTP server, it is very quick and scales particularly well.

On the down side, an exploit against *Tux* could bring down the kernel on the system running it, although it has been in the kernel for a while and all major problems have been fixed fairly quickly. In our tests, *Tux* seemed perfectly stable under a 2.4 kernel, though its inclusion in the 2.6 kernel and how it will perform is obviously still open to question.

Configuring *Tux* is done via the `/proc` interface, so a very simple shell script is required to get *Tux* up and running and to tell it how to serve documents. By giving it specific dynamic URLs, we can have it pass the request for CGI's, PHP and Perl scripts through to another server. Most people will probably just use *Tux* for static



Kernel space is no longer the world of device drivers and filesystems, with *Tux* handling HTTP within the Linux kernel.

content and have dynamic pages handled by a completely separate server on another IP, but being able to forward the content through *Tux* is probably useful to someone.

We really like *Tux*, although its use is somewhat limited. It's fast, flexible, but it's in a totally different market compared to the other Web servers featured here, even *Boa*.

VERDICT

Features	2/10
Ease of use	4/10
Documentation	4/10
Performance	10/10

For serving massive amounts of static content, *Tux* would be a great choice. Unfortunately, it's a little more difficult to install than most other Web servers.

LINUX FORMAT RATING
 5/10

Zeus Web Server

Top performance – at a price.

■ URL www.zeus.com ■ PRICE From £1100

When it comes to Web servers, *Zeus* is certainly considered the best enterprise level server, and comes with an appropriate price for such a task.

Used by high-profile sites, including eBay and Red Nose Day, it's totally rock solid under load and scales extremely well.



For a top-of-the ladder Enterprise Web server, *Zeus* packs a list of popular sites as examples of its scalability.

Zeus has more features than you can shake a stick at, although as it's not Open Source, it's not possible to add your own modules and features without paying *Zeus* a whole load of money. Along with the Web server there are other package which can be used either with, or as a replacement for the Web Server package, including a rather nifty load balancer and a mass hosting server for ISPs wanting to run lots of sites on a shared server.

When it comes to dynamic content, *Zeus* does it all. PHP, Perl, ASP, J2EE. Existing code can be moved from *Apache* to *Zeus* without modification, although the Perl features in *Zeus* are not quite as good as `mod_perl` in *Apache*. Never the less, *Zeus* will happily run highly dynamic sites, making use of the existing performance features in *Zeus*. Even if it's running on a P233, rather than Dual P4 Xeon or even high-end enterprise hardware, *Zeus* will scale to make use of the appropriate hardware and avoid over allocating resources between clients.

Configuration of *Zeus* is totally point-and-click, making it a nice mix between the performance and

features of *Apache*, and the ease-of-use characterised by *Roxen*. Everything is web-based, allowing an administrator to create and manage sites without having to manually edit files. On the plus side, this reduces errors, but for those of us who like to get out the text editor and get the same thing done more quickly, it's a little disappointing.

If you've got the cash to spend on *Zeus*, then it's most likely a no-brainer picking out a Web server. However, when you consider the £1100 that *Zeus* starts at, it may be better to spend your money on improved hardware to run *Apache*.

VERDICT

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

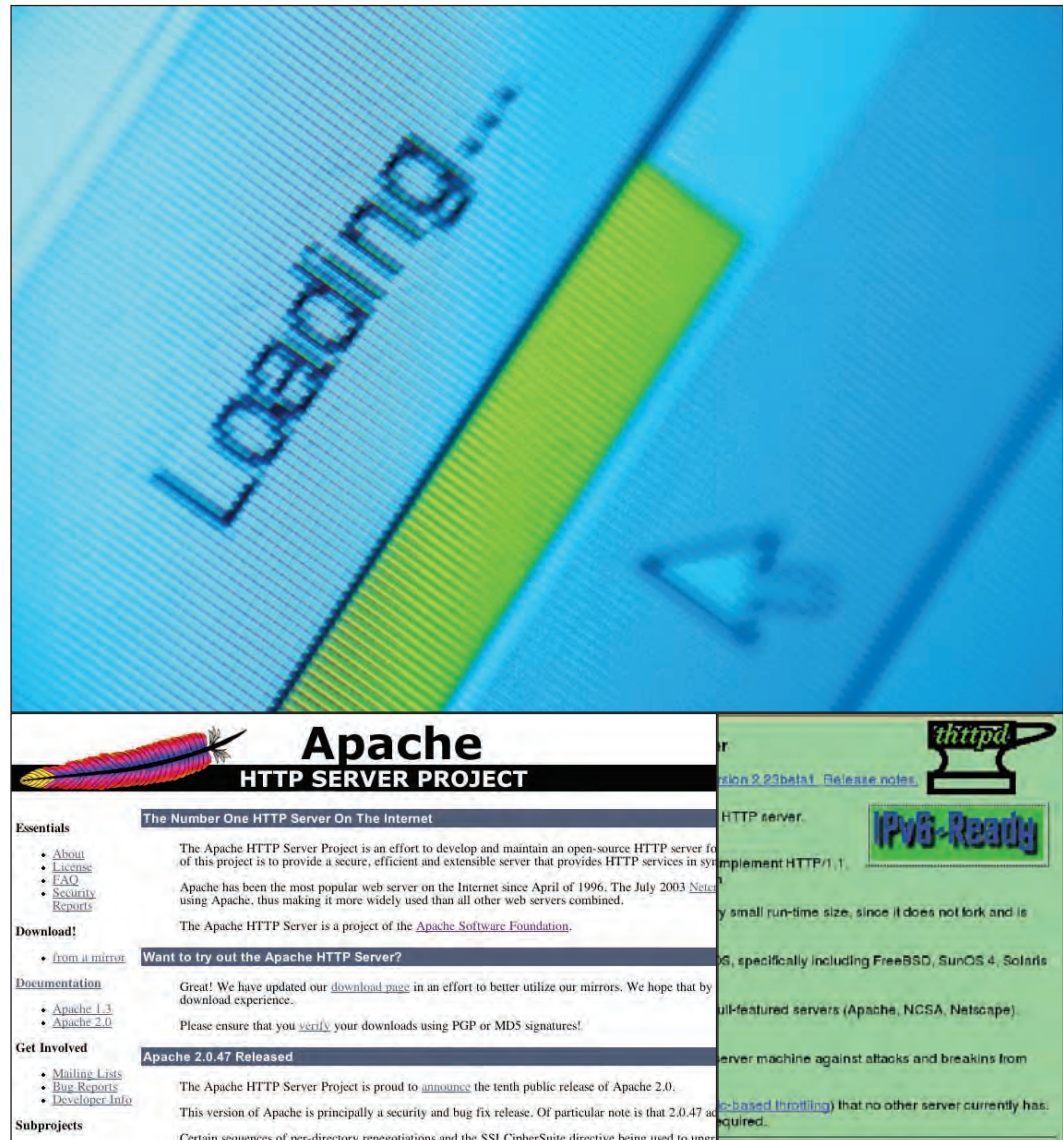
Zeus is a fantastic Web server and is more than able to handle massive amounts of traffic. Sadly, this doesn't come for free.

LINUX FORMAT RATING
 9/10

WEB SERVERS THE VERDICT

We've looked a great selection of Web servers, from those which run directly as part of the kernel, to others which can be clustered across a number of physical systems. It's not easy to pick out a single server which works best for everyone. *Apache* is certainly the most popular server out there, and being totally Open Source, it's trivial to modify the server and tune it to your specific needs. With the tens, if not hundreds of modules available, chances are if you need *Apache* to do something, then someone will have constructed a module to do it already, or it can be implemented in Perl using `mod_perl`. On the down side, *Apache* has far too much overhead for an embedded platform which needs to handle a large number of requests for static content on a very limited hardware platform. *Tux*, *Boa* or *thttpd* are perfect for this, although which of the three is best is a matter of taste. *Boa* is a very stable and scalable system, and there isn't a whole lot of difference in terms of performance between *Boa* or *thttpd*.

On the Enterprise front, *Zeus* is certainly way up there, above *Apache* and *Roxen*, although with the appropriate tuning and additional modules, *Apache* can get very close to *Zeus* on the scalability front. Naturally, without license costs in the mix, one can purchase more hardware and build a more redundant network using Open Source products. *Roxen* is also available as part of a commercial product, but if the end goal is only to



provide web services, then money can be better spent on third-party support

than licenses for software which will never be used.

Whether the budget is zero, or tens of thousands of pounds, Open Source software really can compete with high-end commercial products. One certainly can't complain for a lack of choice and in many ways it is unfortunate that there are lots of different products available which do almost exactly the same things. Of course, it is possible to run multiple Web servers on a single system, one for static content and another for the dynamic pages. Deciding exactly which server to use is entirely up to you.

All of the Web servers we have looked at have a wide selection of capabilities, and implement handling of multiple requests in slightly different ways, as can be seen in the table. **LXF**

Table of features

Web Server Capabilities

NAME	Model	CGI	PHP	ASP	VHosts	Auth	chroot	Throttling
<i>Apache</i>	Pre-fork	Y	Y	N	Y	Y	N	N
<i>Boa</i>	Select	Y	N	N	Y	Y	Y	N
<i>Caudium</i>	Threads	Y	Y	N	Y	Y	Y	N
<i>Roxen</i>	Threads	Y	Y	N	Y	Y	Y	N
<i>thttpd</i>	Select	Y	N	N	N	Y	Y	Y
<i>Tux</i>	Select	N	N	N	N	N	N	N
<i>Zeus</i>	Select	Y	Y	Y	Y	Y	N	Y

Note: Model is the method by which the server handles requests.

- pre-fork – start a pool of processes which each handle multiple requests.
- threads – use threads instead of processes.
- select – use non-blocking I/O and the `select()` system call to handle multiple requests in a single process, single thread.

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

HotPicks at a glance

AbiWord	44
XFce	45
IlohaMail	46
Scribus	46
NaturalDocs	47
IBMonitor	47
Stoned	48
FooBillard	48

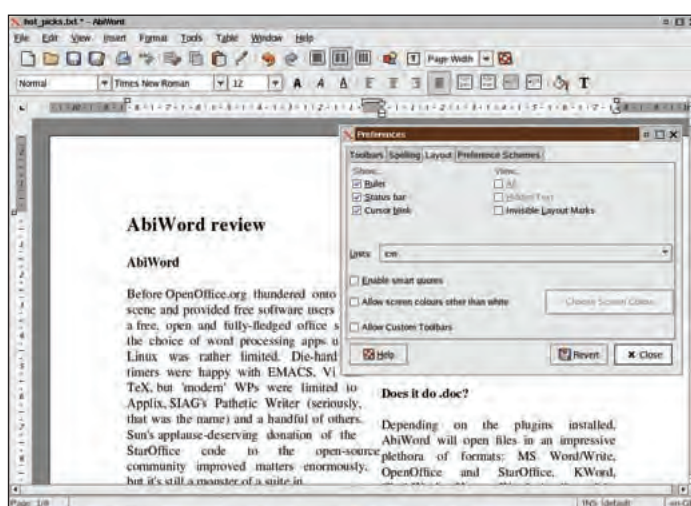
HOTPICKS AWARD

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



WORD PROCESSOR AbiWord

■ **VERSION** 2.0 beta 2 ■ **WEB** www.abisource.com



AbiWord using Red Hat's Bluecurve theme, preferences box in front.

Before *OpenOffice.org* thundered onto the scene and provided free software users with a no-charge, open and fully-fledged office suite, the choice of word processing apps under Linux was rather limited. Die-hard old-timers were happy with *Emacs*, *Vi* and *TeX*, but 'modern' WPs were limited to *Applix*, *SIAG's Pathetic Writer* (seriously, that was the name) and a handful of others. Sun's laudable donation of the *StarOffice* code to the Open Source community improved things hugely, but it's still a monster of a suite in terms of system requirements and the multi-platform *AbiWord* has been filling the gaps.

AbiWord 1.0 was finished at last in April 2002 after a long haul of development. As this issue was being put together, 2.0 was just about ready for release and will be widely available as you read this; as a result, we're covering 2.0 beta 2 here which will be virtually identical to the final version.

Two main build options are available if you're compiling from source: the GNOME version, which ties into the desktop environment, or the pure GTK+

option which lacks a few non-essential printing-related features and other cosmetic touches. We've also supplied RPM packages on the coverdisc, and you'll need the *FreeBidi* libraries installed (also included) to compile it successfully.

AbiWord's GTK2 interface retains the general layout of its predecessor – the two main toolbars and ruler are visible by default and waste little screen space. Many of the icons have been replaced by smoother and prettier equivalents, while the original splash screen is still present. As yet, there's no way to hide disabled menu items à la *MS Office* (though some would say that's a bonus). Overall, the front-end is polished and professional while remaining highly accessible to newcomers.

Does it do .doc?

Depending on the plugins installed, *AbiWord* will open files in an impressive plethora of formats: *MS Word/Write*, *OpenOffice*, *StarOffice*, *KWord*, *ClarisWorks*, *HancorWord*, *Applix* and its own native XML .abw are all supported, along with the usual

RTF, HTML and plain text. The range is more comprehensive than *OpenOffice's*, but the *Word.doc* importer isn't as mature and is also restricted by *AbiWord's* own limitations. We challenged it with .doc files of various complexity; it did a competent job of pulling out the text and basic formatting, but suffered badly with graphic-laden files.

Basic word processing tasks – including paragraph styles, bullet/numbered lists, columns, text alignment, image-insertion, headers/footers, undo/redo, find/replace and spellchecking – are all catered for. New additions in this major release are tables (for basic grid-layout data), mail merge (make letters using pre-defined address lists, but quite glitchy), revisions (highlighting changes from a previous document) and other tweaks and fixes.

Numerous plugins enable the user to access external programs and websites for translating or searching etc, while a few templates are supplied for creating fax coversheets, business reports and the like. Helpfully, *AbiWord* now uses the *Fontconfig* system, which makes managing fonts considerably easier too.

Not a great deal of customisation can be achieved through the settings box though, and toolbar reorganisation would be a welcome extra in further 2.x releases. Nevertheless, the few tweakable configuration options can be saved in separate 'schemes' which is most useful when you're dealing with drastically different types of document.

AbiWord, in all, is a splendid little word processor and an excellent lightweight alternative to *OpenOffice Writer*. It may not be the most featureful program in existence, but a recent study found that 90% of users only actually use 10% of the features in programs like *MS Word* – the real day-to-day tasks of writing letters, reports and notes are handled by *AbiWord* with ease. Its *Word* import code is usable for basic .doc files, the GTK version doesn't tie it to any particular desktop, and it's speedy and user-friendly. Definitely worth a look.

DESKTOP ENVIRONMENT

XFce

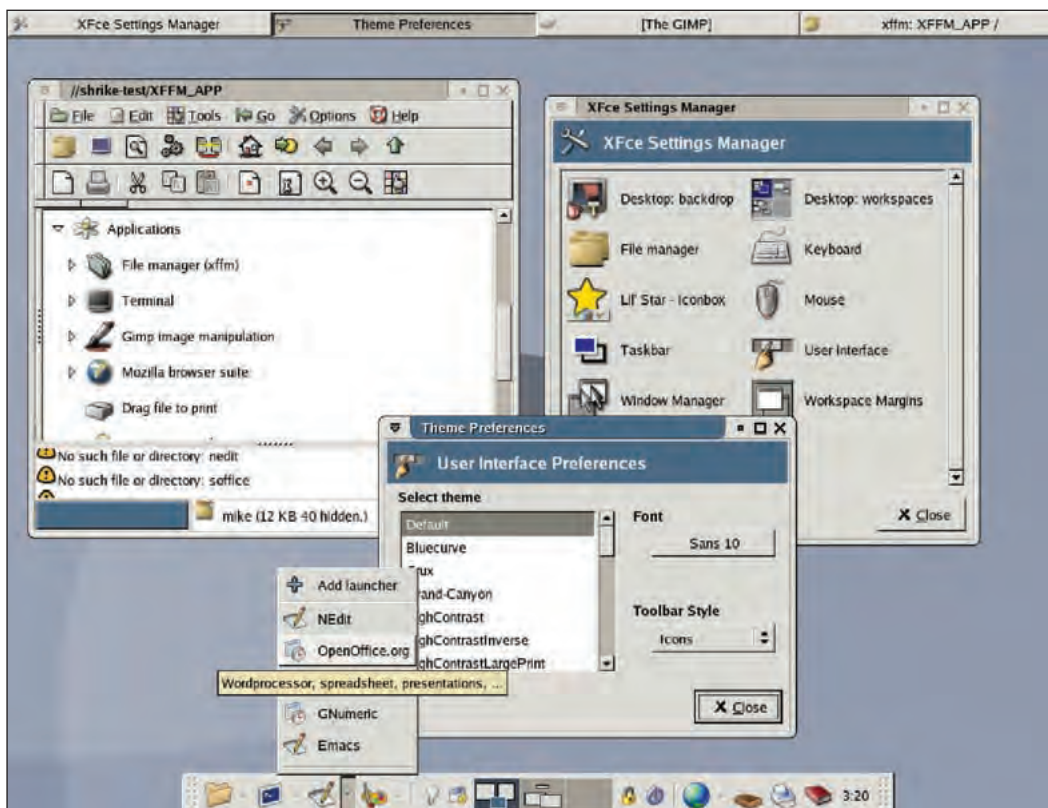
■ VERSION 4.0 RC 2 ■ WEB www.xfce.org

Back in the 1990s, CDE became the most popular desktop environment on UNIX workstations – it stands for *Common Desktop Environment*, and was introduced to provide a coherent and logical front-end over a variety of UNIX flavours. A few companies (including Red Hat) brought the closed-source *Motif*-based desktop to Linux; it didn't take long, though, before KDE and GNOME caught up with it in features and consistency. Then *XFce* appeared, offering a CDE-esque desktop via *FVWM* and *GTK+* code, and we last looked at the 3.x series in *LXF37's Roundup*. In case you were wondering how to get your tongue around the name of this particular program, it's generally pronounced "Ex-face" in conversation!

KDE and GNOME are now the default suites on most desktop-oriented distros, and there's little disagreement that they're immensely powerful and flexible. Unfortunately, size and bloat are the trade-off, and the latest releases of these desktops require a relatively high-spec machine – 128M+ of RAM is the sane minimum for running other apps on top. Recognising that this situation doesn't help in advocating Linux as a svelte and fast OS suitable for older machines, *XFce's* coders have always aimed for superior performance and a slimline codebase.

We've included both source code and ready-made RPMs for Red Hat, Mandrake and SuSE on the coverdisc; they're in *Bzip2* format so you'll need to **tar xfvj** them before installing the packages. (These are RC 2 – the final will be out as you read this and virtually identical.) From there, you'll need to alter your X startup scripts or GDM/KDM login screen properties so that **startxfce4** is run instead of the usual. *XFce 4* is built around *GTK 2.0.6* or above, and as a result you get all the sweets it provides like anti-aliased font support.

Aside from the general cosmetic enhancements over the 3.x releases, the most striking addition is the taskbar – this sits at the top of the screen by default and eats up a little more space,



Default XFce setup, with file manager and settings dialogs open.

but another option is the dynamically-resizing icon box. The CDE-inspired panel remains similar, save for the a graphical pager, and thankfully it snaps into place in the centre of the screen instead of requiring pixel-perfect positioning by hand.

It's XFcellent

For general use, *XFce* is very approachable with a sane set of window management functions and keybindings. The drawer system for the panel works well as a method of grouping related programs together, while stability and running speed are very respectable. As a rough real-life guide, on a fresh Red Hat 9 installation approximately 45MB of memory is used after booting into GNOME, whereas only 20MB is taken up when starting straight into the *XFce* environment. With a lighter distro, *XFce* is very much suited to old 32MB boxes.

As more than a window manager, *XFce* also sports a number of small

utilities: a friendly and attractive file manager, which is capable of browsing SMB shares for Windows networks, alongside a 'Run' dialog, a simple printing manager tool and a session manager for restoring previously running apps when logging back in. The small but well-written manual gives an overview of the desktop using plenty of screenshots.

XFce's configuration window has evolved into a pleasant little bunch of tools, all of which are a breeze to operate. Virtually every aspect of the desktop can be modified through these. It definitely rivals KDE and GNOME in terms of ease-of-use; panel properties, window decoration, theme and icon sets (with Red Hat's Bluecurve and KDE's Crystal icons included) are all here.

Despite the advances, there are still a few rough edges that need polishing up and some extras to be thrown in. Keybindings, for example, can't be configured via the graphical

tools but require hand-hacking of a text file right now. Similarly, the new taskbar module works well but is somewhat basic; being able to right-click on items and perform related window operations would definitely be a bonus.

On the whole, *XFce's* progress since the hastily assembled but workable 3.x series is nothing short of fantastic. Attention has been paid to consistency, functionality and spit-shine – plus it takes the best ideas of its progenitor CDE and brings them into the modern day.

Best of all, though, is that *XFce* delivers what Linux and free UNIX desktops have been sorely lacking: a middle ground between the hulking KDE and GNOME suites and the small and lightweight window managers of *IceWM*, *Fluxbox* and *WMaker* etc. *XFce* gives Linux a clean, fast and easy-going front-end without eating up system resources, and that's something worth shouting about.

WEBMAIL SUITE

IlohaMail

■ **VERSION** 0.8.8 ■ **WEB** <http://ilohamail.org>

Although Linux has a top-notch range of email clients, from the text-console to rich desktop environments, there are still many situations where

web-based front-end is more suitable. Most crucially, accessing email through a browser provides a consistent interface regardless of machine or OS – hugely

useful if you spend a lot of time on the move. *IlohaMail* is one such IMAP/POP3 webmail suite, striving to be easy to install and use, and has already seen large-scale deployment in a Japanese firm (and even won an award!).

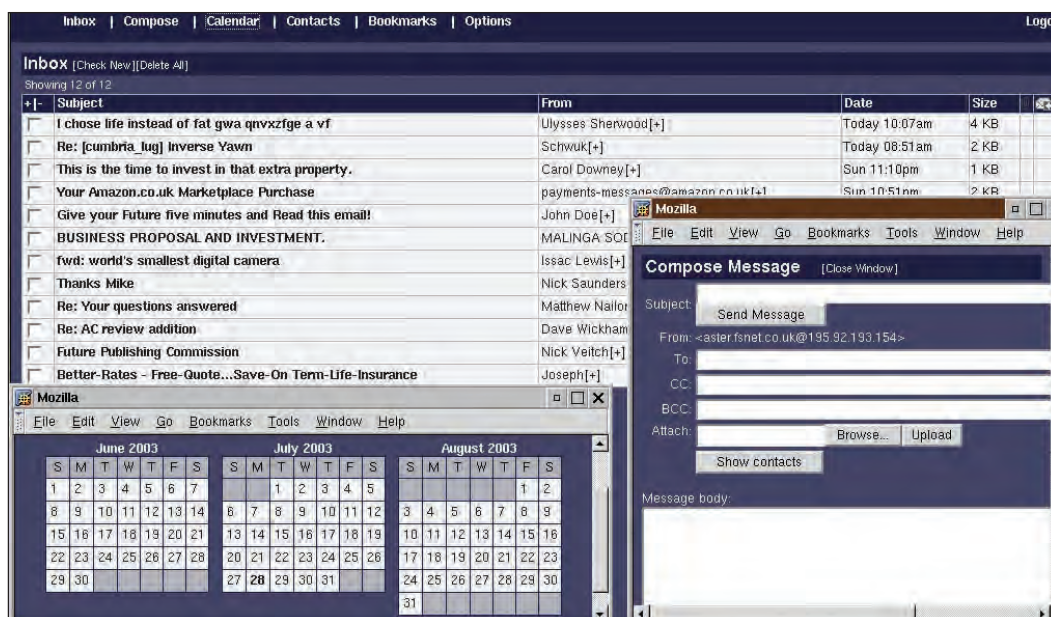
All you'll need to get *IlohaMail* up and running is the *Apache* webserver and PHP4 (the developers recommend version 4.1.2 or newer), and an external database is purely optional. The INSTALL text gives decent instructions on setting

it up; essentially, you simply need to copy the main directory into place, give `httpd.conf` an alias for the location, and make sure file-uploading is enabled in `php.ini` (for adding attachments when composing an email).

IlohaMail's pages consist of a mixture of blues with white text, and fortunately few images to slow things down. After setting some prefs, you're presented with a framed window; the menu remains along the top while at the bottom sits the main action. Intensely brilliant is its performance in \S – working properly in minimalist browsers bodes well for other small-screen and low-spec devices such as handhelds.

Alongside the column-sortable Inbox section, there's a basic message editing pop-up, a rudimentary calendar and a capable address book. It's not the most flexible system webmail system around, but it does the job without any niggles and is easy to manage. We'd like to see more visual customisation options though, as evidently the colour scheme can only be changed by hand-editing the files themselves.

IlohaMail is a fast, highly stable and straightforward webmail suite which has the added benefit of cooperating nicely with even the most basic of browsers. There are naturally more complex alternatives available, but *IlohaMail* fulfils its goals very competently.



Another day, another towering pile of junk mail – sadly, it's the same whichever email client you use...

SOURCE CODE NAVIGATOR/EDITOR

KScope

■ **VERSION** 0.2 ■ **WEB** www.cs.bgu.ac.il/~lahave/kscope/

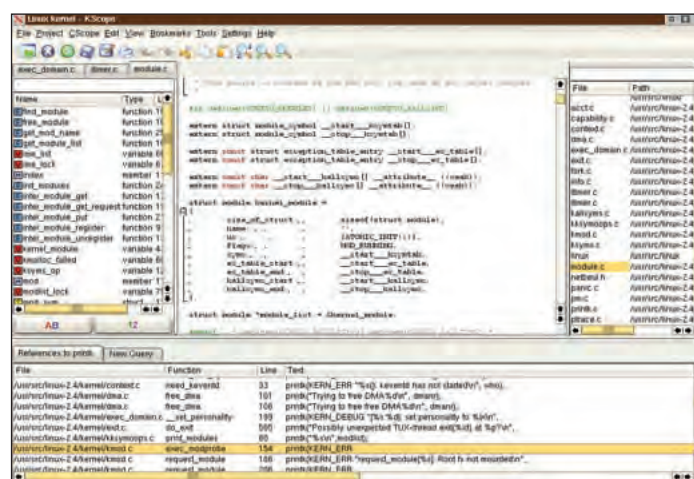
Writing good quality software is always a challenge, even with a small program and limited number of source files. As it grows, the complexity can become unmanageable – the human brain can only keep track of about 7 things at once, so for large projects a more capable system is essential. *KScope* is a KDE-based interface on top of the command-line *Cscope*, and provides robust source code navigation and query mechanisms (but it's not intended to compete with *KDevelop et al*).

To build *KScope*, you'll need KDE 3.1 and Qt, along with the relevant development packages. *Cscope* also needs to be built separately; the author suggests that a special snapshot is used as it provides some extra goodies, and

we've included that on the coverdisc too. Finally, it must be configured for installation under `/usr` at the moment – that's a known bug.

When first started, and after having given the program directions to the *cscope* and *ctags* binaries, the initial step is to create a project and point it at C and header files. These are listed down the right-hand side, while functions and variables for the selected file are on the left. In the middle sits the embedded syntax-highlighted *KWrite* window for editing, and along the bottom lies the query window. It's neat and works pleasing well on lower resolutions.

Alongside the basic editing facilities, several queries can be performed via the underlying *Cscope* engine – these



KScope wrestling with the mighty kernel source tree.

range from elementary *regex* text finds to full-on symbol reference searches across all files. The tabbed interface assists in switching between queries and open files without eating up too much screen space, and the petite docs give a good overview of common operations. *KScope* is respectably powerful and

reliable in its capabilities, and handles huge-scale projects (such as the Linux kernel) with ease. Some work needs to be done on tidying the interface and there's also a lack of customisation options, but otherwise it's a praiseworthy tool and deserves a look if you need a less complex substitute for the big IDEs.

PROGRAMMING DOC TOOL

NaturalDocs

■ VERSION 1.0 ■ WEB www.naturaldocs.org

When creating an Open Source project, documenting the code adequately is essential for two main reasons: firstly, it'll make things easier when you go back to an old untouched-in-months function, and

secondly it assists other developers in contributing to (and thus enhancing) your app. Sadly, many developers seem to see writing documentation as a tedious chore, so *NaturalDocs* has appeared to make the task quicker –

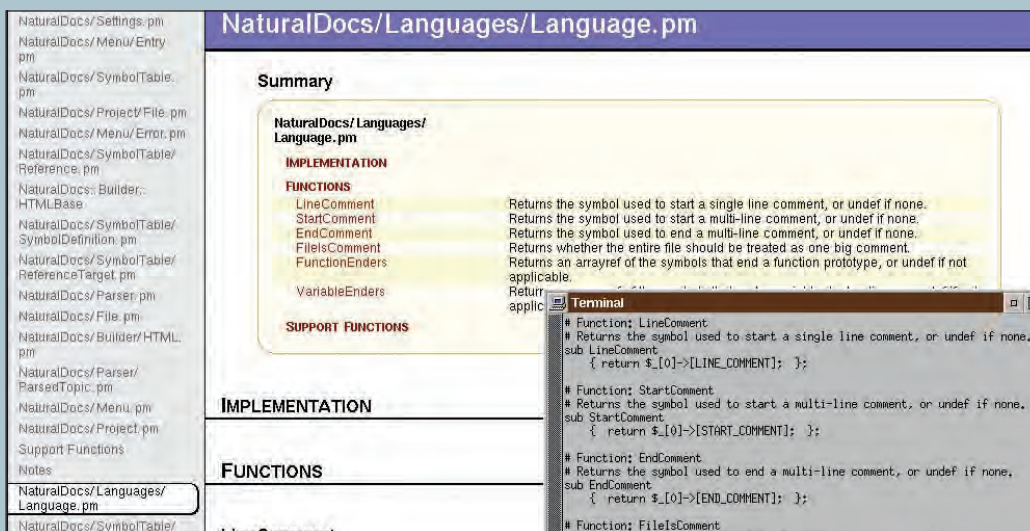
it uses comments written in to the source code itself, avoiding time-wasting and duplication.

NaturalDocs is a Perl program, so you'll need the relevant binaries to run it, but otherwise it's very self-contained. Being driven purely from the command line, the three most important switches you'll need are **-i** for the input directory (ie your code), **-o** for the HTML output of your source comments, and **-p** for the project directory (the latter is a temporary store so that unchanged files aren't parsed every time).

Making full use of *NaturalDocs* involves getting to grips with the commenting system – you have to be quite strict in how source comments are presented, but this leads to consistency in the original code too. For example, in a pre-function description you'd have 'Function: foo' followed by its explanation, then possibly 'Parameters:' and 'Returns:' blocks. *NaturalDocs* recognises a wide range of comment markers – it's not just limited to C/C++ code, and covers Perl, Java and JavaScript too.

Though the parsing and HTML-generating can take some time, the results are attractive, with a navigation list of files down the left-hand side and in-depth description on the right. A few visual changes can be made by specifying an alternate style from the command line (CSS is used), and on the whole it's a doddle to use.

NaturalDocs is worth investigating if you're working on an Open Source project and need well-presented documentation for the Web. It's not too pedantic in its demands and following its style doesn't make the original code look a mess. Meanwhile, the author has some big projections for the next major release – even full HTMLisation of totally undocumented code is one intriguing possibility.



Output from *NaturalDocs*' own code, and the original source in the xterm.

NETWORK BANDWIDTH MONITOR

IBMonitor

■ VERSION 1.0.1 ■ WEB <http://ibmonitor.sourceforge.net>

Running a healthy server involves care and attention – making sure software stays up, crackers stay out and that the machine never gets too overloaded if you're expecting something like a Slashdotting. There are various ways to keep track of network activity – reader Grant Moore suggested LXF looks at *IBMonitor*, a console based Interactive Bandwidth Monitor.

As a small Perl program, *IBMonitor* is a cinch to get running and only requires `TERM::ANSIColor` and `TERM::ReadKey` Perl modules installed on the system (the latter is in the tarball, but most distros include both by default). You should also make sure `/proc` is available as it reads `/proc/net/dev` – again, this is the case in just about every distro.

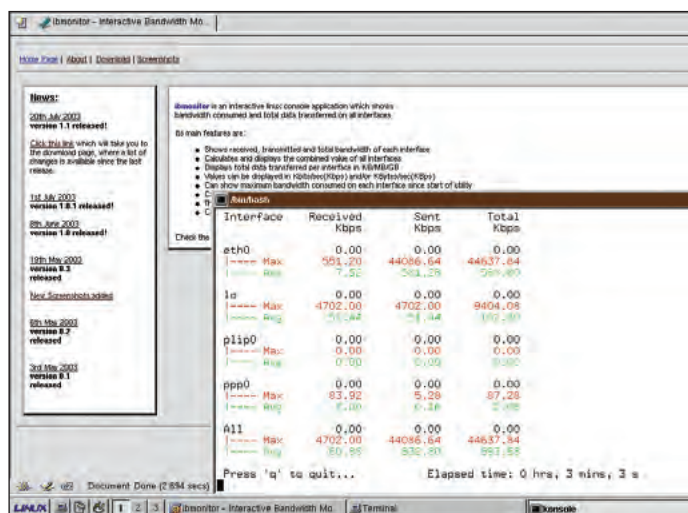
When run without any specific flags, *IBMonitor* pops up a simplistic text-based display which lists the network interfaces down the left-hand side and their current bandwidth levels on the right. Extra command-line options tell the program to display stats for the maximum and average consumption for this session, and these can be coloured to make the display more readable.

IBMonitor works like 'top' in that the information display can be modified through a number of keypresses. Hitting **i** and **y**, for example, alternates between kilobits and kilobytes per second, and the display refresh rate can be altered with the number keys (important on a slow *Telnet/SSH* link, to stop the program eating up your own bandwidth!). Also,

the list of interfaces can be re-ordered, though the method for doing so is not so elegant.

A straightforward little tool without fanfares or pointless frills, *IBMonitor* is ideal for running in an *SSH* window and keeping tabs on a server's performance.

It could do with a way to make the list auto-arrange itself maximum-first, but meanwhile the author has larger goals – external-process communication for HTMLising the stats is one. Plus, the compact and cleanly-written Perl codebase makes it a doddle to tweak.



A *konsole* window displaying today's bandwidth stats.

CURLING SIMULATOR

Stoned

■ **VERSION** 1.1.4 ■ **WEB** <http://stoned.cute-ninjas.com>

Curling: “a game played on ice in which large circular flat stones are slid across the surface towards a mark” as the *OED* defines it. Even Curling fans have accepted that the hi-octane low-temperature sport can get a tad dull – they’ve introduced a ‘free-guard zone’ rule to speed matters up.

Pleasingly, *Stoned* is supplied as a standalone statically linked binary, so you shouldn’t have any major problems running it (and it manages to cram quite a lot into such a small executable). You’ll need to have the *GLUT* library installed; it’ll most likely be on your distro discs or already on the system. From there, running it and hitting space on the title screen will take you to the main play area.

Stoned is network-aware for multiplayer games and, being a slow-paced sport, even the latency of a dialup connection shouldn’t be too troublesome. You’re dropped into a bare 3D arena with ugly wall texturing but happily smooth ice reflections, and an overactive background tune which

works surprisingly well. On the right is some quick help text and a target for

aiming the stones. Simple but effective.

With a target set, hitting both left and right mouse buttons together launches the stone, with a last-second change of direction/pace possible. Your shot then drifts along while the camera follows and dramatically zooms into the target on approach. There’s little in the way of sound effects – just the ‘chink’ of stone collisions.

Stoned is a cute little desktop emulation of Curling, and whether or not it’s totally appealing depends on your interest in relaxed-paced aim-games. Certainly, using the mouse to throw the stone is far more realistic than the tired ‘stop the moving bar’ interface type that golf sims use, and there’s something innately satisfying about objects gliding on ice.



Scotland won an Olympic Gold Medal for Curling at the Salt Lake City 2002 Winter Olympics.

POOL/SNOOKER GAME

FooBillard

■ **VERSION** 2.8 ■ **WEB** <http://foobillard.sunsite.dk>



Rubbishly, my highest Snooker break in real-life is only 26.

It's a colossal shame that Jimmy

White never won the World Snooker Championship, despite many blood-pressure raising attempts. In the meantime, we can give The Whirlwind some moral support by drinking beer and playing computer snooker games like *FooBillard* (it's not a typo – it's a German name).

FooBillard is available in GLUT and SDL versions – the latter is more likely to work on most Linux boxes. Also, the PNG and Freetype2 libraries need to be installed for the game to run properly – most distros put these on the system by default. Once ready, running **foobillard** will throw you straight onto the table, ready to take a shot (there's no title screen or intro).

The basic controls will be familiar to regular Archer Maclean fans, with the left mouse button used for targeting a shot and right button for zooming. Fiddling with the mouse's scroll wheel (or keyboard cursor keys) can set the power of a shot, while pressing the

middle button (or hitting the space bar) finally executes it. Adding spin is a bit bizarre – tap **C** for the external view, hold shift and the right mouse button, and then add spin and/or side.

Four game styles have been crammed into *FooBillard* – 8-ball, 9-ball, Snooker and Carambol. Carambol appears to be a German strain of Billiards. Ball physics are solid and convincing and the computer opposition levels range from disastrous to gung-ho nothing's-difficult brilliance, while the smooth visuals and attractive reflections are backed by simple but appropriate sound effects.

Even with its complete lack of jazzy presentation, *FooBillard* is still an entertaining baize-'em-up with a longevity-increasing variety of game styles and opponents. Undoubtedly the controls need a bit more work (there's way too much momentum and drift when using the mouse to navigate the table) but on the whole it's pretty to look at and challenging to boot. **LXF**

SPEED UP LINUX!

cover feature



Join *Linux Format's* experts in a tour of tweaking and optimising, as we try to get every last clock cycle used to the full. Your journey to a turbo-charged Linux box begins here!



BLUFF! Processor type



Choosing the processor option in the kernel configuration doesn't really make a huge difference to performance, though it does make your kernel more likely to compile properly.

Attempting the installation of any boxed Linux distro onto your machine these days is likely to result in a properly configured, fully working computer system that recognises all of your internal and external hardware (apart from the odd problem with winmodems etc). But the price you usually pay for such ease-of-use is performance. There is no way the installer can tune the system exactly to your needs, mainly because it can't determine what your needs really are.

Over the next few pages, we'll be looking at a range of the things that will increase the performance of your Linux box. Some of the things we detail may make a radical difference, some may result in a small boost, but if you are continually frustrated by the length of time it takes to do the stuff you want, you'll be glad of every second saved.

We'll also be debunking some myths and misconceptions regarding tuning your system, so get ready to tweak!

SWAP SPACE

One subject that is always sure to start a flamewar on a newsgroup or web forum is someone asking how much swap space they should allocate for Linux. In the 'old days', the general rule of thumb was that swap should be twice the size of physical memory. With modern systems being capable of addressing huge amounts of real, physical memory, that could lead to some huge swap partitions, which in truth aren't going to be used much.

So what it's the optimum setting for swap? When considering terms of pure performance, you are better off not having any swap at all, because swap is slower than accessing real memory. Many embedded Linux devices use no swap space at all for this very reason, though obviously, because of rather specific applications, it is easier to profile memory usage and determine an upper limit more successfully.

For a desktop or server situation, it is often advisable to have some swap available, if only to cope with the urgent need to run just one more application, or spawn one more server process. In practice, you will find that your swap space nearly always has some usage, even when there is plenty of free memory. Just because swap space is allocated, doesn't always mean it is actually being used though. Many swap pages, created perhaps from operations in a now, long closed operation, are simply not cleaned up until they are needed for something else.

So, how much?

There isn't a great deal of evidence that having too much swap space is detrimental to performance, it is simply a less efficient way of using the disk space. The reverse is true though – denying enough swap space to



SpeedUp



adequately deal with the processes you are running will ultimately result in something just not working, but your system performance will begin to degrade the moment you start to use swap, and fall off seriously when the kernel has to choose whether to lose cache space, or recover memory from complicated processes.

You can track your memory usage using one of the many graphical tools that are available, or simply by having a look in /proc:

```
cat /proc/meminfo
```

will give you a fairly detailed summary of your usage so far. The important totals are:

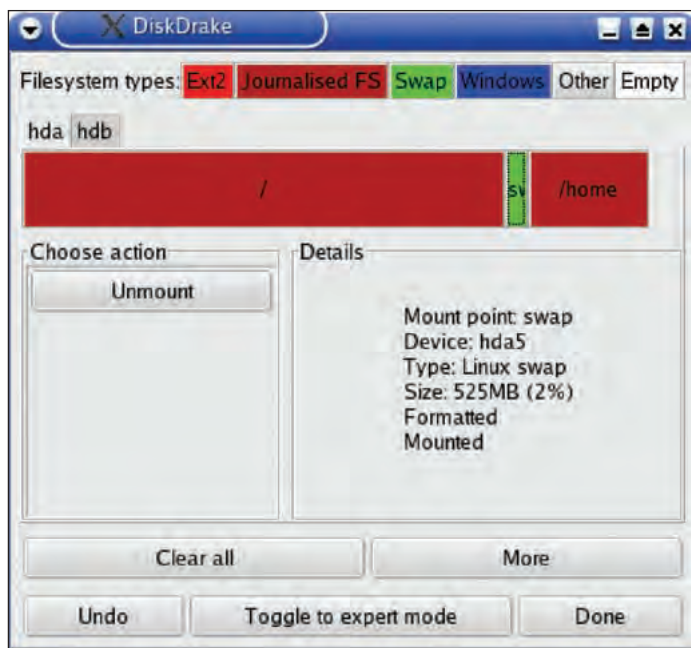
- **MemTotal** – The total physical RAM on the system
- **MemFree** – The total physical RAM available
- **SwapCached** – This is memory that was at one stage in the swap and has been paged back in to physical memory, but is still left in the swap. The reason is, it may need

BLUFF!

Fewer files in your directories is better

We're not sure where this one came from, but it seems that some people out there believe that filesystem performance degrades as you store more and more files in a directory. This is just another myth – the number of files you store in a directory has no effect on how fast your system performs.

Having enough swap space can be critical to overall system speed.



What is swap?

The swap partition on a system is the Linux way of creating virtual memory. In a system using VM, processes, or parts of processes and their data can be 'paged' out to hard disk. When they are required again, they can be paged back into the physical memory of the computer. If memory space is really tight, it may become necessary to actually continue executing the process

while still in the VM system, though it's desirable to avoid such situations because of the performance hit.

Virtual memory arose because of the relative expense of memory in the past. While byte for byte, hard drive space is still significantly cheaper, the cost of memory has fallen dramatically, so in many ways the more effective solution is to buy more RAM if you need it.

Power Postgres

PostgreSQL has come along leaps and bounds in terms of performance in the past few years, but that's not to say you can't give it an extra push to make sure it's working at top speed. Try these simple tips:

- Periodically use the **ANALYZE** and **VACUUM** commands, particularly on tables where you insert, update, or delete lots of values – this helps keep your data defragmented, and also gives Postgres the chance to store info about the data to make its search faster

■ Stored procedures are much faster than sending the same operation across the network each time you want to execute it, because Postgres doesn't have to re-parse, re-analyse, and re-optimize it each time you send it.

■ Like MySQL and other DBMSs, Postgres implements the **EXPLAIN** command – use it before a **SELECT** statement to see how Postgres was planning to execute your query.

to be swapped out again, and if so the system can save some time

■ **SwapTotal** – Total amount of swap space

■ **SwapFree** – Total amount of unused swap space.

Try out your system under different workloads to see just how much swap gets used. Remember that the SwapFree doesn't include the swap space which is just used for cache, and could easily be freed if required, so don't be misled into thinking there isn't much space in swap.

The whole issue is far from straightforward – it would be true to say that there is no real firm answer on how much swap space you should have, since really it depends what you are going to use the system for. On a huge database server with 4GB RAM, you'd probably want another 4GB of swap space. Indeed, Oracle itself recommends the old twice-the-size rule, even for systems with colossal amounts of memory.

On a desktop box, you can easily get away with less swap. 512MB is probably excessive for most desktop use, and it probably isn't worth increasing on this unless memory usage monitoring says the swap is congested.

Multiple swap spaces

As well as determining how much space to have, there are also ways of making the swap space you have work faster. For an obvious start, your swap space should be configured on the fastest drive your system has available. Hard disks are thousands of times slower for storing data on than physical RAM – you don't want to make that figure any worse than you have to.

If you have multiple drives connected to your system, you can also share the swap load. Cunningly, the kernel allows you to easily parallelize swap space in much the same way as RAID striping works – you can use multiple devices to reduce the I/O bottleneck on a single drive.

Imagine you have three connected storage devices – two SCSI disks and an IDE one. Not only can you split the swap space between the drives, you can also prioritise usage, so the fastest drives get used in preference to the slow one, which can act as a sort of emergency backup. A typical *fstab* listing may look like this:

```
...
/dev/sda2 none swap
sw,pri=5 0 0
/dev/sdb2 none swap
sw,pri=5 0 0
/dev/hda2 none swap
sw,pri=1 0 0
```

The **pri** option specifies a priority – the higher the number, the more readily the device is used. The first two entries have the same priority, with the result that the swap space is shared between them. This is analogous to RAID 0 and there is no limit to the number of devices you can use in this way. Obviously, it helps if these devices are identical, or at least of the same spec, otherwise the slower one will drag down the performance.

The final device has a lower priority, so it is only used when the swap space on the others is filled. This means that you can still allocate a huge chunk of space on a slow, old drive to cope with occasional emergencies like runaway processes etc, without sacrificing the normal performance.

MEMORY

You recall all those adverts you

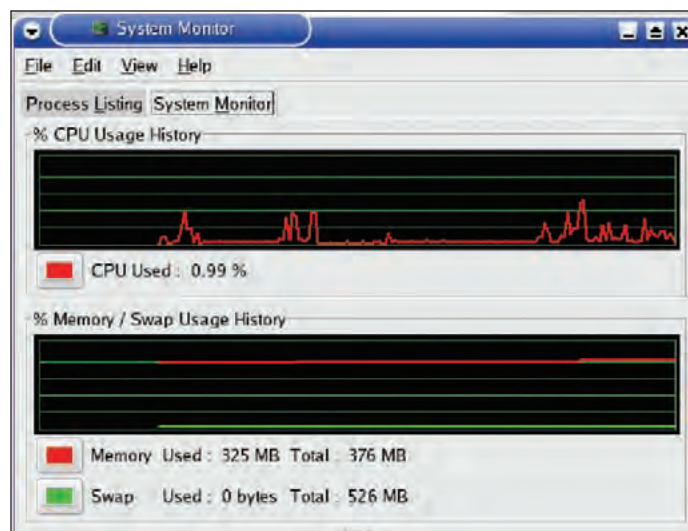
see saying 'buy more memory and your system will run umpty-ump times faster'? Well, though it may sound like it, it's not just a marketing gimmick. Memory is probably the most important factor in overall system speed after the CPU itself. Memory is where the applications live while they are being run, and often more importantly, where the data that they need to run on also resides. In fact, depending on the application, memory can often be more important than the CPU speed – especially for tasks such as running large database servers.

In many respects, memory is closely linked to swap space. By deploying more memory on your system, you can limit the amount of swap access required. As we've mentioned before, any sort of RAM is thousands of times faster than hard disk access, the more the system uses memory and the less swap it needs to use, the faster it is going to be.

Different types

Assuming that you are willing to stack your system with as much RAM as it can handle, what is the best type to use? Memory modules have changed many times over the last fifteen years. In the old days, individual memory chips, in the form of DIP, and later ZIP packages, were individually slotted onto motherboards or RAM expansion cards. Memory modules, starting with the SIMM, made handling memory easier because the (much smaller) chips were already surface mounted onto a card. At the moment, the way memory is added to computers is in a bit of a changeover. SDRAM chips are still widely used, but advances in technology have created the concept of DDR memory. DDR stands for Double Data Rate, because the memory can transfer data on the leading and falling edge of the clock pulse, effectively doubling the speed at which it works.

In practice DDR does double the speed at which RAM can be accessed,



but the system bus and memory controller limit the effective bandwidth, so everything won't run twice as fast, but the differences can be significant. Upgrading to a system board with DDR support would make sense when you need the highest possible performance from your memory.

The GNOME System Monitor can help you keep track of memory usage and is easily understandable. KDE's Control Center has a good memory display too.



Memory FAQ

Debunking the myths

Mark Sanor, Technical Support Manager of Crucial Technology answers your most frequently posed questions on Memory.

Q What effect does parity-checked RAM have on speed?

A You will notice less than 2% of a performance difference. However with the stability of memory modules today, ECC is only needed in mission critical applications, servers, high-end workstations, etc. I'd recommend that you use a configurator like our Memory Selector available from www.crucial.com/uk to ensure compatibility as the two types cannot be mixed.

Q What is the real difference between branded and generic memory?

A Not all memory components are 'PC grade'; some parts are better suited for mobile phones, video game systems, etc. When you are buying generic modules, you are buying from someone who is just getting parts on the open market and making modules out of them. When you buy branded memory (eg from Crucial) you are getting top-grade components and modules. And as we're part of the manufacturer we can ensure quality at every stage – from the raw silicon to the packaging. You can read more about it at the following URL:

http://www.crucial.com/uk/library/quality_page1.asp

Q What do numbers like PC2100 really mean? Are they significant, or just product number designations?

A In SDRAM modules, the numbers that come after the 'PC' refer to the speed of the system's front side bus (FSB) ie PC100 would have a 100MHz FSB, PC133 would have a 133MHz FSB etc. In DDR modules, the numbers that come after the 'PC' refer to the total bandwidth of the module. Those numbers represent the 'theoretical' amount of data it can move each second:
PC 2100(DDR266) = 2.1 GB/sec
PC 2700(DDR333) = 2.7 GB/sec
PC3200 (DDR400) = 3.2 GB/sec.

Q Would you need different memory for a server?

A Yes, some servers may require 'Registered' memory. Registered memory is used to ensure data integrity before the memory is written or read from the module. This type of memory is typically only used only in servers and other mission-critical systems where it is extremely important that the data is properly handled. Since servers can handle multiple clients simultaneously, it is imperative that the data

be correct, speed is not the most relevant part in server memory.

Q How do you know when lack of memory is the cause of poor system performance?

A When you find yourself watching the 'Hour Glass' on your system when you execute a command; opening a program, saving a document, or pauses when running a game, you are a candidate for a memory upgrade. This delay is caused because your system has to move things from memory to your hard drive for temporary storage, and your CPU is being underutilised as far as the functionality it *could* perform.

Q Why is it that RAM prices fluctuate so much?

A The memory market is a commodity market. External influences can dramatically alter its delicate balance between supply and demand. For instance, computer sales have a direct impact on demand, as does the release of more memory-intensive software. Supply can be affected by changeovers to next-generation products, producers entering or leaving the industry, or even natural disasters.

If demand increases with no change in supply, prices will generally go up. If supply

increases with no change in demand, prices will generally go down. As a rule of thumb; the more there is the less it costs... the less there is, the more it costs.

Q I've heard this term used in relation to several different things. What actually is 'dual-channel' DDR?

A The terminology 'dual-channel DDR' is being misused by some in the memory industry, which can mislead the consumer. The fact is there's *no such thing* as dual-channel DDR memory. There are, however, dual-channel platforms which have two RAM controllers that allow you to use two memory banks simultaneously. This means you have more bandwidth to take advantage of more memory at once. You'll notice greater overall system performance with a dual-channel DDR system, but the DDR itself is exactly the same as the DDR used in systems today. If indeed you have a dual-channel platform and you want to take advantage of the performance gain it offers, our advice is to opt for high quality and service over expensive packaging, and simply purchase your DDR memory in pairs. However, be very careful to order two modules with the exact same specifications; the modules must be identical to each other in order to perform correctly.



KERNEL

BLUFF!

Linux can't be defragged

Many 'power users' who come to Linux from Windows hunt around for a command to defragment their hard drive, because this was something they had to do regularly on Windows. We get asked this a lot, so let us make it clear: you don't need to defragment your Linux box, because it naturally keeps itself unfragmented. There's no performance degradation in this, so you needn't worry – Linux can't be defragged, yes, but Linux doesn't *need* to be defragged.



As with many other aspects of

system tuning, the real benefits are to be able to minimise the use of memory (without any speed trade-off) and eliminate unneeded code. Fortunately, the modular design of Linux makes this possible, and not too difficult. If you do want to recompile your kernel though, I'd suggest that you look at the LDP's extensive HOWTO (www.tldp.org).

Modules vs monolithic

One of the early innovations to make it into the Linux kernel was loadable module support. This makes it possible for the kernel to dynamically load modules, such as a device driver, as and when it is required (and also unload it when it is no longer required). The benefits of this are many-fold.

For a start, it means that the core kernel itself can be much smaller, occupying less memory and being generally just that tiny bit faster. It also means that at any one time, the kernel only needs to have the drivers loaded that it actually needs, and you can compile modules for less often used

items, or against the possibility of a system upgrade. You can pretty much compile all the modules for virtually everything supported in the kernel, secure in the knowledge that it won't be used unless needed.

But there are downsides to a massively modular kernel. Firstly, the kernel itself has to manage the modules. This means managing dependencies too, which can slow things down considerably when loading (e.g. the USB-modem module, if you try to load it, may need to invoke the USB-OHCI module and the USB subsystem too). This occurs not only at startup, but whenever you need to load a module. For a more efficient kernel, it is far far better to compile in the features you know you are going to need – e.g. SCSI support or PCMCIA. The latter is you'll usually either have or won't have. It isn't likely to be added to your system. If you use PCMCIA devices, compile it in, if not, leave it out.

You'll find that the default configurations for kernel source supplied with distributions usually includes virtually everything, and in most cases as a module. That's because the distro builders don't know what hardware you have. But it also makes their kernels bigger than they need to be, and often less efficient.

The speed savings from customising your kernel aren't always that large, or even noticeable (except perhaps in the shorter boot times), but they can make a difference – and at least you'll know you have a thoroughly tweaked box.

What do you need?

It's very handy to know exactly what hardware you are using so you can choose the specific kernel components required. Configuring your kernel is probably a feature or two in itself, but here are some pointers:

■ **PCMCIA** – leave it out if you aren't compiling for a laptop.

■ **SCSI support** – if you don't need it, leave it out, but remember that you will need generic SCSI support for CD burning apps, RAID controllers and other devices.

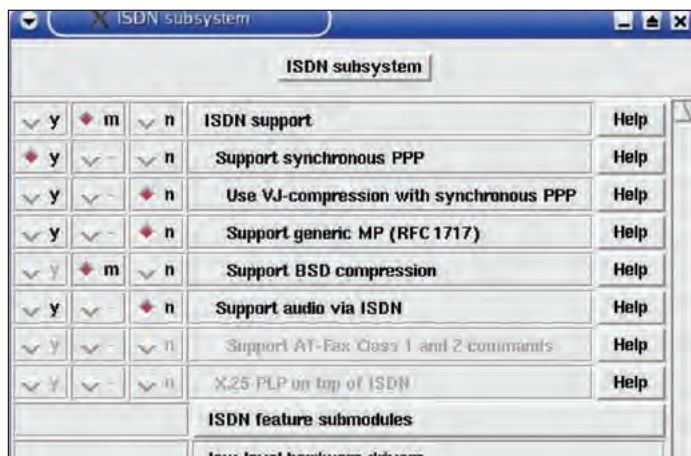
■ **Filesystems** – compile in the ones you use – e.g. *ext3*, *reiser*. Leave anything you don't use as modules (it doesn't hurt). Make sure you select the one your root partition uses, otherwise Linux won't be able to boot.

■ **USB** – probably on of the sections where you do want to leave things modular, expect perhaps the host (UHCI or whatever you use). The modules are only loaded as needed, which will be fine for hotplugging devices. Of course if you *always* use a USB mouse or whatever, you can compile it in.

■ **Ditch it!** – there are loads of sections which are only appropriate for specific uses of Linux, which are often enabled in distro builds. You can ditch things like MTD devices, Telephony, Radio, ATM etc unless you know you will need them.



Correctly configuring your kernel can make a difference.



If you don't use ISDN, don't bother compiling it. The same goes for a lot of sections that your standard distro will have enabled by default.

Speed Start

Many people seem to find the time it takes Linux to boot quite frustrating. In comparison to something like Windows XP running on the same hardware, it will probably take the average Linux distro a while longer to reach the login screen.

This isn't just because the Linux boot code is slow or inefficient – it's to do with different design philosophies. The two main things that slow Linux down on bootup are logging, and discovery. Linux logs almost everything you do, especially the bootup process. The logs are very useful for determining where things have gone wrong, what drivers have been initialised, and if any devices have failed to come up properly. But they are in the form of files written to

the disk, which obviously slows things down considerably. Another significant factor is hardware discovery, where modules are loaded if appropriate hardware is detected. The difference between modules and built-in drivers has been discussed in the kernel section.

One thing that you can easily change to shorten boot time is to turn off unnecessary services. While it isn't true that enabling lots of services eats up CPU and memory (most run in a sleep state, consuming no resources until required), initialising them does take time at the bootup phase. So, if you don't need MySQL, Apache, Samba and NFS all the time, you could save yourself a few seconds right there.

CPU

No matter how much you tweak

the rest of your system, at the end of the day it's your CPU that bears the brunt of the work – whether it's calculating π to the 3,000,000th digit or just handling a mouse button click. Furthermore, it's not really possible to make your CPU run faster without risking overclocking, and while our sister magazine *PC Answers* could probably take you step-by-step through overclocking your hardware, it's certainly out of our remit!

Few people argue against the fact that the best way to get more speed out of your PC is to simply replace your CPU with something newer and faster. Whether that's just ramping up the clock speed, or perhaps even jumping processor generations, eg from Athlon to Athlon XP, changing your CPU gives big gains for all processor-intensive tasks.

Identifying the need

How do you know, other than *OpenOffice.org* taking more than three minutes to start, when your CPU needs to be replaced? The best way to check your CPU usage is using the *top* command, and also the *uptime* command. *top* prints out a neatly ordered list describing how much RAM and CPU time each of your processes are using, as well as a key figure: how much time your CPU is spending idle. You need to watch the output of *top* for a couple of minutes to get a reasonable understanding of what's going on – you may have processes (such as your database server) that only spike every so often.

Using *uptime* gives you the load average for your machine for the last minute, five minutes, and fifteen minutes respectively, which means you can accurately track how much work your machine has had to bear. Low load averages (under 1) mean your machine isn't being used to its full potential by a long shot, with medium load averages (under 3) meaning that your machine is being used well – around 85% of its potential. High load averages (over three) mean that your machine is having to work above and beyond its specification, and usually means you need to upgrade your CPU.

```

Load averages - Konsole
Session Edit View Bookmarks Settings Help
paul@hud-1xf:~$ uptime
 12:57:45 up 1 day, 26 min, 1 user, load average: 54.09, 50.72, 32.13
paul@hud-1xf:~$
  
```

SMP and you

Would you rather have a single 3GHz P4 or dual 1.5GHz P4s? Most people don't seem to know or care about the difference between the two, which is a shame because the difference is indeed quite large. Both SMP (symmetric multiprocessor) systems and uniprocessor machines have their own advantages. For example, the advantage of physically having two processors inside your box is that your machine is capable of doing two things at once. Thanks to the ever-increasing clock speed of uniprocessor systems, modern computers do a very good job of *appearing* to do two or more things at once, but in reality your machine is *either* serving web pages *or* it's reading from your database, etc – it does just one thing at a time.

Moving to SMP, however, your machine becomes able to perform two tasks at full speed, which means that if you're compiling the kernel with one CPU, you can carry on working in X with no slowdown whatsoever. Of course, the extra speed comes at a price: SMP machines are nearly always more than twice the price of their equivalent uniprocessor systems, which means you pay a premium for the hardware.

However, what other way can you get a 6.0GHz box right now other than getting a dual 3GHz machine? The answer is that you can't, which is why SMP is here to stay.

Adopt Opteron lateron?

Despite the deluge of glowing reviews about AMD's new server and workstation CPU, people are still catching on quite slowly. Independent sources now believe that Opteron has already outsold Itanium, despite Itanium having been out for years, which says a lot about the quality of AMD's flagship design. So, why are

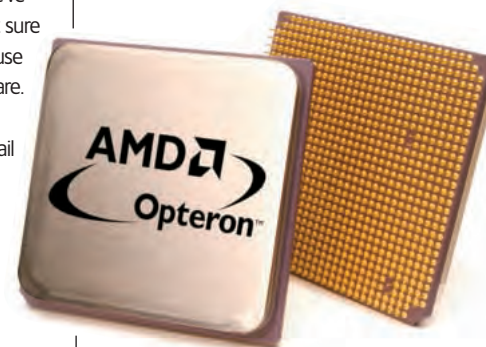
people unsure about taking the leap to 64-bit processors?

The most common reason we've encountered is that people aren't sure they *need* 64-bit hardware because they don't have any 64-bit software. While this point of view makes sense to a degree, most people fail to realise that simply installing a 64-bit OS while still using 32-bit software will make quite a difference to performance – this is one of the key benefits to the AMD64 platform. Add to that distributions such as SuSE's Linux Enterprise Server, which are wholly 64-bit throughout, and you will see a massive performance gain by switching from a high-speed Athlon to even a fairly modest Opteron.

Many pieces of software have made use of 64-bit mathematics for some time – one of the biggest 'culprits' is MySQL, which uses 64-bit integers for most of its calculations. While this is a big problem for 32-bit CPUs, which literally need to break the calculation into two parts then recombine it later, it's a breeze for the Opteron, where 64-bit mathematics is performed natively. While the Itanium series of CPUs from Intel also get a large benefit from the move to 64-bit, they have very poor 32-bit compatibility – rather than executing instructions natively, they have special emulation code that needs to translate 32-bit instructions into native 64-bit Itanium instructions before they can be executed.

Of course, if straight 32-bit CPUs and the Opteron aren't to your taste, a third option will be rearing its head in mid-way through September. The Athlon 64, previously codenamed 'Clawhammer', is set to revolutionise desktop computing by giving home users the power of 64-bit computing at an affordable price tag.

Quake 3 was not designed to run on a 386! A lot of processor-intensive tasks = time to upgrade.



BLUFF!

WINE must be slower than using Windows

Some people we've spoken to in the past often keep a Windows machine around on the basis that it will run Windows programs faster than WINE. This is one of those peculiar myths that has arisen because people just didn't take the time to find out what WINE actually stands for: **WINE Is Not an Emulator**.

What that means is that WINE doesn't read in Windows code, convert it to Linux code, then execute it, like a normal emulator would. Instead, it's a solid port of the Windows libraries themselves, which mean that Windows code running on WINE should run as fast as on Windows. Many people who've taken the trouble to investigate this further have come to the conclusion that Windows programs run even faster on Linux than on Windows, simply because of the more advanced kernel!

Mega MySQL

The Practical PHP tutorial this month is devoted to optimising MySQL – turn to page 84 for more information.





HARD DRIVES



Following on from the

CPU and memory, disk drives are probably the next biggest factor in overall system speed. A lot of software relies on large data files that need to be read and written to disk, never mind the drive's use as a source of virtual memory (see the swap section space).

Performance drives generally fall into two categories – SCSI devices and IDE. In general, SCSI devices tend to be higher performance – especially when you get up to the higher bandwidths offered by SCSI-3.

In performance terms, ATA drives have done a lot of catching up since they were first introduced to the desktop platform, and now offer performance at least good enough for all but the most demanding desktops. The new Serial ATA standard offers the potential for even better, SCSI-beating performance, but it is early days in terms of the development of drives that can make use of this bandwidth. One bonus of SATA that should be mentioned here is that the drives are connected in a point-to-point fashion to the bus adaptor, meaning there is no channel sharing

nonsense going on. This is particularly noticeable when transferring data between drives – if you need a lot of speed copying files between devices, SCSI or SATA are the way to go.

Some of the performance characteristics of drives are determined by their physical construction and properties – spin speed obviously plays a part here, but so too does the number of platters used in the drive – although adding platters, and the read/write heads to go with them, adds to the complexity of the device, it also reduces the time to access any specific piece of data.

Drive your HD harder

There are two simple changes you can make to your system that, together, greatly improve the performance of your hard disk: disable access time for files, and tweak your disk settings. While hardcore Unix users might disregard disabling file access time (*atime*), it's a simple fact that whenever you or a program 'touches' a file by reading or writing to it, Linux has to update that file's entry for its atime. However, this is a rather pointless exercise – few people ever make use of *atimes*, which mean that your system is just wasting its time. Luckily, it's a cinch to disable *atime* for

a given partition, as you'll see. Here's the first line from my `/etc/fstab` file:

```
/dev/hda1 / ext3 errors=remount-ro 0 1
```

If I wanted that partition to not both setting access times, I'd change it to this:

```
/dev/hda1 / ext3 errors=remount-ro,noatime 0 1
```

Once the drive is remounted, access times will be a thing of the past – easier than you thought, yes?

Tweaking your disk settings is a little riskier, because not all hardware supports the latest tweaks. You can do damage to your system by enabling settings that are not supported, and you should exercise extreme caution as well as reading your system manual before you attempt to follow the instructions presented in this feature. Furthermore, as SCSI disks are more advanced than IDE disks, the tweaks presented below will not work with SCSI systems.

The easiest way to test the basic performance of your hard drive is to use the `hdparm` tool, like this:

```
hdparm -tT /dev/hda
```

Note: you will need to change the `/dev/hda` to your own hard drive. Running that on an old test machine, we got the following results:

```
Timing buffer-cache reads: 128 MB
in 1.13 seconds = 113.27 MB/sec
```

```
Timing buffered disk reads: 64 MB
in 16.61 seconds = 3.85 MB/sec
```

3MB a second is hardly a world-beating transfer rate, so we tweaked the drive settings, and tried again:

```
Timing buffer-cache reads: 128 MB
in 1.14 seconds = 112.28 MB/sec
```

```
Timing buffered disk reads: 64 MB
in 7.29 seconds = 8.78 MB/sec
```

As you can see, it's easy to double the speed of your hard drive just by setting things up properly – if you have a newer system, you'll see even larger gains. The parameters we used were `-m16 -d1 -u1 -c1`. `-m16` allows your hard drive to transfer a maximum of 16 sectors per interrupt (some drives even support 32), `-d1` enables DMA, `-u1` allows interrupt unmasking, which will make your system much more responsive, and finally `-c1` enables pure 32-bit mode. It's important that you check your system

BLUFF!

Killing unused processes makes your PC faster

It's a common belief that killing dormant processes that seldom used will make your computer run faster. Sadly, that's rarely the case – unused processes take up no CPU time at all, which means the only case you're likely to get any extra performance by killing them off is where you're very short on RAM. In this situation, killing off unused processes might free up a few megabytes that can otherwise be reused elsewhere – even this is unlikely, though, as Linux is very good at swapping unused data into virtual memory to speed things up.

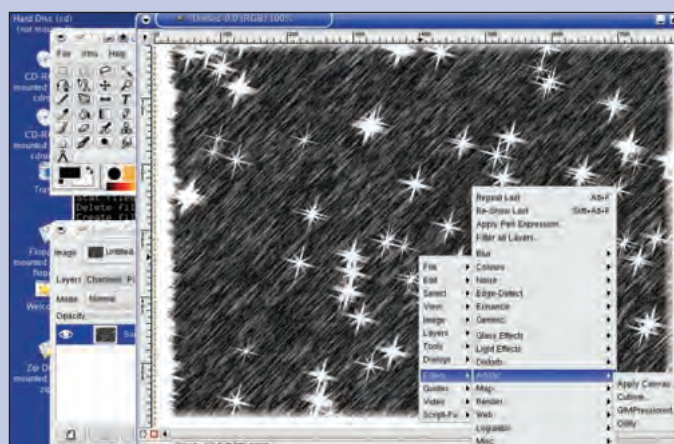
More efficient GIMP

Surely everybody who ever processes images with Linux must use *The GIMP* – there are alternatives, but they're not as capable or feature-mature. There are however, a few settings which you can optimise for your system running *GIMP*.

In *The GIMP's* Preferences Menu, there's an item called 'Environment' where the useful settings are. The first is 'Conserve memory usage'. This implies a performance hit, but paradoxically, if you are short on physical RAM, it will actually speed thing up by not relying on slow swap memory so much.

The 'Levels of undo' have memory implications too. The default setting is five, which is usually generous enough for most people. Reduce it if you badly need to save memory.

The tile cache is important, especially for large images. *GIMP* breaks up images into tiles for more efficient processing and memory usage. When the program



alters the image, it's actually only altering a small block of memory. Obviously, many operations span several of these tiles, and what's more, some operations may keep asking for the same tiles over and over. The tile cache can be

used to speed up operations here. The exact benefit depends on the size of the image and the operations you are performing, but if you do a lot of image processing, it's worth setting this to as large a value as you can get away with.

Grappling with your graphics card

Despite recent advances in 3D graphics performance on Linux, all too many people still use the Vesa graphics driver for X, which is a generic 'catch-all' driver that should work with most video cards. However, it's optimised for *no* video card, which means you get awful 3D performance, and acceptable-at-best 2D performance – hardly what you expect from your shiny GeForce FX!

However, never fear because help is at hand: there are specific drivers for most cards, many of which are well optimised. For example, if you're running an nVIDIA card, you're either crazy or a GPL purist if you're not running their GeForce drivers – they perform excellently in both X Windows and OpenGL games, and are easy to install on any Linux system.



documentation, as well as the *hdparm* man page, to see which options your system supports.

Choosing the right filesystem

Linux offers a number of filesystems for us to choose from, which gives us unprecedented flexibility, but also the burden of choice – which is better? The top three most used filesystems are *Ext2*, *Ext3*, and *ReiserFS*, with the latter two only available from kernel

using *Ext3* is that it is simply *Ext2* with a journal tacked on – if you decide that you want to go back to *Ext2*, either because you're unhappy with *Ext3* or because your system has crashed and you only have a 2.2 repair kernel, then switching back is a simple matter.

For the more adventurous Linux user, there are also the *XFS* and *JFS* filesystems. Of the two, *XFS* is much more stable at the time of writing, but at the end of the day both filesystems

are nothing like as popular as the other three.

RAID

RAID (Redundant Array of Independent Disks) is oddly enough, an array of disk drives which is treated by the system as a single drive. The main advantage of this is for fault tolerance – mirroring data across disks so that if an error occurs on one, the data can be retrieved from the copy. But RAID also allows striping – blocks of data are spread out across multiple devices, which acting independently, can read and write much faster than the bottleneck of a single drive. Simple striping is known as RAID Level 0. You can find more info on RAID in the June edition of *Linux Pro* which came with LXF41. [LXF](#)

“You can double the speed of your hard drive just by setting things up properly – newer systems see even larger gains.”

2.4 onwards. *Ext3* and *ReiserFS* are both *journalled* filesystems, which means they conduct filesystems writes with a high-level of atomicity – if your system crashes, a journalled filesystem is designed to be able to roll back to the previous safe state. *Ext2*, on the other hand, is not journalled – if you pull the plug from your system while it's switched on, you'll need to **fsck** on boot.

Journalled filesystems do have one disadvantage, though, and that's the fact that they need to do extra work to keep the journal up-to-date – if you're concerned about maximum performance above everything else, *Ext2* will be fastest for you. If you want a journalled filesystem while still having excellent performance, you'll find *ReiserFS* wins out – it's not as stable as *Ext2* yet, but it is known to have superior performance. One big advantage to

The joys of GCC

Get extra performance by compiling from source

While it's nice to have software precompiled and ready to install at a moments notice, it's easy to squeeze extra performance out your software by compiling it yourself – not only does this mean you get to choose how you configure it, which usually makes it more streamlined by cutting down on unnecessary bloat, but it also generates faster executables by taking full advantage of your CPU.

You have two options for compiling from source: downloading the source code yourself and compiling it as per normal, or you can use your package manager to compile source packages yourself. The former gives you more control in that you can hand-configure what goes into the finished product, but the latter is easier, and saves you dependency woes.

If you're using Debian, you can download and build source packages

using the following command:

```
apt-get -b source <package>
```

If you encounter errors while building that package because you don't have the necessary libraries and header files on your system, you can get Debian to find all the source dependencies for the package and install them for you, using **build-dep**:

```
apt-get build-dep <package>
```

For Red Hat-based systems, you will need to get the source packages. These are packaged up like other rpms, into a single file. When installed though, you actually get the source rather than the compiled binary. The packages are identified by having *.src.rpm* as the file extension.

Installing from a source rpm is a two-stage operation. First we have to build the source. Locate your package and use the following:

```
rpm -i mypackage2-2.1.4.src.rpm
```



BLUFF!

Using hyperthreading doubles the speed of your PC

People who upgrade their machines to use the newer Intel CPUs to take advantage of hyperthreading often do so in the belief that hyperthreading technology makes one CPU as good as two. Unfortunately, the reality is that hyperthreading often acts to slow your PC down more than it speeds it up. Furthermore, even under the best circumstances, hyperthreading won't make your PC anything like as fast as if it had two CPUs.



```
rpm -bb --target i686
/usr/src/RPM/SPECS/mypackage.spec
rpm -i /usr/src/RPM/RPMS/i686/
mypackage2-2.1.4.i686.rpm
```

The first line installs the source to the rpm directory of the system (by default in */usr/src/RPM*).

The second line actually builds the source, and puts it into a standard rpm package. You should adjust the target here for your particular system (eg *i586*, *i686*, *Athlon*).

The third line actually installs the binary rpm you have just created. You will find it in the */usr/src/RPM/RPMS/* directory, in the subdirectory indicated by the target you chose to build.

Fed up because your programs won't run because they're missing vital libraries? If you are using a package manager such as *urpmi*, you can use this for step one to make sure you get all the relevant dependencies.

What on Earth is... BSD?

Never one to be scared of daemons,
Michael Saunders investigates the world of
alternative Open Source **UNIX** variants...



>> Right, so I've been having trouble with this molar here...

No no, no – you want the *British Society of Dentists*. We're here to talk about *Berkeley Software Distribution*, a family of UNIX-flavoured operating systems which are in wide use and quietly gaining popularity, despite all the media attention that Linux receives at present.

>> Ah, understood. But what do you mean by a 'family' of OSes?

Well, just as 'Windows' refers to a bunch of similar but distinct systems (ie Windows 9x, NT, CE etc.) the same is true of BSD. They all share a common heritage, can generally run the same software and they mix and match ideas and code. They're more separate than Linux distributions though; the BSDs are individual projects with different leaders, developers and users (although there's plenty of overlap between them).

>> Cool, but why are they being covered in *Linux Format*?

Because they're free, Open Source, UNIX-like, stable and fun to explore. In other words, very similar to Linux in many ways – almost all of the software featured in this magazine will run on the BSDs, and many of our tutorials on programming, networking, graphics etc. are equally applicable to these OSes. Indeed, many BSD users buy Linux mags and read Linux news sites because most of the content is about 'free UNIX' rather than GNU/Linux *per se*.

>> So they're UNIXish, open source, and run most Linux software?

Yep. But again, what is commonly described as 'Linux' software is really 'open source UNIX software'; providing the underlying operating system supports many of the basic standards of POSIX and the like, they should have no trouble. For example, KDE advertises itself as a 'Free UNIX Desktop' rather than being Linux-specific, yet the vast majority of KDE users are running Linux.

But why choose BSD over Linux?

Nutshellised: choice and features. The same reasons some prefer Mandrake over SuSE, Debian over Red Hat etc. It's about choice and as we'll see, the BSDs offer some significant differences from Linux – in some areas they're better, in some areas they're lacking, but they all have unique benefits which make them well worth investigating.

» Ah. How did BSD come about, then?

It's all part of UNIX's long and entertaining history. The American company AT&T developed the first versions of UNIX in the early 1970s and licensed out the code. One group particularly interested in enhancing it was the University of California, Berkeley – they added various improvements but eventually AT&T realised that money could be made from the OS, and wanted a 'pure' product of their own. System V UNIX was born, but there was still a problem that AT&T code remained in the 'Berkeley Software Distribution' releases.

After some legal action in the early 1990s to resolve this, UC Berkeley rewrote the offending code and finally had a clean OS. However, the computer market had exploded since UNIX's early days, and a lot of people were interested in getting this BSD code onto modern commodity hardware. William Jolitz began a project called 386BSD to do just this; sadly he wasn't very keen on outside participation and two projects, NetBSD and FreeBSD, decided to continue his work more openly.

That's the basic history condensed down, but it's interesting how Linux popped up when BSD was also presented as a UNIX flavour on PC hardware. There's no doubt that the legal proceedings hampered development somewhat while the new GPLed Linux appeared totally safe, but Linux was also new and 'sexy' with loads to hack on and improve, while BSD seemed 'old-school' and more closed to outsiders.

» Intriguing. So how many BSDs are there now?

Quite a few, but instead of eating up the whole magazine we'll just focus on the main three, the ones closest to Linux in terms of openness, community and software support. These are FreeBSD (on the *LXF* coverdisc), NetBSD and OpenBSD. We'll also have a quick peek at some of the more distant derivatives, but it's this trio that are the big boys of BSD land.

>> So what's FreeBSD? I guess from the name that it's free...

Don't look too deeply into the names – they're all free, they're all developed on the Net and they're all open. FreeBSD (www.freebsd.org) is the most popular of the big three, and is the also most similar to Linux in having a strong x86 slant and focusing on desktops and servers. With a pleasant text-menu

installer program and stacks of superb documentation, FreeBSD is definitely the most approachable and its Ports collection (ready-packaged software, see later) contains an enormous range of apps and tools.

Hardware support for PC devices is quite broad, although not as complete as Linux at present, and the SMP (multiprocessor) code isn't quite as mature either. On the other hand, the kernel and drivers are immensely robust – BSD servers often reach years of uptime. GNOME, KDE, *Mozilla*, *OpenOffice.org* and thousands of other apps run perfectly, and for commercial software there's an impressive Linux 'emulation' layer (which remaps system calls).

And it's ultra rock-hard?

Great emphasis is placed on reliability and performance, and the FreeBSD project runs two separate development trees – STABLE and CURRENT. In a way, this is similar to kernels 2.4 and 2.5, and allows desktop users and developers to try out all the latest bleeding-edge goodies while having a solid release to fall back on for mission-critical boxes.

FreeBSD has gathered a sizable community since the first release in 1993; it's always very difficult to gauge the userbase of any Open Source project, but a recent guesstimate put it at around 2 million users. Not bad at all!

OK, what's the deal with NetBSD?

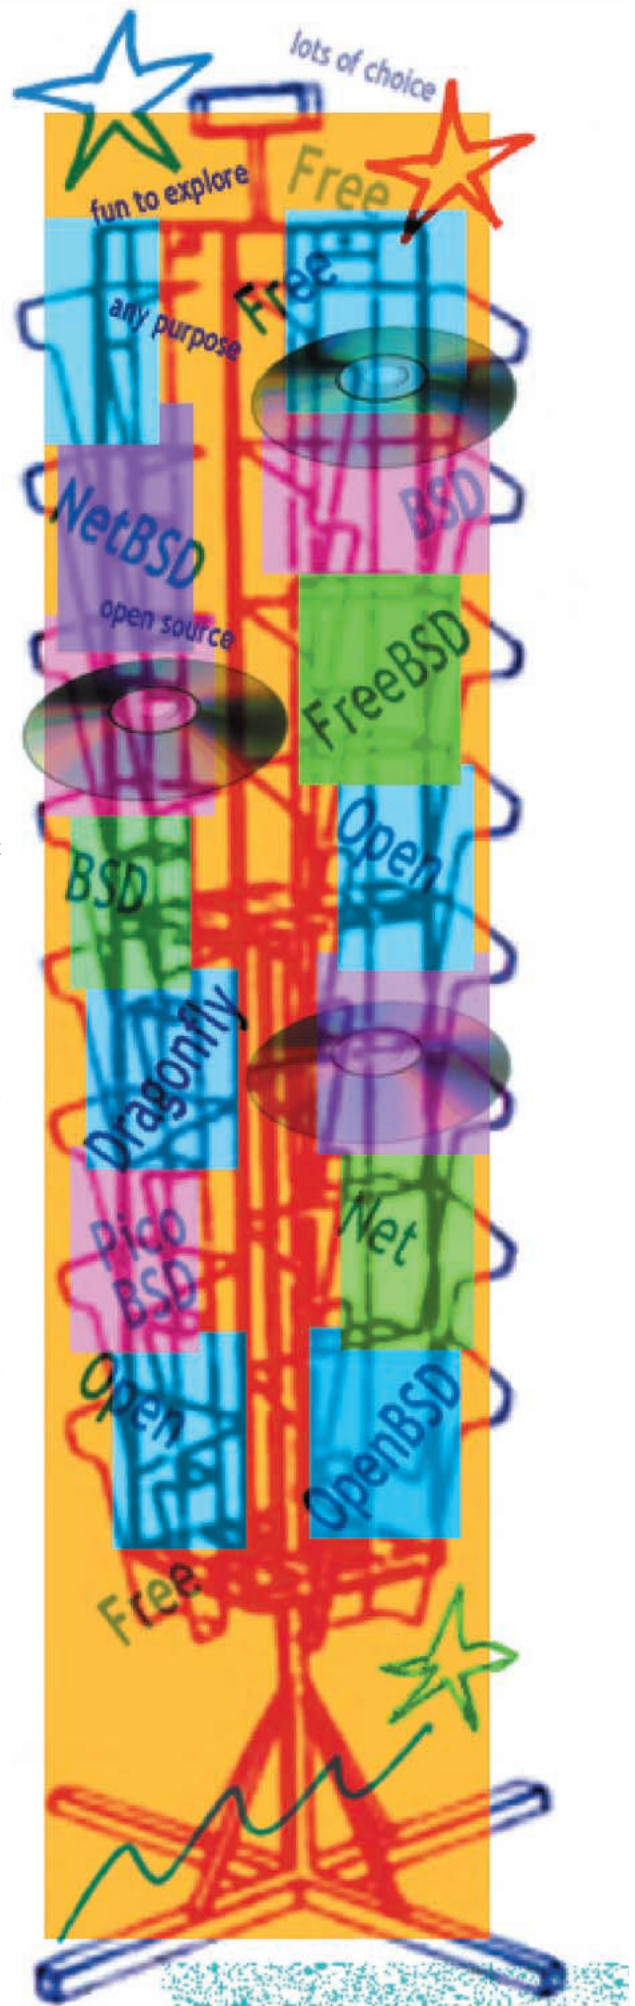
One word: portability. While Linux and FreeBSD can run on various platforms, x86 gets the most attention for those OSes. Not here though – as far as NetBSD (www.netbsd.org) developers are concerned, all platforms are equal and the staggeringly huge range of architectures that NetBSD runs on is second to none – PC, Mac, Amiga, Atari, Acorn, Sun, SGI and VAX are just a handful of the most popular, with development progressing on some truly obscure platforms.

Since NetBSD 0.8 (the first version) appeared in 1993, the OS has become a favourite of researchers thanks to its clean and portable codebase. It even has a commercial company backing it, www.wasabisystems.com, which does embedded work based on the OS.

Wowzers. How unified is the support for all those platforms?

Though Linux can claim support for many of the above architectures, it tends to be largely fragmented with different kernel trees per CPU and individual distros to suit each platform. NetBSD, in contrast, builds for all of these straight from the same source tree – there are no branches or forks or separate distros. Only Debian attempts the same wide support and a consistent environment across platforms, but it's still some way off.

NetBSD now has a rudimentary menu-driven text-based installer, but it's a doddle to operate and doesn't try to be too ambitious or intrusive. On **»**



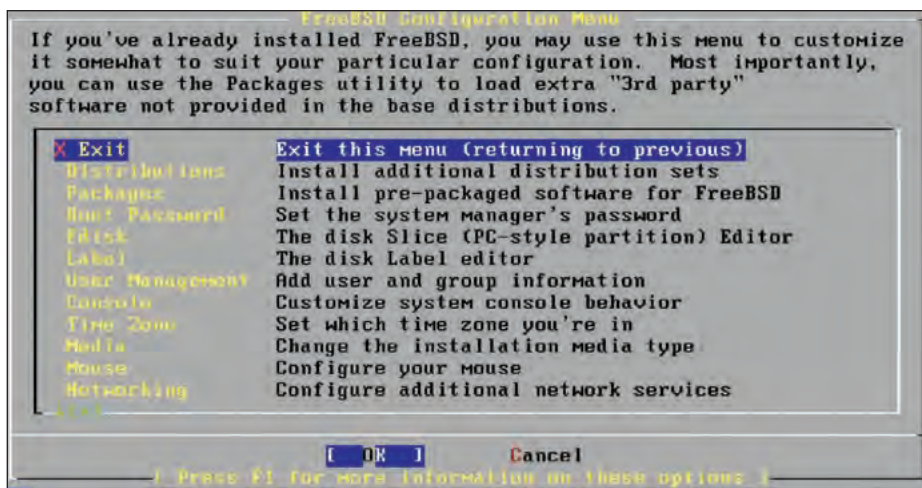
WhatOnEarthBSD

◀ the software front, all the biggies like KDE, GNOME, Mozilla etc. can be found in the *Pkgsrc* bundle (same as FreeBSD's Ports system, but the word 'Port' is used for a hardware platform in NetBSD).

➤ And finally OpenBSD?

OpenBSD forked from NetBSD in 1995, when programmer Theo de Raadt had a major conflict with the politically troubled project leaders. Here, the main objective is security. Flashy new features and desktop frills are all good and well on a home box, but for a firewall or server it's imperative that some script kiddie can't easily gain ownership of your system. OpenBSD's developers are constantly auditing their codebase for any potential security flaws; consequently, in terms of hardware support and the latest snazzy features, the OS isn't cutting-edge but it's a fine trade-off.

Still, OpenBSD has pioneered various developments which are becoming commonplace in other BSDs and Linux distros – the code audits reveal bugs which others can fix in their own code, and they've done some excellent work on privilege separation and 'jailing' processes (i.e. making sure something doesn't constantly have root privileges, and even if security is breached, the attacker has only a small subset of the system to play with).



FreeBSD's solid text-based installer and config tool will soon be replaced by graphical-equivalent *libb*.

➤ Lovely stuff.

Absolutely. By far the biggest contribution to other OSes, though, is OpenSSH – the free replacement to the original Secure Shell and a much better alternative to the plain-text Telnet and friends. Virtually every Linux distro now includes this useful tool for logging-in to remote machines. OpenBSD's installer is a purely command-driven text-based affair, and while it's the hardest of the

three to work with, the online assistance is smart. All things considered, OpenBSD is supreme at being a small server or firewall system; there's no SMP support at the moment which is a problem on larger servers, and even though it has a decent Ports collection and solid variety of apps, that's not where the focus is.

➤ What about the others?

Indubitably, the BSD-variant which gets the most desktop exposure is Darwin, the core of Apple's Mac OS X (developer.apple.com/darwin/), although the Mach kernel and proprietary GUI system make it more of a distant relative. There's also BSD/OS, a commercial system produced by BSDi (now Wind River) which was once popular with ISPs but appears to be slowly on its way out. Dragonfly BSD is a relatively new project from highly-regarded Amiga and FreeBSD coder Matt Dillon, while PicoBSD is a single-floppy version of FreeBSD.

➤ Why don't they merge?

Why don't all the various different Linux distributions merge?

➤ Touché...

The goals are just too dissimilar. There's been some talk of merging bits of the userland – nobody quarrels that the kernel, filesystem layout and admin tools need to be separate, but maintaining individual trees for *ls*, *cp* and *chums* is a touch wasteful. These three BSDs all fold in the external GCC and XFree86 projects so there's no pointless duplication of effort there.

➤ You said the BSDs are developed over the Internet - like Linux?

Like Linux in that the projects communicate and share code through the Net, but they're generally more centralised. CVS is used for managing the code and maintaining branches – competent developers are given CVS commit access after proving themselves. FreeBSD and NetBSD both



have 'core' teams of developers for making important decisions and directing the project, while OpenBSD is closer to Linux in having a single figurehead (de Raadt) in charge.

»» So what exactly is this 'Ports' system you mentioned?

Essentially, Ports is a software packaging and distribution system, with a heavy emphasis on making full use of the source code. Whereas Red Hat, Mandrake, SuSE and Debian etc. use binary archives to install software, Ports acts as a mechanism for downloading, patching, compiling and installing apps. Its basis is a directory tree of package details, in which the user navigates to a program (eg `/usr/ports/graphics/gimp`) and issues a **make install**. From there, the original source is downloaded, patched with any BSD-specific changes and extras, and then built.

It might sound complicated, but it works transparently and because the source code is stored remotely (usually on a project's own site) no disk space is wasted. Naturally, it also handles dependencies well – if an app needs GTK, for example, it'll also download and compile that too. A top-notch and mature system. In fact, the Gentoo distro is enjoying great success by emulating Ports.

»» Doesn't sound ideal on older hardware though...

The BSDs also support binary packages, which are simply pre-built Ports and can be installed with **pkg_add**. It works very similar to **rpm** or **dpkg**, and despite the loss of fine-tuning and optimisations that come from building locally, it's a godsend on low-spec boxes or where time is limited. Thankfully, the database that stores details of packages is aware of all this, so you can easily mix-and-match Ports and pre-built packages.

»» Why does some software need patching?

Because of Linux's enormous popularity, most open source software (especially on the desktop side) is developed on Linux machines, and as a result the occasional 'Linuxism' creeps in. These are basically assumptions in the code that it'll be running on a Linux box – device names, system calls etc. So the BSD packagers sometimes have to make a few minor changes via patches, but as said, it all works without intervention.

»» How do BSD developers feel about Linux's popularity?

Of course, they'd rather it was their own OS making all the headlines (well, apart from the SCO ones) but at the same time Linux provides a constant source of new users. Anyone who removes the Windows ball-and-chain and decides to try out alternative OSes is another potential BSD user, picking up basic UNIX skills along the way. A large chunk of current

BSD fans started off with the easy-going Linux distros of Mandrake/SuSE and the like, moved on to the more demanding Debian/Slackware/LFS, and ended up on a BSD.

Unfortunately for the BSDs, popularity *can* be a big help sometimes. Acquiring specs (to make drivers) from frigid companies grows easier, and the involvement of large companies like IBM has helped recent Linux kernels to scale-up on massive high-end hardware. In all, the BSD communities feel more like Linux circa 1998 – more hackerish and with just a drop of corporate attention.

»» So what license is all this code under? Is it GPLed?

Nope – the BSDs use the 'BSD license', which is quite different from the GNU General Public License that the Linux kernel and surrounding tools use. The GPL demands that any modifications to a GPLed program are also made open if distributed; conversely, the BSD license lets you do just about anything with the code providing you acknowledge the copyright and don't hold the developers responsible if things screw up.

»» That means any old company can just take the code!

That's right, but remember: you can't steal something that's already free. MegaHyperCorp could release their own closed BSD variant tomorrow, but the original source would still be around. Jordan Hubbard of FreeBSD said that their goal is to 'provide software that may be used for any purpose and without strings attached'. Hugely altruistic then, and in the end it leads to wider use of good-quality code.

As you can imagine, there are regular flamewars between BSD and GPL supporters, but all things considered it's good to have two major licenses. The former is perfect if you just want to get some code out into regular use; the latter is more appropriate if you want to guarantee long-term freedom and you're taking a more philosophical approach. Most Linux distros contain a mix of licenses anyway – XFree86 is under a BSD-like license.

»» Are the BSDs better than Linux?

Electric Nauruan gerbil farmers indulging in a toasted celebration of vintage spatula grooming and anti-Nietzschean pole-vaulting!

»» ¿Que?

Ask a silly question... Well, alright, it *could* be a worthwhile (yet flamebait) query from someone new to computing, but as Linux users we're already aware that one OS doesn't fit all purposes. While we know that Linux is stronger than Windows in the crucial areas of stability, security and freedom, at the same time we accept that Microsoft's OS has greater application support and is sometimes easier to use. Similarly, modern Linux distros are friendlier and run more commercial apps than the BSDs, but



that doesn't make the BSDs any less useful. And that's the main thing – the BSDs do their job brilliantly. For a general-purpose x86 server or desktop, FreeBSD excels. If you're looking for a clean UNIX-like to bring an old Amiga back to life, NetBSD is ready and waiting. And if you need an ultra-secure firewall/router or small server on low-end kit, OpenBSD is a very wise choice.


So in terms of which is 'better', it's down to your needs. Certainly, Linux is a more appropriate choice for anyone new to UNIX-like operating systems and the hardware support is considerably larger (especially on x86 PCs). The increasing range of commercial apps is another factor; still, FreeBSD's Linux-emulation code helps out greatly there.

»» Right then, I'd love to give this a go. Where do I start?

This very issue of LXF has FreeBSD on the coverdisc – so it couldn't be easier. If you've installed Debian, Slackware or Gentoo before, you'll know how to deal with problems from the raw command line and should have no trouble – just find some disk space, boot from the CD (or write a floppy) and enjoy. Most commands, file locations and general practices are identical or similar.

Check out `INSTALL.TXT` on the CD first to get an overview of the installation process. If you're a recent Linux convert it's a good idea to have someone familiar with BSD to help out if things get hairy, either in person or online. Some Linux User Groups have been running BSD install-fests, so they might be worth checking out too.

»» What should I do if something goes wrong?

Your first port of call should be the online mailing lists for the BSD you've chosen; find the one for newcomers or basic issues and ask for assistance. The lists tend to be busy though, so it's always best to give an in-depth description of your problem (including OS version, hardware, what you've tried etc.) in order to get the best results. Also try out our online forums. Have fun! 

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 := not
end;
Usually, you'll find the code on
our CD/DVD too.
```

THIS MONTH TEACH YOURSELF...

Beginners' tutorial

We covered KDE last issue, so to be balanced we look at the GNOME desktop this time **p64**

Red Hat Package management

PART 2: Alternatives to Red Hat Network & Packaging Tool **p68**

Compiler writing

Imported from *Linux Pro* – how to use *flex* and *bison* to write our own SKYLang language **p72**

Blender >>

How to use Bezier curves and other advanced modelling features to make more complex scenes **p76**

The GIMP

Multimedia interfaces – better known as 'skins' – and full instructions for building your own attractive viewer **p80**



ILLUSTRATION BY PHILIP – www.blender3d.org

PHP

Optimising your MySQL queries can accelerate the code by applying some simple rules **p84**

Server School

NEW SERIES! An introduction to the ideas and terminology behind running your own Linux server **p88**

TIP OF THE MONTH!

DIRECTORY BROWSING

You should already know how to use **cd** to browse around your filesystem, but there are four special commands that are particularly key.

cd ~

you should already know – it changes directory to your home, just like using **cd** with no parameters.

Slightly less used, although nonetheless as important, is

cd -

which changes directory to your previous working directory. That is, if you were in `/home/paul/sandbox` and you change to `/etc/httpd`, then type **cd -** you'll be placed back into `/home/paul/sandbox`.

The two other special commands that are even more rarely used, are:

pushd

and

popd

These two complement each other, by allowing you to build up a stack of working directories. Rather than using **cd** to change to a directory, you can use **pushd** to change to the directory, and push that directory onto its internal stack. You can browse around all you like, changing directory with **pushd** (or **cd** if you don't want to use the stack), and each time you'll add more to the stack.

To go back to the directory you were previously in, type **popd**. This takes the directory at the top of the **pushd** stack, removes it, then changes directory to it. If you type it again, the procedure is repeated, allowing you to work your way back through your directory stack.

Each time you push or pop something using **pushd** and **popd**, your shell should print out a list of the current stack so that you know what you have to work with. You can get this list up at any time by using the **dirs** command, and you can also pass **dirs** the parameter **-l** to get it to return absolute directories for the stack [LXF](http://www.linuxformat.co.uk)

DESKTOP ENVIRONMENT

Beginners' Guide to Linux – GNOME

PART 4 Linux is all about choice, so it's not surprising that there is more than one desktop environment available. Let's take a look at GNOME, a project formed due to fears about the openness of KDE's original license.

An ancient Chinese proverb says that it's better to light a candle than curse the darkness. So if KDE isn't to your liking, don't let it get you down; give GNOME a spin. Like KDE, GNOME provides a complete desktop environment that can be tailored to your requirements; unlike KDE, it has followed a defined set of Human Interface Guidelines almost from the beginning, so many users find the experience of using the environment more intuitive and refined. The GNOME project is also less determined to cater for every user's needs and so, for instance, it doesn't have a fully integrated office suite or portfolio of communication applications – those functions are covered by *OpenOffice.org*, *Evolution* and *Galeon*.

FIRST STEPS

Once you've logged in to your GNOME system, you should be presented with a screen similar to **Fig1** below. You'll notice it follows the same conventions as most other OS shells such as Windows and KDE, with a taskbar ranged along the bottom of the screen and icons on the desktop.

A. The Desktop

The default desktop is a rather pretty blue gradient, but you can remove this or set a picture as the background by right-clicking

Fig1 Default appearance of the GNOME desktop on logging in.

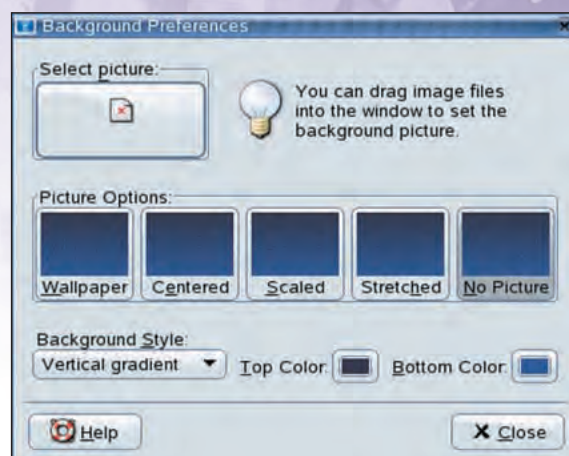
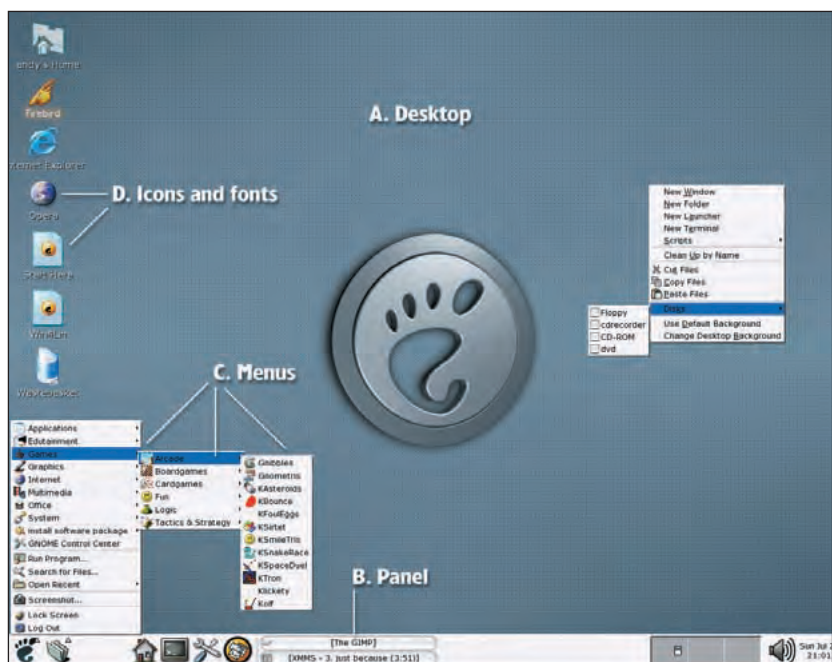


Fig2 Whether you want wallpaper or a power-friendly dark-coloured desktop, GNOME is easily configurable.

anywhere on the desktop and selecting 'Change Desktop Background'. You can also get to this dialog via the GNOME Control Panel as can be seen in **Fig2**. By default changes made here will be reflected across all of your virtual desktops.

The simplest option (and one that is especially useful for notebook PCs with limited battery life) is to have a single colour background. You can select this by clicking on the Background Style button and selecting 'Solid color'. You can then select the hue that suits you best. If you are using a notebook, choosing darker colours (as well as being easier on the eye) will improve battery life. You can, of course, keep the vertical gradient (selecting start and end colours as appropriate) or choose to have a horizontal one.

The more popular option is to put some Wallpaper (family photos, art etc) on the desktop. This not only personalises the computer for you, but can also provide some visual distraction for tired eyes if you spend long hours at the computer. To set a background image you can either hit the 'Select Image' button and navigate to the required picture, or drag and drop the image from the file manager (*Nautilus*) or desktop into the 'Select Image' space.

This latter method is a great way of demonstrating GNOME's dynamic update features; as you drop the image in, the background is changed immediately. There are no 'Apply' buttons to click, and no wait to see the effect of your alterations, and this is reflected in all of GNOME's configuration options. There are a number of options for images: they can be either Wallpaper (tiled), Centred, Scaled or Stretched.

TutorialBeginner'sLinux

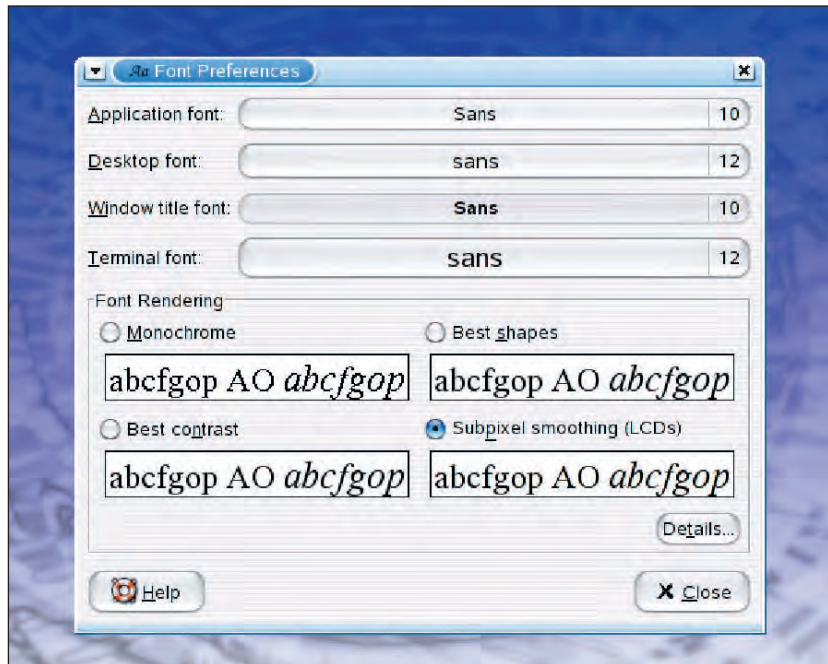


Fig7 Alter font settings to maximise readability for users of both CRT and TFT monitors.

while laptop and TFT screen users will get a better view with the 'Subpixel smoothing' option. Further refinement can be made by hitting the 'Details' button, but for most users the basics will suffice.

GNOME improvement

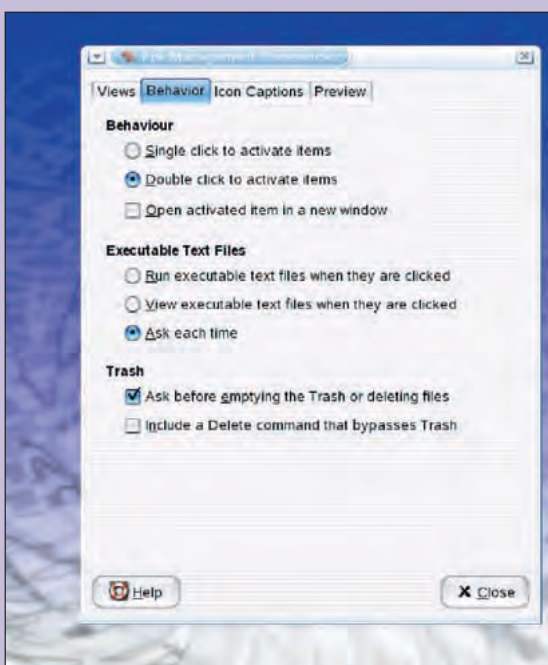
GNOME has a reputation for good looks, a reputation the developers are keen to build on with, amongst other innovations, support for alpha-blended desktop elements and SVG icons (which can be scaled to any size without loss of quality or

Tip!

Single-click or double-click?

If you're used to KDE's single click access to documents and applications (or your fingers just aren't as fast as they used to be), you may have looked in vain at GNOME's mouse menu for an option to change the default double-click. Well, as desktop behaviour is part of *Nautilus*, you need to open the file manager and select 'Edit > Preferences' and find the mouse behaviour options under, you've guessed it, the 'Behaviour' tab. You can set single-click access in here, or change the speed at which double-clicks are noticed, and also change the mouse for left-handed operation.

Fig8 Change GNOME's behaviour with this dialog box for a more comfy user experience.



increased file size). Fortunately these prettifying efforts are not undertaken at the expense of usability.

Almost all of GNOME's changeable elements can be accessed via the *Control Center*, and the most visible aspects of the environment can be changed in the Theme dialog. The three elements of a theme are Controls, which define how widgets such as radio buttons and scroll bars look, window borders and icons. Themes are stored in .tar.gz format and can be either installed via the browse menu or by dragging and dropping from *Nautilus*. Once you've tweaked and teased the desktop to your liking, you can hit the 'Save theme' button and give your theme a name. This approach positively encourages experimentation, so you can download extra icons, borders etc. and, in the tradition of Linux development, mix and match to your heart's content.

One other, often overlooked, part of the GNOME *Control Center* is the Window Preferences dialog. Here you can set windows to become active as the mouse passes over them (ie they become focussed without having to click on them) and have them to spring to the front after a defined time, should this suit your style of working. You'll also find in here an antidote to twin problems of window title bars falling off the top of the screen, so you can't reposition them, and the 'Windows' key on most keyboards being useless.

Storage solution

Where KDE has *Konqueror*, GNOME has *Nautilus*; a file manager with a fine pedigree in both design and functionality. *Nautilus* follows the now standard two-paneled model with the larger space on the right giving access to files and folders and, on occasion, displaying web pages or file previews. The thin pane on the left, which can be removed or reinstated by hitting F9, can display information about the currently selected folder/file, file system tree, browsing history, drag and drop emblems or a short note. This last can be very useful if, for instance, you are pulling together a CD of images (all helpfully called things like dcsf2024.jpg) for printing, as is shown in **Fig10**. Just make a note of the file name and the content and move on to the next folder. The great thing is that your notes are pervasive and will stay attached to that folder until you delete them. You can swiftly access these notes without navigating to every folder by right-clicking a folder and selecting 'Properties'. Look under the 'Notes' tab to view or edit your scribbings.

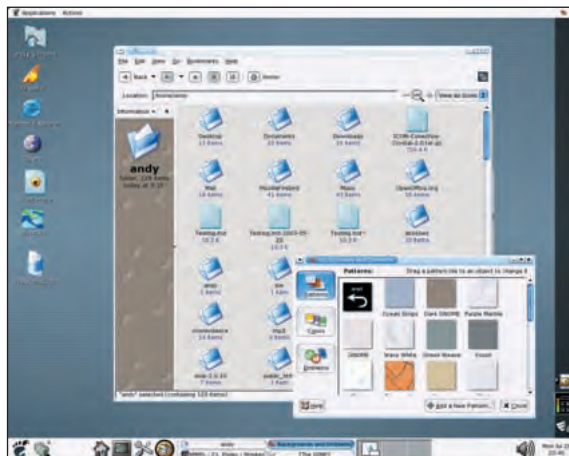


Fig9 Create your own pattern themes on the desktop with an easy drag-and-drop.

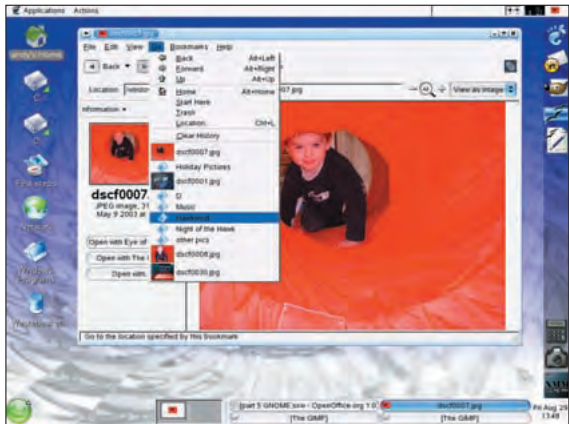


Fig10 **Nautilus** displays info about the recently selected files in the 'Go' menu.

If you regard every control surface of your PC as a blank canvas to be personalised, you can select 'Edit > Backgrounds and Emblems' to get ready access to a selection of drag-and-drop background patterns and colours that beautify *Nautilus*'s panes, like the example in Fig9. You can also drag these patterns onto any of your panels (except the Menu Panel) that have 'Pattern' set as the background.

Unique feature

The other element of this dialog, which can also be accessed via the Side Pane (hit F9 if you can't see it), is the unique Emblems feature. This allows you to associate extra icons to files and folders to denote their content or status. For instance, you may have a Family Photos folder, which would be important, so you could drag and drop a 'camera' icon and an exclamation point onto it. You can add as many as four emblems, and they can be removed or changed by right-clicking to get into the properties dialog. Hit the 'Emblems' tag and select or deselect the pictures you want displayed. The great thing about emblems is that they're just standard .png files, so if your needs are not met by the small selection of included ones, you can hit the net and download some more or even make your own!

Hey, good lookin'!

Having lots of photos and MP3s littering your hard disk is quite a common occurrence, so the file manager is increasingly becoming the centre of users' computer universe, and media previewing can no longer be regarded as a nice addition – it is essential.

If you've opened a directory of photos in *Nautilus* only to be faced with a long list of file names, do 'Edit > Preferences' and click on the 'Preview' tab. Here you can set previews for text, image and audio files. You can choose from a number of options, including previewing local files only and setting a size limit on previews – by default, this latter should be 3MB.

Having done this, you should now see your images as thumbnails, which can be scaled with the magnifying glass icon on the toolbar. You can right click an image for a context sensitive menu which reveals the usual cutting/pasting options as well as an 'Open with...' entry that offers one or two applications that can be used to edit or view an image. If you need a closer look at the picture, double clicking will open the full image in *Nautilus*'s right pane, while the thumbnail moves to the Sidebar, along with the 'Open with...' options so you can, for example, open the picture up in *The GIMP* for editing.

What on Earth are... Permissions?

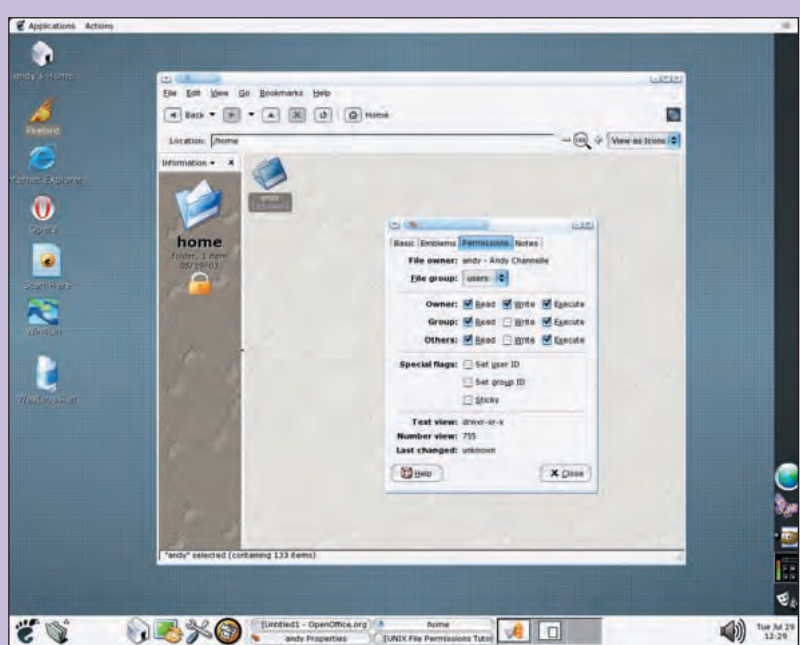


Fig11 Defining the permissions of root and other users doesn't have to be a command-line process.

YOU'RE GOING TO START REELING OFF LOADS OF UNIX JARGON AREN'T YOU?

Not at all! Permissions, at the most basic level, is quite a simple subject and though it can seem complex, we'll keep things at rudimentary level here. When you installed your distro, you would have set a 'root' password and also created one or more users. Permissions merely define what these users can and cannot do on the computer and/or network.

FOR EXAMPLE?

Andy's documents are kept in a subdirectory of /home called /andy. If he navigates his way to /home, right-clicks on /andy and selects 'Properties > Permissions', he ends up with the dialog shown in Fig11. From this you can see that the folder in question is 'owned' by Andy and is part of the 'Users' group. All other users on this machine would belong to that group. Below the owner details is the part of the system which defines who can do what. As the owner, Andy can read, write or execute files within the folder – that is, he has full access. Other users (and other groups) can read documents and execute files, but can't write anything back unless Andy – or root – changes the definition.

WHAT DOES THAT MEAN IN PRACTICAL TERMS?

It means that Andy's files, by default, can be viewed but not changed by other users of the machine. But it doesn't stop there. Navigate up through the file system until you see a folder

called /bin. Open up the properties on this and select the Permissions tab. If your system is set up correctly, this should be owned by 'root' and the group should also be 'root'. Now open a new file manager window (File > New Window), go to your home directory and try to drag and drop a file into /bin.

IT DOESN'T LET ME DO IT.


No, because though you can scour through the directory and execute applications from within, you can't put anything in there. Not a sound file, image, programme, virus or trojan...

AH, IS THIS WHY LINUX HAS A REPUTATION FOR BEING SECURE?

Exactly! A virus or other malicious program (or user!) would have to navigate around your root password to get access to the guts of the system and do any damage. The really important point about all this is: don't run as root for your day-to-day tasks (and don't give your root password to anyone).

WHAT IS ALL THE OTHER INFORMATION IN THE PERMISSIONS DIALOG?

There are various ways you can define permissions, including using text and a numerical system, but they really need a full tutorial to explain. If you want to know lots more about permissions, point your web browser at www.dartmouth.edu/~rc/help/faq/permissions.html for a pretty comprehensive tutorial on the subject.

The only vital feature missing from *Nautilus* is a slide show, but this can be taken care of with a third party app such as *gThumb*, which is part of most standard GNOME distributions. Finally, if you need rapid access to your recent file browsing history and can't be bothered to use the Sidebar, just look under the 'Go' menu entry for a nice visual map. 

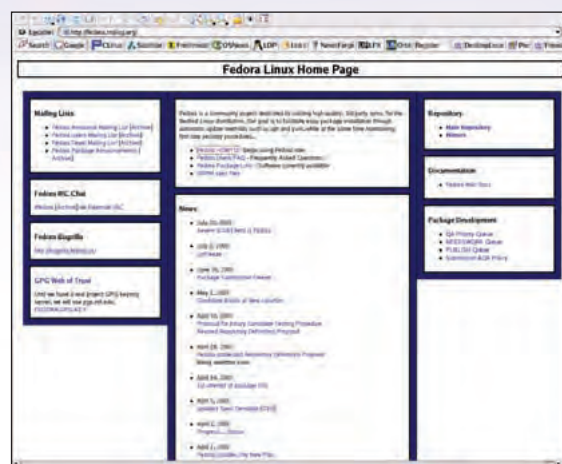
PACKAGE MANAGEMENT TOOLS

Red Hat Alternative Package Management

PART 2 There are many package repository choices for the Linux user willing to acquire some extra knowledge. **Hoyt Duff** examines a few alternatives to the Red Hat Network & the Red Hat Packaging Tool.

Last issue, we looked at the default Red Hat package management system and found that it offered a package installer and uninstaller that restricted your choices of packages to what Red Hat had provided on the two or three official CDs. Updates were provided via the Red Hat network over a client-server model. While this makes great sense for Red Hat (which sells support for this scheme) and sysadmins of large systems (whose employers purchase the support), there are other viable methods of accomplishing the same thing and providing some flexibility.

Perusing APT, URPMI, and YUM, we suggested that you consider using the version of APT modified by Connectiva Linux for that particular distribution because it works with RPM files and provides an automatic way of updating or installing applications from a repository located locally or remotely. A version of APT and supporting repositories are provided by two groups: the Fedora Project and FreshRPMs.Net, shown on the right and below. Because of versioning problems, you should not mix RPMs from both repositories. Either source is excellent, but we selected



The Fedora Project Homepage has many contributors, and a lot of the site focuses on describing the process of the project.

FreshRPMs.Net because of the amount of documentation its site provides and, being controlled by one person, some packages seem to get published a little quicker – it's a big job no matter how many people are involved. As with all things, you should examine all the alternatives before you make a final decision.

This month, we will show you how to set up your own repository for APT and YUM, and show you how to populate them with files to serve over your local LAN via ftp or http. This is useful if you have several machines to admin or will be installing locally. You can also make custom packages available to your machines (or share with everyone) with a properly configured repository. Building custom RPMs is not that difficult once you learn a few tricks, but regrettably it's beyond the scope of this article.

Stocking the Store

To populate our ftp tree, we need to locate an ftp server from www.redhat.com/download/mirror.html, the official list of Red Hat mirrors. We'll use the updates in our example, and looks like ftp.mirror.ac.uk provides the basic os RPMs as well as updates



Excellent documentation is available from the <http://freshrpms.net> homepage.

Finding an rsync Server

Using mirrors makes sense – and is polite!

An operating *rsync* server will provide ‘modules’, ie a directory tree as a symbolic name. To determine if a server provides an *rsync* modules, use the query:

```
rsync fully.qualified.domainandhostname::
and, if running, you'll see a list of available rsync modules. If we try:
$ rsync ftp.dulug.duke.edu::
```

```
mirror.dulug.duke.edu
- This rsync server is currently available to any/all people.
- This is subject to change with little or limited notice.
- We may in the not-so-distant-future restrict rsync access
  to tier2 red hat mirrors.
```

Modules:

```
archive      Everything
redhat-ftp    Red Hat FTP Site
redhat-base   Red Hat FTP Site
redhat-beta   Red Hat Linux beta releases
redhat-rawhide Rawhide FTP Site
redhat-updates Updates FTP Site
redhat-contrib Red Hat, Inc. – Contrib FTP Site
```

and we see that Duke University provides *rsync* services for its Red Hat directories. You may continue to add the modules name to the end of the command until it returns no new information and then use the *-r* command for a recursive listing and further exploration.

We at LXF don't want to cause problems for the Duke University LUG (sponsors of this particular ftp site), so we encourage you to find another *rsync* server by checking Red Hat mirrors at <https://www.redhat.com/download/mirror.html>.

At this server, we further explore using an ftp client (or the *-r* command for *rsync*) and discover that the files we want are to be found at <ftp://ftp.dulug.duke.edu/pub/redhat/linux/current/en/os/i386>, a common Red Hat ftp directory structure. Look for Mandrake sites at www.linux-mandrake.com/en/cooker-devel.php3

and contribs. Conveniently, it provides an *rsync* server as well. Our choice of local directory structure is explained in the *Create A Local Repository* section.

There are several applications that are useful to mirror the RPM files from a public repository or ftp site to your local server. Each has its strengths and weaknesses, but all will eventually get the job done. Our goal is to automate the process as much as possible once we get the wrinkles ironed out. Let's examine a few of the available tools.

CURL

The *curl* application is a client that excels at retrieving specific files from http, ftp, gopher, LDAP, telnet servers, offering proxy support and dealing with authentication, cookies, and resuming interrupted transfers. It can handle multiple URLs and can download a group (such as 1-100 or a-z) of serially named files. To download all three Red Hat ISO images, we use:

```
# curl \
ftp://ftp.mirror.ac.uk/sites/ftp.redhat.com/pub/redhat/linux \
/9/en/iso/i386/shrike-i386-disc[1-3].iso
```

WGET

The *wget* application is a client that downloads files from ftp or http sites. Not only data files can be *wget*-ed, but entire websites can be recursively downloaded and downloads can be resumed. To download the first ISO image, we use:

```
# wget \
```

```
ftp://ftp.mirror.ac.uk/sites/ftp.redhat.com/pub/redhat/linux \
/9/en/iso/i386/shrike-i386-disc1.iso
```

While each of these two tools have their uses (both are used by Mandrake's URPMI for example), others may be better suited.

FMIRROR

The *fmirror* application is a more robust tool for our purposes as it is designed for mirroring ftp directories. It can use regular expression matching for files that are to be included and excluded; it supports authentication and passive mode transfers. For repetitive use, a configuration file (used with the *-f* option) is essential. Since we need to use the full ftp URL with *fmirror*, a little digging around on the ftp server is required to discover the correct path; we used the command-line ftp client and the *dir* and *cd* commands to locate the correct directory.

We also copied the */usr/share/fmirror/generic.conf* file to */etc/fmirror.redhat_updates.conf* and edited it to look like this:

```
# /etc/fmirror.redhat_updates.conf
# All lines starting with a non-letter is considered a comment.
# Also blank lines are treated as comments.

# Fill in with useful values
username: anonymous
password: lfxwriter@hotmail.com
host: ftp.mirror.ac.uk
remotedir: /sites/ftp.redhat.com/pub/redhat/linux/updates/9/
en/os/i386
localdir: /var/ftp/pub/apt/9/en/os/i386/RPMS.updates

# These excludes rules should normally be first, they will
exclude the
# most common junk and uninteresting files. They also ensure
# that you
# don't mirror any temporary files created by fmirror or the
# perl mirror package.
exclude: f ^(\.in\.\.mirror|core$|MIRROR\.LOG|
\.notar\.\.message)
exclude: f ^(\.cache\.\.zipped|lost\+found|Network
Thrash Folder)

# Put your own includes and excludes here
# END
```

You would use a separate config file for each ftp site you target and specify that file with the *-f* command option. Using *fmirror* to download the updates, we use:

```
# fmirror -f /etc/fmirror.redhat_updates.conf
```

RSYNC

See the *What On Earth* feature in LXF44 for an in-depth explanation of the ins and outs of *rsync*. The *rsync* application is an enhanced, secure client/server replacement for *rcp* (remote copy) that uses its protocol to speed up transfers by only sending differences between to files when the file exists locally and remotely. It also supports copying links, devices, owners, groups and permissions. We like to use *rsync* to repair damaged copies of downloaded ISO images rather than re-download the entire 650MB. Using *rsync* to download the updates, we use:

```
# rsync --recursive \
ftp.mirror.ac.uk::updates.redhat.com/9/en/os/i386/ \
/var/ftp/pub/apt/9/en/os/i386/RPMS.updates/
```



TutorialPackageManagement

FTP COPY

The *ftpcopy* application is provided with Red Hat. It will recursively copy an ftp site and delete all local files that were not found on the ftp site. Shell-style wildcards (but not regular expressions) are supported for inclusion and exclusion; authentication and recursive operation are supported. While you may include files modified after a certain date, the entire file will be downloaded even if only the name was changed (unlike *rsync*). Using **ftpcopy** to download the updates, we use:

```
# ftpcopy ftp.mirror.ac.uk \
/sites/ftp.redhat.com/pub/redhat/linux/updates/9/en/os/i386 \
/var/ftp/pub/apt/9/en/os/i386/RPMS.updates/
```

Note that the ftp URL stands alone from the ftp directory structure, as well as the use of the trailing slash on the local directory.

Make it Easy, Mate: Script It

While all these tools have their uses, downloading 'by hand' can be an onerous task. One of the strengths of Linux is its powerful scripting capabilities that can make complex tasks simple to execute. It is no surprise that scripts have been developed to handle these downloading chores.

Kirk Bauer maintains the *AutoRPM* script available from www2.autorpm.org:81/. This Perl script can mirror an ftp site, keep installed RPMs consistent with local or remote RPMs (ie install them), and do the same for multiple local machines. It offers both interactive and scripted modes. The default behavior is to download Red Hat RPMs to a queue and permit to be manually installed by running *AutoRPM* from a command line; this can be modified to install them automatically if you desire. Once you have the application installed, the `/etc/autorpm.d/sample-configs/` directory contains useful tips and suggested configurations; `/etc/autorpm.d/pools/` contains a list of ftp sites available to the program. You may add additional sites plus local directories to the pool. *AutoRPM* implements the concept of 'FTP Pools' or groups of ftp sites that house related groups of files (like several sites that provide update files). *AutoRPM* will, with successive use, rank the sites. The file `autorpm.cron` will run

autorpm nightly and check for newly available updated files, downloading them to the local queue.

Ron Stoddard maintains a Perl script to mirror an English-only local copy of Mandrake using *rsync*; no installation of RPMs is done. The script and supporting info can be obtained from <http://members.optusnet.com.au/ronst/rsync/readme.html>. There are include and exclude lists available as well as a `readme.html` file. The script is well-commented and a few minutes perusal will have you up-to-speed. The selection of which categories of the RPM files you download is determined by which lines you choose to uncomment; the script supports only the current Mandrake version, 9.1. At the bottom of the home page are links to two scripts that are to be called from *cron* to fully automate the process of keeping your local files updated if you are using a dial-up connection. Since there are *rsync* servers for Red Hat as well, it would seem trivial to modify the *troels* script to mirror Red Hat files as well.

If you mirror an official Mandrake site (look at the official list at www.linux-mandrake.com/mirrorsfull.list), you will get the index files used by URPMI by default. If you have custom RPM files, you'll need to generate the index files locally using *urpmi.addmedia* with the `-f` option to force the generation of the `hdlist` file; see the man page for *urpmi.addmedia* for more details.

Create a Local APT, URPMI, or YUM Repository

Assuming that you find it necessary or convenient to maintain a local repository for APT or URPMI, there are a few additional steps that are necessary after you have your local copies of the RPM files.

You can serve the files in one of three ways: from an http server, from an ftp server, or from a local hard drive on a single machine. The choice of http versus ftp is up to you. Red Hat and Mandrake offers packages for both types of servers, and the ftp server can be configured as an anonymous server if required.

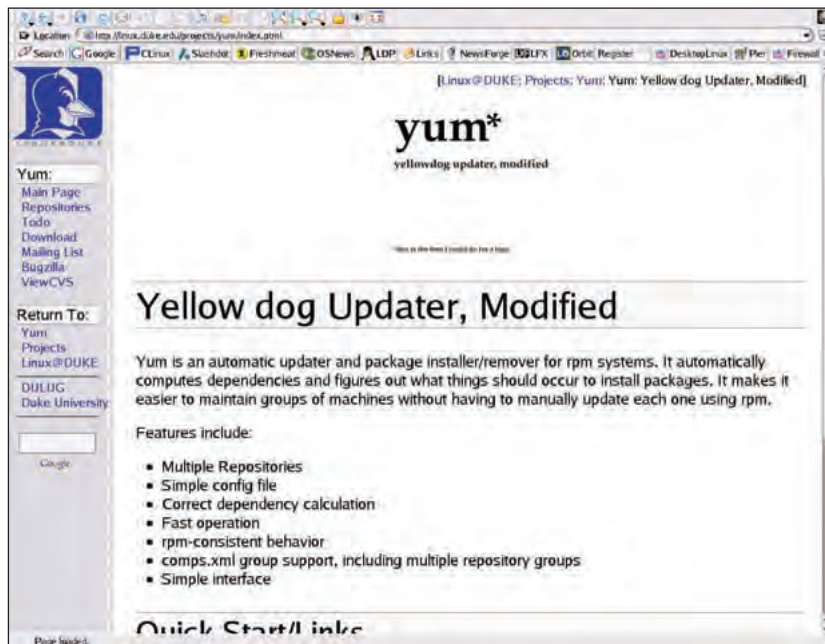
Note: Actually setting up these servers is beyond the scope of this article; bear in mind that security is of paramount concern if your repository will be accessible by the public. The new LXF server tutorial series that starts on page 88 will be addressing this at some point over the next few issues.

When either Mandrake and Red Hat RPM repositories are established, it is best to retain the 'official' directory structure so that the directory structure mirrors that of existing public repositories. This is really more a matter of convention (APT repositories are rigidly structured because of their Debian heritage, for example) but feel free to be a rebel and innovative if you are willing to put forth the extra effort.

For Mandrake 9.1 directories using the *troels* script, the tree will look like this:

```
base_directory +- 9.1-contrib
                +- 9.1-tree
                +- i586
                +- 9.1-unsupported
                +- i586
                +- SRPMS
                +- 9.1-unsupported-MandrakeClub
                +- 9.1-updates
```

The *troels* script uses `/mnt/mandrake` as the base directory. We prefer to modify the script to use `/var/ftp/pub/mandrake/`. For



The *Yellow dog Updater Modified* Homepage has potentially the best tool, but it still lacks a GUI interface.

the Red Hat files, we use a base directory of `/var/ftp/pub/redhat/` and mimic the directory tree found on Red Hat ftp sites as suggested by the excellent documentation found at

<http://freshrpms.net/apt/server/>:

```
base_directory -- 9 -- en -- i386 --
--base
--RPMS.os
--RPMS.updates
```

The “base” directory is used to hold the index files you will create with APT.

FreshRPMs.Net has excellent detailed resources at the previous URL with information as to how the author establishes and configures his own Apt repositories as well as providing a shell script (`genaptrep.sh`) they use to generate their repositories, a directory tree file (`apt-tree.targz`; he uses `/var/ftp/apt` as his base directory, as will we) and a guide to the suggested naming scheme for APT repositories (de-mystifying the contents of the `/etc/apt/sources.conf` file). FreshRPMs.Net also provides an example stanza for *Apache* to make the repository available via http.

Once you have created your local directory tree, populated it with RPM files and run the script to create the indexes need by APT, you can add it to the APT sources.conf file as:

```
rpm ftp://your_localhost.localdomain redhat/9/en/i386 os
updates
```

Creating a local URPMI repository is done via the `urpmi.admedia` command using the `-f` command option to force the creation of a local index file. We would use:

```
# urpmi.admedia -f 9.1_local /var/Mandrake/9.1-
tree/i586/Mandrake/RPMS
```

If you use YUM, you'll need to create a YUM repository. Many public sites create both APT and YUM repositories from the same directories. The `yum-arch` application will create the necessary index files. For our updates directory, we would use:


```
# yum-arch -d -v /var/ftp/pub/apt/9/en/os/i386/RPMS.updates/
```

The `-d` and `-v` commands check dependencies and provide verbose output, respectively. The configuration of the `yum.conf` file would be:

```
[updates]
name=Red Hat Linux 9 - Updates
baseurl=http://your_localhost.localdomain/pub/apt/9/en/os/i386
/RPMS.updates/
```

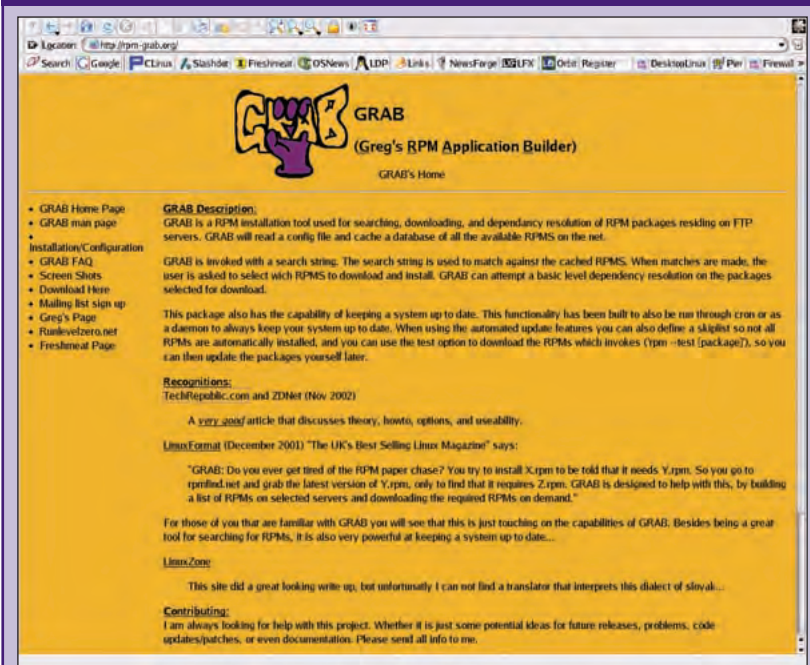
Substitutions for the RHN

While some tools and scripts can be used to obtain and install the base RPMs as well as the updates, two people have decided to re-invent the client-server functionality of the Red Hat Network in GPL software.

NRH-up2date at www.nrh-up2date.org/ is a replacement for the up2date server used with the Red Hat Network. It is Perl code that works with an *Apache* server to mimic the up2date server, allowing the download of updates to the local NRH-up2date server and then permit you to point all your local machines at the NRH-up2date server. The main goal is the elimination of the RHN fees without violating the RHN Terms of Service. Current from <http://current.tigris.org/> is a project in the same vein as NRH-up2date, using an SQL backend database. Both project are still in development. Again, the product that this project is attempting to emulate is the Red Hat Network at <https://rhn.redhat.com/>. 

A Few Quick Links

Supervising your RPMs



Greg's RPMs Application Builder Homepage – no special preparation needs to be done to create a repository, unlike APT, URPMI, or YUM.

The `newrpm.py` script simply compares two sets of RPMs, locally or remotely. It can be obtained from www.logilab.org/newrpm/

The RPM SuperVisor, available from www.klaus.franken.de/rpmsv/, seems functionally similar to APT and YUM, but doesn't seem to have seen any development lately. The HOWTO is in German only, and is too technical to be much use with BabelFish.

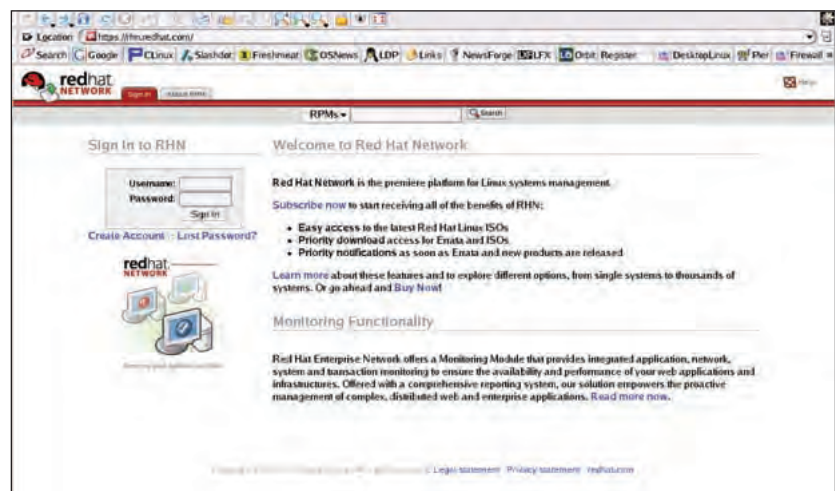
Greg's RPM Application Builder, or *grab*, is available from <http://rpm-grab.org/>. Reviewed in LXF22 (Dec 2001) and functionally similar to APT and YUM, the attraction of GRAB is that no preparation need be made to the repository.

The `rpm-update` script available from www.kleemann.org/rpm-update/ will download and install errata RPMs for Red Hat.

It can also handle multiple local installations.

The Quick Start file is well written and informative. If all you want to do is to keep up with Red Hat errata updates over an Internet connection, this script makes it easy.

As you can see, maintaining a local repository is not a simple task, but can be useful if you need to maintain several machines. Once you get your scripts properly configured, running them as *cron* jobs will take the work out of keeping current. It takes some work on your part to deal with the learning curve, but the results can be well worth it. With quick and easy access to local RPM files, you may find an interest in building custom ISO images populated with non-Red Hat RPMs or custom RPMs you might build yourself. Have fun!



Red Hat Network Homepage is the target that all others aim for. Some focus on enhanced features and ease of use, others simply on avoiding RHN fees.



FLEX AND BISON

Compiler writing

After a brief hiatus last issue for a Q&A and now deserving a move into the main magazine, **Paul Hudson** resumes development of SKYLang, the programming equivalent of a tricycle Rolls-Royce, and finds himself going loopy...

After having a break last issue with the Q&A session, we're straight back to the graft this month by taking our new intermediate code system to new heights. Now that we have intermediate code to work with, there's lots we can do to make our language take off, so, without further ado...

The state of the language

SKYLang reads in a script, parses it into intermediate code, then executes it. It understands a few operators for adding, subtracting, and such, and it also does string to integer type conversion, stores variables, handles complex calculations (**\$a = 5 + 5 * 2**), and outputs variables. In our last full tutorial we had just implemented intermediate code, which opens up a world of possibilities for us – the question is, what do we go for first?

Looping the loop

There are several things to choose from: conditional statements such as **if (\$foo) { echo \$bar; } else { echo \$baz; }**, or user functions such as **function myfunc(\$var1, \$var2) { echo \$var1; }**, or perhaps even built-in functions, like **sqrt()**. Instead of those, we're going to start off with loops – how to make a piece of code execute several times in a row.

Loops are a very basic concept in languages, but in reality they can be a little hard to program. As usual I've tried to make an implementation that's fairly easy to follow and understand, but don't worry if you don't understand it first time, you still may need to re-read this article a few times just to make sure.

If you recall from the first Compiler Writing article back in *LXF41*, the definition of a SKYLang loop is this:

"A loop consists of the reserved word 'for' followed by a variable, then an equals sign and an initial value for that variable, then the keyword 'to', followed by the maximum or minimum value of the variable, followed by the keyword 'increment' or 'decrement', followed by the value to change the variable at each iteration. The code contents of for loops will be enclosed within braces. A valid for loop will look like: **for \$i = 1 to 5 increment 1 { SKYLang Code Here }**"

It's a very basic loop definition, because SKYLang is a very basic language – if you've used BASIC before, you'll notice the similarity. Now, there are several basic properties of loops that I want to outline before continuing, just to make sure we're all on the same level:

- i** The looped code is everything from the opening brace to the closing brace.
- ii** The loop variable is iterated when the closing brace is reached
- iii** A for loop checks its variable first, then executes, and loops. The variable is checked again at the end to see whether to jump to the top for another iteration.
- iv** Loops can be embedded within each other to make inner and outer loops.
- v** Any number of loops can be embedded within each other.
- vi** Loops always close the most recently opened loop.

With those six very basic rules in mind, let's take a look at how loops are implemented. There are four key parts to our loop: the loop variable, the starting value, the ending value, the increment value, and the code to execute inside the loop. Let's break that down into even five key stages:

- 1** Set loop variable to starting value
- 2** Execute loop code
- 3** Increment variable by increment value

4 Check variable is less than max value

5 If it is, jump to "Execute loop code"

We already have our code in place to handle the section "Execute loop code" – this is just lots of other SKYLang statements that may include other loops. The key thing to note is that the main block of our loop code splits the rest of the loop code in two, which will leave us with problems later on. For now, though, let's start implementing the code that is necessary to get started.

New flex matches

We've got some new keywords for the language now: for, to, increment, decrement, open brace {, and close brace }. You need to add the first four of them just below our definition for "echo" in skylang.l, like this:

```
"echo" {
    return T_ECHO;
}

"for" {
    return T_FOR;
}

"to" {
    return T_TO;
}

"increment" {
    return T_INCREMENT;
}

"decrement" {
    return T_DECREMENT;
}

"," {
    return T_SEMICOLON;
}

"{" {
    return T_OBRACE;
}

"}" {
    return T_CBRACE;
}

.\n { /* trim whitespace */ }
```

Inside skylang.y, we need to declare these new tokens. Edit the %token <var> line to this:

```
%token <var> T_ASSIGNEQUALS T_SEMICOLON T_ECHO
T_FOR T_INCREMENT T_DECREMENT T_TO T_OBRACE
T_CBRACE
```

Embedded actions

That's the lexer and parser definitions out of the way, which means we now need to tell our parser to expect these tokens in a certain order. To get started, we're going to implement incrementing – decrementing is the same process. The parser definition for our incrementing for loop looks like this:

Conversion tweaks

Make sure you don't break your code while looping the loop

Previously we did string to integer conversion by literally changing a variable to an integer. This wasn't a problem before, because each statement could only be executed once, however now that we have loops we need to leave our code untouched – we need to change the `sky_strtoint()` function so that it's non-destructive. To accomplish this, it needs to take the variable to convert, copy it, convert the copy, then return the copy, leaving the original as it was.

Furthermore, another problem with the current function is that it doesn't handle variable contents properly – using the SKYLang version from LXF43, try executing this code:

```
$foo = 10;
$bar = $foo + 5;
echo $bar;
```

From that, I get 9 – it sees \$foo, considers it a string, and converts it using `strlen()`, giving 4 + 5. This is clearly no good – we want it to substitute the number 10 in place of \$foo, then add that to the 5. So, the goal of the new `strtoint()` function is two-fold: we need to check whether it's a variable, and if it is we need to substitute in its value, and we also need to work entirely on a copy, so as not to damage variables that may be used again if they are inside a loop.

Here's the new code:

```
SKYVar* sky_strtoint(SKYVar* strvar) {
    SKYVar* working = new SKYVar;

    if (UserVars[strvar->charval] != NULL) {
        working = UserVars[strvar->charval];
        if (working->type != svString) return
        working;
    }
```

```
working->type = svInt;
working->charval = strvar->charval;
char* errors = NULL;
int strconv = strtol(working->charval, &errors,
10);
if (errors == working->charval) { // failed
    working->intval = strlen(working->charval);
} else { // succeeded
    working->intval = strconv;
}
```

```
return working;
}
```

Note that the function returns a **SKYVar*** now, and the first line of code in the function creates a **SKYVar*** called **working** – that's what we'll be returning, and it will contain the converted data. The first step now is to check whether the string we've got is actually a variable – this is simple enough, but the key is where, if we have a variable, we copy it into the working variable, and exit if its type isn't **svString**. If we don't have a variable, we continue on and set the working variable to be an integer, and copy the original string value in too, for the conversion.

The rest is the same, with the exception that we now work exclusively with the working variable, and return it at the end. The calls to `sky_strtoint()` now need to be altered so that they catch the return value – you need to change the relevant code to this:

```
if (op->op1->type == svString) { op->op1 =
sky_strtoint(op->op1); }
if (op->op2->type == svString) { op->op2 =
sky_strtoint(op->op2); }
```

```
T_FOR T_VARIABLE T_ASSIGNEQUALS expression T_TO
expression T_INCREMENT expression T_OBRACE statements
T_CBRACE
```

However, if you remember, *Bison* waits until it matches all tokens before it takes any action, which means that we'll get to perform actions once the **T_CBRACE** is received and the loop is ended. This isn't good enough, because we need to setup the loop increment variable before the statements are added to the op list. Luckily, *Bison* lets us use embedded actions, which break up action execution into parts dependent on the tokens received. That is, we can write our parser match like this:

```
T_FOR T_VARIABLE T_ASSIGNEQUALS expression T_TO
expression T_INCREMENT expression T_OBRACE { do stuff }
statements T_CBRACE { do more stuff }
```

Using this method, we get control immediately after the **T_OBRACE** is matched, then we pass back control to *Bison* so that it can parse *statements* and **T_CBRACE**, then we get control back again after the **T_CBRACE** so we can write the end of the loop.

The problem is that the embedded action works out as two entirely different code blocks once the *Bison* code is converted to C, which means we can't send variables from one part to another – specifically, we can't tell the **do more stuff** section where to jump to without some sort of global variable. While it's possible to >>

Debugging your code

If you use a debugger like *gdb* or *gvd*, you can add debugging information to your compiler code to try to figure out where things are wrong if you have a bug. To get debugging information into your executable, edit your makefile, and change the CC line to **CC = g++ -g**

◀ have some sort of global integer variable that stores the number of the op code to jump to, it would only be able to store one op code at a time, which rules out the possibility of inner loops – we need to be able to stack up loops freely.

To handle multiple loops, we're going to use an STL stack object. Stacks are a very basic abstract data type that works like a tube with one closed end full of tennis balls – you can put a tennis ball in the tube, add another one on the top, take off the top one, then take out the bottom one. You can't put two in, then take out the bottom one, because the top one would be in the way. With a stack, you 'push' items onto it, then 'pop' them off – you can't reference an object by its position in the stack, because you can only read the item at the top of the stack.

Referring to **Fig1** at the bottom of this page, Example 1 shows our empty loop stack. When we hit the start of a loop, we push the line number of the loop code onto our stack, giving Example 2. When we hit an inner loop, that is, a loop within our loop, we push that line number onto the stack, giving Example 3. If we have an inner inner loop (a loop within our loop within our loop), we push that line number on also, getting us to Example 4. Think back to the six basic rules of loops – number six was "loops always close the most recently opened loop", which means you can't write code like this (I've used indentation to show what the programmer wants to happen)

```
for $i = 1 to 10 increment 1 {
  for $j = 1 to 20 increment 2 {
    echo $foo;
  }
  echo $bar;
}
```

There the first closing brace is meant to close the outer loop, and the second closing brace is meant to close the inner loop – crazy, and, thanks to rule six, not allowed. Loops always close the most recently opened loop – this means that the first closing brace will match the most recently opened loop, which is the **\$j** loop, and the second closing brace will match the **\$i** loop. That code, properly formatted, is this:

```
for $i = 1 to 10 increment 1 {
  for $j = 1 to 20 increment 2 {
    echo $foo;
  }
  echo $bar;
}
```

Now, consider the situation in Example 4 – we are inside three loops. This might be from code like this:

```
for $map_x = 1 to $map_xwidth increment 1 {
  for $map_y = 1 to $map_ywidth increment 1 {
    for $i = 1 to 25 increment 5 {
```

```
// we are here, inside the inner, inner loop
}
}
}
```

Immediately after the **we are here** line, we hit a closing brace – which for loop does it close? Obviously reading the source code (thanks to the pretty indenting), we can see it's the **\$i** loop. However, it's just as easy to do in code, because of our stack – the most recently opened loop is the one at the top of our stack, so we just need to pop it off to figure out where the loop was opened.

Enough waffle!

If you're like me, code speaks louder than words, so we need to jump in at the deep end. Below the **std::map** definition of UserVars, add this line:

```
std::stack<int> LoopStack;
```

The LoopStack stack holds integers, which means we can push and pop operation numbers from our OpArray variable.

std::stack is a new data type from STL, so you also need to edit *skylang.h* to include the stack header – add this line beneath

```
#include <vector>:
```

```
#include <stack>
```

Now, onto the hard stuff – the actual parser action to handle loops. Put this next big block of code after the ECHO action – I've included the last line of the ECHO action to help you match it up:

```
new_op->op1 = $2;
}
|
T_FOR T_VARIABLE T_ASSIGNEQUALS expression T_TO
expression T_INCREMENT expression T_OBRACE {
  SKYVar *newvar = new SKYVar;
  newvar->type = svString;
  newvar->charval = $2;
  sky_do_binop(SOP_ASSIGN, newvar, $4);
  LoopStack.push(OpCodes.opcodes.size() - 1);
} statements T_CBRACE {
  int opnum = LoopStack.top();
  LoopStack.pop();

  SKYVar *result = sky_do_binop(SOP_ADD,
  OpArray.opcodes[opnum]->op1, $8);
  sky_do_binop(SOP_ASSIGN, OpArray.opcodes[opnum]->op1,
  result);

  SKYOp *new_op = new SKYOp;
  OpArray.opcodes.push_back(new_op);
```

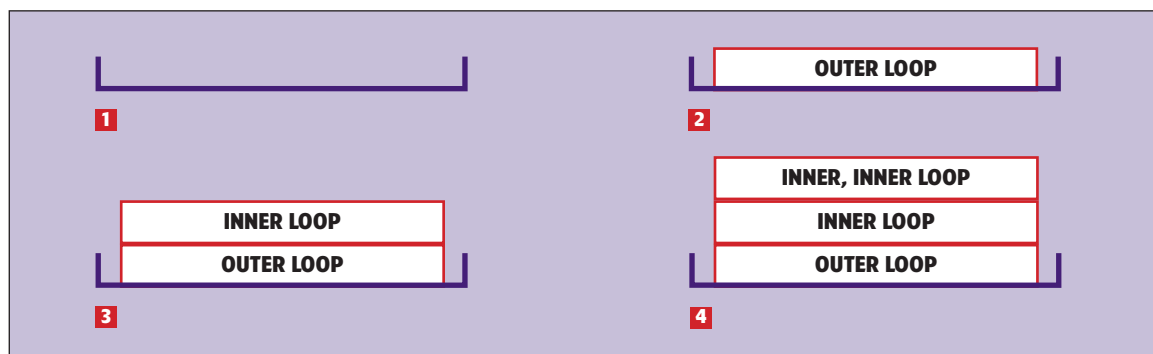


Fig1 Loop stacks and how they work.

```
new_op->opcode = SOP_JMPLTE;
new_op->op1 = OpArray.opcodes[opnum]->op1;
new_op->op2 = $6;
SKYVar *goto_op = new SKYVar;
goto_op->intval = opnum;
new_op->result = goto_op;
};
```

That's quite a chunk of code – don't worry if it looks a little overwhelming at first, because we're going to break it down into easier-to-swallow chunks. Firstly, note that the loop match is split into two, with the first action being performed after the **T_OBRACE** and the second being performed after the **T_CBRACE** – this is essential, because, as described earlier, we need to handle all the statements that come inside the loop separately, so the loop initialisation needs to be done separately from the loop end.

The first few lines of code in the first action are basic and we've seen them before – we create a variable for the loop iterator, and assign it the starting value. The new code, however, is the line **LoopStack.push(OpArray.opcodes.size() - 1);** – this is half of the loop stacking code. Our **LoopStack** STL stack variable stores integers that will correspond to the exact instruction in our **OpArray** of where each loop starts.

Our opcodes vector, which stores each individual operation, has a **size()** function, which returns the current number of operations in there, and as we just added the last line before the loop code actually begins (assigning the loop variable), we can take the current number of operations, subtract 1 because it's a 0-based array, and use that for our jump position – we literally jump to the **\$i = 1 to 5** line. This might sound odd at first, because you might think that that will just leave us setting the loop variable each time we jump, but remember that our **execute()** function increments the operation counter by one at the end of each operation – this means our jump will jump to **\$i = 1** line, and the **execute()** function will nudge it onto the line immediately after, which is the first line of our loop code.

Moving on, the statements get parsed and executed, then we get control back as the **T_CBRACE** (a close brace character) is received, signalling the end of the loop. At this point, we've just completed stage two of five – we've set the loop variable, and generated the intermediate code for our loop statements. Stage three is to increment the loop variable by the increment value as requested by the user, so first we need to grab the top loop from our stack using **LoopStack.top()**. We can then use **LoopStack.pop()** to take off the top value, leaving the operation number of our loop start in **opnum**. Secondly, we send another binary operation in to add **\$8** to the loop variable, and store the result of that operation in a SKYVar result variable – **\$8** is, of course, the eighth part of the *Bison* match, which is the *expression* after **T_INCREMENT**. Finally we add another binary option to assign the result of the add back to the loop variable, effectively incrementing it by the requested amount.

SKYOp new_op signals the start of stage four, where we need to check that the loop variable is less than the maximum allowed variable for the loop. This uses a new SKYOpType (the enum defined in *skylang.h*), which is **SOP_JMPLTE** "Jump if Less Than or Equal". This op type takes a SKYVar to check and a number, **\$6**, to check against. **\$6** is the maximum value from the source code – the *expression* after **T_TO**.

So, the JMPLTE operation will check **op1** against **op2** and jump if op1 is less than or equal to op2 – but where does it jump

to? The last few lines in the action set up the result of JMPLTE operation to hold the line number it should jump to, as provided by the opnum variable we got from **LoopStack** – this is a hack, because the result variable is supposed to hold the result of the operation, but it wouldn't be used for anything else, so we might as well use it.

Is that it?

I'm afraid not – that's just the code to *parse* the loop into intermediate code. We also need to execute it, by adding a new case in the **execute()** function to handle JMPLTE operation. Jumping is simply a matter of changing the **execute()** loop variable, **i**, to another number, which in our case will be the intval inside the result of our JMPLTE operation. Here's the code to handle it:

```
case SOP_JMPLTE:
if (UserVars[op->op1->charval]->intval <= op->op2->intval) {
i = op->result->intval;
} else {
// loop finished, do nothing
}
break;
```

Nothing hard there, really – we just do a direct comparison of the intval of our variable against the intval of the number, and if the variable is less than or equal to the number we set **i** to the new execution point. At the end of the **execute()** function, **i** will iterate up to the next operation, meaning that the next operation to be run will be the first statement inside the loop block.

Conclusion

We've included the full source code to the latest version of SKYLang, v0.3, on the coverdisc – once you've implemented the changes in the boxout *Conversion Tweaks*, check your version against mine to make sure you've got all the latest code, because there are a few extra changes as well as the ones listed above.

At this point, the switch to intermediate code is now paying off – making loops work has been remarkably easy, I hope you'll agree. While the article might seem long for something I consider simple, remember that a fair chunk of it was devoted to the newer string conversion, with about half of it being devoted just to theory on how loops work. The actual code is quite simple, which you'll find is generally the case – we've already done the hard work!

I put a quick hack into my source code to dump out the **op array** list so that you can see how your scripts compile – it's not much good right now, but it is, as I say, just a hack.

Now, thanks to reader *Trickster* who posted on our forum a request that we should end tutorials with some homework, so that's precisely what I'm going to do! The loop code as it stands contains one fatal flaw – take a look over this problem script, read basic loop rule #2, then find the problem in *skylang.y* and fix it.

```
$main = "Main loop";
$inner = "Inner loop";
for $i = 1 to 10 increment 2 {
echo $main;
for $j = 25 to 20 increment 1 {
echo $inner;
}
}
```

I'll be explaining the problem and the solution next month – until then, happy hacking! [LXF](#)



NEXT MONTH

We continue exploring intermediate code, by looking at conditional statements. There are a few bugs in the compiler currently – don't worry, we'll be addressing them in *Linux Format* issue 47.

ILLUSTRATION BY PHILIP - www.blender3d.org

IMPROVE YOUR ART SKILLS

Advanced modelling features

PART 6 Although we have covered an extensive amount of ground, *Blender* has an enormous amount of features still unexplored by us. **Jono Bacon** looks at some of the more advanced modelling facilities.

At this point in our usage of *Blender*, we have been dealing primarily with simple, or 'primitive' shapes such as cubes, cylinders and cones. In most cases, we have been creating these shapes as a basis for our modelling, and then deforming and editing these shapes where necessary. We are now going to take a new route to create more specific and complex shapes – by using Bezier Curves. You may be already familiar with the concept of Bezier curves if you have used them with *The GIMP* or other art packages but for those who are unaware, I will define this first.

Bezier Curves

A Bezier curve is basically a vertex (single point on a shape) that has a 'handle' that can move the vertex and its connecting lines in different ways. The best way to visualise this is with the concept of levers. Imagine that each vertex has a lever with a handle on the left and a handle on the right, with the vertex in the middle. Now, imagine that when you pull the right handle, the curve that is connected to the vertex moves in the right direction and *vice versa*; this is essentially how Bezier curves work. Using these

curves, we can create some pretty advanced shapes, and they are particularly useful for creating metal plate-type objects and logos.

To get us started, create a new project and remove the default plane. Now split your views and ensure that the Front view is the largest view with other views including a Top and Camera. In the front view, ensure the 3D Cursor is in the centre of the view and then click on Add>Curve>Bezier Curve in the Toolbox (remember that you can bring up the Toolbox by pressing the Spacebar). You will then see a Bezier curve and its handles as shown in **Fig1**.

The basic concept of using Bezier curves is to add a new curve each time and then adjust the handles to change the curve to its desired position. Add a new curve by holding down the **Ctrl** key and clicking with the left mouse button where you would like the vertex to be placed. With the two Bezier curves now positioned, you can adjust the connecting curve by selecting either the main vertex or one of the handles (by right-clicking it in Edit mode) and then using the **G** key to move it as you would like any other moving operation.

While you are adding new curves to the shape and changing their visual appearance, there are a few things to bear in mind.

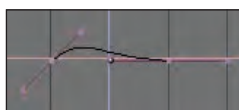


Fig1 Bezier curves form the basis of complex shapes.

Firstly, there will often be a case when you would like to create a sharp corner in your shape. This is simple to do; just select the control vertice (the central vertice) and then press the **V** key. The colour of the handles will then change from purple to green. This little technique is useful for creating smooth shapes with rounded and sharp edges – which is common. Another important thing to remember is that the Blender Subdivide command works with Bezier curves also. To do this, select both control vertices by selecting the first one, holding **Shift** and then selecting the next one. If you then press the **W** key, a menu will pop up and you can select Subdivide. You will then see another Bezier curve is inserted between the two selected curves. You can also remove Bezier curves if needed, by selecting the control vertice and then pressing the **X** key; a menu will pop up and then select Selected.

When you have added a number of curves, and are at the point where you have created a shape that you are happy with, you will want to close the lines to create a single continuous line. You can do this by selecting the two control vertices that will be connected to close the shape, and then press the **C** key. The handles will still be available on the Bezier curves to change the look of the curves if required, and you can still add, remove and sharpen the vertices. When you have created a shape that you like and have connected it, your shape in non-edit mode can look something similar to that in **Fig2**.

Although our shape looks fine, there is one main limiting factor – it does not have any holes in it. If you look at many objects, they have holes and dips in it, whereas our shape is just like a solid sheet of material. If you think about it, adding holes to objects would be a complex function within *Blender*, as the program would need to know what is a hole and what is not. Luckily for us though, *Blender* uses some wizardry to figure all of this out for us, so we can simply create another curve inside out, shape it, and the hole will be automatically calculated.

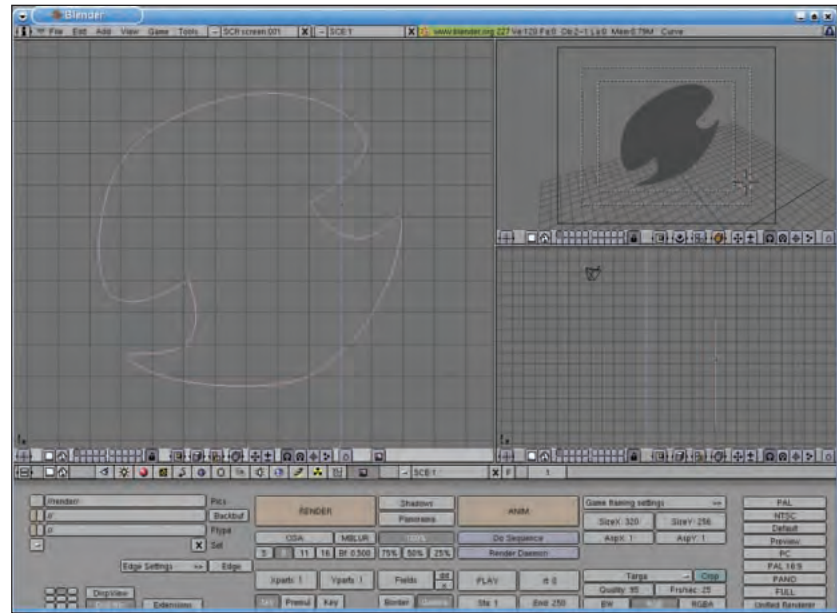


Fig2 A simple shape composed of Bezier curves

Bringing the shape to life

Currently, we have used Bezier curves to create complex 2D shapes. Although this has been useful, the shapes are distinctly... two-dimensional. While we are revelling in the world of 3D, we have to make our 2D masterpiece jump into the third dimension. This is simple and can be done within the Edit buttons (F9). With the shape in Non-Edit mode, to the right you will see some buttons with Width, Ext1, Ext2 and BevResol written on them. These are used for extruding and bevelling our shape. The buttons have the following functions:

Width this sets the width of the shape thickness. This value is

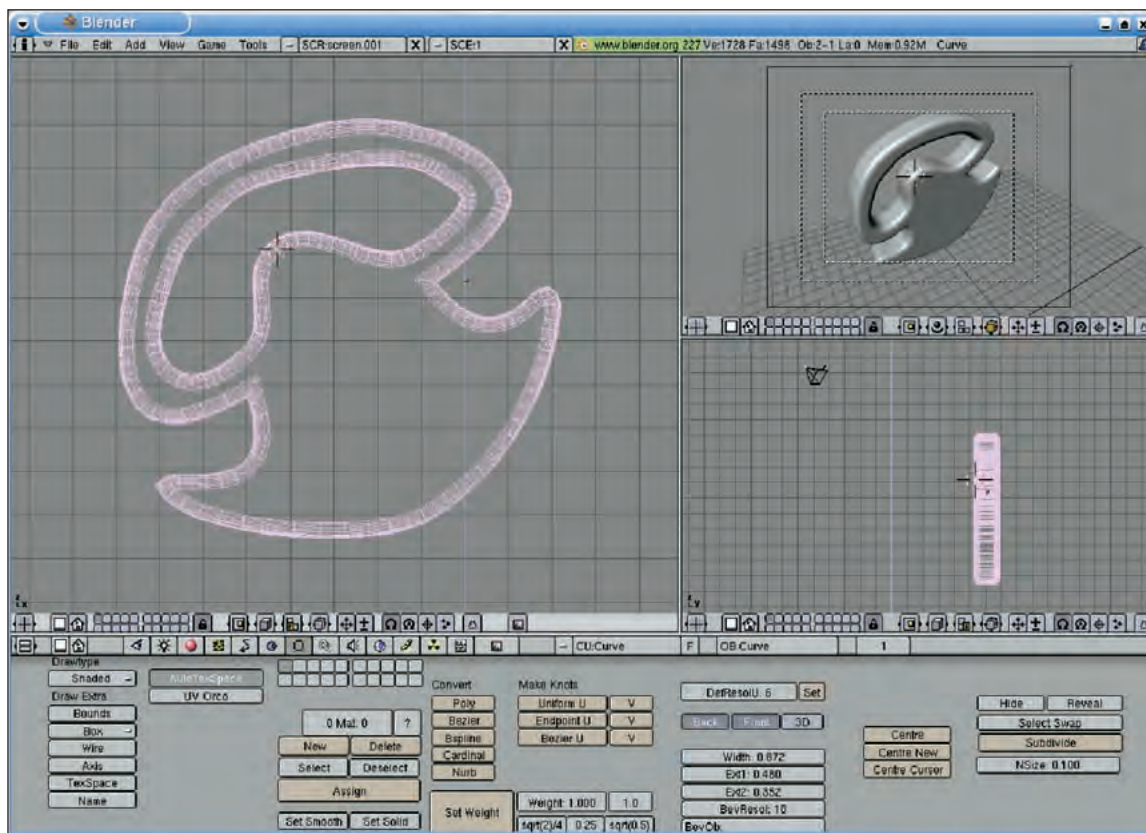


Fig3 Thickening our shape out.

TutorialBlender

◀◀ also calculated around holes in the shape, so the relative width is calculated for different parts of the shape.

Ext1 This sets the extrusion size for the shape along the Z axis (ie making the shape thicker). We can use this button to set how thick the shape appears.

Ext2 This button controls the amount of curved bevel extrusion. You can only use this button if you have already extruded the shape with the Ext1 button.

BevResol This button control the degree of resolution that that bevel has between 1 (very small) and 10 (very high).

With these buttons you can now extrude the shape and give its edges a nice bevel effect. You can see the result of these buttons on my shape in **Fig3** on the previous page. Your results will probably look different – experiment until you are happy.

Spins

Spins are a useful facility for creating repeated sets of objects surrounding a point. The typical example where this can be used is the points on a clock-face. There are 12 points in all (one for each number) and they are all an equal distance from the centre of the clock-face, as well as an equal distance from each other. In *Blender*, the SpinDup function performs exactly this objective, and we will give it a test-drive.

First, create a new project, remove the default plane and add a shape to the scene (I added a cone, as an example). Now, switch to the Front view and move the shape to the bottom of the screen. Now, leave edit mode and select all of the vertices with the **A** key, and then switch to the Edit buttons (**F9**). To the right of the buttons there is a group of buttons with Spin in the middle. We will need to set a few special buttons to ensure our spin is how we want it. These buttons are:

Degr This button sets the amount in Degrees that the shape will be spun. As an example, if you set the button to 360, the shape will be rotated around a full circle.

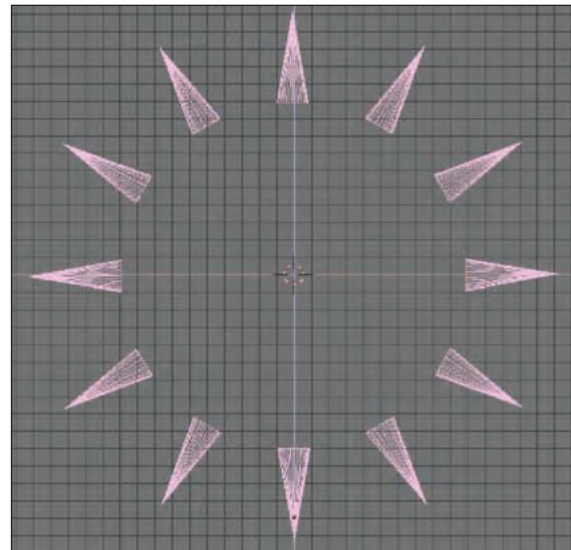


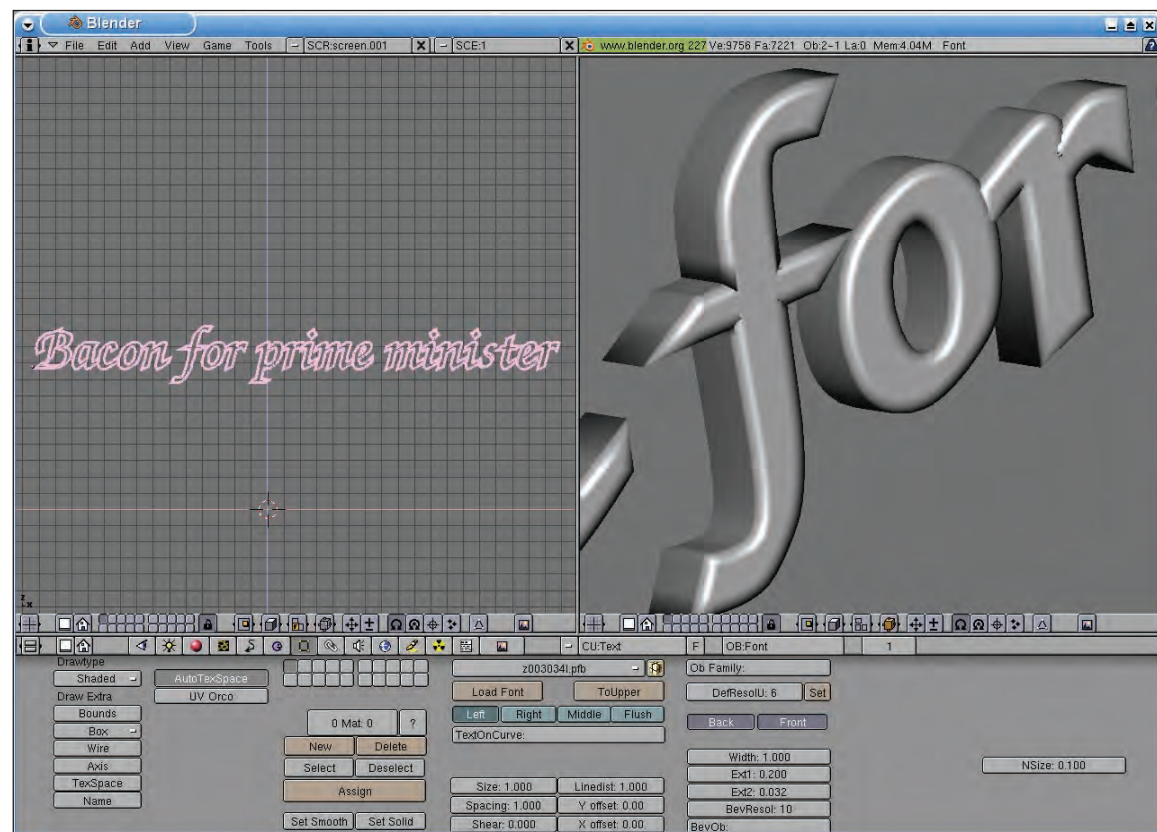
Fig4 Using the spin tool with cones.

Steps This is the number of duplications of the shape that will be applied to the spin, and in how many segments. It is best to think of this in terms of equal divisions of a pie. If for example, I set this button to 12, there will be 12 equal segments in the circle and there will be a copy of the shape placed at each cut in the circle.

Turns This specifies the number of complete turns in the Spin. If you set this to 1, you will get 1 complete spin.

As an example, I have set the Degr button to 360, the Steps button to 12 and the Turns button to 1, and when I then click on the Spin button, I will get 12 equally spaced cones rotating around where the 3D cursor is placed. It is important that the 3D cursor is placed correctly, as it will define the diameter of the circle. You can see the results of this operation in **Fig4**.

Fig5 Extruding and bevelling fonts - a few simple steps makes a dramatic difference to your rendered text.



How to create a Bazaar Jar

Combining the techniques you've learned so far

This month we are going to use some of the techniques we have learned to create a strange little jar as shown in the screenshot. To create this jar we first need a new *Blender* scene. Delete the default plane and then create a UVSphere with the default values. In the Top view, use the **B** key to select half of the sphere, and then press **X** to remove it. Next, select the middle row of vertices and select another half of the sphere. Extrude this half 5 or 6 squares and you will now have the back of the jar. The only thing we need to do with this section is to close the gap where we have removed half the sphere (if you press **Z** to shade the view, you will see that the sphere is hollow). You can make this solid by selecting the front vertices in top view and then press **Shift-F** to connect the vertices as a face.

Get in shape

To create the front shape, switch to the Front view and move the viewport away from the shape we have just created. From the Toolbox, click Add>Bezier Circle and you will see a circle with four Bezier curves inside it. This works the same as normal Bezier curves, although new curves are added the current circle. This way you can add new Bezier curves and create the outside shape of the front of the jar. When you have created the shape, you can then add another Bezier circle inside the shape

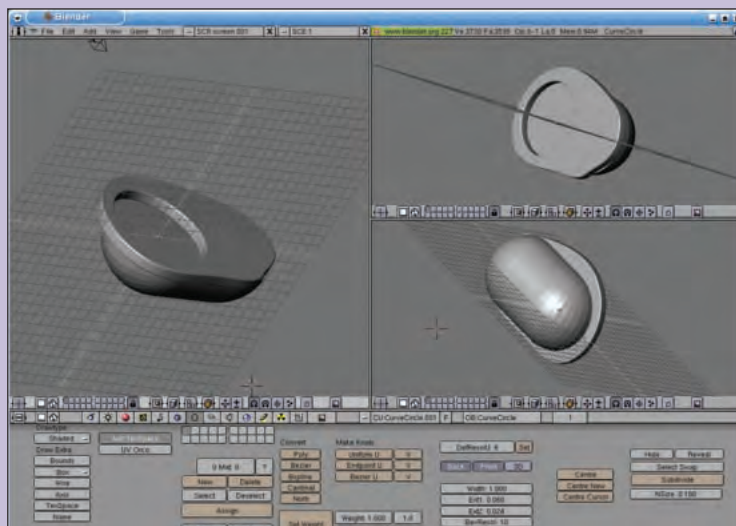


Fig6 Creating simple and strange little objects in *Blender* by uniting two halves.

(remember to do this while in Edit mode or the inner circle will not form a hole in the outer shape). Line the inner circle up as I have and then exit edit mode. You can now use the extrusion and bevel buttons discussed earlier to make the front shape look nicer. When you have finished this you can line both shapes up together, and the bazaar jar is complete.

Although this was a simple project to work on, it is a good introduction of how

these separate techniques all fit together to create more complex-looking shapes. If you have the time, why not add some more features to the jar such as legs, or text labels. Another thing you may want to do is to add materials/lights and render the scene properly.

Finally, you could do a rendered keyframe sequence, and maybe even add armatures to the shape to bring the jar to life. The only limit is your own imagination!

Using text in your scenes

Text is a pretty important part of most computing, and it is certainly no different in *Blender*, and you are probably going to need to have some text in your scenes. Handling text is very easy, and there are some simple yet powerful tricks to deal with text in different ways.

To get us started, create a new project, remove the default plane, and then switch to Front view. Move the camera out of the way and then select Add>Text from the Toolbox. You can now type some enlightening text into the 3D view such as "Jono for prime minister" (yes I know, but it just may happen one day, and ego is a prerequisite of the ambition anyway). When dealing with text, we still have an Edit mode, and you can enter text in this mode. Leave edit mode, and you will see the text turns pink like normal. You can now move the text around and scale it like any other object.

The first thing you will probably want to do is change the font type. At the moment, *Blender* only supports Type 1 fonts, although TrueType font support is planned in a future release. You can change the font by clicking on the Load Font button in the Edit buttons. You can then take a trip over to your fonts directory (such as /usr/share/fonts/type1) and select a font to use. You will see the fonts applied, and the font will be added to the drop down box above the Load Font button.

At the moment, our text is completely flat. We can now make use of the same buttons we used for extruding Bezier curve shapes to extrude and bevel our text. Using these tools alone can be useful for creating authentic-looking text, and when you apply good materials and lighting, it can look pretty impressive.

Another technique we can use with text is using a Bezier curve to determine the direction of the text. In some way, this can be used to similar effect as WordArt in word processing systems, but with greater flexibility. To use this feature, in Front view, create a Bezier curve so it curves how you would like your text to flow. Next, exit edit mode and select the text. If you switch to the Edit buttons, you will see in the centre of the buttons a button called TextOnCurve. If you select this button and type into it the name of your Bezier curve, you will see the text will change to the direction of the curve. You can also Left, Right and Middle buttons to position the text. One other small feature is that the ToUpper button will change all of your text to upper case.

Conclusion

This month we have been taking a look at some of the more advanced modelling facilities in *Blender*. As with any complex application, there are countless other features, and the key to getting to know them is playing with them. One thing that is important with *Blender* is to know that simple techniques are the hallmark of great-looking scenes. When you have a look around the Internet and see some of these impressive looking *Blender* scenes, just try and decode in your mind how the scene was created. Even simple tools such as Extrude are used extensively and in many different contexts. This month we have discovered how *Blender* can create more complex shapes in 2D, but these shapes can then be extruded, bevelled and the other tools can be applied if needed. [LXF](http://www.linuxformat.co.uk)

NEXT MONTH

We are going to continue our look at some of the more advanced concepts in *Blender*, including the particle effects system. This powerful feature will allow you to create effects that are traditionally difficult to model with normal methods. This includes creating things such as smoke, fire, wind and other effects. As well as particle effects we will also look at how we can combine these techniques together to create more complex scenes. Keep on blending!

SKINNING WITH THE GIMP

Multimedia interfaces



MIs – better known as **Skins** – are all the rage with application developers. **Michael J Hammel** wields the tools to create fantastic interfaces, all you have to do is provide the creativity!

The chances are you've played with programs like *MPlayer*, *Xine* and *XMMS* if you're into watching video or listening to music on your Linux system. These tools provide user interfaces that stretch the limits of convention for computer programs.

But skins aren't limited to multimedia programs. The advances in CSS for the Web have made skinning an art unto itself for website design. Take a look at the 'Zen Garden' (www.csszengarden.com/) or Eric Meyer's 'css/edge' (www.meyerweb.com/eric/css/edge/) sites for examples of this technology.

Creating these interfaces is a two-step process. The latter step is trickier and requires understanding how to chop images into component parts and put them back together using a skinning interface, often requiring knowledge of XML or some other text file describing the position of those image pieces. The first part, however, is a breeze: building the interface starting with nothing

more than a blank GIMP canvas and an active imagination.

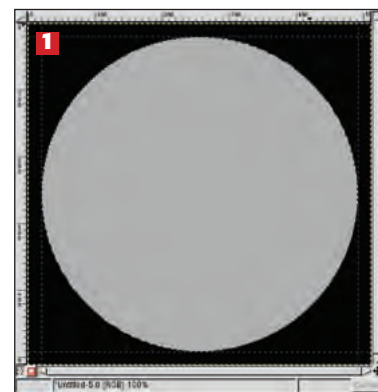
In this tutorial we'll take the steps outlined in *The Complete Adobe Photoshop Handbook*, a compendium of tips and tutorials for Photoshop users published by *Linux Format's* sister publication *Computer Arts*, and convert them into the equivalent steps under *The GIMP*. In this way, you'll not only get to see how easy it is to create these mind-stretching designs, but also how close in features *Photoshop* and *The GIMP* really are!

There aren't many special tricks here that a default installation of *The GIMP* couldn't provide, but you will want to make sure you can find the *Lighting Effects* plug-in (Filters>Light Effects>Lighting Effects) and are very familiar with dragging Guides around the canvas. We'll be using many guides in this tutorial. We also assume you know your way around *The GIMP's* windows and menus, know how to change foreground and background colours, and know how to use selections. If you don't have the *Round* plugin (look for Selections>Round) then skip the part where we round the edges of selections. You don't have to round them to complete this tutorial. Finally, we'll be working with grey colours for most of this tutorial. Color and texture can be added later as separate layers composited with different layer blend modes.

Designing a skin

The design is simple – an imaginary handheld device with a Palm Pilot-like screen, handle grips on the side and a hard-shelled metal structure underneath. We start by creating a 3D metallic plate on which the device is built.

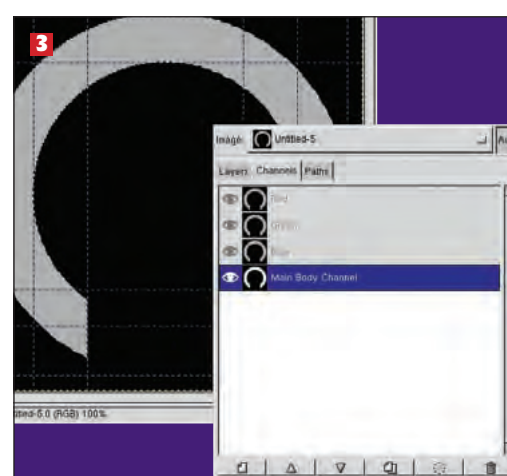
1 Create a new window that is 512 x 512 pixels in size. Fill the background layer with black and then create a new, transparent layer above it. Keeping in step with the original tutorial, we'll name this new layer 'Main Body'. Drag a vertical guide from the left ruler into the canvas area and position it at x=20 pixels. See the lower left of the status bar to see the guide position. Drag another vertical guide and place it at x=492. Place horizontal guides at y=20 and y=492. Make a circular selection by pressing the shift key while you click where the left and top guide meet and dragging to the opposite guides intersection. Fill this selection with equal amounts of Red, Green and Blue (RGB), each set to 175.



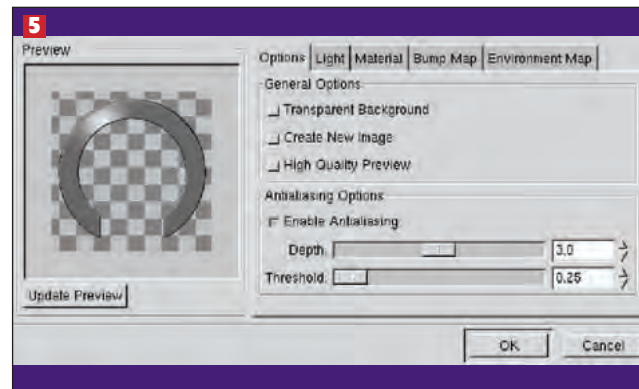
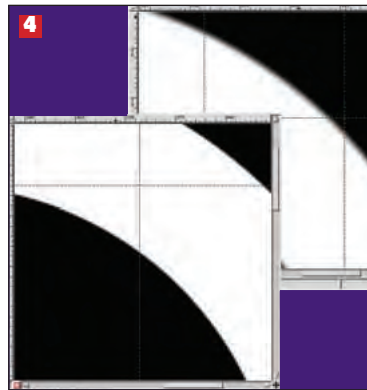
2 Add guides 85 pixels from all four edges and drag another circular selection using these guides. Use Ctrl-X to cutout this selection from the Main Body image. Add vertical guides at 155 and 357 and a horizontal guide at 380 to create a rectangular box cutout at the bottom of the Main Body image. Use the 'Alpha to Selection' option for the Main Body layer to create a selection out of this upside down horseshoe shape.



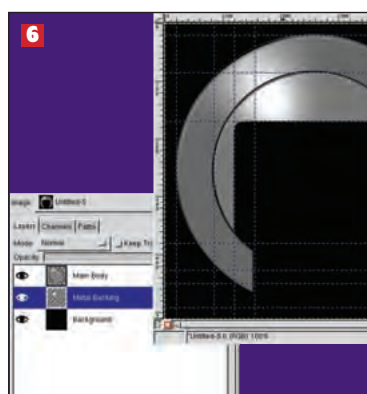
3 Change to the Channels page of the Layers and Channels dialog. Add a new channel and name it 'Main Body Channel'. Set its opacity to 100%. The image should go to black with the horseshoe shaped selection still active. Fill this selection with white, then clear the selection with Ctrl-Shift-A.



4 With the 'Main Body Channel' still active, apply a blur with **Filters>Blur>Gaussian Blur (IIR)** with the horizontal and vertical values set to 3. The blur applied to this channel will be used to create the bevelled edge of this part of the design. The same technique (using different channels) will be used later on other components in the image to add a 3D feel to them as well. When the blur is complete, click on the eye icon to turn off the visibility of this channel (so we can see the whole image again). Go back to the Layers page in the Layer and Channels dialog and click on the "Main Body" layer to make it active.



6 Create a new layer below the Main Body and call it 'Metal Backing', making sure this layer is now the active layer. Use the guides from step two to create another circular selection and fill the selection with red, green and blue values of 200 each. Add vertical guides at 120 and 392 and horizontal guides at 170 and 450. Create a rectangular selection from these guides and round the corners using the **Selection>Round** menu option, with the roundness set to 30. Cutout this selection from Metal Backing using **Ctrl-x**. Create a selection of this layer with the 'Alpha to Selection' option in the Layers menu.



5 Now we'll give the Main Body a little 3D effect. Open the **Lighting Effects** plug-in (**Filters>Light Effects>Lighting Effects**). Set the **Options** to **Enable Antialiasing** with a depth of 3.0 and threshold of 0.25. In the **Light** page set the **Light Type** to **Directional** and the **Color** to white. Set the **Direction Vector** to $x = -0.60$, $y = -1.0$ and $z = 1.0$. In the **Material** page set the **Intensity Ambient** level to 0 and the **diffuse** level to 0.70. Set the **Reflectivity diffuse** to 0.65, **specular** to 0.85 and the **highlight** to 27.00. Finally, in the **Bumpmap** page, enable bump mapping using the **Main Body Channel** as the map and using the default values for the rest of the options. The **Environment** mapping is not used here. Apply this to the **Main Body** layer.

Resources

MPlayer Skin Howto: <http://mplayerhq.hu/DOCS/skin.html>

XMMS supports WinAMP skins. A Howto for this can be found at www.markerrington.com/winampskins/howto/

This GIMP tutorial was inspired by the original work of Kraig Sederquist (www.sedergraphics.com) from a tutorial published in *The Complete Adobe Photoshop Handbook*, published by Creative Arts Magazine, Future Publishing Ltd.

Make the screen Area

Now we'll add an area in which we'll later place display screens. These screens will appear to be encased by a metallic frame set within the empty space formed by the 'Metal Backing' layer. The trick to this part of the tutorial is to be creative with the casing – make it fit in the empty space by drawing a

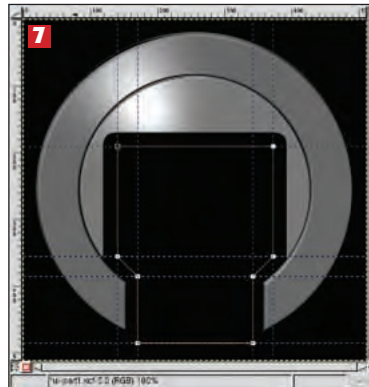
straight and angled edge selection using the Bezier Selection tool, then round the corners with **Selection>Round** again. The fancier you get with your Bezier selection the more high-tech you can make the casing appear.

After the casing is ready, we will spend some

time creating handles to hold our digital device.

Start by removing all the old guides. The existing guides will just add unnecessary clutter and confusion to our work. You can drag them out manually or, if you have the script installed, use the 'Guides>Remove Guides' plug-in.

7 Add a new layer above the Main Body layer and call it 'Screen Casing'. Add horizontal guides at 190, 355, 385 and 490. Add vertical guides at 140, 170, 342, and 372. Draw a Bezier selection using these guides to make straight and angled lines that cover much of the empty space in the 'Metal Backing' layer. Fill this selection with RGB values of 150.

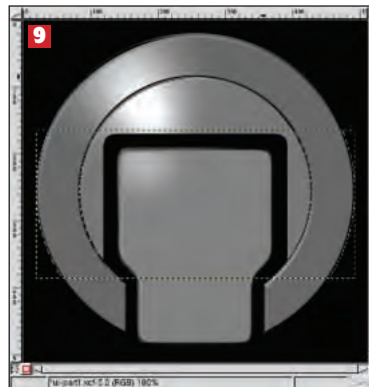


8 Create another channel at 100% opacity, this time named 'Screen Casing Channel'. Fill the selection in this channel with white and turn off visibility of the channel. Return to the 'Screen Casing' layer and apply the **Lighting Effects** filter again using 'Screen Casing Channel' as Bumpmap. To change the reflectivity on this casing (making it a duller metal) you can move the location of the light and change the ambient intensity and the specular value. Don't forget to blur your channels or the 3D effect will not show up.



« Make the screen Area (continued)

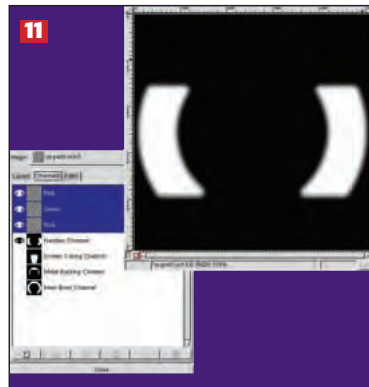
9 Make the 'Main Body' layer the active layer. Make a rectangular selection through the middle of this layer, then copy and paste the selection as a new layer. Name this layer 'Handles'. The new layer will be cropped to the edges of the selection, but we need to grow the selection, so you'll need to choose 'Layer to Imagesize' in the Layers menu to make room for the larger selection.



10 Select 'Alpha to Selection' in the Handles layer – you now have a selection of the two handles we're about to create. Grow the selection (Selection->Grow) by 8 pixels, make sure the 'Preserve Transparency' option for this layer is turned off, and fill with RGB = 112. Keep the selection active for the moment.



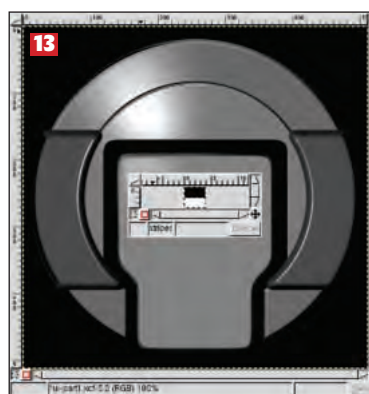
11 Create another 100% opacity channel, this one named 'Handles Channel'. Fill the selection with white. Clear the selection and blur the channel by 16 pixels. Switch back to the layers page and make the Handles layer active.



12 Apply the Lighting Effect filter to the Handles layer using the Handles Channel as the Bumpmap. Change the Reflectivity and Intensity options to change the material appearance of the handles. We'll be adding more texture to them in a moment.



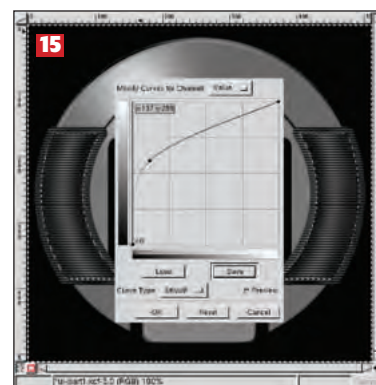
13 Let's add texture to the handles. Create a new, very small 4x4 image filled with white. Fill the upper half with black. Zooming in on this image will make this black fill easier. Save this file in your .gimp-1.2/patterns directory with the filename 'stripes.pat'. When prompted, give the description as 'Simple Stripes'. Once you've saved the file, open the patterns dialog and click 'Refresh'. The new stripe pattern should be listed in the set of patterns.



14 Use 'Alpha to Selection' on the Handles layer, then create a new layer above the Handles layer. Name this layer 'Handle Lines'. Select the Simple Stripes pattern from the patterns dialog box, double click on the Bucket Fill Tool to open its Tool Options dialog. Choose 'Pattern Fill' for the Fill Type. Now click inside the selection. The stripes will fill the selection in the new layer. Change the layer mode for this layer to Overlay.



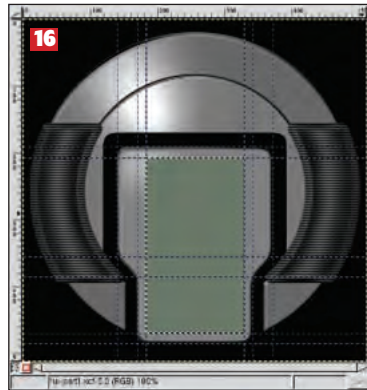
15 The handles are almost complete. We need to add a little more depth to them so they appear more round than flat. To do this, we start with the handle selection once again (it should still be there if you haven't cleared it yet). Add another layer between the Handles and Handle Lines layers and name it 'Handle Shadows'. Select a large, soft edged brush, perhaps 'Circle Fuzzy 76' if you have that one. With the opacity of the airbrush tool set to 50 per cent, paint black strokes along the left and right edges of the selection in this layer. Your strokes should follow the curve of the handle selection. Only add shadows along the outside at first, then add a little to the inside. Just enough give the appearance of a rounded handle. When you're done, shrink the selection by 10 pixels, invert it, and feather it by 7 pixels. Choose Image>Colours>Curves and adjust the curve as shown. This keeps the edges of the original handles while preserving the roundness of the shadows.



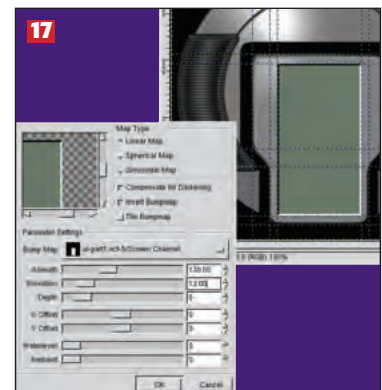
Add the display screen

This can be done in multiple pieces if you want multiple display areas or as one big screen over the Screen Casing layer. We'll use our light mapping trick to give the appearance that the screen is recessed into the metal casing.

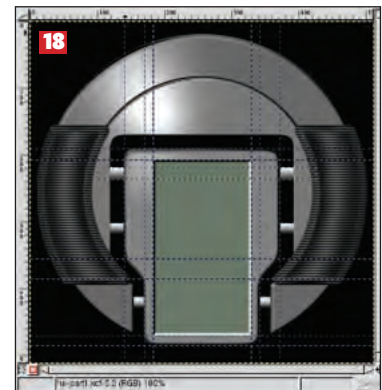
16 Start with new vertical guides at 183 and 329 and horizontal guides at 233 and 470. These provide the outline for our screen. Make a rectangular selection with these guides. Create a new layer just above the 'Screen Casing' layer and call it 'Screen'. Fill the selection with RGB values of 106, 120, 105 respectively. The screen resembles the original Palm Pilot LCD screen.



17 Create a 'Screen Channel' in the channels page filled with white. Clear the selection and blur by 3 pixels. Return to the Screen layer. Select **Filters>Map>Bump Map**. Set the azimuth to 130, the elevation to 13, the depth to 8 and choose the Linear, Compensate for Darkening and Invert Bumpmap options. Make sure the Bump Map image is set to the 'Screen Channel'. Apply this to the Screen layer.



18 You now have the appearance of a recessed screen. At this stage, the basic unit is complete and we can get as creative as we'd like. First we'll make simple connections between the Screen Casing and the Main Body. Add a transparent layer just above the background and below the Main Body layer. Draw a rectangular selection starting on the left inside the Main Body and through to the right side of the Main Body. Set the Foreground to White and Background to Black. In the Gradient Fill Tool Options dialog set the Blend to 'FG-to-BG (RGB)' and the Gradient to 'Bi-Linear'. Then drag from top to bottom of the selection. The result will be simple 3D tubing between the Screen Casing and the Main Body. Repeat this a few times along different horizontal lines in the image. Notice that only the parts of the tubing that are not blocked by higher level layers show through.



In glorious colour

While there are many more features we could add to this interface, so let's turn our attention to adding texture and colour to it as an example.

19 Choose the Main Body layer and duplicate it. Get a selection with Alpha to Selection. Fill this with a pattern - we chose 'Canyon 2'. Desaturate the layer, set its blend mode to Multiply and reduce its opacity to about 50%. Duplicate this layer, alpha to selection the new layer and fill with the Pink Marble pattern. Set this marble layer's blend mode to Color.



20 Duplicate the Screen Casing layer and Alpha to Selection on it. Fill the selection with a blood red - RGB of 102, 9, 9 respectively. Set the layer mode for the red layer to Saturation.

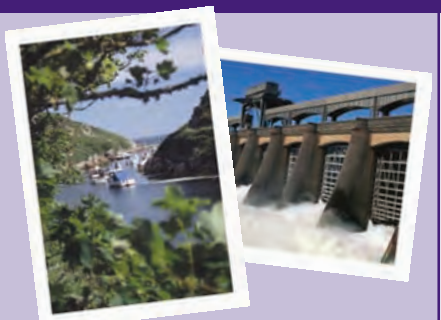


21 Duplicate the Metal Backing layer and Alpha to Selection it. Fill the selection with the Lightning pattern. Set the layer blend mode to Value and the opacity to about 40%.



NEXT MONTH

You'll doubtless have a wealth of summer holiday snaps to tweak and catalogue by the time the next issue of *LXF* hits the shops, so we're going to show you how *The GIMP* can edit them to make them look their very best. **LXF**



FASTER CODE

Practical PHP Programming

Paul Hudson finishes his two-part mini-series into accelerating your scripts with a look at how to optimise your MySQL queries...

Last month we looked at the basics of optimisation, and also went into quite some depth regarding how to make your PHP scripts run faster. However, where's the point in making your scripts run faster if the database that drives them crawls? This issue, without further ado, we're going to be looking at various ways to make your SQL run faster. This is a big area, so you're strongly advised to read up on the suggestions given here to see whether they will work in your given situation.

There are three basic routes to take when it comes to optimising your SQL: change your queries, change your table schema, or change your server configuration. We'll be looking at all three, starting with queries.

One of the many advantages to SQL is that it is a true fourth-generation language – you tell it what you want, and it does all the hard work for you. As such, the actual *implementation* of your query requests is left entirely down to your database server, which means it might not necessarily be doing things in the most optimised way. In order to reach the largest audience, we'll be covering MySQL techniques here, but many of the same rules will apply elsewhere.

Optimising queries

MySQL rule #1:

Select as little data as possible.

Consider the following two queries:

```
SELECT * FROM MyTable WHERE Username = 'Hudzilla';
```

```
SELECT ID FROM MyTable WHERE Username = 'Hudzilla' LIMIT 1;
```

MyTable, we shall assume, contains thousands of rows, with one for each user. The programmer is looking to extract the ID of a user based upon a username, and above are two ways it can be done: select every field from each row and match **Username** against **Hudzilla**, or to select just the **ID** field and match **Username** against **Hudzilla**, and only return one row. I made a simple benchmark to test the exact speed difference between the queries – with 200,000 users (a large number to make the point clear), the first query took 1.31 seconds to execute, whereas the second query took 0.66 seconds to execute.

The reason for the big speed up for the second query is two-fold: firstly, we're not bothering to return anything other than **ID**, whereas there could be dozens of fields being returned in the first query. Selecting as little data as possible means using **SELECT *** rarely, if ever. Secondly, and most importantly, the **LIMIT 1** at the end of the second query will force MySQL to only return one row and, as we don't have any complicated ordering going on, MySQL will simply stop searching as soon as it finds the first row with Username **Hudzilla**. That's quite clearly superior to the first solution, where MySQL will keep on searching the database to find other **Hudzilla** rows – this is rarely the desired situation.

MySQL rule #2:

Let MySQL tell you how it would execute your query.

If you append **EXPLAIN** before any **SELECT** query, MySQL will output a detailed report on how it would have handled the query, how many rows it would have returned, etc – this is a great way to spot problems. The faster of the two queries from above, when

used with EXPLAIN, produces the output that is shown in **Table 1** to the right.

The important parts are **possible_keys**, **rows**, and **Extra**.

Possible_keys shows **NULL**, which means MySQL cannot run the query using indexes – we'll be solving that later. **Rows** shows how many rows MySQL would have had to search through, and **Extra** shows what extra functionality it would have used to perform the search – here we can see it would have filtered the search using the **where** clause.

MySQL rule #3:

Let MySQL tell you about queries that take a long time.

Perhaps you didn't realise this, but simply by enabling one option in your MySQL configuration file, MySQL will keep a log of all the queries it considers slow (usually taking longer than ten seconds to execute), as well as how many rows the slow query returned.

To activate the slow query log, simply start up MySQL with the option **--log-slow-queries = /var/log/wherever/you/want/it**. You can optionally also use **--log-long-format**, which will cause MySQL to consider all queries that don't use indexes as slow queries – this can be a great help if you're trying to add indexes, but you're not sure where they're needed.

MySQL rule #4:

Perform joins carefully.

Joining data from two tables in one query is a very powerful technique – particularly when combined with normalisation. However, if not done right – even by a small margin – you can incur a serious speed hit. Joining data properly in your queries is a complicated and long-winded thing to explain, and we're almost out of the room, so the best I can do is point you to an excellent reference on the topic, the MySQL manual:

MySQL rule #5:

OPTIMIZE TABLE /table/ is a crucial command to achieve maximum performance.

Once you've implemented everything else here, run **optimize table** on your tables every so often – it does a lot of housework to make sure the table is in an optimal state for querying, and will make your queries faster.

Changing your schemata

The second way to speed up your MySQL performance is to alter the layout ('schema') of your tables to be optimised for performance. This can be split into three sub-sections: data types, indexing, and normalisation, ordered by difficulty. Changing your database schema can be considered the same as changing your algorithm – very often you think you're spending hours tweaking things to get speed-ups, only to find you could double the speed by starting from scratch.

Choosing the right data type for your fields can provide a substantial speed boost, but may or may not be good when it comes to space. For the 'select as little data as possible' test above, here's the SQL schema I used:

```
CREATE TABLE lxf45 (ID INT NOT NULL PRIMARY KEY
  AUTO_INCREMENT, Username VARCHAR(255), Age INT,
  JoinDate INT, Homepage VARCHAR(255), Location
  VARCHAR(255), FaveColour VARCHAR(255), Password
  VARCHAR(255), PassRemind VARCHAR(255));
```

It's not a complicated schema, as you can see. However, there are a number of ways it can be optimised. Firstly, **INT** is a vastly over-used data type – it means "store any number between -2147483648 and 2147483647". Now, while it might be possible that vitamin pills and a healthy diet will help you live long, but it's

table	type	possible_keys	key	key_len	ref	rows	Extra
lxf45a	ALL	NULL	NULL	NULL	NULL	199999	where used

TABLE 1: Append EXPLAIN before any SELECT query.

not really likely that we'll ever use all of that range for our age field. Nevertheless, MySQL *doesn't* know that, so we need to be slightly less vague about our Age range. The **TINYINT** data type stores values from -128 to 127, which is probably enough despite it holding minus values. The **TINYINT UNSIGNED** data type stores values from 0 to 255, which is definitely enough for the foreseeable future, and takes up much less room than a full **INT**.

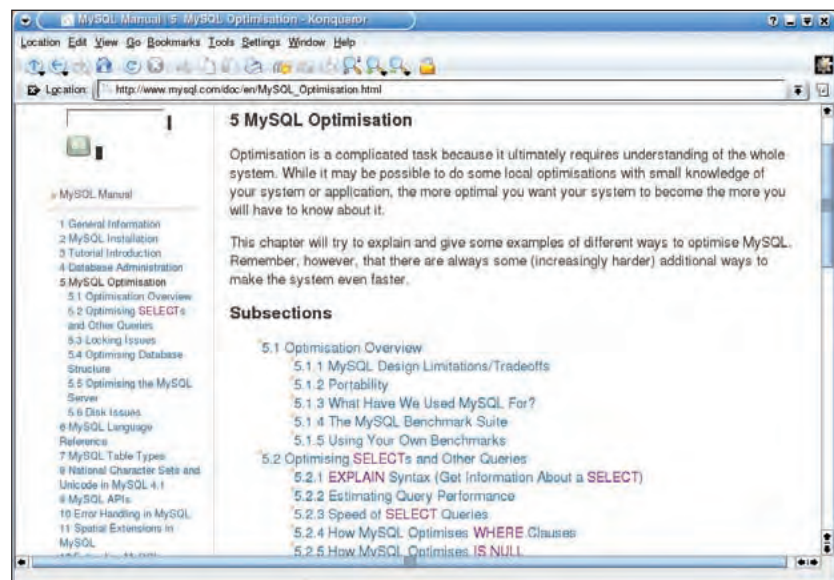
Size vs Speed

Hard drives are much less expensive than just a few years ago, and getting cheaper all the time – you can buy hundreds of gigabytes of space for £200 and still have change, so saving space isn't generally an issue. Processors are an entirely different matter, though – adding CPU power is still expensive, so everything we can do to cut down on CPU usage is a good thing, even if that means taking up more hard drive space, right?

Consider this new schema below, a modified version of the previous one:

```
CREATE TABLE lxf45a (ID INT NOT NULL PRIMARY KEY
  AUTO_INCREMENT, Username CHAR(255), Age TINYINT
  UNSIGNED, JoinDate INT, Homepage CHAR(255), Location
  CHAR(255), FaveColour CHAR(255), Password CHAR(255),
  PassRemind CHAR(255));
```

As you can see, we're now using **TINYINT UNSIGNED** for our Age field, but the key difference is that we're using **CHAR(255)** rather than **VARCHAR(255)**. Both hold up to 255 characters of text, but the difference is that **VARCHAR** only uses as much as space as there are characters for on a row-by-row basis. Variable row sizing might sound good at first – after all, why leave unused space to waste? The problem lies in the fact that, because **VARCHARS** can vary in size, MySQL needs to calculate the length of each **VARCHAR** field in a row so it knows how far to jump to get to the next record. Compare that to using **CHAR** as the data type, where MySQL can just add 255 (using the above example) to get to the end of one field and the start of another. ➤



The MySQL manual is an essential reference whether you are a novice or guru.

« Optimum column types

Another important consideration to keep in mind is that keeping character data short is a great space saver – in **Table1** on the previous page I used **CHAR(255)** for everything, but you should consider each column carefully. A helpful tool in this regard is the **ANALYSE()** procedure – not the same as **ANALYZE TABLE table**, which is like a miniature **OPTIMIZE TABLE**. Use the **ANALYSE()** procedure like this:

```
SELECT * FROM lxf45a PROCEDURE ANALYSE();
```

That will execute the query **SELECT * FROM lxf45a**, then pass the results into the **ANALYSE()** procedure. The **ANALYSE()** results are fairly complicated, but the important column is **Optimal_fieldtype** – this tells you the best field to use based upon the results returned by the query.

While all the above will help you get a big speed boost, the real key to maximum performance in MySQL is mastering indexes. An index in database terms is the same as an index in a book. If you were to take any book, say Hoyt Duff's *Red Hat Linux 9 Unleashed* for example, how long would it take you to find something without the index, and how long with?

Indexing your data

You create an index using the command

```
ALTER TABLE... CREATE INDEX name(column)
```

– the **name** is just a name for your reference, but the **column** should be the name of the row you want to construct an index of. We're working with username here, so we should use

```
ALTER TABLE lxf45a ADD INDEX idxusername(username)
```

It will take a little time to execute, and also subsequent inserts/updates will be a mite slower because of the need to update the index, but it's worth it – the faster of the two queries from earlier on now takes 0.00 seconds. Yes, you read that right – it executes so quickly MySQL doesn't even register the time it took, which is a great deal faster than the 0.66 seconds the next fastest query took, and even faster than the 1.31 seconds the slowest query took.

You can add indexes to as many columns as you want – you can even index every column if you like. However, there are trade-offs, as per usual – indexes take up space, so the more indexes you have, the larger your table will be. Secondly, each time you change a row, you also need to change the index – if

you index every column you'll need to update several indexes to reflect the changes. If you have a table which is written to exponentially less than it is read from, you can go ahead and try adding as many indexes as you think necessary.

We're not finished yet, though – remember, optimisation is the act of improving code, and even though we seem to be running at maximum performance, there's a great deal more we can do.

Normalisation

One key area open to optimisation is data duplication. Very often data is duplicated across rows, for example as row might contain Person A works at Company A, whose address is BigLongAddressHere. Person B also works at Company A, whose address is still BigLongAddressHere. Person C also works at.... etc. In this situation there are many people who work at the same place, but the company address is duplicated in the row of each staff member. What happens if the company changes address? Worse, think of all the drive space that's being wasted needlessly!

This can be corrected through normalisation: a person work at a company, but the address of the company relates to the company, not to the person. Therefore, it's better to create a new table for companies, put the companies in there (with no duplicates) and the addresses in there too. Now you just add a field to the employees table to link each employee to a company.

We covered normalisation in more depth back in the *Practical PHP* article in *LXF31*, but there are also several helpful resources online (just try Google) if you want to learn more.

Configuring your server

No matter how fast your queries are, if you're running cheap hardware or, worse, poorly configured hardware, your performance will always be stunted. The only thing worse than not having the very latest kit available for your servers is actually *having* some big iron and not having it configured properly!

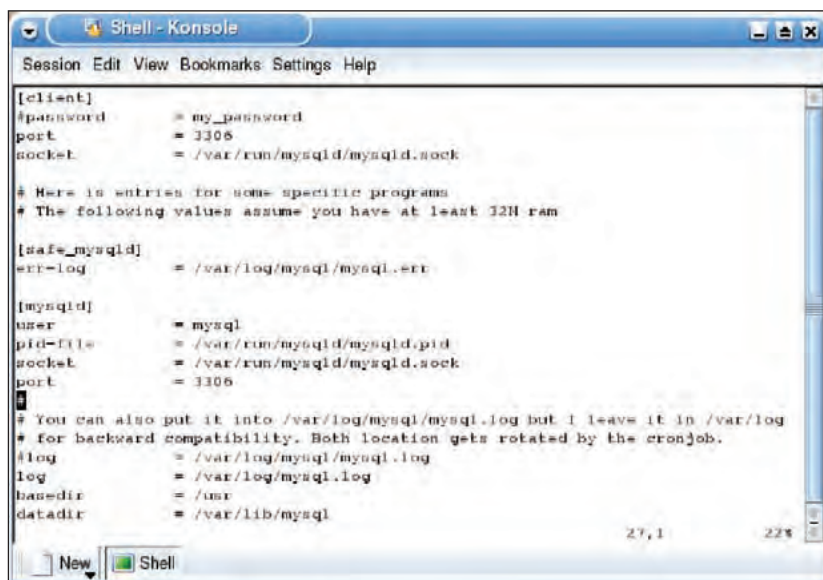
There are several easy tricks to make sure your system is working at maximum capacity. Note that we do not have the space here to print a full 'Optimise Linux' tutorial – however, check out our lead feature this issue for general tips.

There are four steps to database server nirvana: upgrade your hardware, upgrade MySQL, tweak MySQL, and finally keep a watch over your server.

Upgrade your hardware

This might sound like an obvious way to make your code faster, and indeed it is: upgrading your system is likely to make the biggest immediate difference. Attitudes on what you should upgrade are mixed, with much of the conventional wisdom not really holding true any more. For example, MySQL recommend upgrading your RAM, then your hard drives, then your CPU. That may well be true in some circumstances, however the majority of us are likely to find the biggest speed difference by upgrading in the opposite order: CPU, hard drives, RAM.

The reason for this is because the "RAM first" recommendation is based upon the concept that you're likely to be filling up your RAM with your MySQL data. If you are indeed using up all your RAM with your databases, then upgrading your RAM will be a huge help. However, on the smaller of my two servers I have 512MB of RAM, and my databases (large as they are) only take up 200MB – upgrading RAM won't change things at all for me. Naturally if you're running gigantic databases, your priorities will be different!



Be sure to edit your SySQL configuration for maximum performance.

To NULL or to NOT NULL

A little extra performance

As NULL means literally “no value” and not 0 or an empty string, MySQL needs to store one extra bit of data in each field that can have a NULL value to store whether that variable is NULL or not. If you know a field isn’t going to be NULL (because you’re always going to provide a value, or use a default value), then declare the field as NOT NULL – save the space, and get a little extra performance, too.

Upgrading your hard disk helps when committing data to disk and also reading data into RAM. If you can upgrade to a SCSI drive, or, better if you can stripe your data across several disks, you’ll see a noticeable speed up for your server. At the very least, consider using *hdparm* to tweak your settings. On IDE drives we test, I generally use

```
hdparm -m16 -d1 -u1 -c1 /dev/hda
```

NOTE: Using *hdparm* can damage or destroy your hardware – read the manual thoroughly before use.

If your RAM isn’t full and your hard disk aren’t a problem – because they’re fast enough or because they’re not hit often – then it is CPU power that will count the most, and in my experience I’ve found sheer CPU horsepower is the key bottleneck. MySQL uses a lot of 64-bit integer code for its calculations, so if you can afford an Opteron or even an Athlon 64, then you’ll find it makes a substantial difference.

Upgrade MySQL

I’ll make this nice and clear so there’s no confusion: if you’re not using MySQL 4.0.14 or later, you should be. The reason for this, other than for stability and feature enhancements, is that the 4.x series includes a query cache, which stores results for SELECT queries in RAM so that the next time the same query comes in the results can be served up in no time at all. We covered the query cache in some depth in the MySQL 4 review in LXF42, and the basic result is: if you have queries that return the same results all the time or even some of the time, the query cache will raise your performance substantially.

Increase your buffers

MySQL has a variety of buffers that hold information about tables and indexes, and by increasing the size of these buffers you can often gain substantial boosts. Naturally the biggest boosts occur when changing from very small amounts, but if you find that one buffer has used all its memory, the chances are you’ll get at least some performance boost by giving it some more.

To increase the size of your buffers, you need to edit your query cache. Key areas to increase are your key buffer, sort buffer, and query cache. Give your query cache as much as it needs – it’s a big space chewer. If you’re wondering how you can tell how well your buffers are being used, you need to check the next tip...

Keeping tabs on your server

There are two key commands you should use regularly, as they give you detailed information about the running of your machine. Firstly, just type **STATUS**; at the MySQL prompt to get basic information – how many questions have come in, how many were considered slow queries, how many queries per second, etc. Secondly, type **SHOW STATUS**; to get more detailed information – how much of your query cache is being used, how often

```
mysql> status
-----
mysql Ver 12.20 Distrib 4.0.13, for pc-linux-gnu (i686)

Connection id:          552
Current database:
Current user:           lx@localhost
SSL:                    Not in use
Current pager:          stdout
Using outfile:          ''
Server version:         4.0.13-log
Protocol version:       10
Connection:             Localhost via UNIX socket
Client character set:    latin1
Server character set:    latin1
UNIX socket:            /var/run/mysqld/mysqld.sock
Uptime:                 4 days 6 hours 51 min 36 sec

Threads: 1  Questions: 4502  Slow queries: 0  Opens: 62  Flush tables: 1
Open tables: 21  Queries per second avg: 0.012

mysql> show status;
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Aborted_clients | 2 |
| Aborted_connects | 5 |
| Bytes_received | 293335 |
| Bytes_sent | 529552 |
| Com_admin_commands | 4 |
| Com_alter_table | 7 |
| Com_analyze | 0 |
| Com_backup_table | 0 |
| Com_begin | 0 |
| Com_change_db | 538 |
| Com_change_master | 0 |
| Com_check | 0 |
| Com_commit | 0 |
+-----+-----+
```

Closely monitor your MySQL status to make the most of your server.

MySQL has had to trim the cache due to a RAM shortage, how often your key buffer has been read and how much space it has left, etc. Check **SHOW STATUS** once a week or so to make sure your buffers aren’t overflowing.

RAM is cheap, so there’s no harm giving your buffers all they can take and then some extra for good measure. Using the query cache has no impact on performance, which means if you have the RAM you can just allocate 1GB to the query cache to watch your performance shoot up.

Conclusion

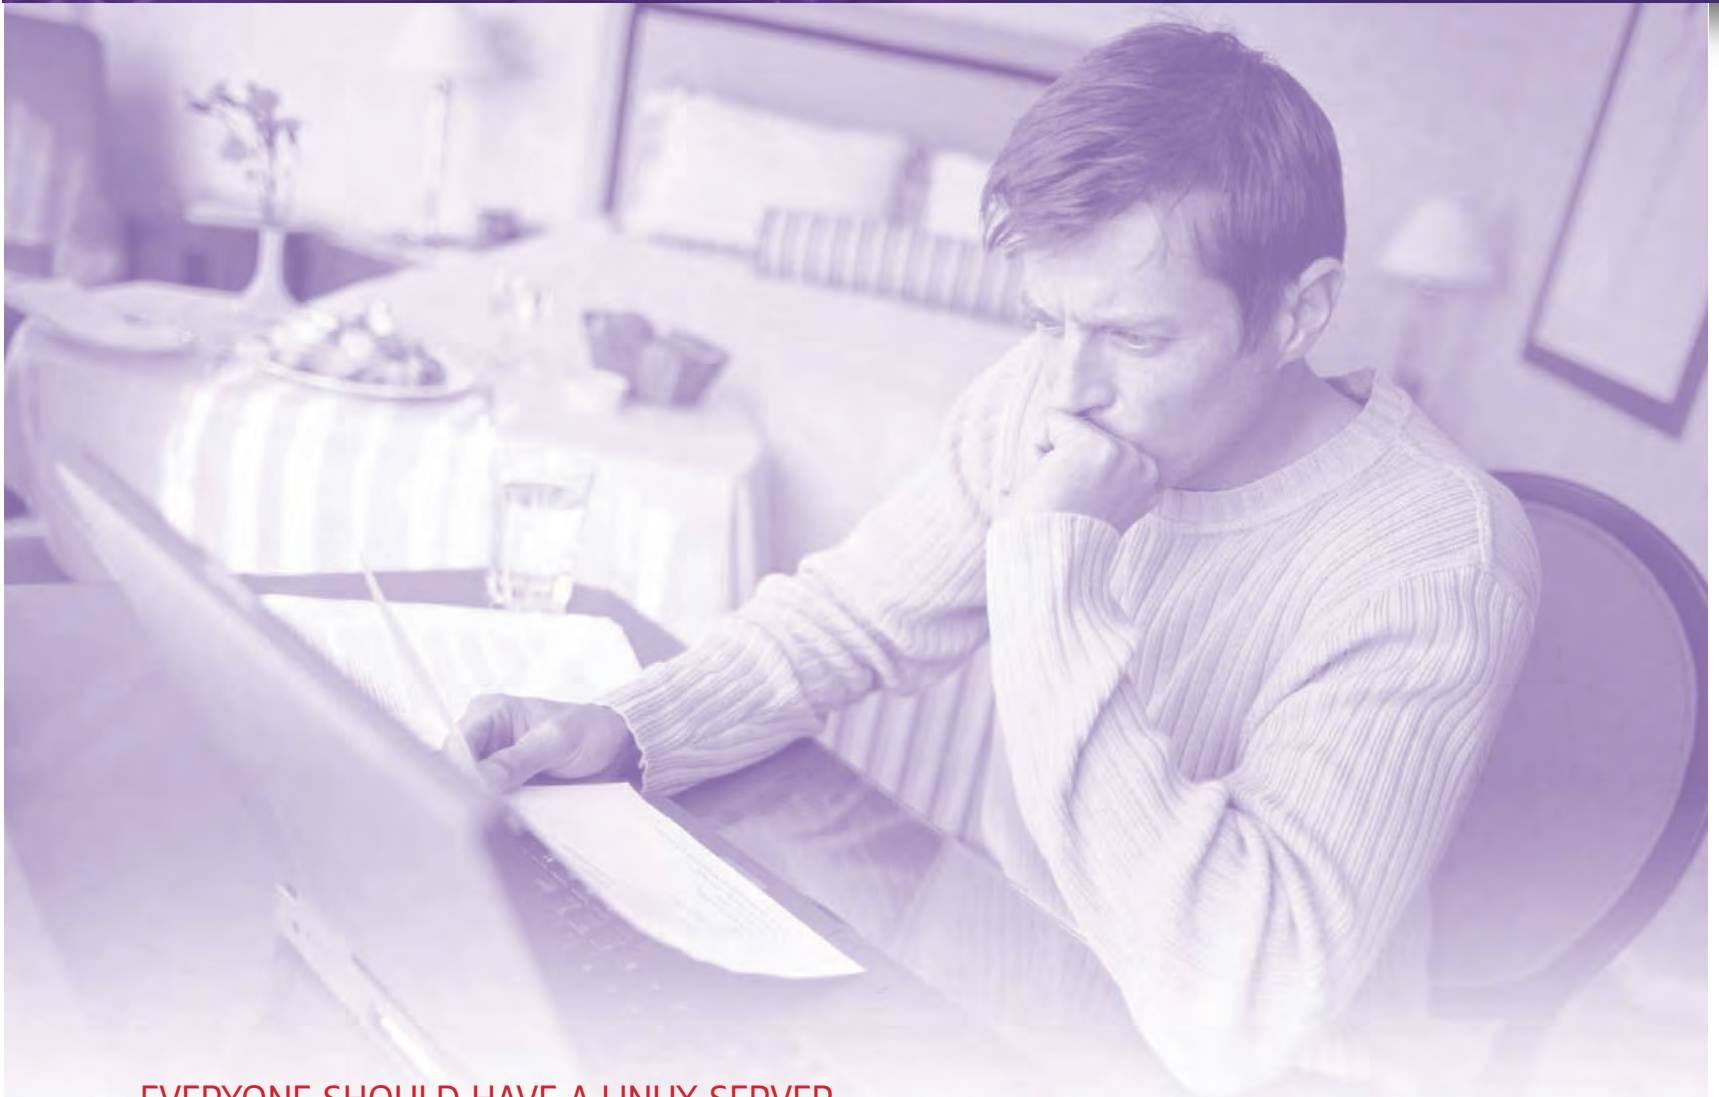
There are a variety of ways you can optimise your MySQL usage for maximum performance, and you can generally spend a lot of time eking out every last bit of speed. However, the chances are you’ll get the best benefit, at least when it comes down to speed improvement vs your time simply by giving using indexes on columns you query by, and giving your buffers and query cache lots of RAM. If you really want to get every last drop of speed without upgrading your hardware, then there’s more than enough above to keep you going for quite some time.

If you really want to take your performance to the max, consider purchasing the second edition of Paul Dubois’ book, *MySQL* (ISBN 0-7357-1212-3, Sams 2003), as he goes into a great deal of depth on optimisation, including much more detail on the various caches and buffers you can configure.

Optimisation is a black art, yes, but it’s lots of fun, and really helps test your coding skills. Very often it’s a skill you learn just by doing, so practice, practice, practice! [LXF](#)

NEXT MONTH

We’ve received an avalanche of mail from readers (keep them coming!) full of ideas for future articles, and next month we’re starting off with the suggestions by covering output buffering – why it’s cool, when to use it, and how to use it in ways you’ve probably not even considered...



EVERYONE SHOULD HAVE A LINUX SERVER

SERVER SCHOOL: First steps

PART 1 Nick Veitch introduces the concepts you need to know in order to get up and running with a Linux server of your own.

Series plan

There are a wide variety of tasks that a server can perform, and we hope to cover them all given sufficient time. Here's our plan of action for the rest of this series...

- PART 2** Web server
- PART 3** Firewalls
- PART 4** Remote access and admin
- PART 5** FTP and filesharing
- PART 6** Mailservers
- PART 7** Databases
- PART 8** Intrusion detection
- PART 9** LDAP
- PART 10** ...

The first thing we should really address is 'What is a server? Many people immediately think of some noisy, rack-mounted box with multiple processors and RAID arrays. While this may be true for the majority of computers used as servers, the reality is that any computing device can act as a server. Web servers are routinely run in embedded devices for example. In order to be a server, the hardware has only to be capable of running the server software.

A server then, is any computing device that 'serves' information. Servers can be used for all sorts of different things – email, directory services, file sharing, databases, applications, web pages, streaming audio – the list of services you may want to use is a very long one. The most numerous type of server in computing is a web server, or to be a bit more technically accurate, an http server. Web servers lurk all around you, from popular public sites like Google and www.bbc.co.uk, to private company intranets and tiny embedded devices – there are even webcams with their own built-in server.

Of course, many other services can be driven through http too, such as file transfer and streaming media. Setting up a web server will be covered in the first part of our server school series.

But we're getting a little ahead of ourselves here, we still have to sort out our actual server hardware.

HARDWARE

As we have mentioned, you can run a server on virtually any computational capable device. Why then is our mental picture of a server a rack-mounted box with lots of flashing lights? Well, mostly because those are the devices most used for this sort of work – it's what they were designed for. You can however, follow this series and run just about any type of server you can think of on an ordinary desktop PC. In fact, you don't even need a separate PC, although it would be handy to have one (serving stuff to yourself isn't much of a challenge).

One area in which dedicated server hardware does have an advantage is that it has been designed for a high demand, always-on lifestyle. Your average desktop, from the PSU to the motherboard and drives, isn't engineered to be used 24/7. This doesn't stop many desktop components from performing admirably with such usage though, and if you are only dreaming of having a part time server, it's likely to be all you need.

Processor

You may be surprised to learn that for the great majority of server tasks, the CPU is probably the *least* important component. For serving static HTML, disk speed and memory play a much greater part in the overall responsiveness of the system than the CPU does, because you are just fetching data from disc and pushing it out the network port.

For databases, the distinction is even more pronounced. Although searching tables of data is more computationally demanding, the bottleneck usually occurs because of the fact that the tables need to be loaded from disk into the memory to be searched in the first place.

Most server tasks are about the movement of data, not processor intensive tasks really. There are obviously cases where the CPU plays a more important part, especially when you are dealing with dynamic web sites with lots of on-the-fly graphics generation, or running application from the server itself. For the most part though, the server can get away with a much slower processor than the one you would be happy using in your desktop. This is the ideal situation for most home users – when you have upgraded your PC so many times that you have enough bits left over to build another computer, it will probably make a really good server.

Memory

While the processor speed isn't so important for most tasks, the reverse is true for memory. When it comes to servers, you can rarely have enough of it, and it can never, ever be too fast. Memory is where the server processes live, and where all the data to be served must also live, at least briefly, on it's way to remote clients.

How much of an effect memory will have on the performance of your system very much depends on what you will actually be using it for. In terms of static HTML, it may not make a huge difference, though if your entire website can be cached in RAM, it will certainly be noticeably faster than if pages are constantly being fetched from disc.

Drives and Storage

Obviously, as an extension of what we have mentioned so far, servers are all about shifting data, so where you store that data can have a big impact. If you really need maximum performance and have a large amount of data to server, you may want to invest



Many thanks to RackSpace Managed Hosting who have donated hardware to help us complete this series. For more info about RackSpace, see page 98.

in SCSI devices. If the integrity of data and minimising downtime are important to you, it may also be worth investigating RAID.

Getting better performance from your drives is covered in some detail in our main feature on speeding up your Linux system this issue, so we won't go into it all again here.

For a basic server, any old drive will do. Usually it is helpful to keep live data on a separate drive from the OS itself, so if you have two drives, that might be helpful.

Power Supply

The power supply is obviously a critical component. You might not think of it often, but obviously, without it, no data is going anywhere. As mentioned before, desktop PSUs aren't really designed to be on all the time, but obviously experience teaches us that most of them are fine for long periods of sustained use.

Critical servers will usually have more than one power supply – either in the form of multiple hot-swappable on board PSUs, or an external uninterruptable power supply. These are usually expensive options, but worth mentioning.

“Our IT practices sometimes circumvent this security profile where we are most vulnerable – from *within* the network.”

DAVID TAYLOR, MANAGING DIRECTOR OF SYMARK EUROPE

Network hardware

It goes without saying that you'll need some network hardware to set up your server. Linux, because it grew up in the server



Post install

Updating and locking down

Whichever distro you have installed, there are a few things you should do when you first start up. Firstly, you should check to see what services your distro installer has activated by default – some distros, such as Mandrake, tend to activate *more* services by default than is considered safe, whereas others, such as Debian tend to activate *no* services by default. The trade-off here is clear: ease-of-use vs security. It is generally recommended that you aim for the practice of being “secure by default”, enabling services once you've had time to configured them properly.

Secondly, if your server is to be available externally, it's worth making sure you are as secure as possible. This takes two steps: updating, and locking down. For

Updating, you should use your package manager of choice to make sure you have the very latest versions of your software. If you have the time, you should also check through recent security announcements from your distribution maker, to ensure you have the latest information at your fingertips.

To lock down your server, there are several tools to help you out. To start with, running *nmap* will scan your local server for open ports – it's worth checking this list to make sure you haven't got services running behind your back. Next, make sure the admin password for your various services that require one is different for each service. For example, your MySQL root password shouldn't be the same as your system root password.

Then, configure your users access levels so that you have complete control as to who can log into where. For example, if you don't ever want guests to access your FTP server, disable anonymous access.

For the ultimate lockdown, consider using a tool like *Bastille Linux*, which can secure many different aspects of your system through a friendly user interface. Note that *Bastille Linux* is only available for certain Linux distributions, most notably Red Hat 7.3, 8, and 9, Debian 2.2 and 3.0, as well as Mandrake 8.x. The reason for this restriction is because *Bastille Linux* needs to be configured in exact harmony with how your system is laid out so that it can accurately analyse and amend your system configuration.

◀◀ environment, has very good support for Ethernet cards, and we'll be assuming that you want to run your server on an Ethernet network. You can pick up network hardware for under £10, as Ethernet cards are almost a commodity item these days. In fact, many motherboards ship with onboard Ethernet these days, so you might not even need to spend that tenner.

For many public server installations, it's useful to have a second Ethernet connection. This is so the 'Internet' can access over one device, thoroughly secured and monitored, and your internal services (like your own remote connection, admin and backup software) can connect over a less restricted and lower demand route. In home setups it often isn't worth the extra setup and hardware costs (effectively you have to build a second network).

How your new server will fit in to your current network setup is up to you. The simplest solution is to set up a local network using an Ethernet hub device. If you have a router for external access, you can connect the hub uplink to this. Network setup is covered in the box, *LAN setup* on the opposite page.

It's important to note that less cutting-edge distros, such as Debian, do not support very new network cards, and have particular problems with gigabit NICs. Other distros, including Red Hat and SuSE, work fine.

Other bits

There aren't really any essential peripherals for a server. Apart from installation, when it can be handy, you don't need a monitor, or even a keyboard or mouse. The idea is that your server just

sits away in a corner somewhere and serves. If you have several servers, you might find it worthwhile to install a KVM switch so that one monitor, keyboard, and mouse can service many servers.

CHOOSING A DISTRIBUTION

Easily as important as which hardware you use is choosing the right distribution for your server. It isn't the same as choosing a distribution for desktop use, because you have different needs and priorities with a server. There's no overriding need to have the very latest KDE apps for example, because it's likely you might not even install KDE, or even X. You also aren't likely to install lots of new software on the server, but you do want it to be easy to update to the latest security patched versions of the software you do have installed. In general, you won't be interested in the most cutting-edge software either; you want reliability more than a few extra features.

Of course, you have to balance this with familiarity and ease of administration. If you are already running one Linux distribution on a desktop, it may be convenient to use it for your server too, simply because you know your way around (of course, this could also be the ideal opportunity to expand your Linux knowledge by trying a new distro).

Red Hat and SuSe

In many ways, Red Hat and SuSE have fairly similar offerings. Both now have a server specific distribution, available at a premium price (Red Hat Advanced Server and SuSE Linux Enterprise Server respectively). The difference between these and your desktop distro is that largely they are using older, more stable and better-tested software, and they have gazillions of support options. They also cost a fair bit of money.

If you are looking to deploy a server for your business, or have other good reasons for needing a solid support contract these may be the best options, but for the more casual sysadmin, this may be a bit over the top.

Red Hat's desktop-oriented distro, Red Hat 9, contains loads of software that you'll never need for a server environment. Since Red Hat has split off the server distro into a completely different package, it has also been more adventurous with the software included – *eg* using *Apache2* as the default web server. SuSE Personal Edition also includes loads of stuff that is of no real use on a server.

Generally though, these distros will work fine as basic server distributions, you just have to be a bit more careful when it comes to choosing which packages to install, and what services and security options are chosen.

Mandrake

Mandrake has long been regarded by many as being the best of the bunch when it comes to a desktop distro, but can it be suitable for a server too? In short, Mandrake is a bit like Red Hat and SuSE, only more so. Mandrake has always tried to include the very latest software, and perhaps has sailed a little close to the edge on occasion – including test releases of the kernel, beta versions of *Apache* and so on. Also, people coming from other distributions find that Mandrake has some funny ideas about where files should be stored.

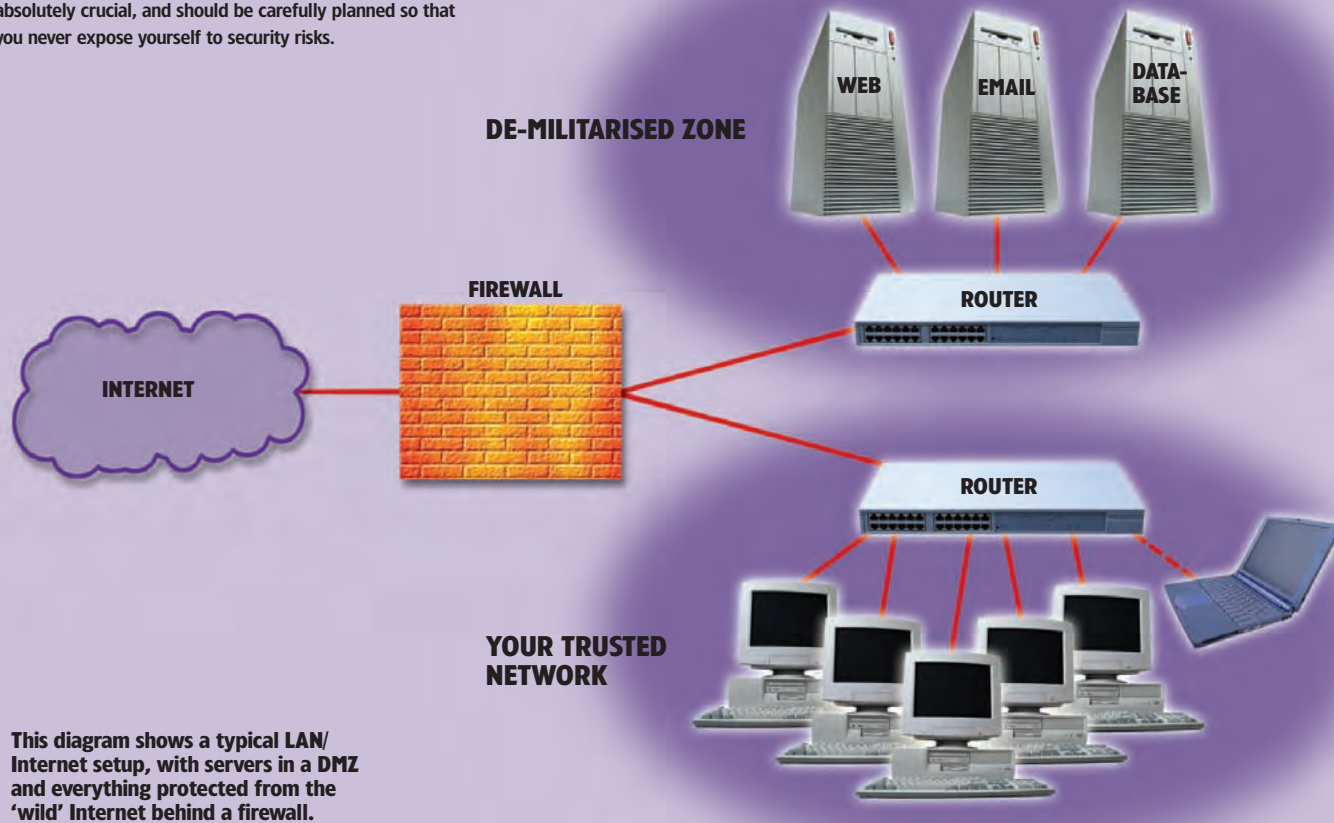
Having said that though, the *urpmi* system is great for package updates on your system – though you probably won't want to use the *cooker* updates on a live server that you want to continue working.

The SGI Altix is the fastest Linux server in the world, but you don't need a big iron to get good functionality. A normal desktop PC would suffice – indeed, Google uses a vast array of unbranded boxes to serve its search engine.



LAN setup

Placing your server in the right position in your network is absolutely crucial, and should be carefully planned so that you never expose yourself to security risks.



This diagram shows a typical LAN/Internet setup, with servers in a DMZ and everything protected from the 'wild' Internet behind a firewall.

That aside, it is easy to install and easy to administer. Some of the included tools, like *diskdrake*, the disk partition manager, are simply the best there are. Even some ISPs use Mandrake as the basis of server offerings!

For more demanding users, Mandrake produces Mandrake Corporate Server, an enterprise-level distribution of their product that is less cutting-edge, but, on the plus side, more secure, and comes with support. Server distros that provide a level of support are of course very attractive to business customers.

Debian

Almost everyone who has ever used Debian GNU/Linux will agree that it's not a very user-friendly operating system, the people who have come to *love* it realise that the reason Debian eats newbies for breakfast is because behind the scenes it's a real power user's OS. While it can be used as a desktop box, Debian particularly excels running on a server, and it does so for several reasons.


Firstly, Debian's package manager, *APT*, offers thousands and thousands of programs that can be downloaded from the web site. These programs are neatly packaged so as to avoid interdependencies between packages – if you've ever used another Linux distro to install one package and found it saying it absolutely needed to install another, seemingly random package along with it, then Debian will be a breath of fresh air. *APT* also allows you to smoothly upgrade your Debian box as new software comes out, meaning that a user on Debian 1 could have upgraded to 2, then 2.2, and now 3, simply by issuing the command **apt-get upgrade**.

Secondly, Debian has a very slow release cycle process compared to some of the other Linux distributions – each new Debian release is tested for months at a time. For example, Debian 3.0, was released just under two years after its predecessor, 2.2. In that time, no new software was added to the stable code tree – the only changes in the later iteration were bug fixes and security updates.

Thirdly, Debian's maintainers keep very tight control over what goes into the distribution, ensuring that the packages all work together well, but also, more importantly, that you have complete control over the packages available to you. For example, during installation of Debian you are asked whether you want to use non-Free software or not, and if you select 'No', Debian will only list packages that are available under a Free licence.

Overall, Debian is a distribution that requires its users to have an understanding of what they want – it doesn't offer much in the way of wizards or helpful documentation, but it is one of the few big, non-commercial distributions left, and also offers unparalleled system stability.

The Undiscovered Country

At this point you should have Linux installed on your server, locked down and ready to take on the various roles you have in mind for it. There's a lot of learning ahead, and more than a few trials and tribulations, too – setting up your server is a relatively easy task, but doing it *right* is where the real effort comes in. This first part has set the groundwork for the following instalments of the series – watch this space! 

NEXT MONTH

The LXF server school is continuing by discussing how you can turn your machine into a webserver using *Apache*.

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.

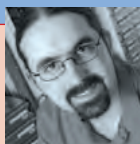
LXF answers guy
David Coulson is a networking and security guru with plenty of sysadmin experience to boot.



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



GCC vs Debian

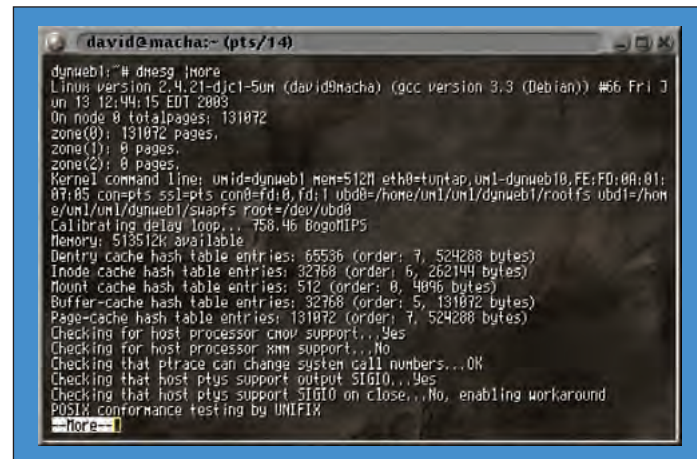
Q I was reading your review of *gcc 3.3* in *LXF42*. It explained why my kernel was not compiling on my Debian unstable box. Do you know how to fix this or when the Linux kernel will work with *gcc*?

Pete, via email

A Debian unstable has both *gcc-2.95* and *gcc-3.x* based packages, so the simplest option is to install the *gcc-2.95* package, which will remove the existing *gcc* package. We tested a build of 2.4.21 with *gcc 3.3.1* from Debian testing, and it built and ran quite successfully. More recent kernels should build with *gcc-3.3*, although those from vendors may be shipped with patches which do not compile correctly under *gcc-3.3.1*.

MDK ISO & directory

Q I have a problem copying the Mandrake 9.1 directory on my hard drive, even generating the ISO. I have a BAD CRC error when I try to copy the



Certain versions of *gcc* work great, such as this kernel built with 3.3, but some don't work quite so well with some code.

program from the DVD to the hard drive under Windows XP (dual PII – 400Mhz, 265MB, NTFS) and also under Windows Me (PIII – 1Ghz, 512 MB, FAT32)... I am lost! Even the image boot returns a 'failed to boot' message...

Any known incompatibility with Mandrake 9.1 and TYAN Tiger 100? Do you know why and what to do to install Mandrake 9.1 on my PC?

I've Red-Hat 6.2 on this PC using the 2CD install already.

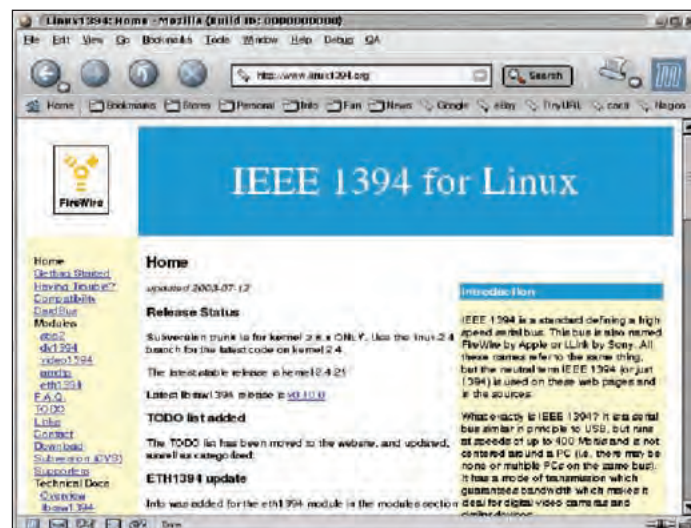
Bruno, via email

A It sounds as if the DVD you have is corrupted, hence why the CRC error is generated. If it fails to copy on two systems and doesn't boot, one can quite easily discount all hardware issues. There should be no compatibility issues with Mandrake 9.1 with the Tiger 100 board. If you take a look in the front of the magazine, there is info on how you can obtain a replacement DVD.

Documentation as to how to install Mandrake 9.1 is distributed on the CD, although there is instructions for 9.0 which is very similar available from www.mandrakelinux.com/en/demo/s/Demo/Mandrake9.0/Install/

Slackware and GCC

Q Reading your *GCC* article, I decided to give Slackware 9.0 another go. *UAE* will not compile with *GCC 3.2* (is someone working on this?), so I imported *GCC 2.95*. *UAE* compiles, but *gtk* doesn't work (is someone working on this also?). I even recompiled my kernel. Is there an organised way to upgrade sources from 2.95 to 3.0



<http://linux1394.org/> contains a large amount of info for users of FireWire on Linux, but the development code may be required for some devices.

and above? *uae* (0.8.19) is about 4000KB and a bit complicated – there's some code generating code etc and 3.2 just stops. So it's back to Slackware 8.1 for now!

J Gray, London

A Certainly things don't compile with *gcc* 3.2 because of problems in the code being compiled, or because of problems with *gcc* itself. *gcc* 3.3 is a little better, but you may experience more success with 2.95, which almost all code should compile under. If you want a nice way to organise code, maybe look at using a package manager or switch to a different distribution. There is no need to recompile everything when upgrading *gcc*, as 2.95 and 3.3 are supposed to be binary-compatible.

When compiling code, one must be aware that if there is a package manager in use, it may install binaries and libraries in different locations than a regular compile would. It's always a good idea to remove any packages before installing something which was compiled by hand, or look at using the package manager to build the



Both Crossover and WineX improve upon the usefulness of Wine for those switching from Windows to Linux.

package for you and install the resulting package.

No acceptable path

Q I just started using Linux (Mandrake 9.1) and I'm a bit new to the *Konsole*, so I could use some help. I get the message while running `./configure` for *Enlightenment* that:

No acceptable cc found in \$Path

What does this mean? Are there additional files I need to download? Would it help if I got the RPMs for Mandrake 9.1? I'm not quite sure what RPMs are used for yet. And, last question, if I do need to install some files, what command do I use?

Michael Gillespie, via email

A Before you can compile anything, you will need to install the development

packages, which include *gcc*, *glibc* headers, and a selection of other packages. These will all be available on your original distribution installation CDs or DVD.

Compact Flash

Q I have Mandrake 9.1 because you put it on your coverdisc. Now I want to end my Windows dependency. I bought the Compact Flash reader because my Ricoh camera does not work with Linux, and it needs a special driver in Windows – in fact, Windows has been behaving a little better since I unloaded the Ricoh programs! I know what I have to do to get another scanner that works. I have read on the net how to get *Nero* running under Wine and will be attempting it soon. The last piece in the Jigsaw is getting my Digital Camcorder working. I know that both the interface card and the camera itself are compatible (from www.linux1394.org). However, I do need to know (in language I can understand!) how to recompile the

A QUICK REFERENCE TO: Zebra

Anyone who works with IP networking a great deal will know how difficult it can get to maintain a selection of systems so that routing remains up-to-date on all systems. As always, it is possible to distribute these routing tables between systems automatically, so all we need to do is add a route on one box and have it distribute across the network to everything else. This functionality is provided by a number of different protocols, the main two being RIP and OSPF.

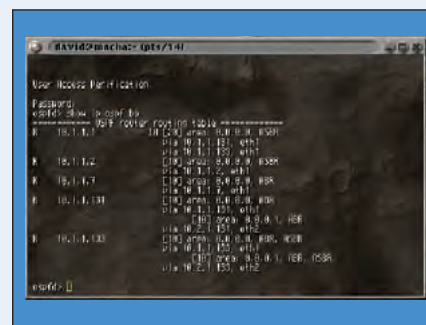
To use either routing protocol, YOU require a routing daemon that can handle the desired protocol. For Linux, Zebra (www.zebra.org) is a great system, which supports RIP, OSPF, BGP4 and the IPv6 variants of each of those. Each protocol has different capabilities and ease of use. RIP is very basic, easy to setup and is designed for a simple linear network. OSPF on the other hand is a little more complex, but is designed for systems that may have multiple routes to the same location. We're going to look at OSPF, since it's a little bit smarter than RIP when it comes to doing things.

Zebra is available as a package for most distributions, so once that is installed we need to have it start *zebra* as well as *ospfd*. *Zebra* is the daemon which interfaces with the kernel routing table, and adds and removes routes

from the system. *ospfd* talks to *zebra* and talks the OSPF protocol to other systems on the network.

The configuration for *Zebra* is very similar to that of Cisco IOS systems, and one can telnet to port 2601 to configure *Zebra* and 2604 to configure *ospfd*. We can run `show` commands to find out what the routing tables are currently doing:

```
ospfd> show ip os
== OSPF network routing table ==
N 10.1.1.0/24 [10] area: 0.0.0.0
    directly attached to eth1
N IA 10.1.7.0/24 [20] area: 0.0.0.0
    via 10.1.1.7, eth1
N 10.1.254.0/24 [20] area: 0.0.0.0
    via 10.1.1.133, eth1
N IA 10.2.0.0/16 [20] area: 0.0.0.0
    via 10.1.1.131, eth1
    via 10.1.1.133, eth1
N 10.2.1.0/24 [10] area: 0.0.0.1
    directly attached to eth2
N 10.2.254.0/24 [20] area: 0.0.0.1
    via 10.2.1.133, eth2
N IA 10.4.7.0/24 [20] area: 0.0.0.0
    via 10.1.1.7, eth1
N 10.4.9.0/24 [20] area: 0.0.0.0
    via 10.1.1.9, eth1
```



Zebra allows Linux to utilise routing protocols made available on a network by Cisco or other high-end networking kit.

```
== OSPF router routing table ==
R 10.1.1 IA [20] area: 0.0.0.0, ASBR
    via 10.1.1.131, eth1
    via 10.1.1.133, eth1
R 10.1.2 [10] area: 0.0.0.0, ASBR
    via 10.1.1.2, eth1
R 10.1.7 [10] area: 0.0.0.0, ABR
    via 10.1.1.7, eth1
R 10.1.131 [10] area: 0.0.0.0, ABR
    via 10.1.1.131, eth1
    [10] area: 0.0.0.1, ABR
    via 10.2.1.131, eth2
R 10.1.133 [10] area: 0.0.0.0, ABR, ASBR
    via 10.1.1.133, eth1
    [10] area: 0.0.0.1, ABR, ASBR
    via 10.2.1.133, eth2
```

FREQUENTLY ASKED QUESTIONS INTERNET CONNECTIVITY

FAQ HOW DO I CONNECT TO THE INTERNET WITH A MODEM UNDER LINUX?

There are many different ways to setup a dialup connection with Linux. Assuming we want a standard PPP connection, then we can use something like *kppp*, *gnome-ppp*, *wvdial*, or one of the numerous distro-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* to make it all happen when we **ppp-on**. The most basic manual configuration would be:

```
/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
OK ATDT08452121666
ogin: username
word: password
ocol: ppp
```

FAQ WINDOWS THINKS MY MODEM IS ON COM1, BUT WHERE IS IT 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 dial-up services, including those who use ISPs which use BT's SurfPort or WebPort dial-up pools, will need to use CHAP.

FAQ I HAVE A WINMODEM; WILL THAT 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 documentation, and links to kernel modules for a number of common chip sets, although your mileage may vary.

FAQ HOW DO I USE MY CABLE MODEM?

```
david@macha:~ (pts/13)
macha:~ # mii-tool
eth0: negotiated 100baseTx-FD flow-control, link ok
macha:~ # mii-tool --help
usage: mii-tool [-v] [-A media,...] [-F media] [interface ...]
-v, --version          display version information
-p, --verbose          more verbose output
-r, --reset            reset MII to poweron state
-r, --restart          restart autonegotiation
-w, --watch            monitor for link status changes
-l, --log              with -w, write events to syslog
-A, --advertise=media,... advertise only specified media
-F, --force=media,... force specified media technology
media: 100baseT4, 100baseTx-FD, 100baseTx-HD, 10baseT-FD, 10baseT-HD,
(to advertise both HD and FD) 100baseTx, 10baseT
macha:~ #
```

mii-tool makes it very easy for a Linux user to check, or change, the current link configuration for a NIC.

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 occasionally some don't like working with certain DHCP servers.

FAQ WILL MY FROG WORK WITH LINUX?

If you're 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 (*pppoe*) 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 comprehensive HOWTO can be found 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

kernel specifically with this in mind. How do I substitute the latest 1394 and dv drivers? How do I compile them separately from general compilations? Is there an alternative to putting them right in the kernel? I have been messing with computers for years, but I find most of the articles and help files I read assume a base knowledge that I don't have. I've tried following instructions line by line but found they did not work. So, as many of your readers will have Mandrake (due to the coverdiscs) why not base a tutorial/project around it, to get digicams working? You could include variations for other kernels and dealing with 3rd party (eg

nVIDIA nForce) drivers. I would be most grateful!

Dave Spagnol, via email

Mandrake 9.1 should come with a recent 2.4 kernel, which will already support the 1394 subsystem out of the box. You should be able to do:
modprobe ohci1394 ieee1394 dv
There are also 1394 drivers available from <http://linux1394.org/>, though those are considered to be bleeding edge kernel modules.

Give the standard kernel modules a try, upgrade to the latest Mandrake packages if necessary, then move to the development code. By running the latest 1394 code, it means that you will have to patch the Mandrake 9.1 kernel

tree, rather than the standard 2.4 tree, which may cause compilation issues. Good idea for a tutorial – we'll consider it for inclusion in LXF in the future!

Grubbing around

How do I install the boot loader in MBR (Master boot record) Grub?

I am about to install Mandrake 9.1 from the CD with LXF41.

Andrew, via email

Grub is reinstalled using the `grub-install` command, referencing the */dev* entry referring to the device from which the system will boot from.

```
# grub-install /dev/hda
```

However, when you install Mandrake

9.1 it will install Grub for you, so there is no need to do it manually.

Winmodem wonky

I have recently installed Mandrake 9.1, running as a dual boot system with Windows XP on a Mesh PC with an AMD 1800 MHz processor. Everything seems to work OK (Epson Stylus printer, Canon scanner etc) but unfortunately the modem does not.

The system has an on-board 'Best Data' Data Fax modem, which works fine under Windows XP, but under Linux seems to be recognised as a Winmodem (which I assume it is not) and then fails to connect.

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 using either *iptables* or *ipchains*, although the prior should be chosen if you don't already have an *ipchains* configuration, or if you have an *ipchains* configuration and don't 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 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
```

Is there a generic modem driver I can use to be able to login under Linux? If so where can I get it? Mandrake software tells me that this modem is unsupported – which means having Linux installed is essentially useless if I have to revert to Windows to connect to the Internet. I have tried browsing the Net (under Windows) but cannot find any suitable solutions.

Chris Goddard, via email

Though your modem's name and designation doesn't suggest such, it is quite probably a Winmodem, so the chances of making it work with Linux are somewhere between slim and non-existent, probably closer to the

latter! The fact that Mandrake recognised it as a Winmodem was probably from the */proc/pci* output, and doesn't mean it actually supports it at all. If it was a hardware modem, Mandrake would simply use it like a regular serial port.

Short of buying another modem, there is probably isn't much you can do, unless you are able to discover the exact chipset and look at <http://linmodems.org/> to find out if it supported by any third-party code. It's somewhat rare to find a hardware-based PCI modem anymore, and many are based on rather obscure chipsets, so unfortunately, you may be quite unsuccessful making it work with Linux.

WineX and Wine

I use a dual-boot machine, with my brother using the Windows XP bit, and me the Gentoo Linux bit. I haven't booted into Win* for more then a year now, but sometimes I just want to play a Windows oldie or something else that's on my brother's partition. The questions I have now are:

Can I use WineX to run applications/games that are already installed on the win-partition AND those that I installed in Linux using WineX?

You said that WineX runs apps that Wine doesn't... So actually WineX should run everything Wine does plus the games and some other apps, or does Wine have something that WineX doesn't?

Matija Suklje, via email

You can most certainly run programs installed under Windows XP with WineX, although the success you experience will be dependent upon how well WineX supports whatever you're trying to run. You can simply mount your Windows XP NTFS or FAT32 filesystem under Linux and configure that to be your C: drive in WineX.

WineX is based on the Wine code, but supports many other applications, just as *Crossover* does. Whether Wine can do more than either WineX or *Crossover* depending how many other modifications were made to the Wine code since it was used as a base for other Wine products.

Generally both WineX and *Crossover* contain fixes from previous Wine versions, so there usually isn't much in Wine that is lacking in something else.



f-prot (<http://f-prot.com>) is a command-line virus scanner for Linux, which can be used with a variety of mail scanners.

Getting Woody

I've installed Debian 3.0 Woody on a new partition and am having trouble getting X to run, giving the message no screens found

On another partition I have Mandrake 9.1 which runs X well (using the nVIDIA drivers for a Riva TNT2 card).

My mobo also has an onboard i810 graphics chip, which I'm perfectly happy to use under Debian, at least until I install the nVIDIA drivers.

My question is: I can copy the Monitor section from the Mandrake XFree86Config-4 into the Debian one and any other needed sections (not the lines which load the nVIDIA drivers, obviously). Will this work? Are there any distro-specific parts of XFree86Config-4 that would prevent this from working?

Also, how do I confirm that the i810 kernel module has been compiled in?

From the LXF forums

The file you are looking for is *XFree86Config-4*, which is located in */etc/X11*. This file is not distribution specific in the slightest. Assuming that you are pointing your X installation to the correct configuration, as *XFree86* will search a few patches for *XFree86Config-4*, then you will be able to copy it between different boxes without any problems. There will be more errors before you get to the point of having no screens, so watching the actual log output should be indicative of what the

problem is. You can do:

```
lsmmod
to find out if the i810 kernel module is
loaded, and if it is not then you can do
modprobe i810
```

The module is not required to use X on such a chipset, only to provide direct access to the video memory of the device to improve performance.

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



Hardware firewall

Q In order to protect my file server (running SuSE 8.1 and *Samba*), two desktops (running W2K plus Mandrake 8.2 and W2K plus SuSE 8.2) and my laptop (Mandrake 9.1 plus W98SE) from outside attacks, I use an SMC Barricade 7004 router with an in-built firewall between the ADSL modem and my home network.

Win2K uses the ZoneLabs Free Firewall, but the log files never mention any attack from the outside. My questions: "How secure is a hardware firewall? Is it not possible to hack/crack it? Could you suggest some basic rules for setting up a hardware firewall?"

Henrik, via email

A I'm not familiar with the SMC Barricade 7004 but there are a couple of things you should be able to check to make sure your internal network is truly blocked. From a little web research, I can see that the router supports network address translation (NAT). This means that all your internal machines use a non-routable IP address (such as a 192.168.x.x) and your connection to the wild Internet uses an IP address assigned by your ISP. Because your internal machines are not contactable directly by IP address, they cannot be attacked directly. Some firewalls with NAT allow you to set up port redirection to explicitly allow access from the Internet to a machine on the non-routable network. If you didn't enable this explicitly, then I wouldn't worry about it. With no port redirection set up, it is very difficult for a casual script-kiddie to attack any of the computers behind your firewall. The fact that you've seen no incoming attempts on your Windows server would mean you'd be pretty safe assuming that the software firewall is recording possible attacks correctly.

Mother-bored

Q I've recently upgraded the motherboard in my server from a Tyan 2462UGN to a

Tyan 2469UGN. Both boards have onboard Adaptec SCSI controllers. From Google I've found that the old motherboard uses the *aic7xxx* module and the new one uses the *aic79xx* module. I've modified my *modules.conf* to this effect but when I reboot the server it still seems to try to load the old module and hence cannot find any of its SCSI hard drives and will not boot. I can verify that this is definitely the correct module as I have an identical server with a SCSI tape drive and an IDE hard drive and I can see the tape drive perfectly.

Josh, via email

A I believe the problem is that you have an initial ramdisk set up with all the modules that your kernel requires to bring the system online. However, since you've changed a module and have not rewritten the initial ramdisk, these changes are of no effect. To create a new ramdisk with the appropriate modules loaded, you'll need to run *initrd* to create the new ramdisk from the information in *modules.conf*:

```
[root@scarlet root]# mkinitrd -f
/boot/initrd-2.4.20-9.img 2.4.20-9
```

This assumes that your kernel version is 2.4.20-9 and your *initrd* file resides in */boot*.

Denial of Service

Q Until recently, I've had had a very smooth sailing when using Red Hat Linux. You see, I run a controversial website and on my previous Win2K system I had many break-in attempts, two of which were successful and caused my site to be defaced on each occasion. Since moving to Linux (and with a little assistance from my friends) I've made the site stable and secure. I've recently been receiving a new form of unwanted traffic, traffic being the operative word. My service provider has informed me that they suspect that I'm at the receiving end of a Denial of Service attack. They say that the source IP address is being spoofed randomly so there is very little that they can do about it. One of the guys there mentioned something about tcp syn cookies. But he was a Solaris administrator and he said he wasn't sure if it was available for Red Hat Linux. If it is, how do I configure and use it?

Simon, via email

A The type of DoS attack that tcp_syncookies would protect you from is a SYN flood attack. This is when an attacker sends

many spoofed SYN packets to your server, your server will acknowledge the request to open a connection and send back a SYN.ACK packet. Under normal conditions the attacker's server would send an ACK packet back to complete the handshake and establish a connection. However, as the source address is spoofed, the ACK packet never comes back, and your server has to keep a list of all the incomplete connections. There comes a time when the server cannot track any more incoming connections and stop accepting connections all together.

TCP syncookies functionality prevents attackers from blocking up your server by filling up the list of connections by using a clever algorithm and a little cryptography in the TCP sequence numbers.

If you're using a fairly new version of Red Hat, you'll be pleased to know that TCP syn cookies support is compiled into the kernel out of the box. To check if this functionality is enabled simply cat out

```
/proc/sys/net/ipv4/tcp_syncookies
```

By default this is disabled but to enable it simply type

```
echo 1 >
/proc/sys/net/ipv4/tcp_syncookies
```

This setting will be reset if the system is rebooted so to make the change permanent just edit */etc/sysctl.conf* and add the following line:

```
tcp_syncookies = 1
```

About Rackspace Managed Hosting

Rackspace Managed Hosting™ provides complete managed hosting solutions with Fanatical Support to small to medium-sized customers as well as large enterprises. All customer platforms include state-of-the-art data centres,

customised servers, burstable connectivity, 99.999% uptime SLA, a dedicated account manager, instant emergency response and access to live expert technicians 24x7 for support of all hardware and core software. Founded

in 1998 with locations in San Antonio, Texas and London, the company manages over 7,000 servers for customers in more than 80 countries. For more information, visit www.rackspace.co.uk and see this month's *Linux Pro*.

The screenshot shows the Rackspace Managed Hosting website. The top navigation bar includes links for Home, Products & Services, Fanatical Support, Data Centers & Network, and About Us. The main content area features a large banner for 'Fanatical Support' with the text 'It's not about SERVERS, it's about SUPPORT.' and a photo of a woman, Jessica, a Rackspace Fanatical Support agent. To the right of the banner, there are several promotional boxes: 'Did You Say 100%?' highlighting 100% Network Uptime, 'Happy Birthday Rackspace' celebrating 4 years, and 'Move to Rackspace and Save' showing potential savings. The bottom of the page lists various services and awards.

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 synchronise 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 synchronise. 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 **Sherina** – your new Zaurus PDA will be with you shortly!

Q Enough of the small talk, how much for the monkey? I'm a network administrator/office manager in a small firm, which uses commercial solutions almost universally. However my boyfriend has slowly, over the last few years morphed into a Linux geek, much to my (sometimes) chagrin and has taken to forsaking all else and others (including me!) at times.

If Windows is Barbie, then Linux is Lego to my boy. So here is my question; while I readily admit that Windows is irritating, it gets the job done (eventually). Erm, well, yes, it does but don't ask me (I'm just a girl!), ask the millions of home and corporate user the world over who use the many handy products that Microsoft produce – Access, Outlook etc to name a few. "Ah," I hear you cry, "What about the excellent StarOffice and OpenOffice.org and others?" True, there are many very capable substitutes out there, but like a lot of the Open Source offerings, they are limited due to their nature of being Open Source. What if Johan P. Supergenius who's working on a new way to remotely access my network gets hit by a bus

tomorrow? Will Wilhelm D. Mastercoder take over or is it dead in the water? To be sure, Bill Gates can easily release another drone from the vats. I'm not knocking Open Source by any stretch of the imagination, but let's face it, besides maybe Red Hat etc, how many Open Source projects inspire the type of confidence which products like Microsoft Office inspire in people due a very human factor – the people who develop them are getting paid for it! Sad as it might seem, people seem to take comfort in the fact that they have paid for something rather than getting it for free. So here are some points to muse:

1. I don't really want to play the latest games (Meet Lego boy!).
2. I'm not a master coder-cum-hacker (Hmm, have you met Lego boy?).
3. I just want to get things done. These are the main points that the average small network administrator/IT manager thinks about when it comes time to decide on what software she needs.

Open Source equals free equals unreliable to the average organisation. Linux needs the

support of not just big business but it also needs to stop treating me like I know what many of the arcane features do. I don't. Nor do I care. I want a small tidy network, with as little growing pain as possible. So when am I going to see true alternatives to the commercial products which dominate the average office that are going to replace Microsoft's dominion over these areas?

Please answer soon as I'll be lectured too until I give up on the ('demonic' – his words) Microsoft, and I jus' cannae do it, Cap'n! Sherina, via email

A Although this page is a System Administration Q&A, since you *do* manage a small network, I'm going to break all the rules and take your bait!

The first point you raise appears to be related to the loose organisation of the open source community. It is true Microsoft and other corporate software vendors have large sums of money to spend on development, and they can just hire a new developer if one should get hit by a bus or have a change of heart. However, while it may seem that Linus Torvalds writes the Linux kernel and Andrew Tridgell wrote Samba, these people are just the

people who co-ordinate often thousands of contributors and have the final say. If they got bored, then someone else (almost always an active developer) will take over. If Andrew Tridgell decided to do something bizarre with Samba and nobody liked what he was doing, the project would fork and somebody else would lead the development of the breakaway project. At a corporate such as Microsoft this sort of thing could never happen. If for some reason Bill Gates decided that he didn't like a certain feature of MS Office it could be removed, nobody would have any power to do anything about it.

There is a dogma attached to software today that "if you didn't pay for it, then it's no good." With the lengthy disclaimers attached to just about all commercial software, you don't really have many more rights than you do with Open Source software. Depending on your administration needs, you will find that there are many useful tools available for Linux. You don't even need to know how to use the command line to use them. A program many people find useful is Webmin – this allows you administer virtually all of your servers' functionality from a secure web page.

Essential disc info

Read this important information before you use your *Linux Format* coverdisc – CD or DVD.
We've collated some helpful info to help you get the most from these jewels of data!

FINDING THE ESSENTIALS

Missing something?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases, the latest libraries and

other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

Package formats

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

Documentation

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

WHAT ARE ALL THESE FILES?

If you are new to Linux, you may find the profusion of different files and extensions confusing. As we try to give as many packages as possible for compatibility, there will often be two or three files in a directory covering different types of Linux, different architectures and usually source and binary versions – so which do you install? They can be identified by their filenames, and usually just by the file extensions.

Someap-1.0.1.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.0.1.i386.deb – The same, but a debian package.

Someap-1.0.1.tar.gz – This is usually source code.

Someap-1.0.1.tgz – Same as the above, tgz is abbreviated form of tar.gz

Someap-1.0.1.tar.bz2 – Same, but uses bzip2 compression instead of zip

Someap-1.0.1.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.0.1.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.0.1.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7.x PPC Linux.

Someap-devel-1.0.1.i386.rpm – A development version.

INSTALLING FROM TARBALLS

A tar ball is a two stage archive. First the files are archived into a single file with tar and then compressed with Gzip or Bzip2. To unpack, cd to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
```

```
tar xvf -bzip2 /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in .tar.gz or .tgz, and the second for Bzipped files, ending in .tar.bz2 or .tbz2. Naturally, you change the paths to suit the location and name of the archive. and replace /mnt/cdrom with whatever is applicable to your system (eg /cdrom). This normally unpacks the archive into a directory of the same name, enter that directory with:

```
cd progname-2.1.0
```

To compile and install the software, type the following three commands:

```
./configure
```

```
make
```

```
su -c "make install"
```

The last line will prompt you for the root password, as this stage must be run as root. If you are already logged in as root, just type **make install**. This will give you a default installation. If you want to change any aspect of the install, type

./configure --help to see the options available. For example, you are usually able to change the default location with the PREFIX argument. When you have finished installing, you may remove the source files with:

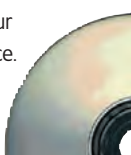
```
cd ..
```

```
rm -fr progname-2.1.0
```

You should also log out as root, before you do anything you may later regret.

DEFECTIVE CDs

In the unlikely event of your disc being defective please email our support team (support@futurenet.co.uk) for further assistance. If you would prefer to talk to a member of our reader support team please call **01225 822 743**.



CREATING INSTALL CDS WITH CDRECORD

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord -scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v  
/path/to/image.iso
```

You can simplify the command by saving some default settings in /etc/default/cdrecord. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the "Burn" tab and the "ISO 9660 Image" gadget near the top of the window. Put the path to the image file in the gadget and press "Combust!". Now put on the kettle while the CD is created for you.

Other OS?

You do not have to use Linux to burn the ISO to a disc. All the 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 does have one, and a DVD drive, and use the CD burning software on their computer. It can be Windows, Mac OS, AmigaOS whatever.

No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions, it is also possible to mount the images and do a network install, or even a local install from another disk partition. The methods often vary between distributions, so check on the distro vendors website for more information. [LXF](#)

Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed DVD – including a Linux distro that runs from Windows on a 'virtual disc'.

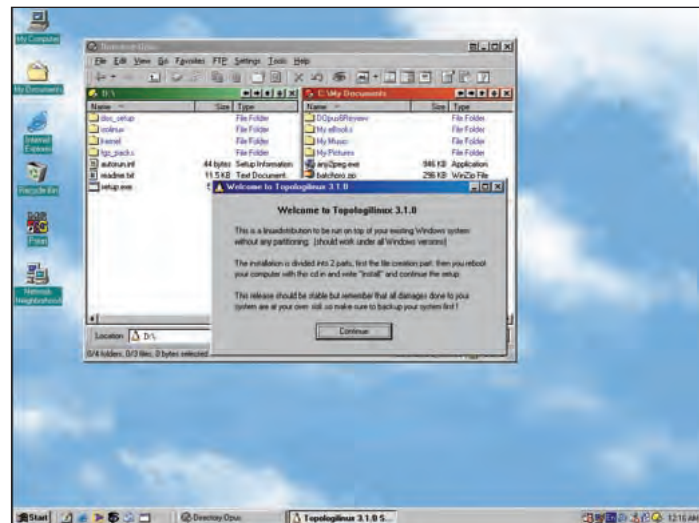
The greatest barrier to trying Linux for the first time is probably the need to repartition your hard drive in order to make space for it. This is a potentially risky operation that often involves the use of relatively unfriendly command line tools. One alternative is the 'Live CD' distribution, like the excellent Knoppix. These run direct from CD, so are ideal for trying out Linux, but sooner or later you still need to resize your Windows partition in order to install it.

Space is not the problem for most people, PCs are usually supplied with large hard drives, set up as a single Windows partition, with plenty of free space. The problem is that this space is not directly usable by the vast

majority of Linux distributions. Topologilinux sets out to change that. Instead of installing into its own dedicated partitions, this distribution uses a 'virtual disk' contained in a file on a Windows partition. The distribution is based on Slackware, so it has a good, solid foundation. Installation is in three stages. First the virtual disk is created, this part is done while running Windows. Then you install the base distribution to the virtual disk. Finally, you boot into your new Linux system to set things up.

Create a virtual disk

I never thought I'd be writing "start by booting Windows" in here, but that's what you do. Once Windows is running, put the first CD (created from Topologilinux-3.1.0-full.iso, see page 107 for information on creating CDs from ISO files) into your CD/DVD drive and press the 'Continue' button on the dialog that pops up. If the CD doesn't autorun, run setup.exe. Now select the location and size of the virtual disk you want to create. There are a couple of limits to consider. You need a minimum of 2.5GB for a full



Using Windows to create the virtual disk used by Topologilinux.

installation, plus some space to save files, 4GB is a good minimum if you have the space. If you are using the FAT filesystem, the maximum size of a single file is 4GB, so your 'disk' can't be bigger than this. This isn't a problem if you are using the NTFS filesystem with Windows XP or 2000. Press 'Create files' and wait, possibly for quite a while.

Once the virtual disk has been created, restart your computer with the CD still in the drive. Provided your computer is set to boot from CD, you should see some text and a boot: prompt. Type **install** and press Return to start the installation proper. This is based on Slackware and is very straightforward, if a little basic looking. After logging in and running **setup**, you just follow the prompts, the default options are correct for almost all users. When the basic installation completes and the system reboots, make sure the CD is still in the drive and press Return at the boot: prompt.

Your new system

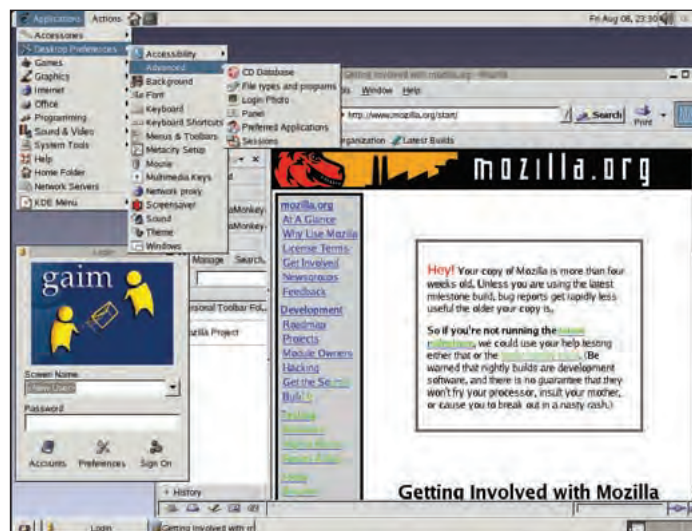
Now it is time to set up your new Linux system, so log in as asked and type **setup** at the prompt. You will be asked a few questions about your hardware, when in doubt accept the default offering. Once setup is complete, you need to set a root password and create a normal user for general use, you only log in as root for system administration. To set a root password, type **passwd** and answer the questions. Then add a normal user with **adduser** and give the information



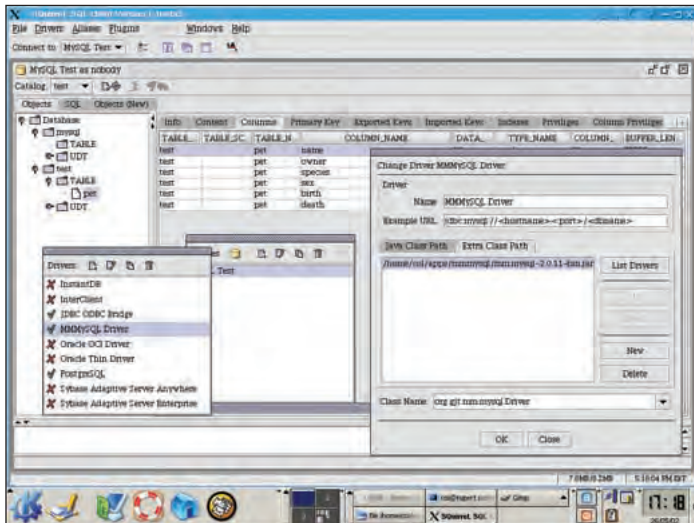
Wherever you see this logo it means there's related stuff on the DVD

IMPORTANT NOTICE

Before you even put the DVD in your drive, please make sure you read, understand and agree to the following: The **Linux Format** DVD is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker on ANY new software. While every care is taken in the selection, testing and installation of DVD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system which may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all important files. Please read individual licences for usage terms.



Topologilinux has a huge range of progs, almost 2GB compressed onto the first CD with plenty more on the second. Here it's running the GNOME desktop.



A Java front-end for SQL is ideal for a cross-platform office environment.

asked for. The login name is a single word, usually all lower case, you can add your real name later. Now you can start the graphical environment with **startx**. If you want to change the settings for the GUI, for example on the computer I used for testing it came up in 640 x 480 resolution, log out of the GUI by right-clicking on the desktop and selecting 'logout' (this is for the default of KDE, there are similar options with other window managers) and type **xf86config**

Once you are happy with the GUI, you will want it to come up automatically when you boot Topologilinux. To do this, go back to the command line prompt and type **joe /etc/inittab**. Find the line that reads **id:3:inittab**, change the 3 to a 4 and press **Ctrl-K X** to save. Now your system will boot directly into the graphical environment whenever you have the Topologilinux CD in the drive. Without the CD, it boots into Windows as before.

You now have a full Linux distribution running alongside your existing Windows setup, without touching your disk partitioning. It will run a little slower than if it were on a proper Linux disk partition, but otherwise it will show you exactly what you can expect from Linux without altering your hard drive's layout. If you ever want to remove Topologilinux – maybe to install Linux in its own partitions – all you need do is delete the *tlinux* folder from your hard disk.

Installing extra progs

The second CD is not used during initial installation. It contains a number of extra packages you can install once you have the system up and running.

Most of these packages are stored in .tgz archives. You install these by opening a terminal window, typing **su** to switch to the root user (remember that you don't normally log in as root) and then you should type **installpkg /mnt/cdrom/path/package.tgz**. For example, to install the kernel source, often needed to install other programs, type

```
su
mount /mnt/cdrom
[enter root password]
installpkg /mnt/cdrom/Kernel-Source-2.4.20-ntfs/kernel-source-t-2.4.20-noarch-1.tgz
```

Don't worry about typing the long path to the package – if you type the first few characters and press the TAB key, the shell will complete the path or file name as far as it can. Then type the start of the next path and press TAB again.

Internet/SendYmail

Many people prefer to use somewhere like Yahoo for their email accounts rather than using the address given by their ISP. Keeping your email address separate gives the freedom to change one without affecting the other. Yahoo tried to spoil this by charging for POP3 and SMTP access, only leaving the advertising-subsidised web access for free. Of course, Linux users are too resourceful to let them get away with that, so we have programs like *FetchYahoo* and *YoSucker* to download the mail ready for reading by a normal email program.

SendYmail takes care of the opposite direction, sending mail. It emulates an SMTP server that then sends you mail through Yahoo as

though you were connected over the web. Now you can go back to using *KMail*, *Evolution*, *Sylpheed* or *Thunderbird* with your Yahoo mailbox.

MOBILE KNOKII SYNC

Before computers helped us get organised, we had names, addresses and phone numbers scattered across various address books, diaries and scraps of paper. Now, thanks to the wonders of modern electronics, we have names, addresses and phone numbers scattered across various computers, PDAs and mobile phones, and probably the odd few scraps of paper too! While the DVD doesn't have a GPL program to manage your scraps of paper (although I expect someone somewhere has attempted to write one) we have had desktop to PDA programs in the past, and now we have *KnokiiSync*.

KnokiiSync was written to sync a Nokia 6310 mobile phone with the KDE address book. Because most Nokia mobile phones use a common set of commands and protocols, it should work with the majority of their phones. It supports Nokia's method of storing multiple phone numbers, as well as email and postal address, for each entry, and will transfer these to and from the KDE address book. Installation of *KnokiiSync* itself uses the standard **./configure && make && make install** process, but before you can begin installation, there are a couple of other programs it requires. These are *gnokii* and *pkgconfig*, both included in the dependencies directory. The latter uses the same

installation method, but installing *gnokii* requires a couple of extra steps. Unpack the *gnokii* archive and read the *INSTALL* file before proceeding.

It goes without saying that you should also have KDE installed to use this program.

OFFICE SQUIRRELSQLCLIENT

Linux has alternatives for most Windows programs, often superior alternatives. One program that doesn't appear to have a viable direct replacement is *Microsoft Access*, the database program. Linux has some excellent relational database backends, like MySQL and PostgreSQL, but is unfortunately somewhat lacking when it comes to front ends suitable for general desktop use.

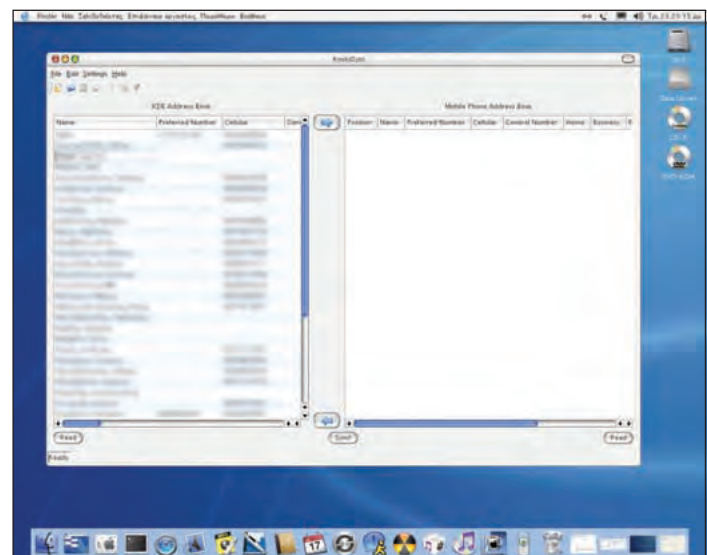
Maybe SquirrelSQL can be just what Linux needs. This is a Java program, making it ideal for use in an office that uses different operating systems, such as one that is in the process of migrating from Windows to Linux. Source code is supplied on the DVD, but the easiest way to install it is from the .jar files, with

```
java -jar squirrel-sql-1.1final1-install.jar
```

or

```
java -jar squirrel-sql-1.2beta3-install.jar
```

depending on whether you want to try the stable 1.1 version or experiment with the new beta of 1.2. The DVD also contains several plugins for SquirrelSQL. The plugin format is different between 1.1 and 1.2, so make sure you install them from the appropriate directory. **LXF**



KnokiiSync synchronises your Nokia phone with the KDE address book.

User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally natter about all things Linux. You can find lots more information online at: www.lug.org.uk

1 HAMPSHIRE

URL www.hants.lug.org.uk
Contact Hugo Mills

2 BRISTOL & BATH

URL www.bristol.lug.org.uk

3 SCOTTISH

URL www.scottish.lug.org.uk

4 OXFORD

URL www.oxford.lug.org.uk
Contact Alasdair G Kergon

5 KENT

URL www.kent.lug.org.uk
Contact Kevin Groves

6 BRIGHTON

URL www.brighton.lug.org.uk
Contact Johnathan Swan

7 WORCESTERSHIRE

URL www.worcs.lug.org.uk

8 NORTHANTS

URL www.northants.lug.org.uk
Contact Kevin Taylor

9 ANGLIAN

URL www.anglian.lug.org.uk
Contact Martyn Drake

10 MILTON KEYNES

URL www.mk.lug.org.uk
Contact Denny De La Haye

11 SCUNTHORPE & DONCASTER

URL www.scundog.org
Contact Shaun Holt – shaun@scundog.org

12 MORAY

URL www.moray.lug.org.uk
Contact Stewart Watson

13 WEST WALES

URL www.westwales.lug.org.uk
Contact Dan Field

14 WOLVES

URL www.wolveslug.org.uk
Contact Jono Bacon

15 PETERBOROUGH

URL www.peterboro.lug.org.uk
Contact Steve Gallagher

16 EDINBURGH

URL www.edinburgh.lug.org.uk
Contact Alistair Murray

17 TYNESIDE

URL www.tyneside.lug.org.uk
Contact Brian Ronald

18 LEICESTER

URL www.leicester.lug.org.uk
Contact Clive Jones

19 GREATER LONDON

URL <http://glug.linux.co.uk/>
Contact John Southern

20 SURREY

URL www.surrey.lug.org.uk
Contact Jay Bennie

21 CAMBRIDGE

URL www.cam-lug.org.uk

22 DEVON & CORNWALL

URL www.dclug.org.uk
Contact Simon Waters

23 FALKIRK

URL www.falkirk.lug.org.uk

24 MANCHESTER

URL www.manlug.mcc.ac.uk
Contact John Heaton, Owen Le Blanc

25 HERTFORDSHIRE

URL www.herts.lug.org.uk
Contact Nicolas Pike

26 WEST YORKSHIRE

URL www.wylug.lug.org.uk
Contact Jim Jackson

27 SHEFFIELD

URL www.shefflug.co.uk
Contact Richard Ibbotson

28 STAFFORDSHIRE

URL www.staffslug.org.uk

29 NORTH EAST

URL www.shofar.uklinux.net/NELUG

30 LONDON

URL www.lonix.org.uk

31 BERKSHIRE & THAMES VALLEY

URL www.sclug.org.uk

32 LIVERPOOL OPENSOURCE

URL http://linux.liv.ac.uk/_liv_linux_ug/
Contact Simon Hood

33 DEAL AMIGA CLUB

Email superhighwayman@hotmail.com
Contact John Worthington

34 CHESTERFIELD

Email spirelug@yahoo.co.uk
Contact Robin Needham

35 SOUTH DERBYSHIRE

URL www.sderbylug.org.uk
Contact Dominic Knight

36 BELFAST (BLUG)

URL www.belfastlinux.cx
Email russell@belfastlinux.org

37 WILTSHIRE

URL www.wiltshire.lug.org.uk
Contact Jason Rudgard

38 SOUTH LONDON

URL www.sl.lug.org.uk
Email edo@perceptiondm.com

39 CHESHIRE

URL www.sc.lug.org.uk
Contact Anthony Prime – enquiry@sc.lug.org.uk

40 NORTH WALES

URL www.northwales.lug.org.uk
Contact Andy Hutchings A-Wing deltaone@virgin.net

41 MIDLANDS

URL <http://midlandslug.port5.com/>
Contact Pete Thompson

42 CUMBRIA

URL www.cumbria.lug.org.uk
Contact Jamie Dainton

43 DORSET

URL www.dorset.lug.org.uk
Contact John Robinson

44 SHROPSHIRE

URL www.shropshire.lug.org.uk
Email shropshire@lug.org.uk

45 SOUTH WEST

URL www.southwest.lug.org.uk
Email southwest@lug.org.uk

46 SOUTH WALES

URL www.swlug.org.uk

47 NORTH LONDON

URL www.kemputing.net/lug/anlug-aims.html
Email jason@voyagercomputers.co.uk

48 MALVERN

URL www.malvern.lug.org.uk
Contact Greg Wright

49 HUDDERSFIELD

URL www.hud.lug.org.uk
Contact Dave Naylor – knocker@caramboo.com

50 NOTTINGHAM

URL www.nottingham.lug.org.uk

51 ST ALBANS & LUTON

URL www.lust.lug.org.uk
Contact Michael Culverhouse – mike@easily.co.uk

52 WREXHAM

Contact Paul Kersey-Smith
Email paul@pkls.fsnet.co.uk

53 PRESTON & LANCs

URL www.preston.lug.org.uk
Contact Phil Robinson

54 DERRY

URL www.derry.lug.org.uk

55 ISLE OF WIGHT

URL www.iow.lug.org.uk
Contact David Groom – info@iow.lug.org.uk

56 SCARBOROUGH

URL www.scarborough.lug.org.uk

57 BLACKBURN

Email matt@consultmatt.co.uk

58 YORK

URL www.york.lug.org.uk

59 LINCs

URL www.lincs.lug.org.uk

LinuxUserGroups

**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 SOUTHEAST-ON-SEAURL www.sos.lug.org.uk

Contact Derek Shaw

97 ORPINGTONURL www.orpington.lug.org.uk

Contact Barry Schofield

YOUNG LINUXURL www.young.lug.org.uk**SCHOOLS**URL www.schools.lug.org.uk

LUGS OF THE MONTH

Goonellabah Linux Users Group (GLUG)

Matt Nailon writes:

When one thinks 'rural Australia', the arid wilderness typified in the *Mad Max* movies is what springs to mind, rather than the lush, verdant greenery of New South Wales rainforests and rolling plantations of sugar cane, nuts, tea, coffee and fruit etc. Goonellabah is an area of Lismore, in the far North Coast of NSW at latitude 28.81 degrees south of the equator and Longitude 153.274 degrees. State capital Sydney is some 860km south by road and

Brisbane is 2.5 hours drive to the north.

The Goonellabah LUG is a group of computer enthusiasts who meet to learn and discuss aspects of GNU/Linux and Free software. Members vary in experience and skills and come from a variety of backgrounds. New members are always welcome! Meetings are currently held every third Thursday at Goonellabah Medical Centre (www.gmc.net.au/). Meetings start at 7.30pm and finish around 10.30pm. Recent topics covered

include: PHP, Lilo, *apt-get*, Wireless Community Networking Project, kernel compiling, *cdrecord*, *gumpi*, *Qmail* and Knoppix 3.2. There is an active mailing list, IRC channel, and if you want to try Linux on your PC, they can help with installations. GLUG also provides free websites and support to community oriented non-profit organisations whose aims are to enrich the physical, spiritual or social environment <http://glug.dyndns.org>



Worldwide Linux User Groups

Free Software users across the globe

Africa

EGYPT

URL www.linux-egypt.org

GAUTENG, SOUTH AFRICA

URL www.glug.org.zaEmail glugmin@revolution.org.za

THE LORD'S ABODE, JO'BURG, SA

Email Andrew Gargan avrin17@iname.com

Australia

ADELAIDE

URL www.linuxsa.org.auEmail mtippet@anu.edu.au

ALICE SPRINGS

URL www.aslug.org.au

MELBOURNE, VICTORIA

URL www.luv.asn.auContact luv-committee@luv.asn.au

PERTH

URL <http://plug.linux.org.au/>

SYDNEY

URL www.slug.org.au

Europe

COSTA DEL SOL (English speaking)

URL www.fuengirola.lug.org.uk

DENMARK

Alssund www.alslug.dkEsbjerg www.eslug.dkFyns www.flug.dkMidt-og Vestjylland www.mvjlug.dkNordjylland www.njlug.dkSkåne Sjælland www.sslug.dkTrekantsområdet www.tlug.dkVest-fyn www.haarby-net.dk/vflugÅrhus www.aalug.dk

EIRE

URL www.linux.ieEmail root@linux.ieURL www.dilu.orgContact 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.orgEmail isam@planet.edu

Asia

HONG KONG (multilingual)

URL www.linux.org.hk

SINGAPORE – SLUG

URL www.lugs.org.sg

SRI LANKA

URL www.lklug.pdn.ac.lk

MYANMAR (formerly BURMA)

URL www.myanmarlug.orgEmail afyde@balug.org

PAKISTAN

URL www.linuxpakistan.netEmail tux@clug.org

HYDERABAD, SINDH, INDUS VALLEY

URL www.geocities.com/slug_pk/

KASHMIR

Coming soon!

China

BEIJING (GB encoding, but mostly written in Chinese)

URL <http://mud.263.net.cn/~linux>

CHINESE LINUX USER GROUP

URL www.linux.org.cn

NANJING

URL <http://jllib.jlonline.com/njlug>

India

LINUX INDIA

URL <http://linux-india.org>

ALIGARH LUG

URL <http://linux.amupost.com>

BOMBAY

URL www.ilug-bom.org.in

CHANDIGARH

URL www.geocities.com/vipinb

CHENNAI AND MADRAS

URL www.chennaiug.org/

CYBERABAD (CLUG)

URL <http://seeknew.freesevers.com/clug/>

DELHI

URL www.linux-delhi.org

KOLKATA

URL www.ilug-cal.org

MADURI

URL <http://linuxmadurai.tripod.com>

NORTHERN INDIA LINUX

URL <http://groups.yahoo.com/group/lug-northindia>

Spreading the word

Free Software is obviously of importance to not-for-profit organisations. Charity begins at GNOME says Jono Bacon.

Let's face it, charities are strapped for cash and need to save money where they can. Although to me charities seem an obvious target for free software, there is surprisingly little info on advocating Linux to charities on the Internet. Part of this problem is finding charities that use free software already, and to remedy this I set up a page to allow charities to show they use free software: www.jonobacon.org/projects/charitiesregister/.

The key to advocating Linux to charities is in telling them what they will save financially and the extra peace of mind they will have with the stability and security that Linux offers. It is also important to stress the future-proof implications of using Linux in a charity – there is no upgrade cycle that they are trapped into. It is recommended that you find the specific requirements of the charities in your area that you are going to advocate Linux to, and ensure that Linux will be a better solution for them. As an example, if there is a small charity that uses a computer or two for office productivity/accounting/Internet etc, then

this could easily be moved to Linux and save them fees on licensing. Large charities may have other requirements such as servers and workstations that can run Linux and save them the all important wedge of cash.

When I was looking around the Internet, I could not find any suitable information to give to a charity to explain what Linux is and how it can help them. Due to this lack of information, I decided to write one myself and developed a 10-page document that discusses what Linux is and how it can help them, at www.jonobacon.org/writing/research/linuxforcharities.pdf. I recommend that this document is sent to charities (contact them first so they don't think what you're sending is spam!), and that you emphasise that they can make use of LUGs for help and support. Next month we will continue our look at charities and identify the key areas in which they can be helped with free software. We'll also identify the key requirements that charities have, and the post-installation support they may require.

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

LINUX PRO

FROM THE MAKERS OF LINUX FORMAT

OCTOBER 2003



COSTING effective storage solutions

Your storage options analysed – find out how the benefits of NAS, DAS and SAN stack up for your business

PLUS

Stop the spammers

Open relays, Viagra, and what the law can do to stop them

Linux Expo 2003 Preview

Get the inside track on the UK's biggest Linux event

Hosting Focus

In the first of a new series, we look at 1&1's hosting solutions

PRACTICAL LINUX SOLUTIONS FOR I.T. PROFESSIONALS

Welcome

Twenty pages of real-world Linux for IT professionals

Every day, the *Linux Format* office gets swamped by more and more emails. Sadly, these aren't all from genuine readers. Most of them are extolling us to take out loans, enlarge various parts of our bodies, order drugs or help in the laundering of monies for various West African former dictators. We even get unsolicited emails advertising software to block unsolicited emails! Yes, 'spam' is a problem that isn't licked yet. But with a lot of companies now receiving more spam than genuine email, now more than ever the law is becoming a further weapon in the fight against spam. But what are the legalities? Who can you sue? And how can you make sure your company doesn't fall foul of the law itself? Our must-read spam epic begins on page 12.

Our *Linux Pro* cover feature this issue is about storage. Always a popular topic with our readers, the storage market has probably never been more active and potentially confusing. The article looks at the supposedly rival technologies employed in NAS, DAS and SAN and looks beyond into the actual time and resources taken to manage these setups. We've finished off our storage coverage this issue with a look at Western Digital's Raptor SATA devices.

Following on from our hosting roundup in the last issue of *Linux Pro*, we've decided to take a closer look at some of the companies vying for your business. This issue we'll be talking to 1&1 and its customers to get a better impression of their services.

Also, we have a small preview of the biggest UK Linux event of the year – the Linux Expo UK 2003. Now in its fifth year, and still at Olympia 2, London, the Expo has become *the* UK event for Linux companies to demonstrate their wares, luminaries to discuss the state of the industry and users to get a chance to meet the above. If you still haven't registered there's still time – more details over the page!

Nick Veitch Editor
nick.veitch@futurenet.co.uk



“A lot of companies now receive more spam than genuine messages – how can the law fight back against these unsolicited emails?”

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
Group publisher
Dave Taylor
Publisher
Kelley Corten
Contributors
David Harris, Robert Jacques

Photography
Powerstock, Photodisc

ADVERTISING

Key Account Manager
George Gill george.gill@futurenet.co.uk
Business Development Manager
Damian Hughes
damian.hughes@futurenet.co.uk
Senior Sales Executives
Diane Clysdale, Giles Crosthwaite-Scott,
Harry Maltby
Classified Sales Executive
Stephen Hall

LinuxPro is a part of
LINUX
FORMAT
magazine

Published by
the future network
Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

OCTOBER 2003

CONTENTS

LINUX EXPO 2003

Preview of the UK's Linux event of the year **p04**

STORAGE

Weighing up the pros and cons of a variety of storage technologies **p06**

SPAM

legalities of unsolicited email – can the law stop spam? **p12**

HOSTING

We take a closer look at 1&1's services and customers **p18**



6

12

LINUXPRO 3

Linux EXPO UK 2003

PREVIEW



London's Olympia 2 is the venue for the best UK Linux show of the year. Find out who's doing what with *Linux Pro's* special preview.

Returning for a fifth time, the Linux Expo UK show is the place to be if you want to get up-to-speed with Linux, investigate new products and software or learn what the future holds for the fastest-growing OS.

Last year's event was a great success for exhibitors and attendees alike. As well as new product launches, visitors, many of whom were completely new to Linux, got to sample systems and products, chat to developers and get a clear idea of what Linux could do for them at the time and in the future. With a great conference schedule and some exciting exhibitors, this year should prove just as popular and worthwhile.

Exhibitors

Platinum sponsor IBM has been a big supporter of this show, and Linux in general, for some years now. There isn't one single product or service to look out for from these guys, they try to do it all from desktops and workstations to the full glory of the new s/990 running multiple instances of Linux for serious tasks. They'll also be keen to explain new initiatives such as 'business on demand', which allows flexible scaling solutions for processing and storage to fit in

better with seasonal business. *Linux Format* is proud to be associated with the event as a 'Premiere Media Partner' of Linux Expo UK 2003.

Linux is part of the strategy of just about every major IT company these days, and HP is no exception. From its early server offerings, HP now does complete end-to-end enterprise solutions on Linux. If it's a single server from the Proliant range you are interested in, deploying a carrier grade solution, or building a high performance cluster, HP has the hardware and the know-how to do it all on Linux.

If you want to get the latest news on what the distribution vendors are up to, and what new versions of Red Hat and SuSE might contain, they are sure to be represented here as well, probably with exclusive previews of their latest products.

As well as the leading lights of the commercial Linux world, don't forget you can also visit the .ORG area of the show to get unbiased information and advice, and meet some of the hackers and personalities that have made Linux great. The exact details of the organisations attending haven't been confirmed yet, but we're fairly sure you can catch up with representatives from KDE, GNOME, Debian and the AFFS. As well as providing information on their

DON'T MISS THE LINUX EVENT OF THE YEAR!

8-9 OCTOBER 2003
Olympia Exhibition Centre 2,
London, UK

OPENING TIMES

Wednesday 8 October 2003:

10:30-17:00

Thursday 9 October 2003:

10:30-16:30

FREE CONFERENCE PROGRAMME

Seminars and The Great Linux Debate

THIS YEAR THERE IS A THREE-TRACK conference programme to suit different types of interest in Linux, though the specific seminars have yet to be confirmed at the time of writing – see www.linuxexpo.co.uk for more.

EVALUATE

Track 1 – Enterprise Linux Case Study Theatre (Managers and Directors Only)

Vendors will demonstrate the benefits of using Linux Technologies by presenting case studies from different industry sectors to an audience of senior IT decision-makers. A telemarketing team will canvass over 5,000 Managers and

Directors from the UK's top companies.

IMPLEMENT

Track 2 – Product Education Theatre

Gives vendors a chance to present their latest products, services and solutions.

USE

Track 3 – Using Linux

Hosted by the .org/user groups, this highly technical seminar programme will be targeted at the Linux developer community.

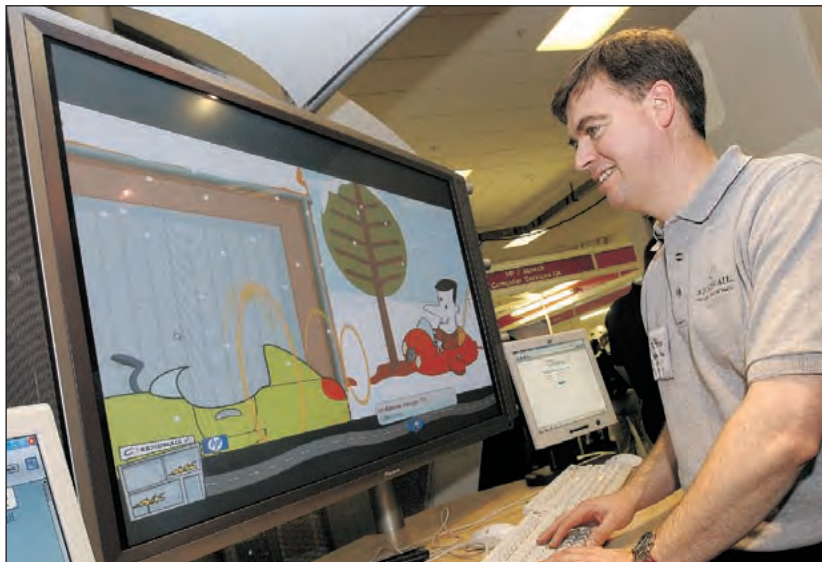
This FREE conference will provide the end user

with all the information needed to evaluate the benefits of Linux technologies. At the same time, it will give you the chance to meet and present to some of the UK's top IT decision makers.

THE GREAT LINUX DEBATE

A regular event since the first Linux Expo, this lively debate involves a panel of experts from all areas of Linux, keen to put across their points of view, predictions and suggestions to questions asked by the compere and the assembled audience. Always popular, arrive early to make sure you get a seat!

It will take place each day from 14.00–15.00.



own projects, these guys can usually give you a better idea of how the whole Linux community goes about things, and why it's great.

Linux Format welcomes you...

Of course, there will also be the chance to talk to luminaries from the world of *Linux Format*. Whether you just discovered Linux last week, or administrate a network

LEFT: Get a first look at some of the exciting developments happening in the Linux world.

RIGHT: The *Linux Format* stand welcomes all existing and prospective readers of all levels of Linux knowledge.



of thousands of users, staff from the magazine will be on hand to poke fun at, give helpful advice or just chat about all things Linux, Free Software and Open Source-related (they tend to be more talkative when plied with beer!)

Please stop by and let us know what you think of the magazine, and give us suggestions for things you would like to see covered in future issues – we always enjoy the opportunity to get to know our readers better. ■

IT'S NOT TOO LATE! – REGISTER FOR FREE

PRE-REGISTER FOR FREE ADMISSION TO THE Expo and Conference events now! From development tools, e-commerce, security and Internet tools to Linux-based training and support; this is a unique opportunity to discuss your Linux issues directly with industry leaders. Visit the Linux Expo website and complete the 3-step easy registration process – just visit www.linuxexpo.co.uk/Exhibition/Visiting/Registration/



COVER FEATURE **STORAGE**

The cost advantage

Storage used to mean filing cabinets, or bulging cardboard boxes of printouts in the cellar. In the days of hardcopy, when every business detail was kept on paper, all this paper was carefully archived into folders and packed into cabinets. And if you had lots of paper, you also needed somewhere to put the cabinets, whole rooms filled with big metal boxes. OK, so that much hasn't changed, today we still have whole rooms full of big metal boxes, but now just one of those boxes can contain the entire British Library on a disk drive the size of a hefty novel.

Every company has to deal with a never-ending and ever-growing mountain of mission-critical digital data, but designing the storage architecture that best fits the needs of your company is not a trivial matter. There's the most obvious method of just adding another server with a big hard drive onto your network, called Direct Attached Storage or DAS. Or perhaps you should look at Network

Storage solutions come in a myriad of shapes and sizes. But if you simply want to know your cheapest route to safe data, where do you start? ROBERT JAKUES reveals your best options.

Attached Storage (NAS) or Storage Area Networks (SANs)? For a more in-depth examination of NAS and SAN, please refer to issue 41's *Linux Pro*. Beyond your choice for the infrastructure architecture, it's then time to deal with the thorny question of storage management.

Definition

Before we wade into all your possible storage options, however, it might be useful to explain what the 'storage' actually means to all concerned. Before the term 'storage' was being bandied about by every hardware vendor and his dog, people used technology, directly attached to computer systems, for external data archiving, backup and restoring. The computer system in question typically consisted of a large, complex mainframe. This method of sticking another server with a big disk drive in it on your mainframe or network whenever you need more disk space has stuck

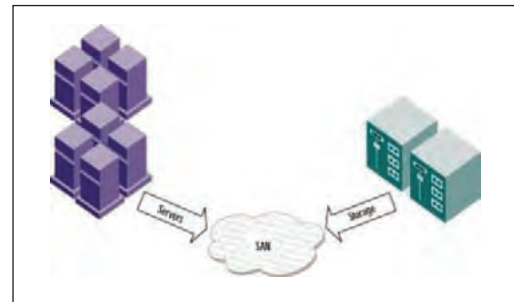
with us right up until today, now more commonly known as Direct Attached Storage (DAS). Its 'old faithful' status and simplicity of setup means that DAS is still very much the storage choice of today. But over the years we have come to realise that this model is inefficient in terms of space usage and requires a lot of time and effort as well as people to manage it.

Today's modern and data-dense landscape, marked by the advent of data centres with huge capacity and availability requirements, and the Internet, with its proliferation of online data that needs to be accessible 24/7/365, demands much more efficient technologies that can be better integrated with business processes. And so out of necessity the boffins have sculpted not just a direct attached architecture, but also two networked storage models, Network Attached Storage (NAS) and Storage Area Networks (SANs), giving us three different technologies with different places within the business. And all companies in the information age, big or small, need to invest in some form of storage for their data.

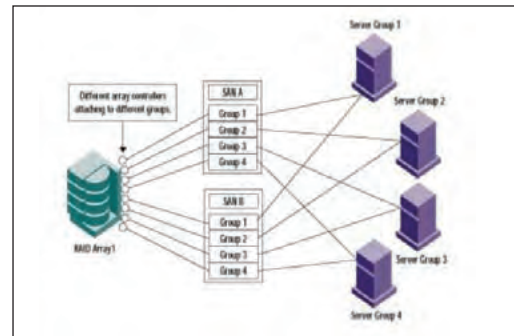
As Trevor Eddells, storage analyst for research firm Xephon, said: "You need storage in the business today, because you need to make sure that the information you have today is still there tomorrow." Developments in storage networking allow companies to deliver data anywhere, anytime, providing a sound architectural evolution away from the model of attaching storage directly to servers. By adding a storage-centric network, users gain the ability to redeploy, manage, centrally protect, and move data between servers and end users easily and efficiently. A storage network also offloads much of the traffic off the front-end network, freeing up bandwidth and improving productivity for users, whilst easing the administration and deployment of data for managers.

But today it's only enterprise scale companies that are looking towards a storage network architecture, although Jon Collins, analyst with research firm Quocirca, said that SMEs would benefit too. "Any systems implemented in the business should have an availability guarantee and any size organisation can benefit from storage from an availability perspective, not just from a cost perspective, because the advantages of external storage are largely based on risk rather than on cost," he said.

The ultra-fast Storage Area Network provides companies with simplified storage management.



By logically structuring a Storage Area Network, data throughput can be maximised between RAID arrays and servers.



Entry-level storage solutions are no bigger than a standard desktop PC.

And because data is of such importance today – indeed the data is often the business – storage cannot be ignored whatever the size of the organisation, be it an upstart dotcom or a global financial institution. As John Vitkus, IBM's worldwide program director for Linux in the Financial Services Sector puts it: "Storage is the currency of the digital economy. You can't work without data; data needs to be protected, saved, made resilient against disaster, and available when you need it."

All aboard the storage bandwagon

According to analyst firm IDC, DAS is still the most common architecture to be found in businesses today, accounting for 65 per cent of new storage implementations. But with a steadily rolling vendor bandwagon and a customer market more conscious of implementing storage strategy, the analyst predicts that NAS and SAN architectures will grow at an explosive rate over the next five years and that all three architectures will be almost equal in market size by 2004.

Quocirca's Jon Collins added that SANs are picking up market share very slowly. "This is because the architecture is more of a strategic buy for the customer and a hard-sell for the vendor in the current climate. Still only about 10 to 20 per cent of storage installations are SAN-based," he said. "NAS uptake on the other hand is growing faster because it's more of a tactical implementation and comparatively cheaper."

"All in all, the market at the moment is hard but fair. There's business to be had but competition is stiff. But the thing about the storage market is that people always need more," added Collins.

Growth in the storage market is being driven in part by the difficulty in managing large distributed databases with frequent updates and a realisation of the benefits of common features found in both SAN and NAS architecture. These are the ability to handle growth and scaling with ease; flexible deployment for any application; »



"...any size organisation can benefit from storage from an availability perspective, not just from a cost perspective..."

**JON COLLINS,
QUOCIRCA ANALYST**

COVER FEATURE **STORAGE**

« separation of storage and server purchase decision; high availability even in the event of disaster; simplified centralised management and administration; and true Unix and Windows file sharing from a single location.

Charlotte Rancourt, storage analyst at IDC, said that networked storage is being driven by the convergence of enabling technologies and needs for increased and more capable storage, reflecting the increasingly networked nature of modern business processes in general. As a result, vendors have several goals for improving their storage solutions including: increasing capacity to handle very large amounts of information; increasing availability via increased bandwidth, as well as meta tagging and improved search facilities; increased security from both physical and virtual attacks; and better business continuity and disaster recovery – networked storage can be used to mirror and replicate data, which allows improved availability in almost all circumstances.

But IDC warns that with such major shifts as these being projected for the future it's increasingly important for users to develop comprehensive strategies to optimise network infrastructure with storage solutions.

With network content expected to continue expanding at explosive rates over the next five years, analysts predict that enterprise storage will account for 75 per cent of all computer hardware expenditure. This will have particular impact on the NAS and SAN markets as stronger standards for NAS result in simpler installation and lower management cost, while increased network speeds can help shrink the performance gap that used to exist between NAS and DAS.

Standards, lies & statistics

With DAS still the most popular type of storage today, disk drives, tape, and RAID arrays comprise the majority of storage devices. These are typically attached to the server over a SCSI or Fibre Channel connection. However, the vendors and committees that established these standards allowed such wide flexibility in interoperability, such as seven variations of SCSI, that standards are weak and numerous.

THE DIRECT ATTACHED STORAGE (DAS) MODEL

A heritage solution

DAS IS THE OLDEST OF THE three models and essentially pre-dates networking. But it's still the most common method of saving and retrieving data today. A DAS device is generally described as one or more spinning or streaming devices attached to a single server via a cable, and contains three software layers, application software, file system software and disk controller software.

During the infancy of storage, islands of DAS servers were implemented in enterprises as demand required, but no forethought was given as to how data would eventually be shared. The evolution and introduction of the network revealed many pitfalls in DAS, such as single points of failure. This means that when a server fails all of the data attached to that server becomes unavailable until the server is restored. These hours or days translate into lost revenue and reduced productivity.



The Sun storage range aims to target the market from SME to enterprise.



IBM attempts to take a bite out of the enterprise storage market with its Shark SAN switches.

Unfortunately, the SAN vendors followed the same strategy and chose also to support many variations of Fibre Channel and SCSI, again resulting in proprietary technologies and weak standards. It should be noted though that the SAN development is still largely in the formative stage and final standards are unlikely to be established until well into the next decade. "The thing about storage standards is that they are so slow moving," said Collins. "You cannot watch them for six months and when you come back nothing has changed."

NAS on the other hand is based on more open networking standards, because it relies on the two true network standards for accessing remote data that have been broadly implemented across Unix and Windows systems. Developed by Sun Microsystems, Network File System (NFS) has become the *de facto* standard for Unix, while Common Internet File System (CIFS), developed by IBM and Microsoft is the standard for all flavours of Windows. Such broadly accepted standards for network data access and devices that serve data directly over a network, are far easier to connect and manage than DAS devices, allowing for universal data sharing between the two file systems.

Open standards means that companies can construct solutions that can theoretically incorporate legacy equipment as well as new and emerging technologies. While open storage networking means that in an age when

PROS AND CONS OF SAN

Pros

- Can be faster
- Easier Backups
- More flexible
- Access to raw devices

Cons

- Can be expensive
- Standards confusion
- More complex to set up

PROS AND CONS OF NAS

Pros

- Relatively low cost
- Easy to maintain
- Scalable
- Easily understood
- No additional software required

Cons

- Higher burden on IP network
- Not the best performance
- Complex backup strategies may become impractical

the opportunities and issues are extraordinarily diverse, complicated, and compelling, companies can tailor their operations and business development efforts to yield the best return.

However Collins of Quocirca warned that "implementation isn't really a standards issue. It will be interesting to see what eventually become the de facto standards, but then again the ones decided on are not always the best," he said. "However iSCSI is definitely one to watch."

Pros and cons of different models

Eddells of Xephon said that most companies have installed DAS as the default, simply because it's the oldest architecture, and when more storage was needed, another server can be added. But he said that more companies are now consciously looking to implement a SAN or NAS architecture if it's economically viable.

"You only benefit from the cost advantage at a certain level, so you've got to be a bigger company to make the investment," he said. "Then there's the question of which is better, SAN or NAS? Well, it's a bit like choosing a car – one architecture is better for one company, and a different one is better for another. It depends on what you want to do with it"

The problem with DAS is that every DAS pool would require its own tape backup system, tape library, and administrator. When multiple DAS pools are involved, the size of a tape library can become ridiculous. SANs and NAS beat this problem by centralising storage and eliminating multiple copies of data, thereby reducing the administrative task of backing up and restoring data. Tape devices can then be pooled into backup farms and mirrored copies of data can be restored in real time from one SAN or NAS to another.

But you also need to consider the types of application that will access the storage and by which mechanisms the data transfer will be initiated. SANs use SCSI over Fibre Channel to move large chunks of data quickly, while NAS over TCP/IP is better suited for moving smaller amounts of data in a less performance intensive environment. This is why SANs have become popular as a back-end storage system for NAS installations, focusing on backing up and restoring solely, without clogging up the enterprise network.

SANs and NAS can scale exponentially because of the multiple number of intelligent switches that can be layered into a redundant network with no single points of failure. While DAS data cannot be re-routed and downtime will reduce productivity. DAS is also attached to only one or two servers without any form of load balancing, so an overworked server cannot access the resources of an under-used server. SANs and NAS on the other hand spread user data over the available capacity and re-assign storage in times of high demand.

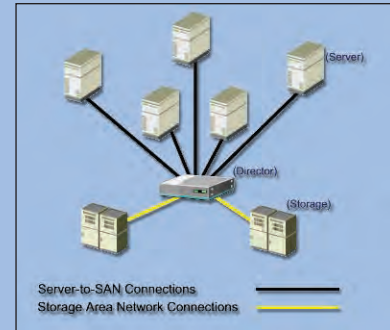
In a report from IDC, analyst Charlotte Rancourt said that SAN and NAS solutions can also be highly specialised and tailored to a company's particular requirements. Rather than buy an off-the-shelf solution, companies must carefully determine their needs. Questions that need answering include: how much and

THE SAN MODEL

Storage Area Network

DEVELOPED LARGELY BY NASA AS AN extension to server cluster technology, SANs use a centralised architecture to connect multiple servers to the same shared-storage device. These systems are usually based on high performance Fibre Channel media and the SCSI protocol, however iSCSI or SCSI over IP is still an emerging technology. High performance networks use switches to connect and route data to and from connected servers and are the primary choice for improving business continuity for organisations that need large amounts of storage.

Although, like DAS, the server performs all accessing and organising of files and data which leaves a high workload on the processor, the file system resides on the application server and a SAN benefits



Server to SAN connections at the heart of a Storage Area Network.

from pooled storage resources and LAN free backup.

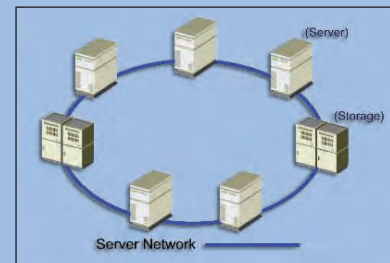
Essentially, a SAN puts a network in between the storage subsystems and the server, adding network latency to the DAS storage model.

THE NAS MODEL

Network Attached Storage

NAS WAS BORN OF THE NECESSITY FOR specialised file sharing, as these islands of DAS servers began to mushroom around the enterprise. NAS is essentially a server loaded with an OS fine-tuned for file sharing, which is easily added to an existing network. NAS benefits from a single user interface, which allows multiple TB (Terabytes) of storage to be managed easily, can be centralised or distributed with no limitations, and consolidates storage for improved disaster recovery.

NAS servers off-load all of the functions of organising and accessing directories



A NAS architecture can be easily created by adding NAS storage devices to a network.

and data, which frees up the server's CPU and reduces potential bottlenecks – a problem suffered by DAS systems.

what kind of data will be stored? Are the users highly skilled or less technically competent – or are they customers and partners? What types of network and devices will be connected to the storage? Are there any specific industry issues that must be addressed?

Collins, of Quocirca added: "When implementing storage, no one architecture is better than another. You can get benefits from one, such as a SAN, which runs on its own fabric, that you can't get on another, such as NAS, which is reliant on the network. The decision is often strategic versus tactical, rather than which architecture is better," he said.

Overall, the benefits networked storage provides are such that it's possible to create a storage system that is not only fast and powerful but also specifically designed to serve the business.



COVER FEATURE **STORAGE**

But DAS is also being displaced by the new model of networked storage because it offers consolidation of storage resources.

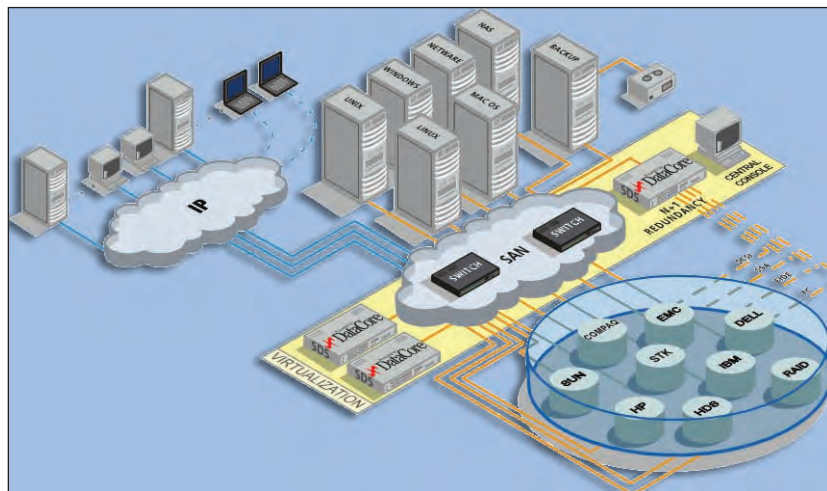
The consolidation question

Storage consolidation can provide substantial cost savings over the more traditional DAS because it uses storage resources more efficiently. When dealing with DAS, the primary cost is the people required to manage the systems, but with NAS and SANs, the actual storage is the primary cost.

The number of people required to administer a DAS architecture will depend on the number of servers installed. But with SANs and NAS, which does not add servers, just storage, fewer administrators are needed overall. According to analyst Merrill Lynch, a full time administrator can manage 1.5 to 5 Terabytes (TB) of data on DAS, but between 6 and 13.3TB on a NAS or SAN architecture. Additionally, the analyst reveals that DAS only uses 50 per cent of available disk space, which increases to up to 90 per cent for SAN and NAS, because more available capacity is used more efficiently, and it's easier to scale up. Improvements in disk utilisation coupled with higher density disk drives helps drive down TCO and improve ROI, two major reasons for implementing storage consolidation.

But the analysts all agree that the supposed cost advantages of storage consolidation have been vastly oversold, the hardware isn't cheaper and the main savings are in going from distributed storage on client workstations to a first level of centralisation at the server. The real value proposition is actually in the manageability of SAN and NAS, boiling down to two key features: the ability to take a snapshot of data that can then be backed up dynamically without downtime; or remote mirroring, which is the ability to mirror data into real time in a geographically diversified location for disaster recovery purposes. But few vendors offer both features, which is why those that do, such as EMC or Network Appliance, charge such high margins for these features.

The key here is management. As an IDC report sponsored by Veritas, a company that makes storage virtualisation software, reads. "The cost advantages of storage consolidation can erode quickly when systems are difficult to manage... storage virtualisation functionality continues to be successful in reducing complexity, customers should enjoy a sharp reduction in the



SAN just requires additional storage to be added, keeping administration costs down.

overall cost of deploying and managing consolidated storage systems. A part of this cost reduction will be due to lower staff training expenses and a reduction in storage management headcount."

DAS boot

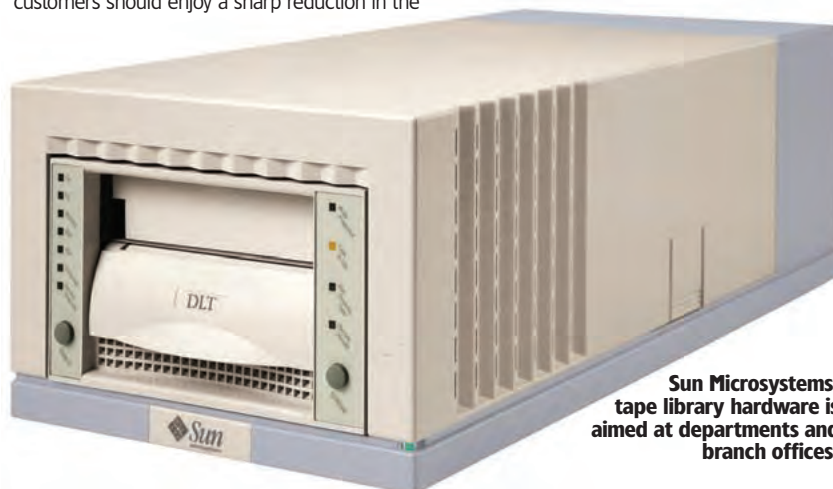
The breakdown of storage costs, according to Merrill Lynch places DAS as the most expensive costing around \$0.84 per Megabyte (MB), with \$0.35 of this spent on administration. A SAN architecture costs about \$0.38 per MB, with most only needing \$0.10 spent on media. NAS works out as the cheapest medium costing \$0.35 per MB, again with the highest expected expenditure of \$0.10 spent on media.

One of the primary reasons for consolidating storage is to improve the resiliency of data, which gives you high-availability for uninterrupted access; reliability for reduced downtime and better business continuity; and better performance for quick transaction turnaround improving overall productivity.

SAN and NAS architectures offer this from the outset. With no single point-of-failure, and multiple paths to data for fast access times, these technologies provide uninterrupted user access to data. But to obtain this resiliency with DAS, you need multiple identical systems with one always sitting idle waiting for another system to fail. And even then if a DAS server fails all of the storage physically connected to that server is unavailable until the server has been restored, reducing performance and availability while making the data more vulnerable to corruption.

So even if you're a small business, which hasn't thought about a networked storage strategy because of cost, it maybe wise to measure the potential benefits in terms of risk instead. Ask yourself these questions: if you loose all your data can your business survive? Even if you can restore, can you survive the downtime it takes to get back on your feet?

Figures from the Strategic Research Institute reveal that companies that aren't able to resume operations within ten days of an IT disaster are not likely to survive. According to Quocirca analyst Collins, businesses should have an availability guarantee, and that's what storage is all about, availability. And the advantages of storage are based on risk negation rather than cost. ■



Sun Microsystems' tape library hardware is aimed at departments and branch offices.

SPAM

Spam

Is there anything you can do?

IT barrister **DAVID HARRIS** addresses some of the legal implications of every email user's *bête noir*.



You'd be hard-pressed to find anyone that thinks well of the daily ingress of porn, Viagra offers and shift Nigerian filling our Inboxes, growing from a small trickle to the tsunami that threatens the very value of email as a communication medium. There are tough problems in even controlling it, let alone stopping it: the ability of spammers to cover their tracks; the technical challenge of distinguishing spam from legitimate email; the inadequacy of legal remedies and the ability of spammers and direct marketers to lobby lawmakers to impede the creation of new controls. This issue is not going to be remedied any time soon. As I outline here, existing legal remedies and prospective legislation fall some way short of full effectiveness.

Economics of spam

Spam is the disreputable cousin of the marketing industry; with its own economics technologies and methodologies. It is subject to the same economic rules as any other business sector – an observation that has ramifications for the strategies to combat it as we shall see later.

The main problem with spam is that it is not entirely unwanted. Indeed, direct marketers say, "Spam is in the eye of the beholder," which is true; it's not spam to the one person in 1000 who wants it, but the other 999 do regard it as unwanted and they did not get a chance to decline it. From the spammers' perspective, perhaps only one in a thousand respond, but it is estimated that each spam costs an average of 300 micro-cents (US¢), so even an infinitesimal response rate makes spamming entirely viable, given the relatively small associated capital costs; a bulk mail program, three or four PCs and a cable modem can be all that is needed. The eighty or so World-class spammers spend substantially more than this; some pay US\$100,000 a month for network pipes and rooms full of servers. However, their income rises proportionately; some clients pay as much as 30% of an invoice to the spammer, others pay \$20/\$30 for a genuine referral. Even if only one in a few thousand respond, a good profit is still achievable. No wonder spam king Ralsky could buy a \$740,000 home with his ill-gotten gains.

The economic downside for spammers themselves is rather limited, few are caught or successfully sued; a rarity was Earthlink, an ISP which was awarded \$25 million in damages in a suit against junk emailer, Kahn C. Smith, though whether they ever get to collect is a different matter.

If the cost to the spammer is small, the same cannot be said for the recipients or ISPs. The torrent of spam requires ISPs to have the bandwidth to receive both it and bounced spams. Once it is received, processing capability and software is needed to handle spam filtering. A Brightlight licence for an ISP with 5000 users can cost up to US\$15000. Using such filtering, AOL rejects approximately 80% of incoming emails. Staff costs are incurred in the form of extra personnel (or man hours) that are required either at ISPs or recipient companies. ISPs can also lose revenue from users who terminate accounts due to dissatisfaction with spam. These have led to estimates of the loss to US industry of around \$10billion. These figures have to be taken with a hefty pinch of salt, since many of them come from anti-spam companies anxious to sell their own products; we disbelieve the RIAA when they grotesquely inflate figures of losses due

to copyright infringement, and so we should be similarly suspicious about these claimed losses. Nonetheless, the analyses are broadly accepted by many independent groups and indeed user experience: we can all see our Inboxes fill with rubbish on an hourly basis.

It is little wonder then, that some predictions are that global spam will rise to 15 billion emails a day by 2006.

Legal attacks on spam

The difficulties

The underlying legal difficulties surrounding defences against spam relate to the nature of the relationship between someone sending email and someone receiving it. Email, like the rest of the Internet, was born in an era of trust. Designed and used by academics who either knew each other or occupied the same peer space, one would not send another an advert for pornography. However, as the circle of users widened, so this relationship between users lengthened and senders no longer cared about the opinions of the recipient – they were just punters. When this resulted in the first spam being sent is unclear (see box at end). Brad Templeton has asserted that the first documented one was sent in May 1978 from the old DEC minicomputer company to promote its new DEC-20 machine to academic users. DEC did at least anticipate and discuss the problems that they might suffer. The second high-profile spamming case departed from even this consideration for users, and set the pattern to date. Cantor and Siegel, which was a tiny (now defunct) law partnership in the US, cross-posted spam relating to US green card immigration services to 6000 Usenet groups, ignoring the etiquette of many of these groups. To add insult to injury, they wrote a book *How to Make a Fortune on the Information Superhighway*, encouraging others to do the same. That they did so was, unlike DEC, merely a consequence of their lack of interest or connection with the recipients. What is the relevance of this lack of connection? It hints at the difficulties that lie in finding an existing legal remedy: *prima facie* there is usually not a sufficient proximity in law between spammer and spamee to convince a court to prevent such abusive behaviour.

Contract and quasi-contract

One idea occasionally mooted, more-or-less seriously, is to reply to a spammer with the invitation to read their spam in exchange for some sum of money. Lawrence Lessig has indulged in such a scheme, though how seriously I'm not sure. Under it, the spammer is invited to accept the terms of this 'contract' by either replying and explicitly accepting the terms or accepting by conduct – that is, by *sending more spam*. This latter mode of acceptance of a contract is permissible in many situations. Rewards are a classic example of this; when a dog is lost, and the owner offers a reward for its return, it is not necessary for the finder to have previously contacted the owners and accepted the contract verbally or in writing. However, while the owner can accept an obligation, he cannot impose one on a stranger. He may not add to his reward notice a clause saying "...by reading this notice you accept the obligation to search for my dog for at least half a day..." You may have read the notice out of curiosity, greed or sympathy, but probably not because you intended to enter

DISCLAIMER

NOTHING IN THIS ARTICLE is intended as or should be construed as advice or acted on without seeking your own lawyers' advice. The author cannot give personal legal advice.

into an agreement. Likewise, even where a spammer reads your complaint email and sees purported 'contract', the mere act of reading it or resending spam will be very unlikely to be enough to form that contract; it is extremely unlikely he will want to pay you to read his spam, since that would undo the economies of spamming. Therefore a court would be likely to say there was never any intent by him to form a contract, and this intent to form a legally binding relationship is an essential condition to the existence of a contract.

Nuisance

There are remedies to restrain people from spoiling one's enjoyment of property. Usually this is used where a property owner is being subject to a nuisance: such as a neighbour who creates a smell, loud noise or pollution. Some lawyers view spam in a similar light; an external



JUNK MAIL

And other insidious forms of direct marketing

OF COURSE, OTHER FORMS OF DIRECT marketing pre-date spam and inspire it. Prior to the advent of spam, householders would moan about forests of junk mail, and in the US after a certain time of day there are the almost unavoidable telemarketers calling. Various legislators and industry bodies around the World have responded to complaints about these intrusions by creating laws and codes of practise that reflect their free market sentiments. In America and many European countries, enforceable preference lists exist which if breached lead to civil and/or criminal penalties.

In the USA, there are 'do not call' lists in many states. However, the list of exemptions for these is so long as to render them almost useless to householders: political parties, credit card companies, banks; oil, electricity and gas companies, insurance companies, financial services companies, charities, people doing surveys, companies who you contacted in the prior 6 months and so on – it would seem easier to specify those unable to make such calls. Violations can result in fines. A slightly stronger system of controls applies to unsolicited faxes and the reason for the stronger sanctions seems to be because the waste of resources in fax broadcasts is greater: time, toner and paper, and because they block legitimate business faxes. The criminal sanctions for unlawful faxes are higher than for telemarketing calls.

In the UK the regulation of direct marketing is self-regulatory and,

unusually, more business friendly than the US. Schemes such as the Mail Preference Service have no legislative bite or civil or criminal sanctions. The stiffest sanction a rogue UK marketer will suffer is a 'tut-tutting' from the Advertising Standards Authority (ASA). In practise this seems to work rather better than one would expect, in part due to the ability of the ASA to instruct its large number of members to reject advertising from those who are declared to have breached the rules. See the list at www.dma.org.uk/Shared/Consumer.asp The Direct Mailing Authority provides links to UK Mailing, telephone, fax and email Preference Services, allowing you to register your wish not to receive unsolicited marketing. There's also Baby MPS – a service which allows parents who have suffered a miscarriage or bereavement of a baby in the first weeks of life to register their wish not to receive baby-related mailings.

The equivalent of the Mail Preference Service in some European countries is known as the Robinson List (from *Robinson Crusoe*) since it often has the effect of creating a desert island isolated in a sea of surrounding junk mail. With the twisted values of marketing people, they regard this as a great disadvantage. That one should be deprived of the offer of yet another credit card or a voucher for ten pennies off a new toothpaste seems to be regarded by them as a screaming horror. Spammers frequently justify themselves on the same specious grounds.

SPAM

« nuisance that electronically enters one's property through a phone or broadband line and detracts from one's enjoyment of the use of a home computer. In the UK until recently, there was support for this view in case law. A daughter being harassed by an ex-boyfriend was able to restrain him from harassing her by phone at her parents' house. Since an email, like a phone call, is a variety of electronic communication, the view was that this may be a basis for acting against spam. This analysis suffered a knock-back when the House of Lords (in a separate case) ruled that that analysis was wrong and that nuisance actions were confined to the property owner alone in relation to actions which directly affected the enjoyment of the property. The nuisance calls were from a stalker and aimed at the girl, not the property; she could receive these nuisance calls at work and on her mobile as well as home. I suggest same is true for spam, which is targeted at an email address, not property, and has no direct effect on the property.

Outside the UK, the position may well be different. The UK decision was based on a Canadian court ruling in a similar case, which was much less narrow in its interpretation of the meaning of nuisance. The UK until recently has not had an anti-stalking law, and it was felt the initial UK decision was an attempt to fill in the gaps. With the passing of UK anti-stalking laws, it was felt that a return to a narrower definition of nuisance was desirable, but outside the UK this may well not be so, and spam may be a nuisance in those locations.

Spam haiku

Spam haiku has been used in an intriguing blend of Japanese poetry and intellectual property law to combat spam – the brainchild of the Habeas corporation. The technique is to add a haiku to headers in outbound email either by the user or perhaps an ISP, for example;

X-Habeas-SWE-1: winter into spring

X-Habeas-SWE-2: brightly anticipated

X-Habeas-SWE-3: like Habeas SWE (tm)

X-Habeas-SWE-4: Copyright 2002 Habeas (tm)

X-Habeas-SWE-5: Sender Warranted Email (SWE) (tm).

The sender of this

X-Habeas-SWE-6: email in exchange for a license for this Habeas

X-Habeas-SWE-7: warrant mark warrants that this is a Habeas Compliant

X-Habeas-SWE-8: Message (HCM) and not spam. Please report use of this

X-Habeas-SWE-9: mark in spam to

<http://www.habeas.com/report/>.

Compliant software such as *SpamAssassin* recognises these headers as certifying that the mail is not spam, and also compliant with Habeas' terms and conditions that require double opt-ins for commercial email. As such, it can pass straight through any filters with confidence. The haiku is both copyrighted and trademarked, and the legal theory is that it acts as a certification mark warranting the authenticity and *bona-fides* of the message. The legal theory is broadly sound; copyright law will protect the haiku since it is an original work of literature that could not usefully be reproduced by a spammer in a lawful way. The header line **X-Habeas-SWE-3: like Habeas SWE (tm)** will likewise function effectively a

trademark. There are questions relating to whether people downstream of a lawful user of the scheme, *eg* an ISP using Habeas, can be compelled as a sub-licensee of the system. Inevitably, spammers will want to use this haiku to circumvent filters, but Habeas, founded by a lawyer, has said it will litigate against transgressors aggressively, a promise they seem to be keeping. Like most of the legally-based schemes, this may share the weakness of getting hold of the spammer. Habeas however, says that, uniquely, this scheme allows them to track down the hiding spammer by targeting the spammer's client with a vicarious or contributory infringement claim. While I do not consider this to be an unarguable allegation, it is very good practical law: 99 per cent of spammers' clients, when told "we will either sue the spammer or you," will quickly hand over the spammer's name, address and trouser size. Habeas says that in these circumstances that its record is two hours in obtaining details that would be otherwise unobtainable!

Data Protection Act

Within the European Union, the *Privacy Directive* will have an effect. This directive provides that so-called data controllers process data on data subjects in accordance within certain principles. These principles are that the data should be:

- 1** fairly and lawfully processed
- 2** processed for limited purposes
- 3** adequate, relevant and not excessive
- 4** accurate
- 5** not kept longer than necessary
- 6** processed in accordance with the data subject's rights
- 7** secure
- 8** not transferred to countries without adequate protection.

Personal data comes in two broad flavours: that which is identifying and that which is called 'sensitive personal information'. The latter would include things such as trade union membership, voting behaviour, racial origins, health, sexual orientation and any criminal record. This category is tightly controlled and gives data subjects powerful rights.

The average spammer will stomp over most of these principles. Email addresses will often be personal data of the former type since it may identify an individual:

david@ukitlaw.com. This is not necessarily so:

webmaster@geeklawyer.org does not necessarily identify any individual and may not fall within the Act.

In most cases, spammers grab email addresses from webpages or Usenet postings, a method which would not be regarded as a fair method of collection under the Act: the email address would have been posted to permit communication with people for the purposes of that webpage or discussion group and not receiving spam. While people might be aware of spambots, it would be unreasonable to say they must have therefore consented to the risk. Not posting an email address – and so not being able to actively participate in many forums – would be a result that would damage societies' public interest in freedom of speech and freedom of association. Likewise, the address is likely to be used without any relevance to the recipients interests and in an excessive manner; one week Viagra would be promoted and in another week marketing tips, cheap home loans or whatever else the spammer is contracted for.

RESOURCES

US Do Not Call lists:

www.privacycorps.com/pages/state_donotcall.htm

US Telephone Consumer Protection Act:

www.privacycorps.com/pages/law_1991.htm

Spam haiku:

www.habeas.com/

Anti-spam organisations:

www.spamhaus.org/legal/answer-03-80295.html

www.cauce.org/

UK Government spam consultation:

www.dti.gov.uk/industries/content/dti_consultation_on_implementation_of_directive_on_electronic_commerce_200031ec_summary_of_responses.html#unsolicited

Another principle of data protection is that it should be possible to stop the processing of data on good grounds. Most spammers will be unlikely to give any real option for this, indeed the 'Unsubscribe' option is an invitation to confirm that the email address is valid and so receive even more spam. Email address lists are also unlikely to be kept securely; spammer gangs will often swap or sell among themselves CDROMs with millions of addresses. This would offend against the requirement to keep data secure and to process it fairly.

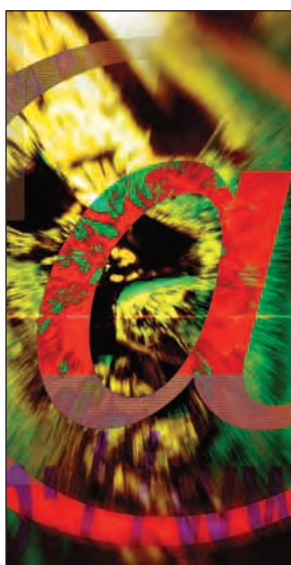
In reality, data protection legislation is aimed at responsible companies rather than spammers, and hence limited in its protective scope. The EU has recognised the weaknesses in the fit between email and data protection by introducing both the *E-commerce Directive* and the *Directive on Privacy and Electronic Communications* that will specifically address this issue. These should come into force in the UK sometime towards the end of 2003. Typically, an opt-in to unsolicited commercial emails is required except where there is a prior commercial relationship in relation to similar products, even then however it must be possible to prevent the receipt of spam. The EU has recognised the weakness of the legislative approach but with increasing political and legislative hostility to spam, it may act as a template to US and Asian legislatures.

Criminal law

The observation that spammers frequently offer goods or services that contravene the law is almost trite. Inboxes routinely bulge with offers sex with under-age girls, unlawful prescriptions such as Viagra, and stock market information that if acted on would amount to insider trading. Others are criminal deceptions, such as the comical Nigerian 419 scams that promise millions of dollars to the gullible in exchange for the use of the victims' bank account. Such emails are clearly deceptive at best and as such, subject to more civil and criminal actions than one could shake a stick at.

Additionally, many spammers increase their chances of getting past spam filters by forging headers. This would be an offence under the *Forgery and Counterfeiting Act 1981*, as the utterance of a false instrument intending that the recipient should take it as real. Normally, forgeries are on paper and directed at humans. However neither the Act nor case law require this; an email and its headers are an instrument for this purpose, and as many people convicted of using forged ATM cards have found, one need not attempt to deceive a human, a machine suffices. This tack will work in most countries with counterfeiting laws. As we mentioned earlier in the case of Habeas spam haiku, it would also be an intellectual property law violation. When the EU directive on privacy comes into effect, it will also be a violation of that.

The latest trick spammers have latched onto is to use trojans such 'jeem' or 'proxy-guzu' on infected clients to send email. Paradoxically, this may be encouraging news. The effectiveness of spam filtering and legislation is forcing spammers to indulge in outright criminality. The act of infecting someone with a virus is a criminal offence in most Western countries; in the UK it is an offence under the *Misuse of Computers Act* as an act of unlawful access to, and modification of, the victims' system. This Act initially targeted crackers, but the acts of dissemination and operation of a spam trojan fall squarely within the Act's prohibited behaviour.



Another line of attack that exists against spammers is to allege criminal damage but there are problems with framing an indictment; it may work in some jurisdictions but not others. Where a spam flood is initiated to such an extent that it takes a recipient server off-line the utility and function of the server is damaged albeit temporarily. Criminal damage normally arises where permanent damage is done to property, such as putting a brick through a window, and it can lead to a prison term of up to 10 years in the UK. It is however not necessary for the property to be broken or deformed for it to be damaged within the terms of the *Criminal Damage Act*. An often-used example is the disabling of a car brake hose by uncoupling it. It is said to be sufficient damage even though it could be easily fixed by a spanner and 5 minutes under the bonnet. On this argument, denying legitimate users access to a mail server by flooding it with bogus connections and unlawful emails is a mere variant of a Denial of Service attack that amounts to tampering to

“99% of spammers’ clients, when told ‘we will either sue the spammer or you’ will hand over the spammer’s name.”

prevent proper server functionality. However attractive as it might seem to put a spammer in jail for a decade, this is an untested proposition. English criminal damage usually requires a degree of permanence in the damage inflicted on something that is physical in character. Merely denying others the use of a server for a week will probably not have that degree of permanence. It is the requirement for physical property damage that also prevents criminal damage being used in patent or copyright cases (so, it has a plus side!) This is, however, the position in English law only, and other jurisdictions have a different requirements that render this approach workable. In a similar vein, it should often be possible to sue a spammer for DoS attacks as being a trespass to property where the use of the server is unlawful, e.g. because it is criminal or in breach of contract.

The obvious difficulty with all of this is that getting hold of the spammer is typically a very hard job indeed. Spammers have resorted to hijacking entire dormant /16 subnets, with the suggestion they may try to hijack live ones. Many other spammers route their traffic through series of open relays from hacked accounts or disposable user accounts. Most of it appears to come from various Asian sub-nets, to the point where some sysadmins block these entirely. This is becoming something of a political issue, with the Chinese government complaining. Whilst sysadmins in this region often bring the problem on themselves by setting up mail-servers incorrectly that is often a consequence of the poor availability of Chinese documentation. One could sue the owner of such an open relay, either in their own country or one's own. When one considers the devastating consequence of a spam attack, this is not at all a fanciful option. Telewest recently had mail outages lasting several weeks. This would have been expensive and it originated, I understand, from a small block of addresses owned by one operator. Were I Telewest's lawyer, I would be advising them to seek



HISTORY OF SPAM

ORIGINATING IN THE DEPRESSION AS food for the unemployed, SPAM™ was made famous in the Second World War, when its long shelf life made it a staple of army food. Its full name is Hormel Spiced Ham, a name that was chosen by visiting New York radio actor Kenneth Daigneau in a drunken competition where the winner the winner got \$100. Marketed as ideal for making 'spamwiches' it went from strength to strength.

The email usage comes from 'spam chorus' in the Monty Python sketch, where the chorus drowns out a conversation where customers try to avoid getting spam, so much like real life!



Hormel have said that it won't object to the use of the word spam to designate unsolicited commercial email, since they say it doesn't tarnish their mark; so long as use is all lower-case and the spam graphic isn't used (Slashdot take note!).

to recover those losses from that ISP. Whether that ISP permitted the client to spam, or acted negligently in administering their mail server to be an open relay, I would allege that they were liable for the loss. I'm not aware of such a case, but the risk for sysadmins is clear. This may seem harsh since the relay operator is not the guilty party, but in reality in commerce, much loss is discussed in terms of risk and loss allocation: if there must be a loss, it is best to try to pass it on, even to another innocent person.

Spammers striking back

It isn't just spam trojans that signal that spammers are beginning to feel pain. Some are now resorting to lawsuits against anti-spam activists. In Florida, it appears that marketersamerica.org (memberships \$3000/year!) is representing various purveyors of Viagra and teen-sex in complaining that these are "legitimate businesses wrongly pushed to the precipice of extinction" by activists acting in ways that are "...Ruthless. Stopping at nothing – even stomping on good taste...". One has to admire theirchutzpah if not their optimism. But are the anti-spam activists in real danger? After all, many are small groups of volunteers, operating without funding and certainly not having the very deep pockets of spammers with which to defend litigation. marketersamerica.org has initiated what can only be described as a very speculative action against activists – including the well-regarded UK-based Spamhaus project (www.spamhaus.org). Their accusation is that Spamhaus is blocking legitimate US businesses from sending email to Internet users by posting information about blocks of IP addresses, interfering with the spammers business and third-party contracts, defamation and (get this!) invasion of privacy. The 'Realtime Blackhole List' operator MAPS has been sued on similar grounds on several occasions, but none have reached a final hearing, most settling out of court. Media3, who hosted 17 spammers as well as 1500 non-spammers, was put on the RBL but it was denied an injunction against MAPS. Given the scattergun approach of MAPS that involves adding non-spamming as well as spammers IPs in order to coerce ISPs into compliance, this



was a welcome, if rather lucky decision that could easily have gone the other way. That said, the creation and reliance on RBL seems to be desirable on public policy grounds if it is properly structured. Whether the MAPS approach goes too far in shooting innocent bystanders is a matter of diverse opinion. MAPS would no doubt say (correctly) that to do otherwise would not provide a sufficient incentive for ISPs to aggressively police their networks and AUP. My own ISP, Telewest, regularly probes my mail server for open relay behaviour in order to protect its IP addresses. Nildram does the same and I suspect so do others. Would they do so without the threat of an RBL? In many cases, one suspects not. On the other hand, if you are the operator of a small web-based business, you could be bankrupted by blacklisted email. Factors a court might weigh would be the speed with which the RBL operator rectifies the list on conformance by an offender, the ease with which the innocent can get an alternative provider, the difficulty of an ISP rectifying the abuse and the extent of the damage done to society by spam and to innocent parties or the ISP being blocked.

Politics & law

Telemarketing does give rise to a partial analogy with spam, in that there is a diversion of resources and expenditure. Some of the proposed legislative changes utilise the analogies with telemarketing in their proposed solutions, though commendably many legislatures are beginning to see that spam is a problem of a much larger magnitude. Even in the US, champion of the free market, various committees are showing signs of losing patience. While previous efforts at legislation have foundered on the lobbying of spammers and their apologists, there appears to be a recognition that the status quo is unacceptable in the face of increasing UCE/UBE.

That notwithstanding, they are seriously considering proposed changes that would make the problem much worse rather than better; eg it would prevent private actions against spammers and only state officials would be able to act. These officials state their unwillingness to undertake policing, because of funding and resource problems. The opt-out proposals would do particular damage, as it would undermine much of the current opt-in requirements of some states and voluntary schemes, as well as putting the burden on the recipient. Each and every spammer would only need an opt-out, and even that would lapse after three years. This seems to be an area where pro-spam/direct marketing lobbying has been effective. On the plus side, many proposals will make it unlawful to forge headers or provide an invalid return address. Nonetheless, according to Brightmail, 90% of spam originates from just 180 or so American spammers routing their junk through other countries. It is clear that only an effectively policed American law will resolve spam.

The European Union is being altogether more switched on. The *E-commerce Directive* has mandated that UCE is "clearly and unambiguously" identifiable as such on receipt. For ordinary commercial mail, it is necessary to identify it as commercial and on whose behalf it is being sent, as well as all the term of any offer. Also Article 13 of the *Directive on Privacy and Electronic Communications 2002/58/EC* requires opt-in for UCE (but not for unsolicited non-direct marketing emails; eg religious political etc) and also prevents forged

sender address and requires an effective unsubscribe address. However, the UK government has disabled the honest labelling requirement that requires spam to be "...clearly and unambiguously identifiable..." This was intended to require spam to have an subject line indicating it as spam: eg "[ADV]..." The UK Government is trying to render this into a form that would make automatic filtering impossible. It has said it would "limit entrepreneurial and commercial creativity in marketing and scope". They appear to be of the view that making it clear it is spam in the text of the body is enough. By this time of course, you have already downloaded it and wasted time reading it. And by not requiring standard form labelling, it makes it hard for filters to catch it with absolute certainty; though honest marking should make Bayesian filtering easier. This damaging hack will probably be implemented in the UK by then end of October 2003.

My view is that an additional effective solution would be to drive a wedge of mistrust between spammer and client. Not all spammers operate on their own account. Many act under contract to businesses such as insurance companies, anti-virus vendors and the like. A useful legislative change would be to render such contracts unenforceable; this would increase the uncertainty of spammers that they were going to get paid. For example, some spammers take payment by credit card, but on my proposal a client could pay to get the services, get the spamming done, and then get the payment clawed back by the card company; ie legally renege on the deal. Of course, this client should not get away scot-free either so I would also go further and make the beneficiaries of spam advertising liable to those affected by spam where it could be shown, on a balance of probabilities, that they had employed the spammer. This would undermine the incentive of people to use spammers while preventing miscarriages of justice to innocent websites who are, for whatever motive, maliciously and falsely accused of employing spammers.

While it may seem futile to debate criminal and civil sanctions over spam when the spammers are apparently untouchable, I would take a less pessimistic view. Spam has moved in time from a system where the breach of standards was a social one between scientists, to one where it is a civil liability, to one where the damage to society is such that it is becoming criminal. So far this escalation of penalties has not led to a reduction in spam, indeed it has risen – though I suggest not for connected reasons. However, as this

"A solution could be to drive a wedge of mistrust between spammer and client... and undermining the incentive for both."

behaviour becomes increasingly legally (as well as socially) unacceptable, the expense and risk rises; its one thing to have a user account terminated for spamming, but it is quite another to risk jail-time. To evade such consequences, it then becomes necessary to move operations to a spam-friendly jurisdiction outside of Europe, America or other Western states, which is inconvenient and expensive. Ideally, the remaining countries can be persuaded by diplomatic pressure and international treaty to crack down too. This is a technique deployed effectively by lobbyists to nefarious ends; it is the manner by which American media companies have foisted the DMCA and clone legislation on the rest of the World. None of this would be quick or absolutely effective, but at worst it would have the effect of localising the source of spam that might make it easier to police. Those countries famed for originating spam risk being blocked, as happens with China at some ISPs at the moment, and to be disconnected from the Internet would be very damaging indeed for any emerging economies. ■

ABOUT THE AUTHOR

DAVID HARRIS IS AN IP barrister who wiles away his days as www.ukitlaw.com waiting for an interested client to ask him about Linux rather than software patents. When not suing people, he writes 'depressingly poor code'. www.ukitlaw.com www.geeklawyer.org

THE EARLIEST SPAM

A piece of email history

BRAD TEMPLETON HAS UNCOVERED THE first documented piece of spam sent on 1st May 1978 (although it seems from the complaints by other users adverts had been seen before but not whether they were spam *per se*). The spammer was the old Dec minicomputer company of Vax fame. It looks rather tame but the furore it generated then was just as fierce as that now. Her it is in all its commercially innovative marketing glory:
"DIGITAL WILL BE GIVING A PRODUCT PRESENTATION OF THE NEWEST MEMBERS OF THE DECSYSTEM-20 FAMILY; THE DECSYSTEM-2020, 2020T, 2060, AND 2060T. THE DECSYSTEM-20 FAMILY OF COMPUTERS HAS EVOLVED FROM THE TENEX OPERATING SYSTEM AND THE DECSYSTEM-10 <PDP-10> COMPUTER ARCHITECTURE. BOTH THE DECSYSTEM-2060T AND 2020T OFFER FULL

ARPANET SUPPORT UNDER THE TOPS-20 OPERATING SYSTEM. THE DECSYSTEM-2060 IS AN UPWARD EXTENSION OF THE CURRENT DECSYSTEM 2040 AND 2050 FAMILY. THE DECSYSTEM-2020 IS A NEW LOW END MEMBER OF THE DECSYSTEM-20 FAMILY AND FULLY SOFTWARE COMPATIBLE WITH ALL OF THE OTHER DECSYSTEM-20 MODELS. WE INVITE YOU TO COME SEE THE 2020 AND HEAR ABOUT THE DECSYSTEM-20 FAMILY AT THE TWO PRODUCT PRESENTATIONS WE WILL BE GIVING IN CALIFORNIA THIS MONTH. THE LOCATIONS WILL BE:

**TUESDAY, MAY 9, 1978 - 2 PM
 HYATT HOUSE (NEAR THE L.A. AIRPORT)
 LOS ANGELES, CA**

THURSDAY, MAY 11, 1978 - 2 PM

**DUNFEY'S ROYAL COACH
 SAN MATEO, CA
 (4 MILES SOUTH OF S.F. AIRPORT AT BAYSHORE, RT 101 AND RT 92)**

A 2020 WILL BE THERE FOR YOU TO VIEW. ALSO TERMINALS ON-LINE TO OTHER DECSYSTEM-20 SYSTEMS THROUGH THE ARPANET. IF YOU ARE UNABLE TO ATTEND, PLEASE FEEL FREE TO CONTACT THE NEAREST DEC OFFICE FOR MORE INFORMATION ABOUT THE EXCITING DECSYSTEM-20 FAMILY"

Many consider the infamous 1994 Cantor and Siegel Usenet Green Card Lottery spam to be the genesis of mass spamming.

Why didn't someone patent unsolicited commercial email? Sometimes the strangling of innovation by patents is a good idea!

HOSTING

WEB HOST FOCUS

1&1 Internet

With over 2,500,000 customers and 3,500,000 domains registered, 1&1 is among the biggest web hosts in the business, and it's not surprising – the company has over 10,500 Mbit of connectivity spread over nine different carriers, as well as owning their own European backbone. To make extra sure that there's absolute reliability, 1&1 has spread its servers across four data centres around Europe, costing over £15 million.

Targeted mainly at home users and small- to medium-sized companies, 1&1's hosting options start at incredibly low prices for shared servers, at just £1.99 per month, but

In the first of our series taking a detailed look at individual web hosts, PAUL HUDSON investigates what 1&1 has to offer...

1&1 has a sliding scale of services. For example, its starting dedicated server price is £29 a month, which gives you much more flexibility over your system as you have almost absolute control over it.

With just over 15,000 servers deployed around their networks at the time of writing, 1&1 has come a long way in the eight years since it was established back in 1995. With 99.9% uptime, and amongst the best bandwidth we've ever seen (see the issue of *Linux Pro* that came with LXF44 for the full test and comparison with other UK providers), 1&1 has an above-average offer at incredible prices. Furthermore, 1&1 has a money-back guarantee in

MARK BEYER SPEAKS

We spoke to Mark Beyer, Director of 1&1 Internet, about the company, how 1&1 use Linux, and their plans for the future

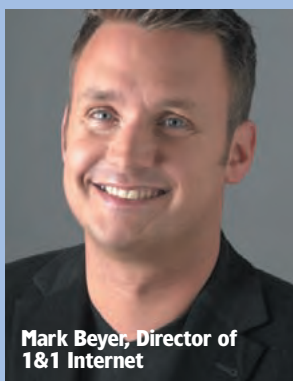
LINUX FORMAT: What are 1&1's unique advantages over its competitors?

MARK BEYER: We have quite a few key advantages over our competitors. Firstly, price – we've tried hard to provide the highest specification at the lowest prices, and, thanks to our investment in R&D, we can offer full-featured packages at these excellent prices.

Secondly, we offer a one-stop shop – rather than have products and services from different providers, we offer customers all they need for the development and growth of their web needs, from start to finished. We're not a marketing company with lots of outsourcing; we do everything in-house, even the development of our new features. We believe this goes hand-in-hand with providing cutting-edge technology.

Finally, there's the fact that we offer the largest connectivity in Europe to ensure the best possible performance for our customers.

LXF: What are 1&1's key target markets?



Mark Beyer, Director of 1&1 Internet

“We can formally confirm that 1&1 Internet is the fastest-growing web hosting provider in the world”

NETCRAFT, APRIL 2003

MB: At the smallest level, we cater extensively for consumers, and SoHo businesses (1-10 employees). Of course, our services also scale up to SME level – we've got a lot to offer companies with up to about 500 employees – and beyond.

LXF: What kind of hardware does 1&1 deploy, and what are your hardware plans for the future?

MB: Rather than relying on a few very large mainframe computers to host our customers' sites and emails, we host on a number of smaller machines. We've found that this gives our customers a much higher reliability rate because the chances of their machines going down are much slimmer. Also, if one machine goes down it doesn't take the whole company with it, which counts for a lot! This is a proven concept which has

place, that means you can test drive 1&1 for 60 days, and if you're not entirely satisfied with all aspects of the service that you experience, you can cancel the account and 1&1 will refund your money in its entirety with no questions asked.

Perhaps 1&1's biggest claim to fame is the fact that when readers of *PC Format*, *.NET*, *Internet Works*, and other digital magazines voted for the *Future UK Internet Awards*, 1&1 came away as the Best Business Web Host and the Best Consumer Web Host awards for both 2001 and 2002.

1&1 in action

What do their customers think of them?

Simon Ruckledge, a partner at Studio 49 Website Design (www.studiofortynine.com), a company which relies on 1&1 for web hosting, had this to say...

"Our search for a web host first brought us into contact with 1&1 over five years ago. New to the web design business, we were looking for a reputable company that would offer easy administration and set-up linked with value for money. 1&1 offered all of these, and also the added knowledge and confidence that we would be using a Linux operating system.

"We have always chosen Linux servers over Microsoft. Described in a few words, the 1&1 Linux packages offered us reliability, security, performance, stability, flexibility and compatibility with all our clients' demands and our existing systems and software. Our projects often required integration with PHP and MySQL, and this became a breeze with Linux described as 'the operating system for businesses and web applications'. With the addition of SSL/SSH, we were able to provide the ultimate

security we promised our customers and future users. The 1&1 Linux packages also provided us with one major advantage over other operating systems, the Linux community – millions around the world using Linux and conversing to provide free, friendly advice and solutions.

"The interface within 1&1 is truly intuitive, to set up and integration is achievable by even the most modest of operators. We were reliant on both 1&1 and Linux to deliver all this to our expanding and demanding customer base with high traffic sites and were not disappointed with the continuing results.

The three-way relationship between 1&1, Linux and ourselves has provided us with a successful business with 90% of our clients originating from referral business, 'happy customers with happy customers.' " ■



Simon Ruckledge, Studio 49 Website Design.

been a huge part of Google's success, and we have worked this way from the very beginning.

LXF: How much does Linux figure into your hosting plans?

MB: Although we do offer Microsoft hosting, we are primarily and foremost a Linux company. All of our internal systems, email, databases etc are all running on Linux systems. The fact that Linux is Open Source has enabled us grow the company to the size that we are now, which would have initially not been possible with a product that requires massive licensing fees and limited options on changing configurations.

1&1 hosting started in 1995 based on a very basic Linux setup, which has been refined and tuned over the years to an extremely powerful setup with 15,000 servers up and running and the capacity for 25,000 in our data centres. In this process, we have contributed numerous patches and bug fixes to the Linux community and have also our own GPL projects running.

Our Data Centre capacity is due to be at a

maximum by end 2004, so our engineers are already planning investment into an *even bigger* data centre project.

LXF: Where do you see the hosting market being in three years time?

MB: There will be a considerable shift towards dedicated hosting because of the larger bandwidth being available to home and smaller users. This means that customers can use the Internet for not just websites, but file sharing and file hosting.

Private users would potentially use the Internet more for sending more pictures and video clips to share; for example holiday pictures, family movies. A customer will soon be able to take pictures via their camera phone or video phone and upload this to their website, then email their friends inviting them to view their pics, even while they are still away on holiday.

For the business users, it is becoming increasingly important to have a website, and we envisage that more of the smaller business users will join the web community, even if it's just to have an online

business card. The higher-end websites for the larger companies will move more and more towards dedicated hosting, simply for the reason that for an affordable price they can utilise powerful servers for a lot more than dedicated hosting. Well, you can with 1&1!

Companies can move even more towards 'mobile offices' by making all of their information electronic, making them accessible anywhere in the world – we're talking more than just emails; it can be your entire database information system. It would be like accessing your hard drive from wherever you are in the world.

With search engines advancing in their accuracy, Google is becoming more and more important – more so even than the Yellow Pages. Nowadays I think that more people are turning to Google and similar search engines than the Yellow Pages – if Google can't find you, then you're nobody! So it's even more important to get a website.

Of course 1&1 will monitor the market very closely and will continue to develop innovative products for the online community.