

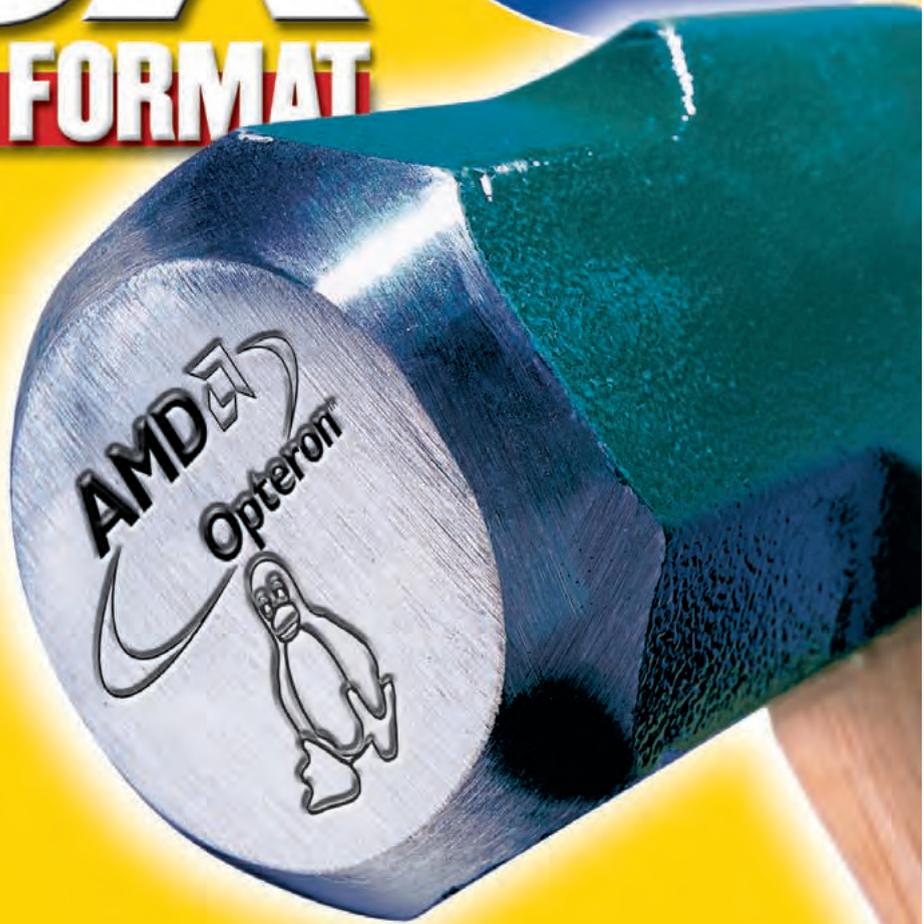
LINUX FORMAT

www.linuxformat.co.uk

AMD'S LATEST,
MOST AMBITIOUS
PROCESSORS EVER –
RUNNING LINUX!

HAMMER TIME!

The inside track on AMD's revolutionary Opteron CPUs – will 64-bit x86 computing live up to the hype? **PAGE 50**



LAPTOP MEGATEST!



Seven of the latest laptops tested with Linux **p43**

JCODEBOX

J2EE RAD development tool reviewed in depth **p28**

HOWTO:

- Build your own RPMs **p62**
- Use cable/DSL modems **p64**
- Blend images with GIMP **p70**
- Make PHP run faster **p74**
- Control flow in Python **p78**
- Create objects in Blender **p82**

DISTRO WARS PART II

Mandrake 9.1, Red Hat 9 and SuSE 8.2 all on test inside **p17**

WHAT ON EARTH IS THE HURD?

p56

CD issue also available Printed in the UK
LXF41 JUNE 2003 £6.49
 the future network 06
 Media with passion 9 771473 968012
 (£6.49 outside UK and Republic of Ireland)

THE UK'S BEST-SELLING LINUX MAGAZINE

Taking the lead

Politics may or may not be on your mind following the recent local elections in many parts of the UK. But politics isn't limited to Whitehall – it permeates most areas of life, including the world of computing.

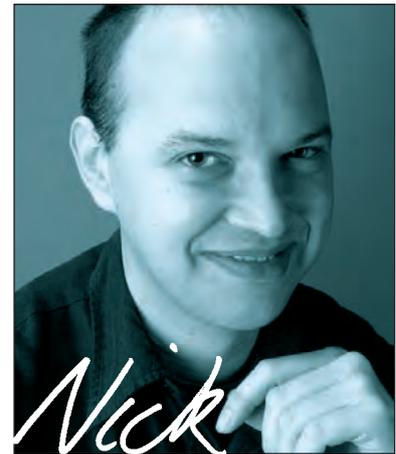
In this issue we take a look at AMD's Opteron processor, launched at much expense in New York at the end of April. AMD partners, server builders and the like who are bringing out products based on the new 64-bit processor, were there in force, but many suggested that politics had stopped some people turning up.

Inevitably, much of the hype surrounding x86-64 may come down to marketing differentiators. Though the follow up Athlon-64 line will no doubt create processors with greater performance, there is no evidence that suggests average desktop users need that much more power – which is why Intel went down the Itanium route. But it indicates that AMD is capable of taking a technological lead and following it through. The real win for them will be if Opterons sells in big numbers.

Of course, as Linux is one of the few operating systems to support the Opteron at launch, it is already a win for Free Software. While others were no doubt debating the politics, Free Software developers just got on with the job. AMD's Opteron is the focus of our lead feature this month, where politics aside, we'll be taking a look at how it performs.

Thankfully, there's no complicated politics concerning laptops, the focus of our roundup this month. We had thought this might be a bit tricky, with endless fiddling to get them to work properly with Linux, but they all pretty much worked out of the box. Take a look at the results (and see our new benchmarking system in action) on page 43.

There's plenty more in this issue too – from reviews of all the latest distro releases, a new server from NEC and some updates to old favourites. With a bumper selection of tutorials too, plus Mandrake 9.1 on the CDs and double-sided DVD, it'll make you wonder how on earth we manage to cram so much in!

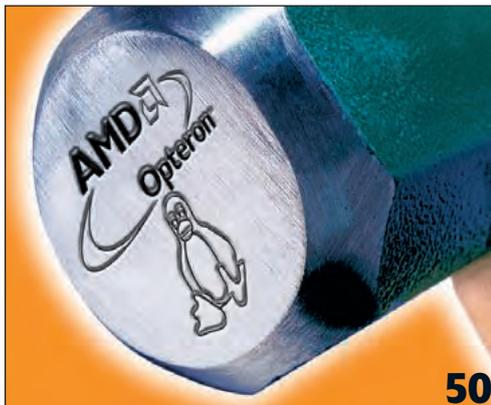


Nick Veitch EDITOR

AMD's Opteron technology explored and explained **p50**

Making a bid to be best of the current new distros – SuSE 8.2 **p22**

LAPTOPTASTIC! – seven big-name laptops benchmarked within an inch of their Linux lives **p43**



LINUX
FORMAT

AIMS OF THE MAGAZINE

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

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

MEET LINUX FORMAT'S TEAM OF WRITERS...



Andrew Channelle
A great nose for Linux news, and a distro expert as well, Andy's reviewed Red Hat and SuSE this month.

Michael J Hammel
Professional GIMP artist Michael is penning (or pencilling) our current Open Source graphics tour-de-force.



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

Hoyt Duff
Fishing pier proprietor Hoyt spends his spare time installing Linux on anything that stays still long enough.



Rich Drummond
Only in USA for a few months, Richard is already spelling 'proper' English words the American way!

Maurice Kelly
When not coding, Maurice spends his time downloading more code. Broadband anyone?



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

Biagio Lucini
This Italian is fanatical about efficiency, which is why we gave him the new iteration of Mandrake to pore over.



Paul Hudson
Formerly our PHP guru and a web designer, our Reviews Editor is 'Mr Workrate' – making the rest of us look bad...

Patrick O'Brien
Published by O'Reilly and IBM, the author of *PyCrust* has a very deep affection for Python scripting.

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

LXF41 June 2003

Contents

LINUX
 FORMAT

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

COVER HIGHLIGHTS

HAMMER TIME!

Extra capability and backwards compatibility: AMD's new 64-bit CPU is out in the wild!

50

17 Distro Wars Part II

Mandrake 9.1, Red Hat 9 and SuSE 8.2 all on test

28 JCodeBox review

Frustrated with J2EE? This'll cure your Java blues

43 Laptop MEGATEST

Seven portable PCs benchmarked with Linux

56 What On Earth is The Hurd?

A history lesson or OS of the future? You decide...



'Not exactly poetic, is it?
Let's see if we can't make our love counter more refined to keep Miss Browning from groaning'

78 Patrick O'Brien's still rather infatuated...

'These cretins attempt to make you and I feel guilty, as they claim we are stifling development of "free" Linux...'

09 Nothing riles Hoyt Duff like Linux snobs!

LINUXPRO

E-voting Democracy and Linux
.NET Linux Use Mono and DotGNU
Thin Client computing get rid of your desktop box
Writing compilers flex and bison
Storage NAS and SAN
Qtopia Embedded applications





- 06 News**
Latest Linux stories and comment
- 12 Mailserver**
Got an opinion? Have your say!
- 18 Mandrake 9.1 review**
Read our review, use the coverdisc
- 20 Red Hat 9 review**
We rate the "Microsoft of Linux"
- 22 SuSE 8.2 Professional review**
Just don't mention the distro-war...
- 26 AC3D review**
3D modelling in many file formats
- 27 Arstaro Security Linux v4 review**
Full-featured secure firewall distro
- 28 JCodeBox review**
Generate Java code in seconds!
- 29 OpenAuth review**
Security enhanced Debian PC
- 30 NEC Express 5800 server review**
1U iron with off-beat configuration
- 33 Book reviews**
PostgreSQL, LDAP, UNIX, MySQL
- 37 HotPicks**
Thunderbird gets the **Hottest Pick**
- 43 MEGATEST: Linux laptops**
Seven top portables benchmarked

- 50 Opteron**
AMD's new Itanium-beater
- 56 What On Earth is The Hurd?**
The great OS that Linux usurped
- 62 RPMs tutorial**
Creating your own packages
- 66 Cable modems tutorial**
Speed up your Internet connection
- 70 The GIMP tutorial**
PLUS nearly 400 plugins on disc
- 74 PHP tutorial**
Run your scripts 326 times faster!
- 78 Python tutorial**
Containers, namespace and wx
- 82 Blender tutorial**
Getting started with primitives
- 86 Runtime Revolution tutorial**
PART 2: Menus and images
- 88 Answers**
Got a problem? Let us sort it out
- 98 Coverdisc pages**
Essential reading for installation
- 110 User Groups**
Local support, *The Lord's Abode*
- 114 Next month**
A glimpse of tomorrow, today...

22



56



COVERDISCS

A double-sided DVD or 3 CDs bursting with Linux apps **98**

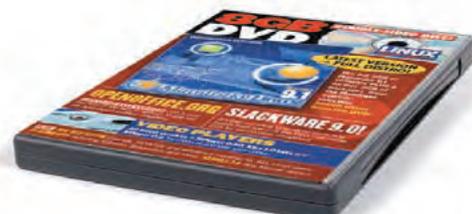


CDS A AND B

Mandrake 9.1 Latest download version of this popular distro

CD C

NEW HELP SECTION! Tutorials, HOWTOs and reference material, for every level; **The GIMP** Latest release (1.3.14) along with nearly 400 plugins; **NVIDIA Drivers** latest versions with a simpler install



DVD

SLACKWARE 9.0 Installs direct from the DVD on anything from a 386 to the latest Pentium 4 and AMD XP+; **OpenOffice.org** Latest release version (1.0.3) and beta of the new 1.1 version; **Samba 2.2.8a** Share files with other OSes; **Video Players** take your pick from three of the best

Please read all the coverdisc instructions on pages 98 to 107 before installing software from coverdiscs!



SAVE MONEY!

SUBSCRIBE TO LINUX FORMAT

See page 96 or phone 0870 4448645



Newsdesk

- Photoshop and OfficeXP on Linux ● Linux fights disease ● KOffice database integration ● iPronto
- WineX3.0 brings 250 games to Linux ● VisualFox Pro on Linux annoys MS ● OpenBSD funding cut

DUAL KERNEL STRATEGY

Novell embraces Linux

Novell has signaled its wholesale entry into the Linux market with the upcoming release of NetWare 7 which, for the first time, will run on either the company's own kernel or Linux. CEO Jack Messman said the strategy would help solidify the diminishing band of NetWare buyers by providing a direct upgrade path for those looking at 'Linux on Intel boxes'. Novell would, he told Computerworld, also add much-needed maturity to the operating system.

"Linux is an immature operating system right now. It hasn't had someone like Novell making it robust, reliable and scalable for very much time. We think we can bring that to the Linux kernel," he said. Not a remark targeted to ingratiate himself with the Linux community. Fearing his initially rather tactless choice of words was too disparaging, Messman penned a hasty retreat and apology.

"Clearly Novell wouldn't be taking this bold step if we didn't feel Linux was a solid OS with tremendous momentum in the marketplace. In fact, we believe Linux is quite stable and scalable. If we didn't, we would not commit to using it with our NetWare 7.0 release." He added that Linux would obviously continue to grow with or without Novell. "The Open Source community is a model Novell endorses. It is the talents from the developers in this community that attracted us to Linux. We are not experts here, we need your help.

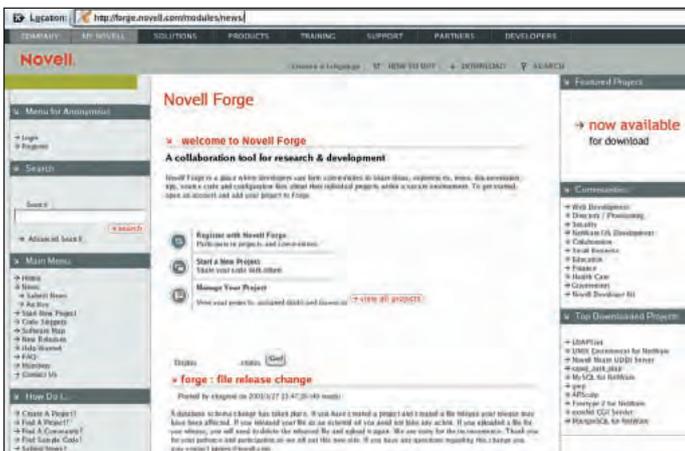
"Novell wouldn't be spending the tremendous time, money and resources to make this strategy a reality if we didn't believe in the present and future of Linux. After building and enhancing



Novell CEO Jack Messman proved adept at removing his foot from his mouth in recent interviews.

NetWare for 20 years, this is new territory for us. We simply ask for your patience along the way."

The apology seemed to do the trick, allowing executives to turn attention back to their product. Vice Chairman Chris Stone, at Novell's Brainshare conference said that embracing Linux wasn't an easy choice, but was the right one. "The Linux operating environment isn't slowing down. It will clearly be the most pervasive of all the kernels out there. I can't find anybody that disagrees with that one. So it became pretty obvious." He said that the dual-platform strategy would give existing customers 'a reason to stay' on the NetWare path. "What we may find is that because we did this, customers will renew their maintenance, just knowing there's a new way to go," he says. But maintaining existing customers is not enough. New users have to be recruited, and Linux will play a vital role.



Novell Forge is a new website dedicated to Open Source development.

Novell Forge No 'incompatible license'

As well as giving Linux a central role in its operating, Novell has also unveiled a new website which aims to give developers a direct path to the company's Open Source initiatives. The Novell Forge site <http://forge.novell.com> will build on the contributions the company has made to projects such as *OpenLDAP* and *Apache*. The launch was accompanied by the donation of the source code for the *Novell Nsure UDDI Server*, which adds identity management capabilities to the Universal Description, Discovery and Integration (UDDI) standard.

Kris Magnusson said Novell would not repeat the errors of other major software vendors who had dipped their toes in the Open Source waters. Novell, he said, would not be creating its own 'incompatible license'.

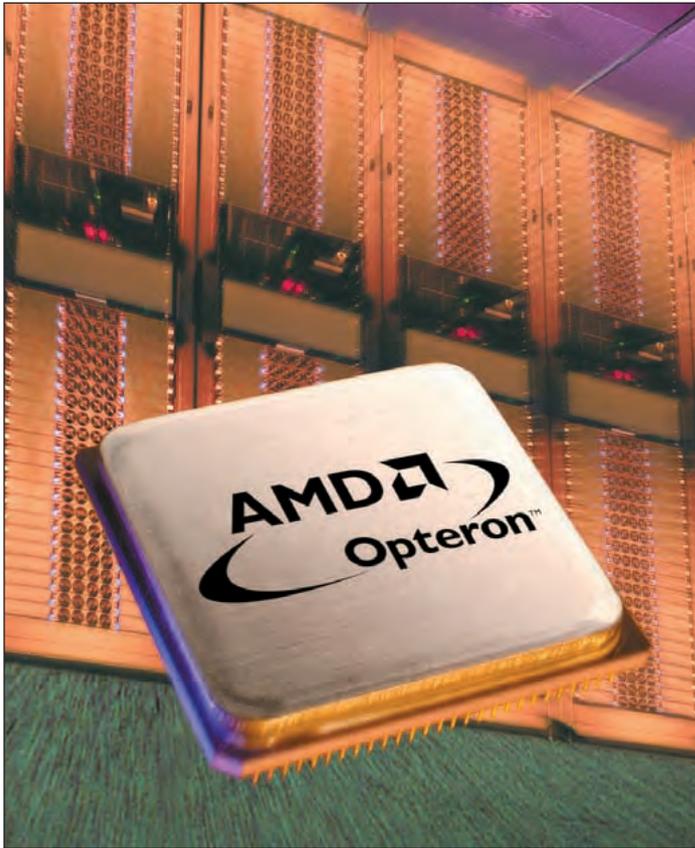
"We've done an extensive review of the GPL - we're not at all scared of it, we

know what it's about and how it works, and we know how to use it in conjunction with our own proprietary services that will run in user space. So you should see 'Novell' and 'GPL' in the same headline in the next 18-24 months."

Bruce Perens, who advised Novell on its Open Source strategy, said the company was genuinely attempting to be "act as good citizens of the Open Source developer community."

"Recently, Novell contracted with MySQL AB to provide a commercial MySQL license to its users with every NetWare system. Since MySQL uses the same code base for its commercial-licensed and GPL versions, Novell's license payments support the continued development of MySQL under the GPL license," he said.

"Now, Novell is taking the next step, by producing its own free software under OSI-approved license."



AMD's new 64-bit chip should drive down the cost of 64-bit computing.

64-BIT ENTERPRISE SERVER

AMD and SuSE

Hot on the heels of the official launch of AMD's 64-bit Opteron processor, SuSE have announced the availability of SuSE Linux Enterprise Server 8 for the architecture. Richard Seibt, CEO of SuSE Linux, said the company was the first off the ground with a tailored solution for the next generation CPU.

"Only SuSE has a 64-bit enterprise-ready server product available for AMD Opteron," he said. "SuSE Linux Enterprise Server for AMD64 enables customers to combine the stability and security of Linux with the performance enhancements available only through the 64-bit architecture."

SuSE has been instrumental in adding Opteron support to the Linux kernel as well as enhancing a range of development tools to leverage the additional capabilities of the system. AMD's Marty Seyer said the close cooperation between the two companies ensured that enterprise customers would have a simplified upgrade path to 64-bit power while

retaining the ability to use legacy 32-bit applications. "AMD Opteron processors provide high performance, flexibility and scalability for demanding enterprise applications in both 32- and 64-bit environments, while also preserving existing investments in 32-bit hardware and software."

One of the first major projects to be built on the pairing of SuSE and Opteron is a new super cluster ordered by Texas A&M University. The system will utilise 128 dual Opteron servers fitted with a total of 384GB RAM and will be used by the University's College of Science to run bioinformatics, mathematical and physics applications. University mathematician Steven Johnson said the new cluster, and the additional memory addressability 64-bit computing provides, would be a tremendous asset. "At a university, price-performance is a major factor in our computing purchases, and the biggest benefit of Intel and AMD getting into the 64-bit market is to drive costs down."

NEWSBYTES

■ The 'freeness' of the GNU projects own GNU Free Document License (GFDL) has been called into question by Debian's legal team. The problem is contained in a section of the license covering 'invariant sections' the ambiguity of which, says Branden Robinson, "isn't good for us or our users." Robinson is calling for a comprehensive critique of the current GFDL (1.2) detailing the perceived problems with the license.

■ A collaboration between India's Council for Scientific and Industrial Research (CSIR), the Department of Biotechnology (DBT) and Tata Consultancy Services (TCS) leads to the development in a new end-to-end bioinformatics software suite for Linux called *Biosuit*. The first Alpha release will out as you read this, with the first official release in May 2004. The project will also involve utilising the 'domain knowledge' and pre-existing tools of the 20 leading academic bodies in India.



■ California's Axis Systems have completed the task of porting their entire software range to Linux, creating what the company claims is the first complete Linux-based acceleration and emulation system for the creation of electronic system and system-on-a-chip designs.

■ Oracle is spending an estimated \$150 million to extol the virtues of 'Unbreakable Linux' to ISVs with the intention of encouraging them to embrace Oracle on Linux. Oracle will also offer porting and developing assistance as well as software tools.

■ IBM's annual DevelopWorks conference stressed its pleasure at the speed with which developers are taking to Linux. Linux strategy manager Adam Jollans said IBM systems had registered more than 5,000 Linux apps. "We're seeing an increase in momentum in developers using Linux and writing for Linux."

■ In the midst of its lawsuit with IBM, SCO has released its first Linux distro optimised for Intel's Itanium 2 processor. The damaging spat with IBM grew out of plans to jointly port SCO's UNIXWare to the architecture before the acquisition of SCO by Caldera and the collapse of the then named Project Monterey.

■ Opera is making some of its voice tools available for free download: www.opera.com/products/verticals/multimodal/

Jono Bacon

Founder of UK Linux and KDE developer, Jono's neighbours are clubbing together to buy Jono headphones for his electronic tuba.



COMMENT

UN and Linux

“ Whatever your views on the United Nations, there should be a collaborative force that brings countries together for commonality in thinking, strategy and for upholding human rights. Linux has been the subject of much discussion regarding its symbiotic nature with human rights; freedom, innovation and availability are key to Linux and as such Linux has been used in poorer countries where such technology would not be typically available.

Linux as a force and culture could mean big things for developing nations and cash-strapped educational communities. Wouldn't it be great if African schoolkids who have little access to computers could have old PCs running Linux with access to learning content on the web? It is contexts such as this where the true power of Linux comes through: technology that helps those in need. Schemes such as this are indeed in operation such as the Redundant PCs project in Sheffield by James Wallbank and the A2RT project in Birmingham, and see page 112 for info on *The Lord's Abode*. The human element of Linux use is unquestionable, and every chance to make 'the powers that be' fully aware of the scope that Linux and Open Source offers should be taken.

Whether the UN will mandate this is anyone's guess, and I am guessing that Microsoft would have a few words to say about this. Irrespective though of what the UN or Bill's army says, we as a community can make it happen. Speak to your local schools and charities and tell them about Linux. Linux could make a real difference.

As a community in times like this where much of the world is in an uncertain state, community spirit should be celebrated and put to good use.

CROSSOVER OFFICE

OfficeXP and Photoshop without Windows

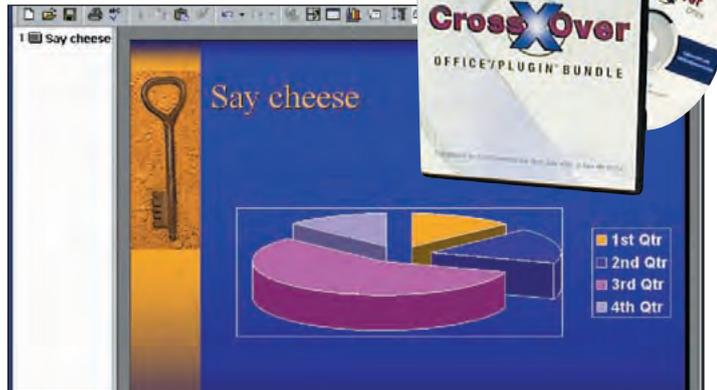
CodeWeavers' popular 'Windows-to-Linux' app *CrossOver Office* has had a makeover, gaining the ability to run most of Microsoft Office XP as well as *Adobe Photoshop*, Intuit's *Quicken*, *Lotus Notes* and *Internet Explorer* without recourse to any Windows native code. The new release also introduces handling of languages such as Japanese, Chinese and Korean but, as yet, doesn't work with Outlook XP. CodeWeavers promise this will be available by the next release.

Jeremy White, founder and CEO said the latest version adds support for many of today's most in-demand professional applications. "*CrossOver Office* continues to prove its value at many of the world's most progressive companies due to its reliability and ease of use—not to

mention its ability to lower costs throughout the enterprise."

White said that, just as SuSE's *OpenExchange Server* was providing a migration path for users of *Microsoft Exchange Server*, *CrossOver Office* would give users of *MS Access*, the database application which has come to dominate small and medium sized enterprise, an upgrade path. "We see *Access* as a key piece of corporate infrastructure that needs an Open Source portability option. By making legacy database applications available under Linux, we remove barriers to broader Linux adoption in many workplaces"

The same could be said of the long anticipated to run *Photoshop* under Linux. Being able to work in *Photoshop* is, according to some pro photographers, a "fundamental requirement" and many



CrossOver Office 2.0 is available now.

advertising agencies, special effects houses and digital production outfits who had previously sought out solutions from *VMWare* or *Netraverse* would now get access to the industry standard photo editing suite without the need to

purchase Windows licenses. Support for double-byte languages will, according to white, open the door to the large and 'potentially very receptive' Linux market in the Pacific Rim.

www.codeweavers.com

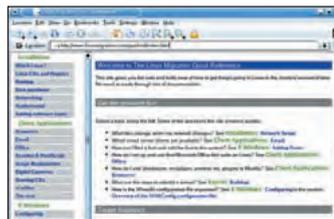
Linux Web Watch/



www.cdc.gov Don't cough while visiting here.



Google's zeitgeist tells us what the connected world is thinking.



linuxmigration.com Good advice for users on the cusp of adoption.



Hyperion's new deal should see more games coming to Linux.

Save the world

Don't 'severe' and 'acute' mean exactly the same thing though?

As we go to press, world medical opinion can't make up its mind whether the SARS virus has peaked or whether we're all doomed. If the former is true, why not turn some redundant computer cycles to good use by helping find a cure for the next time there's an outbreak? www.d2ol.com/ has clients for its distributed computing effort available for Linux, OS-X, Windows and Solaris. Unfortunately, we couldn't get a

screenshot due to the high traffic on the site. It should be better by now – or not, should the 'crisis' continue...

On a connected matter, the Centers for Disease Control SARS advice page (www.cdc.gov/ncidod/sars/) is not the best place for hypochondriacs to visit. Neither is www.who.int/csr/sars/en/, which is the World Health Organisation's take on the crisis. If you want to know how much the virus is dominating the

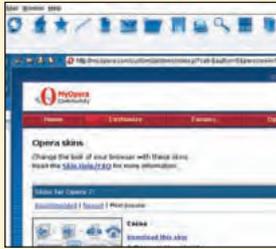
minds of web users, making a quick visit to Google's popular Zeitgeist page www.google.co.uk/press/zeitgeist.html should fill you in. At the end of March, SARS was the fifth most popular query. It will probably be number one for April. And what's the top sponsored link if you search for SARS on Google? Not surprising really, it's face masks!

If you've dipped your toe in the Linux waters, and now want to dive in head

first, www.linuxmigration.com is a good place to start with short tutorials on everything from choosing a distribution and adding software to burning CDs.

Hyperion Entertainment (www.hyperion-entertainment.com) have inked a deal with IncaGold which will see a number of new games, including the lovely looking *Midnight Racing*, coming to Linux.

NEWSBYTES



■ The time lag between Windows and Linux releases of the *Opera* web browser are getting smaller with every version. 710 (beta) has just been released for Linux, bringing features such as the fast-forward button, dynamic skinning and the rather nice M2 mail client to the masses. There is also the intriguing Notes feature that lets you keep track of your thoughts as you surf. Get it at www.opera.com

■ Apple has updated its *Konqueror*-based *Safari* browser to include the latest features, including tabbed browsing. The company has also released the first Apple-d version of *Shake*, the industry standard compositing software. This is the first version not available on Windows, but features an interesting price feature: a slight difference between the cost of the OS-X version (US\$4,950) and the edition for Linux and IRIX (US\$9,900).

■ As featured last month, *TextMaker* for Linux, the new word processor from SoftMaker, is now out to buy.

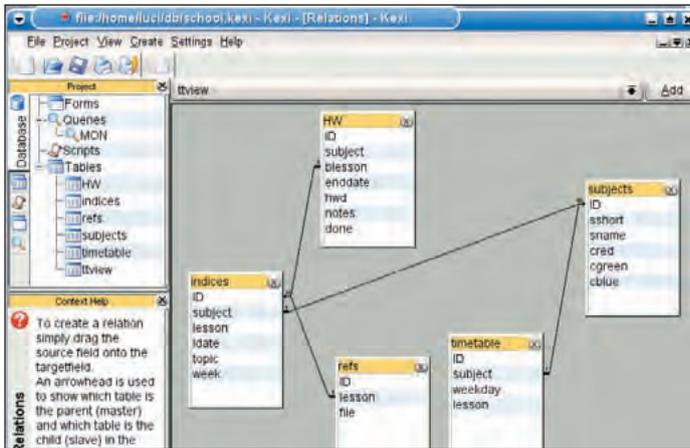
■ A new version of Gaim, the multi-platform, cross-protocol instant messaging application, has been released. Pick it up at <http://gaim.sourceforge.net/>

■ HP and Intel have equipped BP's innovative seismic research facility with a 1000 Itanium 2 processor cluster complete with four teraflops of storage and 8,000GB of memory. BP spokesman Keith Gray said that using Linux, in tandem with a range of Intel and HP apps had achieved a significant increase in computing power and "produced better-defined results in a fraction of the time."

■ Red Hat hopes to build on its enterprise successes with the launch of a new Portal and Content Management System.

■ OpenBSD has received a funding cut from the American government after one of its lead developers made anti-war comments.

■ Mozilla's *Phoenix* browser has been rechristened *Firebird* (which didn't please the makers of *Firebird DB* one little bit) while the new stripped down mail client is now called *Thunderbird*. See our HotPicks section on page 37 for more about the latest version of *Thunderbird*.



KOffice gets integrated access to MySQL and other DB backends at last.

PRODUCTIVITY KOffice gets database friendly

The latest beta release of the KDE integrated office suite has been released complete with a brand new graphical database frontend, as well as the usual collection of squashed bugs and file filter improvements. The latest addition to the *KOffice* family for version 1.3, which is scheduled for a September launch, is *Kexi*, a database client which integrates with both the office suite and KDE itself. While *Kexi* is designed to play well with most database backends, this initial release

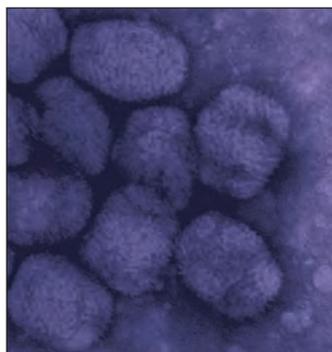
only supports MySQL fully, though there is partial support for CQL++. Close cooperation with other elements of the suite means *Kexi* can be used to create, for instance, mail merges in *KWord* and reports using *Kugar*.

KWord itself has an improved *MS Word* importer and the RTF filter has been rewritten to support pictures, tables and footnotes, while *KSpread* has gained the ability to import *OpenCalc* files and over 100 new formulas. www.kde.org

GENOME MAPPING Linux fighting disease

Linux Networkx has been contracted to construct an Evolocity supercluster by the US gov't's Centers for Disease Control (CDC) to study the genomics of the smallpox virus.

Current Smallpox treatments often cause recipients to become sick and, in one or two in a million cases, will result in the death of the patient. Though the risk is small, the Bush administration (which decreed earlier this year that the US military and members of frontline medical services should be inoculated against the disease) is anxious that new safer methods are developed as soon as possible. Linux Networkx Vice President Eric Pitcher said that the price/performance benefits of Linux made it very attractive in the life sciences, especially vaccine research, "where



Linux can be used to fight physical as well as virtual viruses.

speed is of the essence." In the UK's life sciences field, Screensaver Lifesaver is looking for volunteers to port their efforts to Linux and Mac OS: www.chem.ox.ac.uk/curecancer.html

Hoyt Duff

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



COMMENT WINE-ing?

Both CodeWeavers and Transgaming continue their success with WINE because they've focused on facilitating the use of particular apps, not evangelising Linux. They preach their Linux gospel obliquely: "Linux is nice, but do you want to use a favourite office app or play a favourite game? Then use our product."

They continue to be the brunt of objections to their marriage of Linux and MS Windows, coming mostly from self-proclaimed purists who eschew all Redmond-related code solely because it was spawned there. These cretins attempt to make you and I feel guilty as they claim we support an effort that will stifle development of "free" Linux games and office apps. Give that canard a rest!

Linux gaming will continue to develop slowly until it becomes feasible for major developers to invest in it; likewise for office apps. You can help by politely informing everyone you spend money with that you would appreciate more abundant Linux support and thank them for any support they do offer.

But until Linux conquers the world, the zealots would do everyone a favour by redirecting their efforts toward productive things like improving Linux docs and beta testing new code. There are several orphaned documents at the Linux Documentation Project that could use their help. What about them volunteering editing or writing skills at an online mag like www.linuxnetmag.org or writing a product review for the *LXF* website? They might take part in a newbie mail list to actually help (rather than flame) novices. Or perhaps they could present a few HOWTOs at their local LUG instead of trolling those mail lists and dissuasion groups. Stop whining, start WINEing!

VISUALFOX PRO DB ON LINUX ANNOYS MICROSOFT

Violate EULA, go to jail?

VisualFox Pro developer Whil Hentzen had his plans to demonstrate the Microsoft database development environment running under Linux thwarted by the intervention of Ken Levy, Visual FoxPro marketing manager. Hentzen had planned to show his work to the Bay Area Association of Database Developers, but postponed the demo after being told it would contravene the End User License Agreement (EULA) of the software.

Under the EULA, developers can use VFP to create an executable, which can then be distributed (along with a support library .dll file), to an unlimited number of end users. The EULA does not limit VFP to running on just MS operating systems; but only says you can't use VFP to design, develop and

test programs that are not designed to operate on MS products. That's a critical distinction: if a product program runs successfully on Windows, then surely it's OK to run VFP on Linux to design, develop and test the product on Linux?

The Microsoft position was confused further when Levy posted his "last comment on the issue" essentially stating that VFP is a Windows-only tool: "VFP was designed and tested for use in creating applications that run on the MS Windows platform; the same applies to the components that are provided to developers for redistribution with VFP-based applications," Levy wrote, "If a developer wishes to distribute the VFP runtime with an application, the runtime may only operate in conjunction with a Microsoft Windows platform"

Rather than the episode reinforcing



the strictures of the Microsoft EULA, opinion in the DB community suggests that it merely serves to highlight deliberate ambiguities that MS write in

to its EULAs in order to interpret the agreements as it wishes after release. "MS makes EULA rules up as it goes along," as one poster succinctly put it

MORE LINUX GAMES!

WineX 3.0 is here

While Office users revel in the features of CrossOver Office, Gamers can now rejoice in the release of WineX 3.0 from TransGaming Technologies. The latest version supports a significant number of new games including *BF1942*, *SimCity 4*, *EverQuest* and *Medal of Honor: Allied Assault* as well as, under the hood, improved graphic abilities with implementations of Vertex Shaders and Cube Maps, better copy protection support, extended installer functionality and a number of game speedups. However, the most obvious improvement

is the addition of a GUI, Point2Play, to simplify use of the system.

Point2Play replaces the interface of previous releases, allowing multiple WineX installations, saved games and store options without going anywhere near a command line.

Gavriel State said this was the most significant release in the project's history. WineX now boasts the ability to run 250 titles and, for users running the 2.5 series development kernel, it now supports Force Feedback joysticks.

www.transgaming.com



WineX 3.0 brings EverQuest and many other games to Linux.

Embedded Linux News



- Electrical giant **Philips** has launched what must be the world's most expensive remote control. The Linux-based iPronto retails at a wallet-busting \$1,699 but goes well beyond the scope of ordinary controllers. In fact, Philips says it "combines audio/video, home automation, and home network control capabilities in an easy-to-manage device that offers consumers a new solution for controlling a wide range of devices in the connected home." Essentially, using either IR, RF or Ethernet (wired or wireless) the iPronto will facilitate access to every home device from a humble light switch to a broadband Internet connection from one handsome looking tablet. It also has a USB port and SD/MMC card for expansion, microphone and stereo speakers for audio tasks and a touch sensitive screen. The device is built on MontaVista software and, Philips says, should be available to buy in the second half of the year.

- MontaVista** has also inked a deal with Panasonic to bring audio and video over broadband to a television near you (if you

live in Japan). MontaVista's embedded Linux system will power the Broadnow set-top box, which features traditional video/audio, features as well as Internet access and a selection of TiVo-like PVR tools. The Broadnow device is the second major product launch to grow out of Panasonic's investment in MontaVista, following the January announcement of a Linux-based videophone.

- After the jailing of a US retailer for selling Xbox mod chips, the desire among many users for a way to put Linux onto an unmodified machine was great. Fortunately, someone has done it thanks to a method called the **habibi_xboxes 007:Agent Under Fire game save exploit**. Find more at www.xemulation.com.

- AMD has created a reference design around its Au1500 MIPS32 architecture system-on-chip. The system is built for access systems such as routers, wireless access points and self-contained Internet Access Devices. The full spec is available at www.amd.com/connectivitysolutions/au1500



Software AG has completed the job of porting their entire product line to Linux.

SERVER

Software AG gets SuSE certified

Underlining its recently stated commitment to Linux, Germany's Software AG has completed the process of porting its entire product line to Linux. The full range of database solutions and XML development tools is certified for use with SuSE's latest Enterprise Server

product. R&D specialist Dr. Peter Mossack said Linux would play an increasingly important role in the server sector, and this was a trend that could no longer be ignored. Citing a report by The Meta Group, which claims Linux will count for some 45 per cent of all new Intel-based server

installation by 2004, Mossack said the company was reacting to the needs of customers. "By porting its products to this open-source operating system, Software AG is meeting the emerging market demands and allowing its customers the freedom of choice they are looking for in platforms."

TYPOGRAPHY

Nice fonts for GNOME

The GNOME Foundation has struck a deal with font design behemoth Bitstream to donate a collection of ten serif, sans serif and monospaced fonts to the open source community. The Foundation says the move will fill the "remaining gap in font display for Open Source software. Foundation board member Jim Gettys said the fonts would benefit not just GNOME, but the wider OSS community.

"The donation of these fonts to the free software community is the final piece that will give full functionality to projects like Freetype, XFT2 and XRender extensions of the XFree86 project, Pango, KDE and Trolltech's Qt, among many others," he said "These fonts will be available to all developers and users, giving GNOME and other



Bitstream's Vera fonts will bring high-quality text display to Linux and even embedded Linux 'right out of the box.'

Open Source programs a great look right out of the box that has been lacking until now!

Bitstream's Bob Thomas said the company had been interested in the idea of open development for a long

time, and he hoped the fonts would help generate new fonts for the community. "We're very happy to contribute these fonts to the open source community and we hope that developers build on them."

Paul Hudson

Paul is Reviews Editor at *Linux Format*. The creases in his shirts are ironed so crisply that the LXF office has no need whatsoever for a pencil sharpener.



COMMENT

New kernel!

“ Yes, fellow Tuxters, at long last and after over two years of waiting, the new Linux kernel is approaching a sufficient level of stability that I will deem it safe to upgrade. So, it will be with great anticipation that I will type **apt-get install kernel-image-2.4.18-686**. What? You were expecting, perhaps, that I'd be installing a 2.5.x-soon-to-be-2.6 kernel? I think not. My own boxes run a 2.2 kernel, and have done so since late 1999. Performance- and security-wise, it's been a great release and I've been very happy with it.

All too often people coming from Windows hear about Linux being 'uncrashable' and 'rock solid'. From that, they seem to think that the two together are paramount to allowing them to install whatever they wish, even if it's hot off the press from CVS, without Linux batting an eyelid.

Usually, these same people call me up late at night to tell me that the fonts in KDE have gone crazy, or that they get "Oops" messages from the kernel, and – pretty-please-with-sugar-on-top – will I fix it for them?

Sadly, these cases of "if it ain't broke, 'fix' it 'til it is" are becoming more common amongst newer Linux users, and I think perhaps the community as a whole needs to do more to encourage users to stick with what works until the new version is stable.

So, when 2.6 is finally released, think twice before you rush to download it and get it running.

PS: www.thinkgeek.com sell a T-shirt with "No, I will not fix your computer" printed on the front. If some kind soul wishes to save me hours each week answering calls from 'friends', the LXF address is at the back of the mag ;)

Mailserv

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

Linux not kids' stuff

Being new to both Linux and your mag, the move to both has been a very enlightening experience. My computer, run by water-cooled dual AMD processors, has moved me into the realms of needing an OS with the appropriate support; due to the work involved in running a site, the need of reliability and consistency have also now become paramount. The move to a dual system followed – I opted for a bay caddy – one tray with my Win2K and one with Red Hat. I don't trust Win2K enough to go for multi-booting and this is an easy option.

I went down this road to some extent to find an alternative to Windows and have to say I have not found it, but Linux is perfect for my wife who needs to email and web browse. Though I miss some of the slick packages on Windows, their departure has not caused too many problems and I am even beginning to like *The GIMP*. With programs like *GAIM*, many people don't know which system I am on and I now run it on my W2K install as well.



Linux isn't for everyone – LXF condones symbiosis with closed source.

However, for the children it is a disaster. None of the educational packages such as *Reader Rabbit* and their games such as *Action Man* will work. They are 3 and 6 years and need a system that does not use passwords, usernames etc. One day

they will, but as yet, its not appropriate. I read a comment in *Linux Format* that intimated that Linux will never make it as a family option until it has full game support. I have to agree, as this game support also brings with it

the educational and fun packages for the smaller computer operators. I cannot only consider my needs but also that of my whole family.

I know lots of people will start to say how their 2-year-old loves raw coding, but my kids love being kids, and inacey wincey spider is more in the league of my youngest.

I will not give up my Linux drive and believe that it is not an alternative to Windows, simply a better OS. As a family OS, it sadly cannot compete with what is clearly an inferior alternative. Hopefully, as they get older I will be able to wean my kids off Windows, but that will depend on how well Linux, as an OS, will start to consider the whole family as its user and not just geeks like me.

Andrew Willmore, *via email*

Well, there are plenty of games for Linux – perhaps not the ones that their friends play though. It is also possible to set up Mandrake for example, to login automatically and therefore not need a password (though the login system has now been adopted by Windows too).

A lack of good educational software for Linux has been a point raised in the mag quite a few times. There are a number of projects set up to address this. For early learning it's fairly easy, but differences in teaching methods and focus on particular curricula, make it difficult to create appropriate material.

Linux isn't for everyone. I think it's good to point out though that the nature of Linux doesn't prevent such applications being developed. It's merely the mindset of 'everyone uses Windows' and economic forces to a varying extent that cause proprietary software to be developed for that market.

Image Scan!

I recently picked up *LXF37* at a *Barnes & Noble* store in the US. As a relative newcomer to Linux (I tinker with it at home and, for the moment,

★ Letter of the month

This month's winner receives a copy of **SuSE Linux Professional 8.2**

Video death

Re *LXF39* Answers: drivers needed I'm a 'silver surfer' returning to UNIX after a 15-year gap – I can offer some help. I had problems with SuSE8.1, ATI Radeon 9000 Pro, Hitachi CM751ET (monitor).

All installed fine, but got a blank screen after the video card stage.

I had not expected a full 3D support just a basic 1024x768 16-bit, which any old card and a simple driver can deliver. So I tried expert install, this lets you test a

video mode before committing to it. My monitor had been correctly recognised, however the scan frequency ranges were a bit out (too wide)! Correcting them fixed all the 2D problems that I had. I will wait for 3D.

I have since discovered someone else with a similar problem and the same solution applies in their case. So try any default driver, but get the monitor settings right. A few kilohertz can be video death...

Roger Wilkins, *via email*



It would seem then that the monitor definitions are incorrect. Thanks for pointing out this particular problem, I hope you can also give feedback to SuSE to correct this in future releases if it hasn't been done already.

To thank you for your trouble, we'll award you this month's star prize, which by handy convenience is SuSE Linux Professional 8.2, reviewed on page 22 of this issue.



Epson provide a range of drivers that work with Linux for their scanners and printers.

run one server on Linux at work), I had not heard of your mag before and thought I'd give it a shot. Great stuff! The accompanying DVD was impressive, and the mag's packed with good stuff. I really like the fact that, at least in this issue, I'm finding there is at least as much information

“The key to Linux’s domination is to lull the Windowphiles out of their rut by making the setting familiar, not scaring them with loads of technical info.”

that I would use at home as there is information I would only use at work as Technology Director at Gardner-South Wilmington High School.

I wanted to respond to one of the letters in this issue. In the Letter of the Month, another reader asks about scanner software provided by manufacturers. There actually is one maker that supplies scanner drivers and software for all their scanners, and that's Epson. Its *ImageScan!* and it reputedly performs as well as its Windows and Mac counterparts (I've not tested it myself). Try visiting: www.epkowa.co.jp/english/linux_e/isd_e.html. Links to *Image Scan!* as well as some drivers for their inkjet and laser printers can be found at www.epkowa.co.jp/english/linux_e/inux.html.

Mike Oliveri, *via email*

I'm glad you find the mag as informative for your home use as it is for work. We try to cover all areas of Linux use, which means helping out new users and covering desktop stuff as well as servers. Thanks for the Epson tip too!

Trained snobs

I left the military three years ago, with fairly modest IT skills so when offered the normal re-training package I went for the MCSE path (a guy's gotta start somewhere). So now I'm a competent engineer earning a meagre wage and it doesn't take a chimp to know that

the way ahead is Linux. My main problem is my wage allows me to learn all of the Microsoft skills bit-by-bit, the average Microsoft skills book is about £38-£80. I've contacted Red Hat about training and the training course may be excellent value for money but if I'm training I ain't working and if I ain't working I ain't earning.

So how on earth do I gain the skills? Well, it is not easy because I've found that the average Linux professional has an attitude towards Windows users and this has made it difficult to learn or search (thank god for the newbies site). All my present study is done online – but despite having broadband I don't want to spend all my life in front of a monitor. Can anybody free me from my monitor and help me become a Linux engineer (either RHCE or any other qualification available) and are there any readily available study guides? I only want to be able to study on the train, hill-walking and even in the bath!

May I also finally add to the Linux snobs out there that they are now holding up the world domination of Linux? The key to this domination is to lull the Windowphiles out of their rut by making the setting familiar to them, not scare them away with loads of technical information.

Eugene Meenan, *Merseyside*

There are a number of training companies that offer Linux courses, mostly based on the vendor-neutral Linux Professional Institute Certification (LPIC) scheme. The exams themselves are relatively cheap (about £60). Companies like Training Camp www.trainingcamp.co.uk offer intensive courses if that would suit you better. And of course, there's plenty of material in the mag, and our new Help section on the coverdiscs!

Red Hatty Hatty Joy

I have to tell you 'what a joy' to install *Star Office 6.0*. I simply loaded the CD onto Red Hat 8.0, found an icon called 'setup', double clicked and away went the installation. I then restarted my session and found the *StarOffice* applications and away I went. Productivity within 15 minutes!

I then think, hey, wouldn't it be nice to upgrade GNOME to Version 2.2 and while I'm about it I could also upgrade KDE, all of which are free on your cover DVD.

Ah, a slightly different story, I went wrong somewhere along the complex line and have to re-install Red Hat! Productivity resumes after about 3 hours. I know Linux is meant to expand our knowledge and increase our learning curve and I'm happy for it to do so, when I have the time. But why can't developers make installing and upgrading software a little bit easier?

If this was the case, I would have no problem getting the rest of my family into Linux, as well as encourage my colleagues at work who are all sick of the big MS and are keen to change. Hey, you never know, the company I work for might also change over. (though Linux's cheapness is beaten into a cocked hat when you have to employ highly skilled IT engineers to implement it!)

I have recently been looking into Mac OS, the look and feel seems akin to Linux. I bet you don't have to ./configure-make-make install on that? I like the look and feel of Linux, its stability and its functionality. I persevere with installing software. But if someone would take out all the hard work I would be willing to pay for the OS as I'm sure would a lot of other people?

I tried the *SME* server – excellent! Could I get it to dial my ISP? No. As stated in your review, there is a lack of facilities. Could do with more access to Dial Up, PAP/CHAP passwords etc. Still, have made good use of two redundant PCs, one with Smoothwall and the other with *SME* as a file server. Thank you very much for the excellent software bundles that come with the magazine

Adam Cusack, *Shropshire*

Thanks for your interesting letter. We're sorry your GNOME upgrade didn't work out for you. In some ways, when you look at the complexities of upgrading a major piece of the systems software, while the system itself is running, perhaps we should not be too surprised when sometimes it goes wrong, though obviously it would be preferable if it didn't. I'm not an apologist for lazy install routines – you should expect a commercial software



Mailserver

« package (which usually have no real interdependencies) to install with a couple of clicks. Because system software like GNOME and KDE is complex, and builds extensively on the work of others in the form of libraries etc, it is a slightly different proposition to create a foolproof, runs-every-time installer.

I'm glad you have found a use for *SME* server. Development continues on all areas of the system and I'm sure it will be greatly improved by the next release version.

Bad behaviour?

I have been reading a lot about failed distributions lately. I would like to recommend that the sales departments of companies distributing Linux take a close look at how they handle customer relations, or more importantly don't handle customer relations. I would like to think that I represent a vast

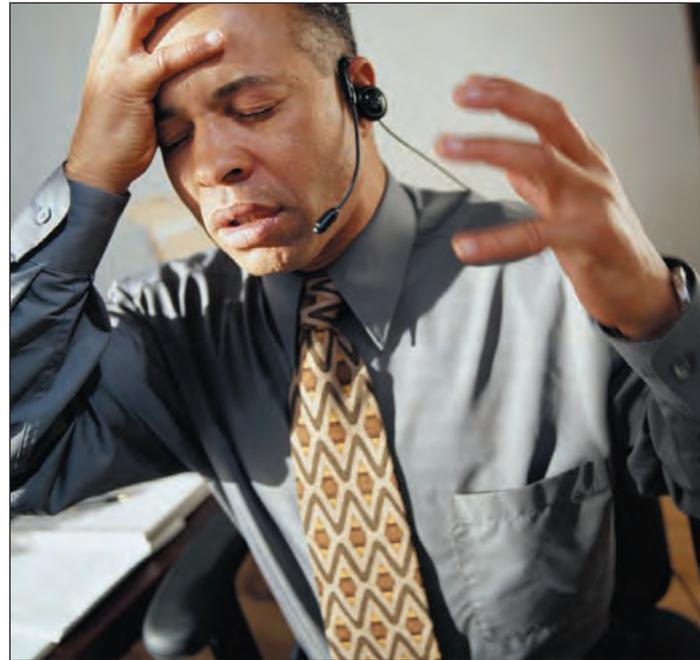
majority of people that are willing to spend the money but refuse to if I can't even get a few simple questions answered.

For example:

Mandrake Linux: refuses to even post an email address to the sales department for pre-sales questions. I have attempted to contact these people for detailed questions about product component level support prior to purchase, without success. No answers means my money goes somewhere else.

Red Hat: They just flatly refuse to respond to the individual user for pre-sales inquiries. I have tried more than once. Heck, I have even tried to get information out of them for my place of employment, an enterprise level type of situation. I do see how they get compared to the Redmond guys. Ditto about my money.

SuSE: My hat is off to these guys. A



If your company gives away distros, funding support can be a headache.

person may not get the perfect response, but they do acknowledge you exist and they do try to help. SuSE is probably the only reason I have had any success with Linux, pre-sales or after. I have a copy of SuSE 7.3 Pro, and am most likely to purchase another version once I get a home wireless network sorted out. The bulk of my most recent questions have been answered.

Debian: Again, my hat is off, and they don't even sell anything. I personally have had little success getting a version 2.x to cooperate, but that should reflect on me as an absolute beginner. I have had many responses to inquiries from their support base.

Thanks to SuSE and Debian. Please keep up the good work. John R. Klaus, *via email*

“I have been reading a lot about failed distributions lately and recommend that the sales depts of distro companies look at how they handle customer service.”

Helpdex

shane_collinge@yahoo.com



I think you might have to excuse Mandrakesoft – it has been going through a lot of reorganisation and a few money woes, so I think its priorities might be somewhat skewed at present. It is a good point though – how can distro vendors answer the many legitimate questions that customers may have before someone even forks out for their product? I think you'll find that the answers are usually more effectively found from the community as a whole. This isn't to say that the commercial distro vendors shouldn't be trying to address these areas though.

The 3D World

Being a recent Linux convert and reading the mag I have found pleasure in Linux, tinkering and learning about kernels etc. Being a designer, Linux conflicts with my career somewhat! Yet Linux is gaining recognition as a powerful 3D and CG OS. It would be a good idea for you to feature some Open Source 3D and design software on your cover discs maybe an article too. For example softimage has a demo of xsi 3.0 for Linux, that would be a nice addition!!

I know many people who lament the lack of macromedia software available for Linux. At present *vmware*, *win4lin* and WINE are the best options (besides using Windows, but we dont want that!)

I'm not sure what companies like Adobe make of Linux's growing popularity, but it's something they can no longer ignore.

J. Stars, Leicester

Linux is indeed becoming more and more important in the 3D and graphics industries. With standard software such as *Maya* already ported to Linux, not to mention free software projects like *FilmGIMP* which already has many Hollywood credits to it's name, there's no debating that Linux is a powerful tool at both the design and rendering ends of 3D graphics.

Adobe has dabbled with Linux in the past. *Framemaker* was available for some time as a free, beta download, but they discontinued it due to 'lack of response'. My opinion is that the graphics/DTP/3D scene was still a little nascent in that period and *Framemaker* didn't compare too well with tools that were already available.

I think the news that *Photoshop Albums* has been created using the Qt

toolkit is an interesting move – the effort to port such software to Linux would be fairly trivial, so it's just a question of the companies having the desire to do it. I think largely the support for high end tools is there or getting there. User-space tools might be a different matter, but they are fairly well served currently by free software such as *The GIMP*, *Blender*, and other software we have actually featured in the magazine.

Gphoto2 drivers

In the January 2003 issue of *Linux Format*, someone on the forum was having problems with getting their HP Photosmart working with *gkcam*.

I use KDE with *gphoto2* and a Kodak DC-280 Zoom and had problems with it as well, but I did get it to work. What I found out was that *gphoto2* includes its own driver for my camera and the kernel driver interferes with it. When I had the kernel driver loaded (compiled in actually) and tried to connect to the camera, I would get the error "Could not claim the USB device". Once I removed the camera driver, everything works flawlessly.

I would recommend that the user checks for any loaded kernel modules for their camera and, if they're present, remove them.

Steven Elling, via email

Yes, this can be a problem. The kernel module will automatically be loaded when the camera is plugged in if present. If you are using any software that uses its own drivers, it's best to leave such modules out of a recompile. In the latest SuSE and Mandrake iterations, most digicams automatically appear as an external device. [LXF](#)

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 questions – direct those to our Answers pages!
- Spam
- Nonsense rants and random abuse
- 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

»There is no reason for any individual to have a computer in their home.«

Ken Olson, former President, Digital Equipment Corp., 1977

»Linux is only for computer science students.«

Another myth disproved !



SuSE LINUX is for everyone – anytime – anywhere!

SuSE Linux requires no prior training or special hardware. Simply install and start using a multitude of applications that allow you to work and play as you wish. An ideal replacement for your old operating system, whether on your desktop PC or your laptop, it is just as easy – but faster, more secure, and much more economical.

SuSE LINUX 8.2 PERSONAL

- The Concise Linux for beginners and home users
- Easy working with OpenOffice suite (compatible with MS Office)
- Games and handwriting recognition programmes
- Multimedia applications (e.g. CD recording, MP3 and MoviePlayer)
- Secure e-mail and Internet access

SuSE LINUX 8.2 PROFESSIONAL

- The Complete Linux for advanced and SOHO users
- Automatic network identification for your mobility
- Multimedia and video editing to professional standard
- Optimised installation, operation, and help functions
- Numerous development and programming tools for the experienced user

@ info@SuSE.co.uk

+44(0)208 846 3918

+44(0)208 748 0538

SuSE LINUX Ltd.
Appleton House
139 King Street
Hammersmith
London W6 9JG · UK



Reviews

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

LXF verdict explained

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

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

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

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

Documentation / Value for money: Whichever score 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...

Mandrake 9.1

Past releases have been criticised for not enough improvement – does this iteration cut the mustard? **p18**

Red Hat 9

Unarguably important to the business market, but whether RH 9 is as good for home use remains to be seen **p20**

SuSE 8.2

It's not often that LXF lavishes a particular product with so much praise – find out why we're so excited **p22**

AC3D

Many Linux 3D apps don't support enough proprietary filetypes – this one does! **p26**

Astaro Security v4

Easy-to-upgrade software firewall with great scalability to protect your network **p27**



JCodeBox

If generating lots of repetitive Java is annoying you, you'd best try this code generator on for size **p28**

OpenAuth

Following Evesham's lead, this desktop box comes with Linux pre-installed and ready to go **p29**

NEC Express 5800

1U server with two internal gigabit network cards provides a whole lot of connectivity **p30**

Book Reviews

MySQL Second Edition, *LDAP Directories Explained*, *Unix Power Tools*, *PostgreSQL* **p33**

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**

LINUX DISTRIBUTION

Mandrake 9.1

Biagio Lucini puts Mandrake's latest release through some heavy testing to assess the scale of the improvements over previous versions.

BUYER INFO

Competition: The other two 'big' distros, Red Hat and SuSE have both been recently updated.

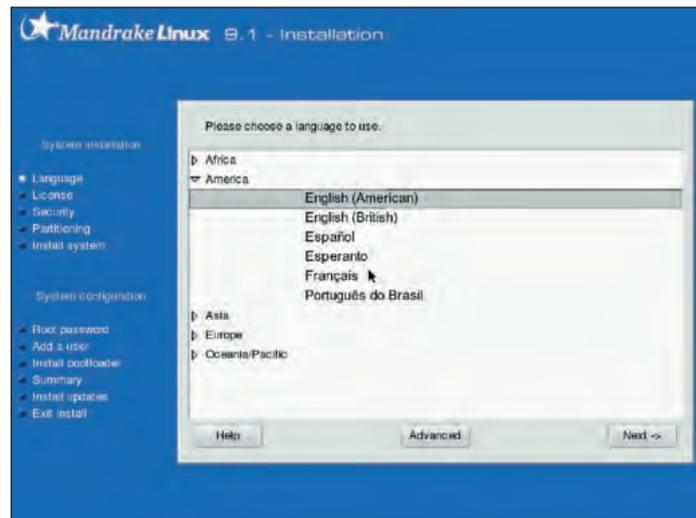
- **PUBLISHER** Mandrakesoft
www.mandrakesoft.com
- **PRICE** ProSuite £199
PowerPack £69
- **WEB** download from www.
mandrakelinux.com/en/ftp.php3

The previous couple of Mandrake releases haven't quite met users' general expectations in terms of the traditional MDK focus on ease of installation and use, arguably showing little improvement over preceding versions. Because of recent financial troubles, if Mandrakesoft wants to stay in business, this 9.1 product must be very solid if it is to heal cashflow problems. Indeed, the company has promised in the past that Mandrake 9.1 (alias *Bamboo*) would be its best release ever. How does this claim actually pan out now that the new iteration is here?

Installation Process

Specific instructions regarding installation from our coverdiscs is covered on pages 98 to 107. For those unenlightened souls who don't read *LXF*, Mandrake 9.1 can be downloaded as three ready-to-burn CD ISO images from one of the ftp mirrors. Other installation methods (eg ftp install, more on which later) are also available. The available packages have been divided in such a way that only the first CD (Installation CD) is required, while the Extension and International CDs contain optional packages. The third CD will be handy though, unless you decide that your new system won't need the kernel source and other development packages that for some obscure reason ended up buried among the i18n rpm!

After booting from the first CD (alternative boot kernels are available on the second CD, should the default one not work for you), we chose the 'Expert' installation option, which as usual does



A new simplified layout and great antialiased font for the MDK installer.

not require a higher degree of expertise other than having heard words such as partitioning, boot loader etc before.

The installer

The 'Expert' method also offers the possibility to load third-party modules, which can be skipped unless you know what you are doing and why. The next step fires up the graphical installer. Though the installation steps are somewhat similar to those of the previous editions, the installer itself has been greatly simplified and (thanks to underlying *gtk2* libraries) appears altogether more attractive through the use of antialiased fonts. The layout of the installer looks much cleaner, with a reasonably sized light grey main window on a blue background. As before, installation steps are listed on the left of the screen, with the step being performed highlighted. The task list looks much shorter than before, but the simplification is only apparent, since the highlighted task often does not reflect exactly what it is going on.

For instance, you'll have to configure your mouse under the section Language and choose the packages under the section Install! Many advanced features are kept hidden by default and can be reached only by pressing the Advanced button,

which is available for almost all steps. Whether these options are really advanced or not is questionable, but it does show that Mandrake is becoming ever more 'ready for desktop'; the day that stores like Dixons start stocking boxed Linux distros draws closer..

Same but different...

The install process is pretty simple and self-explanatory, and the Help button will definitely come very handy, should this not be the case.

So what are the main differences between 9.1 and earlier iterations? Very welcome for those who will have to shrink the latest release of that 'other' OS, is the ability of the partitioning tool to resize NTFS partitions on the fly. Another long-awaited feature is the ability to choose among different journalised filesystems during the install, with ReiserFS and XFS joining EXT3. There is also a claim that upgrade should now be trouble-free, but we haven't tested that properly yet due to space restrictions.

After copying the files to the hard disk, you are prompted to set the root password, to add users to your system and to install the boot loader. It comes as a surprise that the installation of X is not a step in the default install process. Once you have installed the boot loader,

a summary screen will appear, showing what are the components configured so far and – importantly – highlighting in red what still needs to be done, including configuring X. In our test, many of these configurations were performed automatically, so no complaints there. This is also the place where you are offered the chance to go back to some previous sections and change some configuration options. With no other input than confirming the autoprobbed hardware, next the installer configured and brought up our wireless connection and our NIC card and picked the correct printer from a list of more than a thousand. If you try to quit before configuring X, a friendly reminder pops up to ask you if this is really what you meant.

We then hit our first glitch – the following step should have been the download of bug-fixing packages, but the installer failed to complete it.

We also tried an ftp install. Here we were somewhat surprised that, having used the Net to install the system, in the summary the network was listed as Not Configured. The fact that this time the installer wasn't able to identify the wireless card despite our hints towards the module to use led to the conclusion that the different degree of success may depend on the installation method, the auto-booting CD probably being the safest one.

Back to the Desktop

The new default login manager is *mdkdm*, a reworked version of *kdm*. Originally engineered to look similar to the previous versions of *kdm*, *mdkdm* has evolved to a very competitive alternative to the more traditional *kdm* and *gdm*, which are still available should you want an alternative.

At the first login, the user is presented with the by now well-known MandrakeFirstTime wizard, which at the end offers the possibility to join the Mandrake Club (see box above right).

Bamboo comes with a range of desktop/window managers which will satisfy all users, from the 'just-point-

Mandrake Club

www.mandrakeclub.org

The Club is the natural complement to the (100% free) Mandrake Download Edition. For a reasonable monthly fee, Club members get bonuses like voting for packages that will be in the next main distribution, deciding which applications have to be recompiled for the current distribution, accessing online docs and exclusive rpm packages (among which there are some commercial apps) and others.

and-click' user to the die-hard 'give-me-a-terminal' geek. However, not the same level of attention is paid to all offered environments: while (not so strangely) lots of effort has been put in the look-and-feel of KDE and GNOME, the other environments ship with little (eg *IceWM*) or almost no customisation (eg *WindowMaker*).

What really shines in this release is the new and wholly unanticipated MandrakeGalaxy desktop, a unified and fast theme for GNOME and KDE. With much of the hype surrounding Red Hat's BlueCurve, Mandrakesoft has wanted to show with this new environment that even if it is now widening its interests, they can strike back quickly in the desktop arena. MandrakeGalaxy is not an imitation of BlueCurve: while the latter tends to bring the two desktops as close as possible to each other (same look-and-feel), the former offers good visual consistency but keeps much of the individual identities of the two desktop managers. Of course, this is not clear-cut; since for instance, on the look side there is the lack of a unified icon set, while on the feel side users can choose if they prefer single- or double-click.

In the default desktop configuration, all the ads pointing at Mandrake websites have been removed. Instead, a Welcome icon has appeared, which acts as a container for links to *drakconf* and various Mandrake support sites. Mandrake developers have also done a great job in configuring nice-looking antialiased fonts for many of the available applications.

Irks and quirks?

In our test system all hardware not configured during installation (CD-RW, USB and IEEE1394 controllers) had been identified correctly and was ready to be used after the first boot. On the plus side, we would like to stress that

despite these annoyances all our hardware was recognised and ready to use either automatically or (as in the case of a scanner) by the corresponding wizard. Removable media are managed via the hotplug package and the mandrake dynamic desktop: when you plug-in your zip drive for example, a suitable entry is added to the *fstab* and when you try to mount it, an icon appears on the desktop, providing easy access to that device. Actually, this is not terribly useful: if you know how to mount a drive, you don't need an icon.

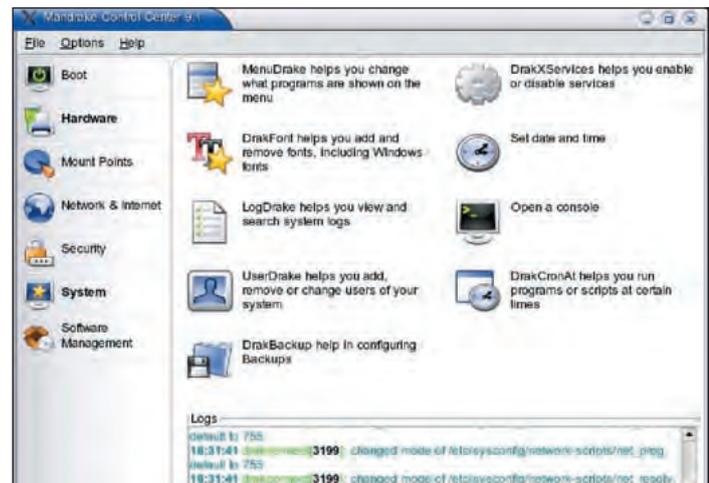
The same holds for a USB camera – seen as a storage medium – though here things are slightly worse: the entry in *fstab* is added after you try to mount the device giving as a path its full hierarchy in the SCSI tree. This can be tracked down to a problem in the configuration of *devfs*; should more attention have been paid to this by developers? Again on the dynamic desktop side, an icon for an automatically recognised webcam only appeared after installing *gnomemeeting*, and again this isn't terribly useful: if you know that for using a webcam you have to install *gnomemeeting*, you don't need that icon on the desktop.

Traditionally a peculiarity of the MDK distribution is the collection of utilities grouped under the *Mandrake Control Center (MCC)*, alias *drakconf*. This offers access to almost every aspect of the system, from setting up X to configuring new hardware; from organising scheduled backups and sharing a connection, through synchronising the system time with an external server to managing users and groups, and many others. As with the installer, the new *MCC* has been largely ported to GTK2. The main window has a three-pane structure, with the main categories listed on the left, the utilities in the working category taking up much of the space in the top right and a log of the operations performed displayed at the bottom right. Apart from the nice-looking icons and antialiased fonts, new for 9.1 is some explanatory text describing each utility. Before, you had to guess what an application here was meant to do, while now you get at least some little explanation that can help you in getting an approximate idea.

If this isn't enough, the long-awaited official *drakconf* documentation has been released and is available with the distro. Still on the documentation side, surprisingly in this release there aren't accompanying guides, the link in the



Konsole, Konqueror and The GIMP on the default KDE desktop with the new MandrakeGalaxy theme.



There is still some work to do on the Mandrake Control Center, but progress since Mandrake 9.0 is certainly encouraging.

MandrakeGalaxy page pointing to a doc website where at the time of writing only documentation for Mandrake 9.0 is present. Official guides are available online only for Mandrake Club members, but this may have changed by the time you read this.

There are a few details that still keep *drakconf* from perfection: embedding of the various apps is not perfect (eg *userdrake*); some wizards do not work as expected (eg the network wizard insists in setting wireless parameters for the wired Ethernet card); and *drakcronat* has occasional crashes.

Conclusion

After almost one month of use on a production machine, we did not feel any need to get back to the previous MDK distro. With excellent hardware detection and configuration tools, Mandrake 9.1 is the best answer to those criticising Linux on this point. Traditional Mandrake tools like *supermount* and *urpmi* have grown to

full maturity. Younger features like the dynamic desktop and MandrakeGalaxy need still some work. The reworked *Control Center* and the new installer are other strong points of this distribution, though both (especially a few components of the former) need polishing. Overall, Mandrake is a reasonably solid distribution up to the job for many common tasks, as a Desktop as well as a Server OS. **LXF**

VERDICT

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

Perfection is an unreachable ideal, but Mandrake 9.1 would have got very close if the developers had paid more attention to some of the little details.

LINUX FORMAT RATING
9/10

LINUX DISTRIBUTION

Red Hat 9

The USA's most popular distro has jumped a whole number in just seven months. But is this leap, usually reserved for a really significant changes, justified? **Andy Channelle** gets on the case.

BUYER INFO

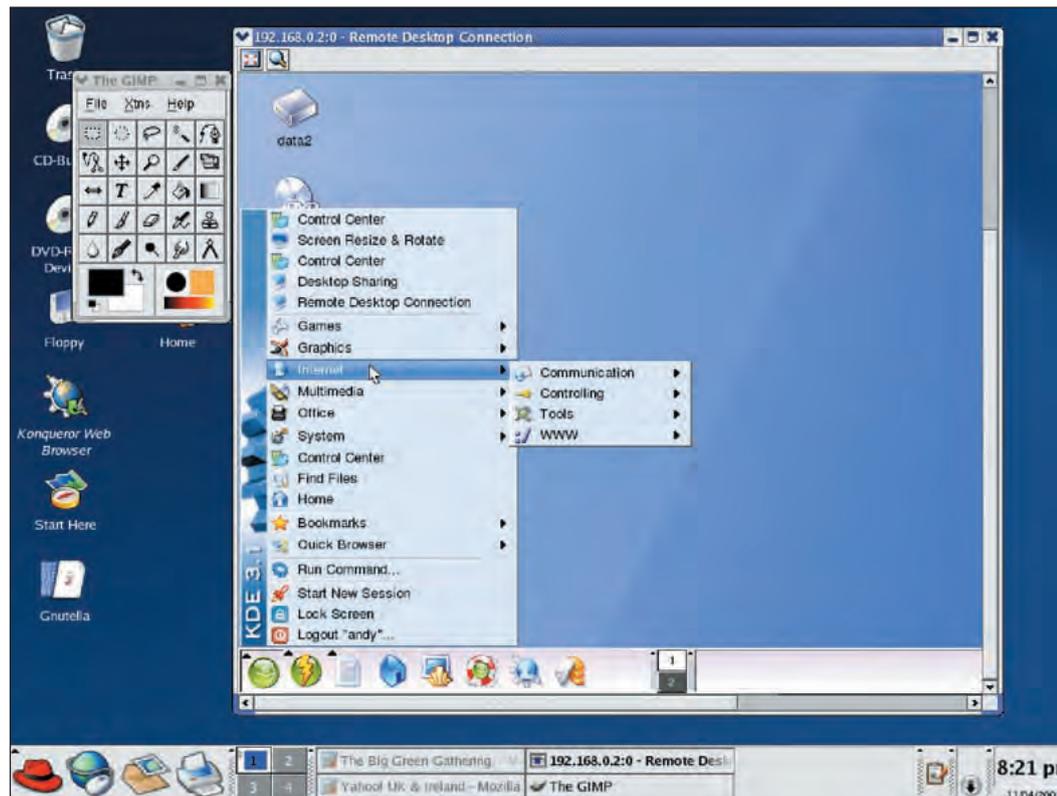
Competition: The other two 'big' distros, Mandrake and SuSE have both been recently updated.

- **DEVELOPER** Red Hat
- **PRICE** Personal: £35 inc VAT (£29 exc) Professional: £116 inc VAT (£99 exc) from www.dabs.com
- **WEB** www.redhat.com

Red Hat has, over the past few years, worked hard at becoming the byword for Linux in the corporate sector, dealing with the threat from the UnitedLinux initiative with aplomb and forging links with most of the major hardware vendors. The flipside to this conservative reputation is a degree of antipathy from the wider Linux community, where Red Hat, the company, is often derided as 'the Microsoft of the Linux world' and Red Hat, the distribution, castigated for (a) unifying the two main desktop environments and/or (b) breaking KDE – delete according to taste.

With such a mature product and the heroic efforts of Open Source device driver writers, you would expect installation on 'basic' hardware to be as simple as choosing a few options, hitting the 'Next' button a couple of times and changing the CD twice. And you would of course be right. *Anaconda* – Red Hat's installation routine – makes hard disk preparation and package selection very easy, and even if your needs extend beyond the basic, there is plenty of hand-holding available. It is not foolproof though, as we'll see later, especially when you stray from the desktop.

New to this release is the facility to upgrade a previous Red Hat install which should save time and reduce the potential for lost data if you are installing over an older version. In this case, however, we completely wiped the required partition and went for a full install. The first machine to be 'hatted' was my regular test machine, which in addition to Windows and



The remote desktop wizard would be useful in a networked environment, but only the client is configured by default.

Mandrake has a 2GB partition set aside for installing new distributions. Now 2GB isn't cavernous, but it has managed standard installations of all the main Linux version, with data held on a separate partition. *Anaconda*, however, thought I was being tight, despite claiming that the Personal Desktop version needed 2GB. After selecting the installation class – you have a choice of Personal Desktop, Workstation, Server and Custom – we move on to the detailed package selection screen which offers intuitive access to a lot of detail on what each installation class includes and, more importantly, the option to change it.

Enter the laptop

Notebook computers with their integrated components, reliance on external expansion and specific requirements are often seen as problematic for Linux users. But as

sales, and computing power, are constantly creeping up on their desktop cousins, it would seem just as important for the future to ensure that mobile users are well catered for.

There was a problem though: while Mandrakesoft is trumpeting the ability to resize the NTFS format partitions used by Windows XP/NT/2000, a Red Hat spokesman has gone on record as saying he would be wary of trusting anything but the MS implementation of the format at the moment. So in order to dual-boot with WinXP, I had to use *Partition Magic* to first resize the XP space and create the partitions needed for Red Hat. With that done, it was just a case of booting from the CD and kicking off *Anaconda*. The Toshiba's fixed resolution TFT screen caused the installation screens to be off-centre for some reason, meaning that while selecting options was easy enough, I had to guess where the 'Next' button

was. Not impossible, but a little annoying. Fortunately, once the second stage of installation (adding users etc) started everything was as it should be. Bluecurve here we come!

Bluecurve, Red Hat's controversial one-size-fits-all reworking of GNOME and KDE, is actually quite lovely, and comprehensively envelops the whole environment, with the whole thing being anti-aliased to within an inch of its life. By default I was dropped into GNOME, but apart from small details, for instance the format of the splash screens, it hardly makes a difference whether you're using GNOME or KDE. Icons, backgrounds, Window decoration and usability are all consistent. I really like it and the attention to detail means that, if you're short of space, you could easily drop one of the environments and free up some 300MB. And if Bluecurve doesn't stir you, both KDE and GNOME are easy to reconfigure to

their 'intended' states. These two desktop environments, by the way, are the latest stable versions and both do a good job of feeling instantly familiar and usable while still retaining the hidden depths that long-term Linux users will certainly appreciate.

Logical menus

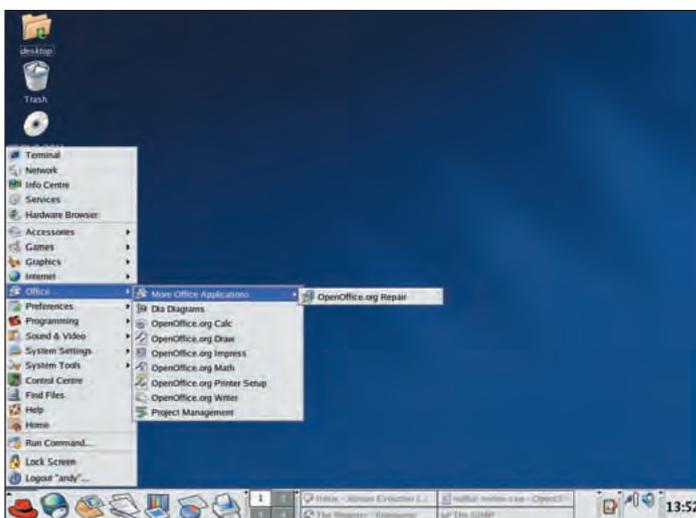
The menu system is logically set up into areas of use such as Programming, Games (of which there are many) and Office, though I do miss the Recent Documents folder that almost everyone else uses. The taskbar has short cuts to a range of default tools including *Mozilla 1.2.1*, *Evolution 1.2* and the main applications from *OpenOffice 1.0.2*. The initial desktop is very spartan, featuring links to removable storage devices, the home directory, Trash and a Start Here option which opens up some basic user options in *Konqueror*.

Surprisingly I've only had a few problems with the laptop hardware: Red Hat reports a removed Belkin PCMCIA network card on boot up, though it is usable once the machine is fully running; Power Management is non-existent; the glidepoint thing doesn't work and *Klaptop* maintains that the machine is plugged in regardless of whether it is or not. The most significant of these problems for mobile users is the lack of power status feedback: it makes for very paranoid usage. Fortunately my Toshiba notebook has a hardware power monitor, but this isn't a ubiquitous feature. One thing I have noticed, though, is that there is a lot less disk access and fan noise coming from the machine than when using Windows,

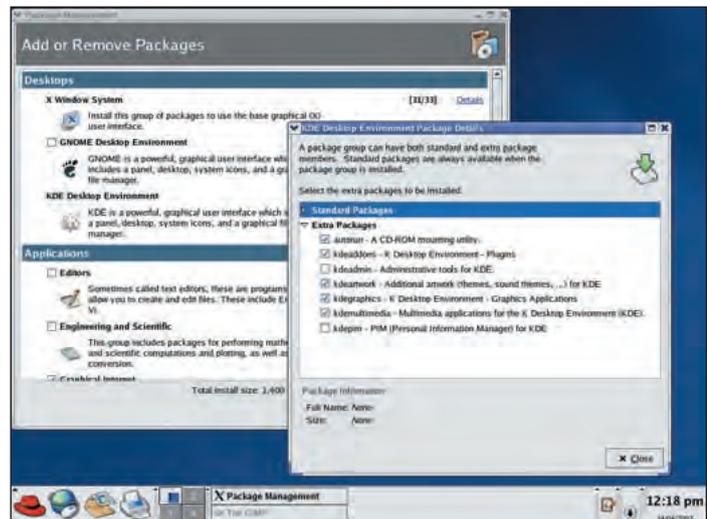
which should translate into longer battery life, though this may be mitigated by the lack of options when it comes to configuring Sleep and Suspend modes. So what is given with one hand is taken with the other!

There are a couple of big changes which potentially justify the big number change: this is the first Red Hat release where CUPS is the default printing system which makes, for instance, setting up network printing a breeze. If you were using *LPRng* before an upgrade, your previous choice will remain the default, but CUPS is so much more useful; *LPRng* is likely to be dropped altogether in future iterations. Version 9 also implements the Native POSIX Thread Library (NPTL) which, in theory, should offer a speed boost and better scalability to i686 processors and beyond. In practice the benefits are not as visible on the desktop as they would be in a server environment, applications like *Mozilla* and, more usefully, *OpenOffice.org* showed an improvement in load times, but this is traded off against incompatibility with applications such as older versions of Sun's and IBM's Java Runtime Environments. It is possible to run these using the old LinuxThreads implementation.

Fears over licensing issues again prevent Red Hat from shipping MP3 functionality. Ogg is a useful format but doesn't cut it when my portable MP3 player only plays MP3s. Fortunately the developers of XMMS have created a very small RPM which adds the appropriate plugin. Less understandable is the decision to not mount Windows partitions automatically. Rectifying this



The logical menu structure looks quite Windowsy and makes navigation intuitive, but a Recent Documents facility would be very welcome.



Package selection within Red Hat is a painless experience.

involves wandering into `/etc/fstab` and adding a few arcane lines to the file; this is fine if you know what's going on, less fine if you're used to distributions doing this, admittedly small, task for you. And if a Windows partition doesn't get mounted automatically, what are the chances of having something like a digital camera automounted? If the words 'absolutely none' spring to mind, give yourself a pat on the back.

Conclusion

With so much effort having been expended on making installation routines as idiot-proof as possible, distribution developers now have to find new ways to differentiate themselves. Though Red Hat 9 is aimed as much at home as well as corporate buyers, its natural habitat is business, where buyers will be attracted by the comprehensive Bluecurve makeover, which will make desktop training and administration more manageable; the range of support options and the backing of big hitters like IBM and Sun. Hotplug support, NTFS partition resizing, automounting Windows partitions etc. would make this a far more attractive prospect for casual users looking for a migration path or an alternative to a bundled OS. More critical is the new 'errata support' rules and Red Hat's decision to give these products a twelve-month shelf-life. In the area of security, this just seems stingy, and doesn't help foster the concept that Linux is the best route off of the incessant upgrade path that Microsoft inflicts on users.

While there is an installation class covering the main server tasks, Red Hat is obviously keen to push the new Advanced Server line of products.

However, for small time web, mail, print and file serving, this has all you will need; the addition of NPTL is one notable example of the way development efforts on the Advanced Server will continue to filter down to the basic products.

Red Hat 9 is a very good product; mature, stable, and extremely well-presented. However, with the likes of SuSE (see page 22) upping the ante to such a degree in the field of use and usability, it is unlikely that Red Hat will come to dominate the home user market in the way they do in the corporate sector at the moment. Let's hope the competition spurs them on to greater things. If the ethics of Anaconda and Bluecurve can be extended to the rest of the setup – home users need Flash, Real and MP3 as standard, not as an additional download which, for modem users at least, is an extra cost – then Red Hat 9 would be fantastic. That said, corporate buyers will find much here (errata support notwithstanding) to inspire. [LXF](#)

VERDICT

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

Easy to use and maintain, Red Hat 9 is a good choice for corporate users. Home users though may be better served looking at SuSE (especially if you're mobile) and Mandrake for the whole package or Xandros, Lycoris or Lindows for general usage.

LINUX FORMAT RATING
8/10

LINUX DISTRIBUTION

SuSE 8.2 Professional

With the broadest smile, Andy Channelle tests and falls in love with the latest version of SuSE LP. Can it really be this good? Oh yes!

**BUYER INFO**

Though its main competition is MDK and Red Hat, SuSE 8.2 has the potential to dominate the Windows/Xandros end of the market too.

- **DEVELOPER** SuSE
- **PRICE** SuSE 8.2 Professional Price: £55 inc VAT (Personal Edition: £25)
- **WEB** www.suse.com

The pro version of SuSE costs a shade under £55 for the boxed pack, falling somewhere in between the likes of Red Hat and Lycoris. If you were to venture beyond the Linux world, you could easily expect to pay ten times that and still come nowhere near the breadth and quality of this package. Value for money is nothing new with SuSE – just take a look at their manuals – but from installation to network set up and maintenance to daily use, SuSE Professional 8.2 has been a revelation. Such a gushing opening obviously demands some justification, so here goes.

There have been a number of featurettes written recently in the UK press and on various websites in which a knowledgeable Windows user attempts to install Linux before

grudgingly admitting that things were “okay... ish” or giving up and rejoining the Microsoft path. It is understandable, as the system is playing catch-up, that Linux has to elevate the art of installation far beyond the multiple reboot, device driver hunt that most people never get to experience with pre-installed Windows. And *YaST* (Yet another Setup Tool) is the Leonardo of installation systems. You could – and I did – simply rattle through, selecting the appropriate language and changing CDs (there are five) when requested, and 15 minutes later you have a perfectly useable system.

But of course there is the depth to deal with other eventualities. For instance, on a subsequent install I selected a preformatted FAT32 partition, and *YaST*'s suggestion involved creating the usual partitions for the base system (with a little extra headroom) and then shrinking the original FAT32 space for data, so that whatever was saved there would be accessible from Windows or Linux. This was then automounted as ‘D’ alongside the usual ‘C’. It is just a small

thing, but this attention to detail is repeated all through the installation and into usage of the system. Going even deeper the partitioning tool allows you to reorder things as you please and has a very usable GUI. Moreover this, along with every other element of the set up system, is available post install which makes light work of adding or reconfiguring hardware. The joy of it is that almost everything has a graphical front end which is the what the modern computer user has been reared on. Not everyone has caught on. Sharing folders across a network with Red Hat (see page 20) involved wading through host lists, *fstab* and a few other configuration files. In contrast, *YaST* needed a few details, the host name or IP address and shared directory – entered into the NFS client dialog and that was it. NFS server setup is equally uncomplicated.

And it gets even better. For instance, you can plug in a digital

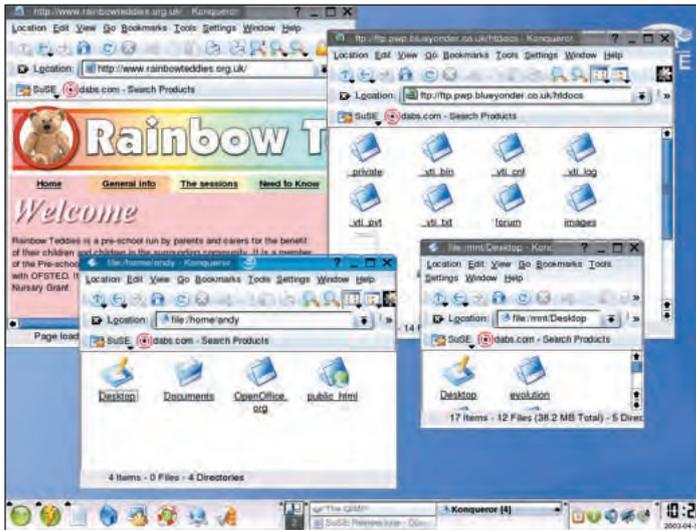
camera and, if it is designed to be used as a USB Mass Storage device, an icon labelled *sda* pops up on the desktop after a few seconds. Open it and you have access to your pictures, and as this is all happening in KDE, you can use *Konqueror* or even *PixiePlus* (LXF's Hottest Pick in issue 38, which isn't installed by default but should be) to browse, manage and even publish your image collection. This is just as it should be! No creating directories in */mnt*, no *fstab* editing and no need to fire up one of the ten shells (why so many?) on offer to type in long and winding commands. Call me simple, but I find this level of abstraction makes the job of actually maintaining a working system easier and a lot more fun.

Multimedia

For new Linux users the sort of things mentioned above will soon fade into the background, and usual statistics – kernel 2.4.20, XFree86 4.3, GCC 3.3 etc. – will be of little concern, it's what they see



SuSE provides a complete multimedia studio system as well as a great OS.



Konqueror handles local, network, FTP and web browsing. As well as the install, YaST handles post-install configuration and package management.

that matters. It's probably a psychological effect, but masking the 'verbose' boot sequence with a pretty picture and a progress bar makes the whole thing seem more human and less intimidating. The basic graphical system only installs KDE, though all the libraries needed for GNOME apps are present and correct and GNOME is available as an option. Curiously Konqueror has been chosen as the default web browser so Mozilla has to be chosen explicitly if you want to see it in the menu. This is fine, Konqueror is a useful browser with many of the best features of Mozilla (tabbed browsing, popup management etc.) and a smaller footprint; it also integrates nicely with the rest of the system and, for once, comes with Flash, RealPlayer and Adobe Acrobat reader ready to go. The other benefit of using Konqueror, I find, is the blurring between remote and local files, so updating a website can simply involve dragging and dropping from /home to ftp://path.com completely transparently. After struggling with Windows XP's awful 'My Network Places' system and Web Folders, this is magic in the Arthur C Clarke sense.

The default email application is KMail, though Evolution 1.2 is also installed to provide more extensive groupware functions. The reliance on KDE may annoy GNOME zealots but it offers uniformity in both form and function and improves real-world use.

After noticing the eye-candy and remarking on the elegance of the user interface, it's time to get to work. Again, SuSE has done a remarkable job in acquiring a remarkable range of applications, and nowhere is this more apparent than in the area of multimedia.

Beyond the brilliant but standard issue applications (GIMP, XMMS etc.) you'll also find high quality applications to deal with audio editing, MIDI sequencing, 3D design and rendering, DTP, HTML, image creation and video production. In fact, with the appropriate hardware, you have all the tools required for a complete production studio, and while the included applications don't yet offer the facilities of, say Adobe Premier or Cubase SX, it all comes in at about 10 per cent of just one 'professional' packages.

MainActor

What really sets the Professional edition apart from the Personal version (which is nearly half the price) and every other distribution on the market is the addition of MainActor, a professional production suite covering the capture, editing and sequencing of broadcast quality video. While this (version 3.70) is not the most

recent release of the software (v5 for Linux is imminent) it does have all the tools, transitions, effects, titling and output options you would expect of a product aimed at broadcasters. Due to having an unsupported USB capture device, I was unable to test the capture element of the system, but messing with a few pieces of stock video gave me the opportunity to play with the editing and sequencing tools. The application is very similar in use to Premier; and blending between cuts was achieved using the same two track timeline with a transition strip in the middle. The transitions range from the basic to more sophisticated options, and the final output was very smooth indeed.

MainActor can also handle basic audio editing, but if you get beyond its capabilities you can always export the audio to Audacity, a very good Open Source editor, for polishing. If your video production demands a full soundtrack, SuSE 8.2 is the best way yet of getting hold of Rosegarden4, one of the best audio/midi sequencers currently available for Linux. Particularly ambitious directors could also wade through the interface of the recently open-sourced Blender to add 3D animations or titles to your project.

For those with a print bias, the limited desktop publishing tools from the word processor in OpenOffice.org (which is obviously included) can be augmented with the likes of Sodipodi for vector/SVG creation, Scribus for layout and the ubiquitous GIMP for still image editing/creation. And just out of interest I attempted, via the installed CodeWeavers WINE implementation, to launch Adobe Photoshop 5.0 from the Windows partition – and it worked. Not

only that, but all the filters and tools worked too!

It's not all perfect

In conclusion I would probably say I liked this, a lot, and there is very little to detract from this relentless positivity. The ACPI power management only half works on my notebook, meaning the screen will power down after a set time, but there is no battery monitor, no warning and no power-dependent screen brightness control. However, battery life has been extended over Windows XP by, I'd estimate, about 30 percent. There's less disk access going on, and the only time the CPU fan has been called for was when using MainActor. The Thesaurus doesn't seem to work in OpenOffice.org, I've had occasional trouble accessing the integrated audio device in Audacity,

I'm not keen on the SuSE window decoration and, for some reason, the EuphoriaGL screensaver is not included. But really I'm struggling to find negative things to point out, even the silly glide point and its attendant buttons work as you would expect, with the third button simulated by hitting both together.

Doubtless over the coming months both LXF staff and our readers will uncover some annoying bugs in the course of their daily use, but as of today this is the best Linux OS package I've used. It's stable, logical and above all, comprehensive. The selection of apps, robust WINE implementation and the dynamic USB device support are the cherries on a very well-made cake. And there has not even been space here to give the the online update client, the superb documentation, font management and any of the hundred other things that will make you smile the mention that they truly deserve. And if you have no need for video editing, the Personal edition, if it follows the style of its big brother, looks like the bargain of the year at £25.00. **LXF**



Blimey, we weren't expecting that to work. It's Adobe Photoshop, folks!

VERDICT

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

Raises the bar for kitchen sink Linux distros. OEMs looking to preinstall an alternative to Windows, this would be it.

LINUX FORMAT RATING
9/10

Reviews Comparing Distributions

LATEST INFO

DistroWatch

How do the latest Linux OS distributions compare?

With a reappearance of our Distro Watch section, last sighted over a year ago, you'll find a comparison of the software and features of the distros tested this issue. Obviously there are far too many packages to include complete lists, so we have merely highlighted the more important software here.

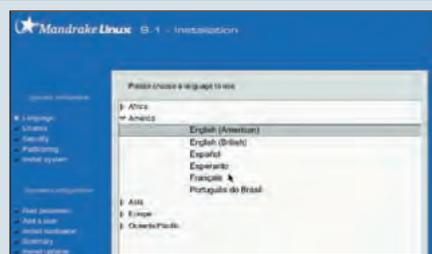
Note that the higher numbers are not always the best choice. For example, Mandrake has a habit of including tested but still officially beta packages – you get slightly more up-

to-date versions, but potentially with some problems. The other distros tend to play safe, although, as you will also see from the table, some packages tend to lag behind the others for unexplained reasons.

The final point of interest is new or optional software. *Apache* is a good case in point. *Apache* 2 has been around for a while now, but the 1.3.x series is in more widespread use, because of fundamental changes to the code. Most of the new distros offer a choice in these cases, though they may choose different options to be the default. **LXF**

By coincidence, the name we gave this section when it first appeared is the same as a website which performs a similar function. With frequent updates and much more comprehensive lists of software, it is a great resource if you want to know what's making a distro tick. Check it out at www.distrowatch.com

Mandrake 9.1



Price	£53 (powerpack), £159 (Prosuite)
Website	www.mandrake.com
Architecture	586 (PPC planned)
Default Desktop	KDE
Download version	Yes, ISO/FTP

Kernel Version	2.4.21pre
-----------------------	-----------

Linux Filesystems	ext2/ext3/JFS/XFS/ResierFS
--------------------------	----------------------------

SOFTWARE

Package format	RPM
KDE	3.1
GNOME Desktop	2.2.0.1
OpenOffice.org	1.0.2
Samba	2.2.7a
PHP	4.3.0
Apache	1.3.27/2.0.44
Mozilla	1.3
Sendmail	8.12.8
MySQL	4.0.12
glibc	2.3.2
gcc	3.2.2
Xfree86	4.3.0

SERVICES

Support	Installation (web) 30/60/90 day phone 5 incidents (Prosuite only)
Updates	URPMI

Red Hat 9



Price	£35 (Personal), £116 (Professional)
Website	www.redhat.com
Architecture	586
Default Desktop	GNOME
Download version	Yes, ISO images

Kernel Version	2.4.20
-----------------------	--------

Linux Filesystems	ext2/ext3
--------------------------	-----------

SOFTWARE

Package format	RPM
KDE	3.1
GNOME Desktop	2.2.0.1
OpenOffice.org	1.0.2
Samba	2.2.8a
PHP	4.2.2
Apache	2.0.40
Mozilla	1.2.1
Sendmail	8.12.9
MySQL	3.23.54
glibc	2.3.2
gcc	3.2.2
Xfree86	4.3.0

SERVICES

Support	Installation (email) 30/60 days Further Support options
Updates	RedHatNetwork (60days sub)

SuSE 8.2



Price	£38 (Personal), £59(Professional)
Website	www.suse.co.uk
Architecture	586/others to follow
Default Desktop	KDE
Download version	None as yet

Kernel Version	2.4.20
-----------------------	--------

Linux Filesystems	ext2/ext3/JFS/XFS/ResierFS
--------------------------	----------------------------

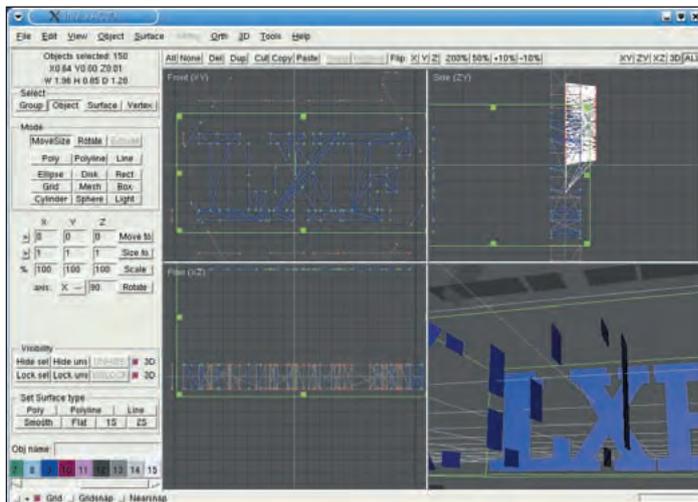
SOFTWARE

Package format	RPM
KDE	3.1.1
GNOME Desktop	2.2.0.1
OpenOffice.org	1.0.2
Samba	2.2.8a
PHP	4.3.1
Apache	1.3.27
Mozilla	1.2.1
Sendmail	8.12.7
MySQL	4.0.12
glibc	2.3.2
gcc	3.3pre
Xfree86	4.3.0

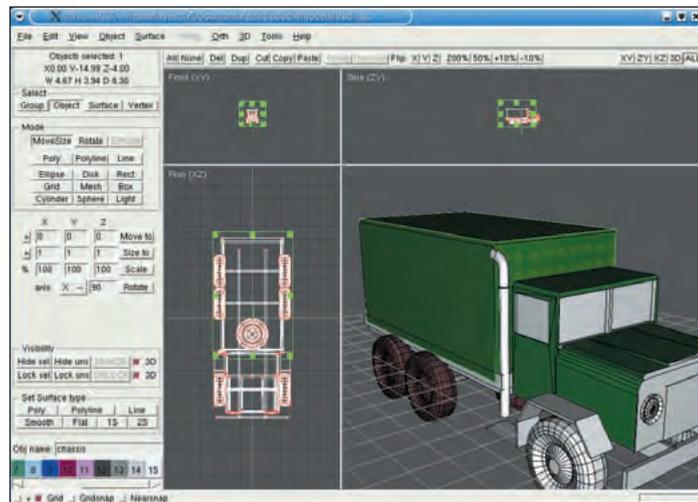
SERVICES

Support	Installation (email/phone) 60/90 days Further pay support available
Updates	Via website

ReviewsAC3D



Plugins include TrueType font support, as well as various 3D file formats



Simple texturing is handled with image maps or flat colours.

3D MODELLER AC3D

Famed for being a chain-smoking stick insect with attitude already, **Nick Veitch** considers a career in modelling with this long-standing 3D tool.



AC3D is a pretty good modeller, but lacks some of the rich features that would make it truly great. It does excel at low-polygon count models, such as you might require in game design or visualisation, but isn't going to compete with more feature-rich modellers commercially available for the professional user. There are a few User interface issues, mostly to do with the requestors, which could be a little more intelligent and/or helpful. Overall though, this is a fine modeller. Currently, *AC3D* is actually a shareware product, not Open Source, but price-wise it does compare favourably with shareware modellers that are available on other platforms. **LXF**

BUYER INFO

3D modelling software. Also consider the Open Source *Blender*.

- **DEVELOPER** Invis
- **PRICE** \$49.95
- **WEB** www.ac3d.org

Linux has made great inroads into the professional 3D industry, with plenty of software now ported to the OS, and loads of film credits to its name. But while the big commercial tools support Linux, there is still some way to go for Open Source modelling and rendering tools. *Blender* is powerful, but a little difficult to get to grips with (unless you follow the *LXF* tutorial series, of course!).

AC3D has been around since the late 1990s, so it is now a very stable, solid and reliable tool which has found a great deal of satisfied users in the Linux community and on Windows, for which the software is also developed.

Features

Primitives are simple shapes that are used to form the basis of most. There are nine primitives supported, plus three line drawing modes to help you build your objects. This is a decent selection and should make it easy to build basic objects fairly quickly.

Scaling and transformation of objects is easy enough, and the distinction between vertices, surfaces and objects is a great help when dealing with complicated transforms or operations.

On the downside, there really are only the basic tools for modelling. You can revolve objects, group, split, fragment and even use a limited beveling function, amongst other things, but there are no, what you might like to call 'higher level' functions, such as boolean operations, fractal subdivision or fancy deformers.

There are useful and easy to use tools for surfacing, including a quick colour palette, and texture-mapping tools. Surfaces can easily be switched between types (single or double sided, smooth or flat rendering etc), and it isn't difficult to see why *AC3D* is a popular tool for building simple 3D models for games and flight sims.

On the file format side, there is plenty of support for model formats. This includes the ones you might expect (3D Studio 3ds files, Lightwave, the ubiquitous DXF, wavefront .obj) and some which are lesser-used but very welcome, such as MilkShape files.

There is a similar level of support for output file formats, which also includes various VRML options.

Most of these filesystems, and a few other features, are supported through the plugin system. This makes it easier to add functionality, particularly file format support for texture images or specific rendering platforms, and most of the plugins have been created by an active and knowledgeable community.

Surfaces and objects can be grouped, named and even ordered in a hierarchy. This makes sense when there is the need to support other file formats for modellers which can use this data, but is also useful for modelling itself, as it is also possible to select objects by name.

Rendering

Apart from the 3D preview window, *AC3D* won't actually do full renders of your objects, so you'll need to rely on external renderers for that. There is a built-in to render directly to *POV-Ray* should you have it installed. This is fairly unintelligent though, and we had difficulties with texture image files, which worked okay in *AC3D*, but failed to be configured properly for *POV-Ray*.

Of course, you can export the models for use in any number of different rendering engines thanks to the great breadth of file format support that *AC3D* provides.

Versions

Register for full functionality

AC3D is shareware, but can be run without registering. The main limitation is that the unregistered version doesn't support the plugins – which are pretty much essential for importing and exporting model files to and from other packages. The unregistered version is on this month's coverdisc.

VERDICT

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

Competent and easy to use – could do with some more editing tools though.

LINUX FORMAT RATING
7/10

FIREWALL DISTRO

Astaro Security Linux V4

Ever the belt-and-braces man, Richard Drummond secures his setup with ASL.

BUYER INFO

Competing firewall products include *Smoothwall Corporate Server* and *Securepoint Firewall Server*.

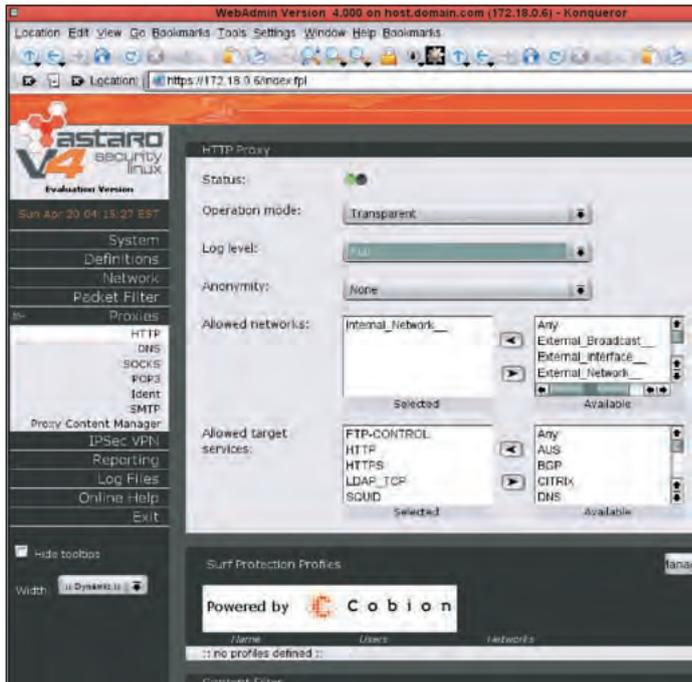
- **DEVELOPER** Astaro
- **PRICE** See box 'Pricing' opposite
- **WEB** www.astaro.com
- **PHONE** +49 (0)721-49 00 69 0

How do you protect your local network from the risks of being connected to the Internet? Many small- and medium-sized business are finding that the most cost-effective and flexible solution to this problem is to employ a software firewall – a package which when installed on a regular PC turns it into a fully-functional firewall. This approach is often cheaper, easier to upgrade than a dedicated hardware device, and offers better scalability.

One of the many such software firewalls on the market is *Astaro Security Linux (ASL)*, a Linux-based integrated solution offering stateful inspection, routing functionality with up to 20 network zones, NAT and VPN, and a range of application proxies. It combines ease-of-use with powerful optional extra features such as load-balancing, high-availability, bandwidth management and virus- and content-filtering. The ASL install CD is available as a free download from various mirror sites. It is free for home use, and per-client commercial licenses can be purchased at a range of price points depending on your network's size and the features you need.

ASL is based on a security-hardened Red Hat 7.3 core – with a tweaked 2.4.19 kernel and intrusion-detection. It employs a familiar set of open-source applications, such as *Apache*, *BIND*, *Exim* and *Squid* – and most of these services run in a chroot jail for added security. Astaro adds to this base its own middleware and a web-based management interface called *WebAdmin* to create an easy-to-use, reliable and integrated solution.

The new V4 release brings a host of new features including better hardware support, a faster web-based



Besides firewalling and routing, ASL includes a range of application proxies such as the web cache and content filter shown here.

interface (by employing *FastCGI* and HTTP stream compression), heuristic spam filtering, a transparent POP3 proxy, LDAP user authentication (in addition to the existing internal database and RADIUS and SAM authentication), and more. An interesting new addition is support for Prism2-based PCMCIA wireless network adapters. This means that your ASL firewall can work as a wireless access point for your network, obviating the need to buy an expensive hardware device.

Firewall management

One of ASL's strengths is its ease of use. Installation is also simple, and the only significant pre-install configuration required is to supply details for your internal network interface. ASL takes over your entire hard drive, creating a series of *ext3* partitions for the installation. IDE and SCSI disks are supported, as well as a variety of RAID controllers. Software RAID isn't supported however. (For more on RAID, see this month's *Linux Pro*.)

The rest of the configuration is done by booting your firewall, and accessing *WebAdmin* remotely with a browser. The *WebAdmin* interface is well-designed and navigation simple thanks to the logical layout. The initial learning may be steep due to the complexity of the system, but there's plenty of context-sensitive and online help available via Javascript pop-ups. An excellent 220-page manual can also be downloaded from <http://docs.astaro.org>.

ASL offers ease-of-maintenance too. Its *Up2Date* service can manually or automatically check for software updates and download them for you. You can then read the change log for each update and decide whether you want to apply it. Backups of your configuration can be easily made and restored via *WebAdmin*, and it can automatically email back-ups at a chosen interval. New in V4 is support for encrypted back-ups. ASL is well-equipped with monitoring features too. It can email the administrator with various alerts, such as a detected intrusion. Extensive controls are also provided for

Pricing

You get what you pay for

Prices start at \$390 for a 10-IP licence including 10 VPN tunnels, 3 network zones, and 5 mail domains, 90-days of web-based support and a 1 year *Up2Date* subscription. A 100-IP licence with 200 VPN tunnels, 7 zones, 25 mail domains and load-balancing is \$1495. Virus protection starts at an extra \$295 per 10-IPs per year, and surf protection (content-filtering) starts at \$275.

configuring and viewing various system logs, and remote logging is supported to an external syslog daemon.

While there are always more features you could ask for – for example, support for software RAID, UPS support, and perhaps a wizard to speed up initial firewall setup – the reality is that ASL provides everything that most users will need. Advanced features such as load-balancing, high-availability and bandwidth management set it apart from its rivals and enables ASL to serve large and sophisticated networks. If you compare ASL's pricing with some of its competitors, you will realise that in addition to all this, you get remarkable value for money. Its closest rival is the less-expensive *Smoothwall*, and, while ASL is more time-consuming to set up and lacks *Smoothwall*'s support for dial-up networking, it does offer more features as standard and has better monitoring support – though a new version of *Smoothwall* is due soon... [LXF](#)

VERDICT

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

V4 confirms ASL's position as one of the most full-featured and cost-effective firewall solutions around.

LINUX FORMAT RATING

JAVA CODE GENERATOR

JCodeBox

Work with J2EE and looking for a way to avoid having to redo the same old code each time you start off a new project? Paul Hudson reviews a new possibility: automatic application creation.

BUYER INFO

Java-based J2EE code generator. Also consider Borland's *JBuilder* or Sun's *Forté*.

- **DEVELOPER** Michael Campbell Associates
- **PRICE** £176 (inc VAT)
- **WEB** www.jcodebox.com

Writing database J2EE code has never been particularly easy to get started with, as there's usually a lot of housework required to get started – handling database connections, creating the JSP web pages, setting up a Java-enabled web server, and so on. Sometimes it's possible to alleviate these pains by reusing bits and pieces of your code, but that can often end up being more hassle than it's worth!

But now, almost out of thin air a solution has arrived in the form of *JCodeBox*. Confidently touting itself as “an affordable Java application and component builder that dramatically accelerates the development of J2EE”, you would be forgiven for thinking the product is getting a little ahead of itself.

However, you would be wrong – *JCodeBox* does precisely that, and it does it in such a way that you almost feel like you've *cheated* because it

The default end product. From start to finish this took about one minute to create, and is fully functional.

takes so little time to create a fully working product. Here at *Linux Format* we love Java, and this product looks almost able to end our Java-writing career opportunities simply because you *can* point-and-click to write your Java application.

When the program starts up, you're presented with simple two-pane view. The left pane contains all the SQL drivers you have available to you, including Oracle, MySQL, PostgreSQL,

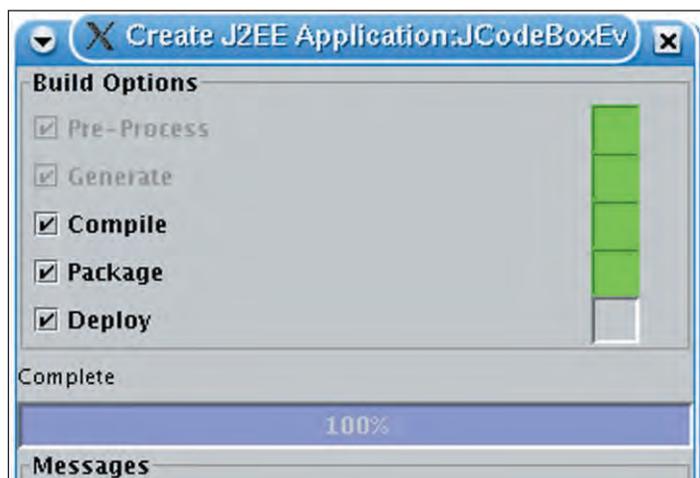
and others. From there, you can define connections to SQL databases of each type, for example I defined a connection to the *LXF* website's MySQL server. With me so far? Good – here's where it gets its most complicated. Once you've connected to a server, you right-click on the database you're interested in, and select 'Generate from Database'. That's the connection stuff all done away with – all too easy.

Does it all for you

To continue, you select from the GUI interface what table you'd like to be used in your application, how you would like it to operate, and so on – *JCodeBox* visually breaks things down into the EIS Tier, the Business Tier, the Web Tier, and so on. Once you're finished defining the application (and really this requires little more skill than using the mouse), you simply click 'Create'. Here you get one final dialog box that allows you to select whether you want to compile the application, package it, and/or deploy it – if you select all three, *JCodeBox* will do everything for you, and leave you with a working web application.

Is it really that easy? Sure it is – and there's a lot more it can do besides. However, it's sadly not all a rosy story – the application, for us at least, was remarkably hard to get working correctly. Firstly, you need a very specific version of *libc* installed for Java to work, although that is a Java problem not a *JCodeBox* problem to be fair. Secondly, you seem to need a very specific version of the driver for your database installed if you're trying to use MySQL or PostgreSQL. Thirdly, the product doesn't come in a box, which means no printed manuals – a big mistake, in my book. Without any sort of 'Quick Start' or 'Getting Started' guide to get me going, I admit I was pretty much clicking around randomly to begin with – not a good sign.

However, once you're past these problems – if you *manage* to get past these problems – the product is a major time saver. Furthermore, it's available for an incredible price – no company would balk at the cost of this thing. If you're using it with a commercial database, the driver problems are non-existent, and if you don't mind using the online documentation you'll be fine there too. If the developers release a new version that fixes these problems (hint: friendly error messages go a long way), then they'll have a truly hot product on their hands. In the meantime, *JCodeBox* is merely smouldering – a great product, but not *perfect*. Yet. [LXF](#)



Compiling your application is, uh, *incredibly* difficult and fraught with problems. Damn it, I'm out of a job as a Java programmer!

VERDICT

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

The way forward for J2EE development. If they offer a cheap upgrade to the new version when it's released, buy it now.

LINUX FORMAT RATING
 9/10

DESKTOP HARDWARE

OpenAuth

Linux is renowned for being 'as safe as houses' so why change something that works fine? Paul Hudson takes a look at a new offering to desktop users looking to migrate to Linux.

BUYER INFO

Debian-based, security-enhanced Linux PC pre-configured for ease of use. Rebrands software already available under the GPL such as *snort*. See also Evesham PCs.

- **DEVELOPER** OpenAuth
- **PRICE** £914+VAT
- **WEB** www.openauth.co.uk



Modern Linux distributions often work under the premise that Linux should be given away freely, with the value-added element that justifies a price tag being a support contract. OpenAuth, however, are seeking to sell Linux with a new twist on the idea, by providing Linux boxes that come pre-configured for maximum security.

The units themselves are fairly innocuous-looking micro towers that internally come with some pretty hefty hardware – our review box, for example, came with a 2.4GHz P4 inside combined with 256MB of RAM. The monitor, a 15" Relisys flat-panel, is of very good build, and as it has quite a wide optimum viewing angle, can be viewed from a large range of angles while remaining easy to read. One slight niggle with the monitor as supplied was that the top line of the terminal was off the top of the screen, however this was easy to correct once spotted – although I wonder how many new users would have realised that there was a problem.

The keyboard and mouse supplied are on the cheap side, but not too bad. On the back of the machine, there are two USB slots available, which is probably adequate for most people – no Firewire, though. Hidden away in the front is a CD-RW/DVD combo drive, which makes for a nice extra at this price. The machine runs very quietly, which fits in well with its low-profile chassis – this kind of

machine can sit under a desk and pretty much not be noticed.

On the software front, the machine came pre-configured with a customised version of Debian, and booted by default into a well-configured X-Windows GNOME desktop. The usual favourite software products are there as standard, including *OpenOffice.org*, *Evolution*, *Gnumeric*, and "*The GIMP* (unstable)" (*sic*). It's a shame, though, that there is little choice of window manager by default – certainly no KDE in sight, although I have been told this is corrected in other machines. The general setup of GNOME is a little messy, though – *OOo* is under 'Other' rather than 'Office', there's a KDE menu at the bottom, despite KDE not being available, and there's even a set of working *StarOffice* 6 icons under the defunct KDE menu that aren't available in the main GNOME menu itself!

The system performed adequately for a machine of its class, although we've seen better. We were unable to compile the MySQL *Super-Smack* tests due to a poorly configured GCC setup. Overall, the hardware is more than good enough for business machines, and the low-profile style is very easy on the eye.

One particular high point in the system is the ability to allow authorised external users to remotely access and

configure your system in order to solve otherwise fatal problems, and we got this working just fine. This is further enhanced by the fact that the system comes partitioned so that a backup copy of your system is kept separate to the live system so that in the even of any problems the system can be restored easily. All user data is, by default, stored on its own partition that's made available to both boot images. Combined together, this technology means that external support people can log-in and repair operating system-level problems without hassle.

Help meeee!

With regards to documentation and support, the system is a little flimsy – two sides of A4 and a telephone number. However, there is an email address you can write to, which presumably will reply with some help. Nevertheless, that's hardly sufficient for users at any level, and we hope this is rectified soon.

The system can be obtained either through an outright purchase or on a leased basis, which allows you to spread the cost of your systems over regular payments and is therefore ideal for companies. The prices themselves are actually quite good, especially so considering the pre-configuration that is done on the boxes – our review

model is available for just £914 + VAT, or £47.60 on a 24-month contract.

Overall, the system is rather a good commercial idea that is seeing its first serious outing in the hardware marketplace. Hopefully many of the problems can be addressed to better streamline the user experience. As it stands, you get a fair amount for your money, and the leasing plan will make this system an attractive option for corporate users. **LXF**

BENCHMARKS

hd	1.93
apache	1.22
mysql	FAILED
compile	1.16
oggenc	1.74
Overall	1.51

VERDICT

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

A good idea, let down by some glitches across the system. More could be done to hone this, which would make for a killer setup. We're looking forward to seeing the next release from OpenAuth.

LINUX FORMAT RATING
8/10

**1U SERVER****NEC Express 5800**

NEC is renowned for big irons, but what do you get for your money, and does Linux take full advantage of the hardware? **Paul Hudson** has a butchers...

BUYER INFO

Powerful, mid-range 1U server with a lot of muscle. Also consider the less expensive SR113 (reviewed LXF 40)

NEC EXPRESS5800/120RC-1 1U RACKMOUNT SERVER

- **DEVELOPER** NEC
- **PRICE** £1761 (inc. VAT)
- **WEB** www.nec-online.co.uk

How much horsepower can you shoehorn into a 1U server? 1U of space is tiny, if you didn't know – about 19 inches wide and 2 inches high is all the space available to cram in as much as possible. With enterprise customers demanding more and more from their machines, the battle to produce the highest spec

1U and 2U system is ongoing, and price isn't always the biggest consideration. To compete in the medium-range to high-end server class (departmental servers and above), machines need multiple CPUs, fault-tolerant RAID support, and RAM to spare.

It's into this market that the Express5800/120Rc steps in. Our machine came with 2 Xeon 1.8GHz CPUs inside, 512MB of RAM, two NICs built-in, and an 18GB hard drive. As you can see, this machine is very much processor-heavy – that is, with so little RAM and hard drive space, the computer isn't designed to handle huge databases or act as a file server. With the supplied hard drive being connected to a 160MB/s Ultra-SCSI adapter, this box is capable of shifting files around

exceptionally quickly – perhaps the ideal role this machine is as a web server that handles threads well to make proper use of the dual-CPU layout.

Being Intel Xeon processors, both the chips inside were hyperthreading ready, and the BIOS came preconfigured to enable this functionality. If you don't know, hyperthreading works by creating virtual CPUs able to be given independent tasks – for example, this machine reports four CPUs when you run **cat /proc/cpuinfo**. Modern x86 CPUs use instruction scheduling to non-destructively re-order commands in order to maximise performance. Hyperthreading, by creating virtual CPUs, allows CPUs to increase execution throughput by juggling threads as part of

instruction scheduling. In the perfect scenario, one physical CPU could do most of the work of two similarly clocked CPUs. The usual effect, however, is generally only about a 20% speed improvement, and then only optimistically – mostly because perfect scenarios don't exist, and also because each virtual CPU has to share physical on-die cache which limits the possible performance improvement.

What docs?

So, you've laid down your cash for the machine, the box arrives (and, my – is it a big box!), and you rub your hands in glee as you imagine all the incredible things this machine will allow you to do at long last. Peering inside the box, you find the machine itself (not the most attractive design in the world, but we've seen worse), some wires, some rails to allow you to affix the unit to your rack, a slip of A3 paper, and a 120 page book marked 'NEC: "Excellent," you think, "at last a company that provides decent printed documentation with their hardware". Sadly, this isn't the case – the book, smartly printed and bound, uses its pages to tell you under what conditions NEC will be prepared to repair your machine, and it does so in twelve languages – presumably to make sure the points are quite clear.



Pathetically enough, it's the sheet of paper that's the printed 'documentation' for this server. Even the low-cost SR113 reviewed last issue had a short booklet included so you could handle basic troubleshooting issues like "It doesn't power on" and such, and that machine was nowhere near the price of this one. There is some documentation included on the CD, but I doubt my wife would be too happy if I dragged my monitor into the toilet with me to read at my leisure!

As long as you don't mind the slight inconvenience, the documentation provided online is rather good, covering what each piece of the hardware does, how to change parts, how to perform basic maintenance tasks, and so on. Given that SuSE give away over 900 pages of printed documentation with their £60 SuSE Linux, it astounds me that NEC supply next to nothing with this hardware given the price.

Connectivity

The unit comes with four USB slots in total – two at the front and two at the back – which should be more than enough for anybody. Also as standard is the low-profile CD ROM and floppy drive. One particularly important feature is the inclusion of two internal gigabit network cards, that work at 1000/100/10Base-T – this is the Rolls-Royce of networking right now, and provides enough connectivity for anyone's needs. One slight problem was

that I was unable to get the network cards working under Debian, and had to install Red Hat instead. Given that there is a wide availability of Linux-compatible hardware available for servers, this isn't terribly good – many people like to have the choice of distribution left to them, myself included.

One neat feature is the inclusion of two RJ45-type serial ports to allow terminal output – one on the front and one on the back. The combination of the USB ports and the serial ports being on both sides of the machine make this unit quite flexible in how you use it – NEC have shown quite a bit of tact with the design.

Go faster stripes

With such meaty hardware pushing this monster, it's not surprising it did so well in the tests. Unsurprisingly, it exceeds on the tests where multiple CPUs can be used – both the Apache and MySQL Super-smack results are excellent, which makes this machine just perfect as an intranet/Internet server. Also of note is the hard disk result, which is also incredibly high – clearly NEC have equipped this machine with top of the range equipment internally, because the throughput seen here is incredible. Performing at almost 4 times the speed of our test unit, the Express5800 would do very well for a file server – if you pay the extra for more hard disk space.

The server *didn't* perform too well on kernel compilation, returning a score only(!) 1.45 times faster than our test machine. However, it's important

to remember that compilation runs on only one of the two CPUs, and so the poor performance on this front can perhaps be forgiven.

Overall, the machine returns a score of 2.78, with particular highs on *Apache*, hard disk, and MySQL performance. If you were to use higher-spec CPUs, I imagine that the *Apache* results will increase fairly linearly. Perhaps the best addition to make would be to double or even quadruple the amount of RAM, and do the same for the hard drive space – that would make this server sit happily in any environment.

Usage

This is a machine built to be rack-mounted, and NEC have taken no pains to make it more attractive for non rack-mounted situations. It's noisy and heavy, but it does the job. Moreover, the noise can be forgiven on the grounds that the fans are obviously pulling their weight – the server is cool to the touch even while taking a hammering from our benchmarks. In nearly every environment, operating temperature is far more crucial than noise, particularly as most server racks are housed in isolated rack cabinets, out of the way.

This system comes with such an unusual configuration that you're almost certainly going to want to order your own options. With such speedy connectivity (when it works), this is definitely a machine that is asking to be given some hard network tasks – the ideal role for this machine is as part of a front-end/back-end web server/database-server system, where there's a solid storage device holding a large database, with a speedy NEC web server able to make most use of its multiple processors (virtual or otherwise!).

This NEC Express5800/120Rc-1 is certified to work with Red Hat Linux, and as such Red Hat will work perfectly on the machine first time. This is unsurprising, as Red Hat have gone to a great deal of trouble to get support from hardware vendors for its products. However, it *doesn't* mean that the server will work smoothly with other Linux distributions – as mentioned, Debian encounters major problems trying to get a networking driver working. While SuSE Enterprise Linux does have some support for the server range the 120Rc is in, it isn't specifically certified for the hardware and may encounter problems.

NEC's UltraCare warranty comes bundled with the server, which means buyers are guaranteed three years of on-site technical support and maintenance as part of the agreement. With so many companies cutting back warranties to the minimum accepted level, it's good to see NEC are sticking with more traditional service levels.

Final Thoughts

This a durable, reliable, and powerful machine that can take quite a beating when it comes to intensive processing. With two reasonably good CPUs, both capable of using Intel's hyperthreading technology on fairly 2.4-based Linux distributions, this is certainly a great server to use if you're willing to pay the going rate. While it's definitely not in the 'budget' class, you get a lot of performance for your money – I don't think anyone will be disappointed with the performance figures of this machine. Furthermore, it all comes packaged up in a neat 1U rack configuration that belies the performance of the hardware.

To conclude, the Express5800 isn't going to revolutionise the server arena – it offers no surprising changes or evolutionary designs when compared to other products at the same level. However, it does offer slightly above average performance for a machine in its class, backed up by a warranty suitable for SOHO users and enterprise users alike. Furthermore, while certainly not the cheapest available, you won't find all that many alternatives at a better price for this kind of kit. **LXF**

BENCHMARKS

hd	3.67
apache	3.67
mysql	2.58
compile	1.45
oggenc	2.54
Overall	2.78

VERDICT

Features	9/10
Performance	9/10
Ease of use	6/10
Value for money	8/10

Solid server with off-beat configuration and no printed documentation – a good performer for the money, though.

LINUX FORMAT RATING

8/10

MySQL, 2nd Edition

Paul Hudson looks at the second edition of the authoritative MySQL handbook.

BUYER INFO

- **PUBLISHER** Sams
- **AUTHOR** Paul DuBois
- **ISBN** 0-7357-1212-3
- **PRICE** £36.50



MySQL is acclaimed as being the world's most popular open-source database management system, and, co-incidentally, *MySQL* is also the name of the book that is acclaimed as being the world's most popular open-source database system! If you haven't heard of the author, Paul DuBois, then you're missing out – he's a true MySQL guru who regularly answers questions of all types on the mailing list. Not only is he friendly and informative, but he also actively cares about the MySQL community – so, who better to write the definitive MySQL book?



If you bought the first edition and you're thinking you can probably skip this new release, you're almost certainly wrong – MySQL 2nd Edition. is over four hundred pages longer than its predecessor, and is entirely revised to focus primarily on MySQL 4. This is not just a minor revision by *anyone's* standards.

If you didn't own the first edition and use MySQL regularly, what on Earth are you waiting for? MySQL 2nd Ed is packed with information for both newcomers and experienced DBAs.

Looking for a basic MySQL tutorial? You'll find it here. Want to know how MySQL optimises your queries? Just turn to chapter four. Want to know about disaster recovery, the MySQL C API, or MySQL security? Check, check, and check. Authoritative is Paul DuBois' middle name, it seems.

As well as being an excellent teaching book, there's also about 350 pages near the back just for reference alone, and covers the perl, C, and PHP interfaces extensively, as well as also providing general discussion of SQL

syntax. Yes, the book is very heavy, but you certainly get a lot for your money – £36.50 (\$49.99) is a bargain price for a tome such as this.

Furthermore, the book isn't just about MySQL 4 – DuBois obviously had a very early release of MySQL 4.1 available to him because you'll find discussion of 4.1's sub-select functionality, and also, very cleverly, *how to work around 4.0's lack of sub-selects*. Marvellous.

This is a book is so engrossing and so power-packed that you can read it straight through first time – DuBois makes it that easy. If you use MySQL at any level, you should keep this book close to hand – I certainly do.

VERDICT

The only MySQL book that is better than DuBois' *MySQL 1st Edition*.

LINUX FORMAT RATING
10/10

LDAP Directories Explained

Nick Veitch swots up on the database alternative.

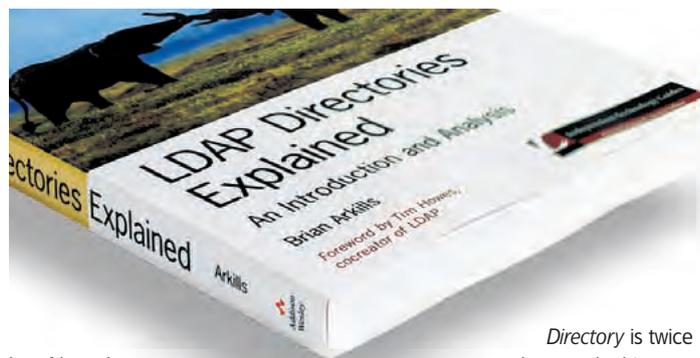
BUYER INFO

- **PUBLISHER** Addison-Wesley
- **AUTHOR** Brian Arkalis
- **ISBN** 0-201-78792-X
- **PRICE** £24.99

LDAP has really found its niche in the computing world. The LDAP directory protocol is commonly used for, well, directory services, whether that be contact lookups (as in mail clients etc), login authentication or even parts catalogues. This book starts out explaining why that is, and the difference between a directory and the more omnipresent database.

Subsequent chapters explain the use of hierarchies, the LDAP namespace, object structures and common client operations which might be performed against such a directory.

Along the way concepts are explained with useful black and white diagrams and illustrations, and quite a



lot of boxed out material explaining concepts or answering common questions. These actually work pretty well, and seem to cover most of the concepts a newcomer to the subject might have difficulty understanding.

There is a good 20 pages dealing specifically with OpenLDAP, which deals intelligently with the pros and cons of this implementation of the LDAP protocol and its drawbacks. The corresponding section on *MS Active*

Directory is twice as long, and a bit more glowing, but as I can't say I have personally used it, I can only accept the authors opinion about how fantastic it is.

The one drawback with this book in my opinion is that it doesn't really go into enough practical detail, and some important topics aren't covered in great enough depth. For example, the topic of distributed directories is covered in a mere six pages. There is some discussion of the pros and cons, but the subject is closed with a line

about co-ordinating modifications being challenging, with no real clue as to whether it is possible, or what tools and systems you might use to effect such modifications.

As this book is supposed to be an introduction to LDAP, rather than a working manual for using it, perhaps some of the criticism here might be a bit harsh, but some more practical examples might have helped get across some of the more difficult concepts in the book, plus make it more handy as an oft-used reference while you're working. If you are looking for a more practical guide, Wrox's *Implementing LDAP* (ISBN 1-8610-0221-1) is probably a better bet.

VERDICT

A bit too much explanation, a bit too little practical help.

LINUX FORMAT RATING
6/10

Unix Power Tools

Power up your Unix experience with **Nick Veitch**.

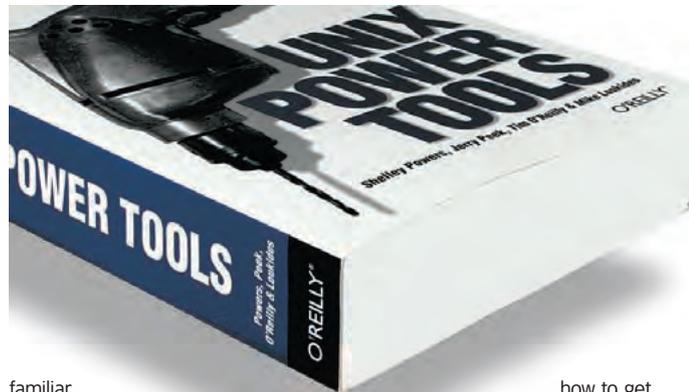


BUYER INFO

- **PUBLISHER** O'Reilly
- **AUTHORS** Shelley Powers, Jerry Peek, Tim O'Reilly, Mike Loukides
- **ISBN** 0-5960-0330-7
- **PRICE** £34.99

One of the great things about computing, and perhaps more true of Unix and Linux specifically, is that there's always more for users to learn. It's great then, to find a single resource full of all the things you didn't know, and that is almost precisely what Unix Power Tools is.

The first chapter gives a basic introduction to Unix, and while much of the book might be a little over the heads of beginners, there is a lot of valuable reference material for everyone here. Experienced Linux users may find sections they are



familiar with, but there is almost certainly some tip or trick you won't have come across in almost every section – for example, did you know that using `$11` in a `sh` script won't return argument 11 of the previous line, but actually the first argument with a '1' appended to it. This problem,

how to get round it and thousands of other tips are what make up the substance of this book.

An impressive 79-page detailed index means you'll always be able to find what you want, but in the office we've found it fun to open the book at random now and then, just to discover another great tip. Fifty-one individual

chapters group together the information together loosely by subject matter. You'll find sections on *Vi*, *grep*, files, Python, shells, security, printing, networking and more. Each individual topic is usually accompanied by examples, cross-references to other parts of the book and occasional 'gotchas' pointing out potential errors you might want to avoid.

If you're looking for one single book that presents the distilled wisdom of hundreds of years-worth of Unix use in an easy-to-read and simply navigated single volume, then this is it.

VERDICT

This really is an essential book for all Linux/Unix users. Top marks

LINUX FORMAT RATING

10/10

PHP And MySQL Web Development, 2nd Edition

Paul Hudson looks at the second edition of the popular web database development book.

BUYER INFO

- **PUBLISHER** Sams
- **AUTHOR** Luke Welling and Laura Thomson
- **ISBN** 0-672-32525-X
- **PRICE** £36.50

At the time of writing, *PHP And MySQL Web Development* is the best-selling PHP book on Amazon, being their 42nd top-selling book in any genre – for some, that probably says it all! The popularity is with good reason, too: the first edition of this book was the first book to combine PHP and MySQL together, despite the two having gone hand-in-hand for so long.

This new edition is now updated and expanded to include coverage of PHP 4.3. There are still one or two layovers from the first edition, such as "A new edition to PHP 4 is the reference operator", which is old news, really – PHP 4 has been out over three years now, so I think references to PHP 4 being new are pretty irrelevant.

On the topic of "new", there is no entry in the index for the *new* operator



– the closest there is is a reference to the *New York Times* website (no, really!) The index, as you can guess, is sadly less than perfect, and as this is the second edition it would have been good to have these kinks worked out by now.

The MySQL coverage is quite weak, being strongly tailored towards end results rather than how MySQL actually works. For example, the 'section' on optimisation is a little over a page long, and barely scratches the

surface of what tweaking and streamlining is actually possible. Additionally, if you're looking for MySQL 4-specific information, you're out of luck, as the book references 3.23.x throughout.

Be wary when reading, because often the authors seem to mention things that 'make sense' as opposed to being correct. For example, one of the five (yes, only five) points offered on how to optimise your PHP code is "use single-quoted strings instead of double-quoted strings where possible [because] PHP evaluates double-

quoted strings looking for variables to replace". Rubbish! Yes, PHP evaluates double-quoted strings to replace variables, but users will speed /little to no speed difference/ in changing double-quoted strings without variables inside over to single-quoted strings, and it merely acts to misinform.

However, the book isn't *all* bad – far from it. If you're looking for the most hands-on manner of learning PHP, this is definitely the book for you. There is little time wasted on theory – in its 872 pages you'll find a direct, solutions-based problem solver that tackles real world scenarios with no fluff added to chew up space. It has its flaws, but they're not terribly fatal – the book makes for an interesting read regardless of the oversights.

VERDICT

Not really an improvement over its predecessor, and you can probably pick up 1st edition on the cheap.

LINUX FORMAT RATING

7/10

PostgreSQL

Paul Hudson looks at the long-overdue treeware reference guide to PostgreSQL.

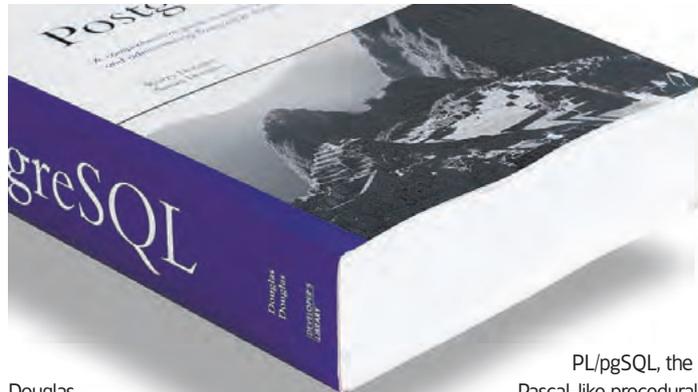
BUYER INFO

- **PUBLISHER** Sams
- **AUTHOR** Korry Douglas, Susan Douglas
- **ISBN** 0-7357-1257-3
- **PRICE** £36.50

Shocking as it may seem to some, MySQL is *not* the only Open Source relational database management system (RDBMS) out there, indeed MySQL barely counts as a RDBMS as far as some are concerned, simply due to its lack of functionality.

PostgreSQL has been around for a long time, and is a very feature-complete and mature RDBMS, including such functionality as views, stored procedures, and triggers – powerful functionality that MySQL users are unlikely to see in the near future.

As such, it's a surprise that a comprehensive book on PostgreSQL has taken so long to be launched. However, *PostgreSQL* was certainly worth waiting for. Korry and Susan



Douglas have done a marvellous job creating a mix of discussion on basic functionality such as how to select data, all the way up to expert discussion on how to use the libpq library to write your own Postgres-enabled C applications.

Furthermore, PostgreSQL fans will be delighted to see a thorough guide on how PostgreSQL optimises queries, and how it stores data internally. Particularly of interest will be the large chapter on

PL/pgSQL, the Pascal-like procedural language built into Postgres, and there are almost forty pages devoted to the topic.

There are one or two minor niggles with the book, particularly the fact that some parts of the book lack sufficient explanation and therefore are left hanging. For example, page 41 discusses the use of "ORDER BY.. USING", but doesn't give any examples of it. Furthermore, page discusses MIN() and MAX(), then gives an example of how to use the two functions when

dealing with string data. However, it doesn't explain the results – that is, when it says the "minimum" customer name in the example table is "Funkmaster, Freddy", is that because it comes first when ordered alphabetically on the first letter, because it comes first when ordered alphabetically on the average of all the letters, or because it has the minimum number of letters?

To conclude, it's an impressive book written by authors who know a lot about PostgreSQL and also write well together. It's well structured, well edited, and has something for all levels of readers. Yes there *are* niggles, but they're small, and still aren't enough to drag it down. With books like this available, PostgreSQL may get a serious popularity boost. **LXF**

VERDICT

The best available tome for PostgreSQL fans, and a cracking read to boot.

LINUX FORMAT RATING

■■■■■■■■■■ 9/10

EASY TO BUY • EASY TO SET UP • EASY TO SEE

DOMAIN NAMES
WEB HOSTING
E-COMMERCE
VIRTUAL DEDICATED SERVERS

log on

and be seen

Pick up a new domain name from £3.50 or a web hosting package from £30 per year. Backed with robust support and technology, Easyspace is one of the largest ICANN accredited registrars in the world, and the UK's favourite domain name registrar.

log on today at: www.easyspace.com

Established in 1997 with over 1/2 million customers. Accredited ICANN registrar & nominet member. Prices exclude VAT.

Easyspace 
your perfect partner for the web

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

Thunderbird	37
Ratpoison	38
Monkey HTTPD	38
JuK	39
LNx-BBC	39
Jay's Firewall	40
Endeavour II	40
Search and Rescue	41
Emilia Pinball	41

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!

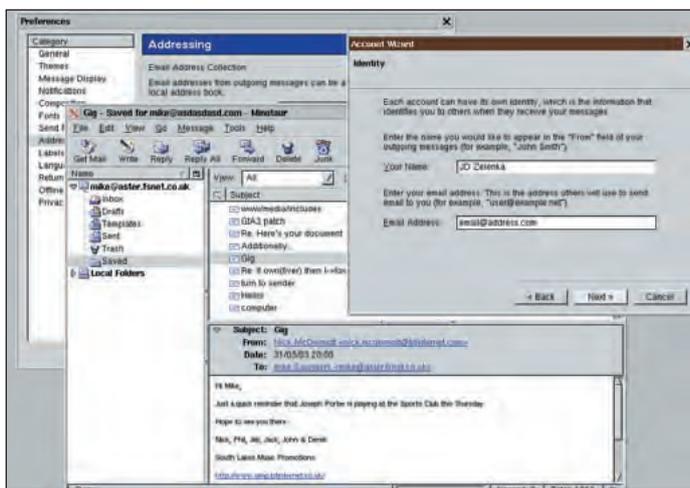


MAIL/NEWS CLIENT

Thunderbird



VERSION 040403 WEB www.mozilla.org/mailnews/minotaur/



Thunderbird's main window, preferences box and account wizard in action.

Mozilla, the Open Source Internet suite derived from Netscape Communicator, was confronted at first by problems with the kludgy original codebase and difficulties in getting community support. After a slow start, though, the 0.9.x releases brought about a stable and featureful program; additions like tabbed browsing, advert filtering and popup blocking have made it a favourite of power-users and Linux newcomers alike. However, many are still unsatisfied with the suite's sheer size and performance – it's a resource-hungry app and not ideally suited to low-spec machines.

Thunderbird has just been renamed – in development, and until the time of writing, it was known as *Minotaur*. In issue 37's Hot Picks we looked at *Phoenix* (now also renamed as *Firebird* – not to be confused with the database of the same name, of course!), a smaller and lighter browser-only variant of *Mozilla*, and now we have *Thunderbird*: a similar project but providing just the email and newsreader components. Gone are the browser, IRC client and page creator – the focus is much tighter

Current *Thunderbird* development is based on the 1.3 trunk of *Mozilla*, and there has already been work on the size issue – this snapshot release is a download of 10MB, in contrast to the 14MB of *Mozilla*. Not a huge reduction then, but it's early days and we expect to see more trimming as time goes on. Equally, we found that startup times have improved by around 25% (cached) in our tests – not massive either, but it's still a good sign.

Thunderbird presents a familiar three-paneled main window, with the folders on the left, message list on the right (sortable by subject, sender, date, priority etc) and message viewer at the bottom. A wizard appears when the program is first run – this steps through the process of setting up the mail or newsgroup settings, and is suitably simple to understand. Along the top sit the usual bunch of icons for replying and forwarding mail (this is now configurable), and the interface as a whole is easy to navigate – *Outlook* users won't be too fazed.

Plain text and HTML composition is catered for, along with return receipts,

online/offline settings, audio notifications, graphical emoticon substitution, an attractive addressbook, and more. The main window layout, font sizes and colour coding (for message priorities) can all be fine-tuned, as well as the security-enhancing options of disabling Javascript and remote-image loading too. Viewing the full source of a message and switching themes appears to be broken in this release, but *Thunderbird's* growing featureset is decent enough for all but the most heavyweight users.

Hit mail on the head

Other highlights include a comprehensive searching facility, with pattern matching for any variety of header fields, alongside an equally powerful filtering system for automatically organising new messages. To combat spam, some junk mail options are available too – this is achieved by the user flagging obvious spam messages, and from then on Thunderbird attempts to catch other suspicious mails in further sessions. Potentially dodgy messages from people already in the address book won't be trapped, though, helping to avoid false-positives (genuine mails being marked as spam).

Those who use *Mozilla's* mail and news components won't find any drastic changes in *Thunderbird*, but development is moving along steadily and the new features (like the customisable toolbars) and improved performance are great to see. Stability is decent too, with the client withstanding large message folders well. For later releases, the coders plan to use the *Phoenix* toolkit and a *Qute*-like icon set and theme, further improving the speed and interface.

Thunderbird's real strength, though, is in providing non-*Mozilla* users with a flexible and feature-packed mail client – if you prefer *Konqueror* or *Opera* for browsing but you're still looking for a good mailer, give *Thunderbird* a try. It's still a beefy app and we hope to see more work on the size and startup time, but like *Firebird*, the project is attracting a lot of attention and development.

KEY-DRIVEN WINDOW MANAGER

Ratpoison

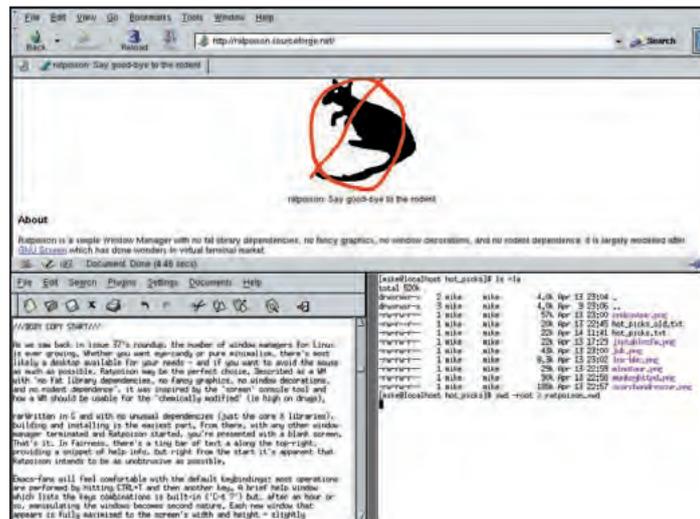
■ VERSION 1.2.0-beta3 ■ WEB <http://ratpoison.sourceforge.net>

As we saw back in issue 37's Roundup, the number of window managers for Linux is ever-growing. Whether you want eye-candy or pure minimalism, there's most likely a desktop available for your needs – and if you want to avoid the mouse as much as possible, *Ratpoison* may be the perfect choice. Described as a WM with 'no fat library dependencies, no fancy graphics, no window decorations, and no rodent dependence', it was inspired by the screen console tool and the idea that a WM should be usable for the 'chemically modified' (ie drunk as a skunk or high on drugs). Hmm...

Written in C and with no unusual dependencies (just the core X libraries), building and installing is the easiest part. From there, with any

other window manager terminated and *Ratpoison* started, you're presented with a blank screen. That's it. In fairness, there's a tiny bar of text along the top-right, providing a snippet of help info, but right from the start it's apparent that *Ratpoison* intends to be as uncluttered and unintrusive as possible.

Emacs fans will feel comfortable with the default keybindings; most operations are performed by hitting CTRL+T and then another key. C-t and O-9, for example, switches windows. A brief help screen which lists the key combos is built-in ('C-t ?') but, after an hour or so, manipulating the windows becomes second nature without too much thought. Each new window that appears is fully maximised to the screen's width and



No borders, no menus, no buttons - just framed to perfection.

height – slightly intimidating at first, but it's easy to split them horizontally and vertically.

An example commented config file is provided with the source, which explains how to add new bindings. The man and info pages are detailed and well-written too, while stability and performance were excellent during

our testing. In all, *Ratpoison* is a curious little WM which either instantly repels or immediately attracts; functionally, a keyboard-only WM may seem odd, but with some patience it's enormously fast to use and doesn't waste a single pixel of screen real-estate. If you're sick of the rodent, spend some time with this.

SMALL WEB SERVER

Monkey HTTP Daemon

■ VERSION 0.6.3 ■ WEB <http://monkeyd.sourceforge.net>

Apache is undoubtedly one of the most popular Open Source projects, at present eating up a whopping 63% of the Web server market (according to Netcraft) and

leaving Microsoft's *IIS* in the dust. However, in some cases the features and complexity can be overkill; as a result, various other projects are aiming to provide better alternatives

for simpler jobs, and here we have *Monkey HTTPD*, which aims to be small, fast and easy to configure.

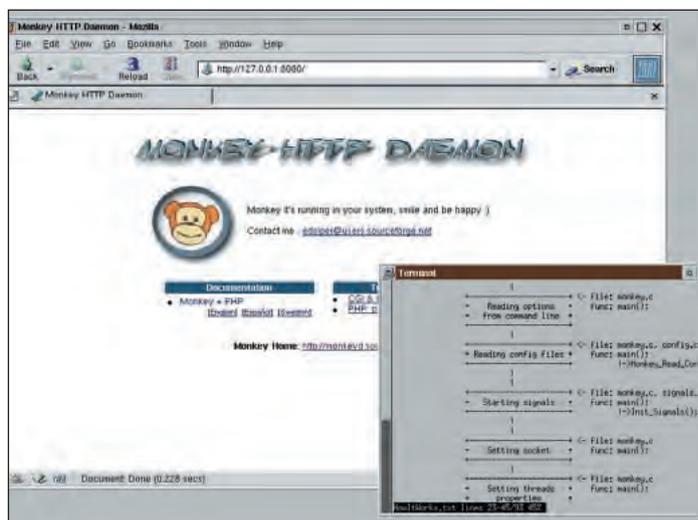
Monkey includes CGI and PHP support, virtual hosts and MIME types, and at only 63k for the source, *Monkey's* C code (with Spanish comments) is clean and lean – taking up only half of the total tarball size. Also included is a detailed explanation of how the daemon starts up and accepts connections, along with a stock directory tree containing an example config file and HTML front page. Despite the docs, **make install** won't work and the automatically produced *monkey.conf* will have hard-coded paths in (ie where it was compiled), but a quick sed job can sort that out.

Monkey's config file has the appearance of a compacted *Apache* equivalent – many of the keywords are identical, and although the translation may not be too lucid in places (the author Eduardo Silva is Chilean) it's well commented and easy to get started with. Together with the

basic name, port, document root etc. settings, the location of logfiles can be changed, timeouts and max connections tweaked, and version number omitted from a served page for security. It doesn't approach the flexibility of *Apache*, but for its diminutive size it does a lot.

Also in the conf dir is *monkey.deny*, where URLs can be blocked if they match certain patterns, as can IP ranges using wildcards. *Monkey's* log output is almost the same as *Apache's* default – the log-parsing tools should only need minor tweaks to work. On the whole, *Monkey HTTPD* is a compact, friendly and functional webserver which makes a perfect alternative to *Apache* for smaller tasks – yet it still has a decent range of features. Improvements to the docs and installation process would be welcome to attract more users, though the quality of these still manages to surpass some apps whose authors claim that English is in fact their first language.

As with so many of the programs featured in *LXF's* Hot Picks section, work on *Monkey* is ongoing, so if you have particular skills relevant to web servers or Spanish language, why not get in touch and contribute?



The outcome of a successful installation, with the tech docs on display.

KDE JUKEBOX

JuK

■ **VERSION** 1.1 ■ **WEB** www.slackorama.net/cgi-bin/content.pl?juk

Last month we looked at *Rhythmbox* in Hot Picks – the new GNOME media player aiming to displace *XMMS* as the best music tool for that desktop suite. Naturally, KDE's

huge band of coders don't just sit back and watch, though; they've been working on their own jukebox app, *JuK*, which is due to be included in the upcoming KDE 3.2 release.

Pronounced "Jook", the project's goal is to build a media player akin to *iTunes* or *RealOne*, according to the hackers behind it.

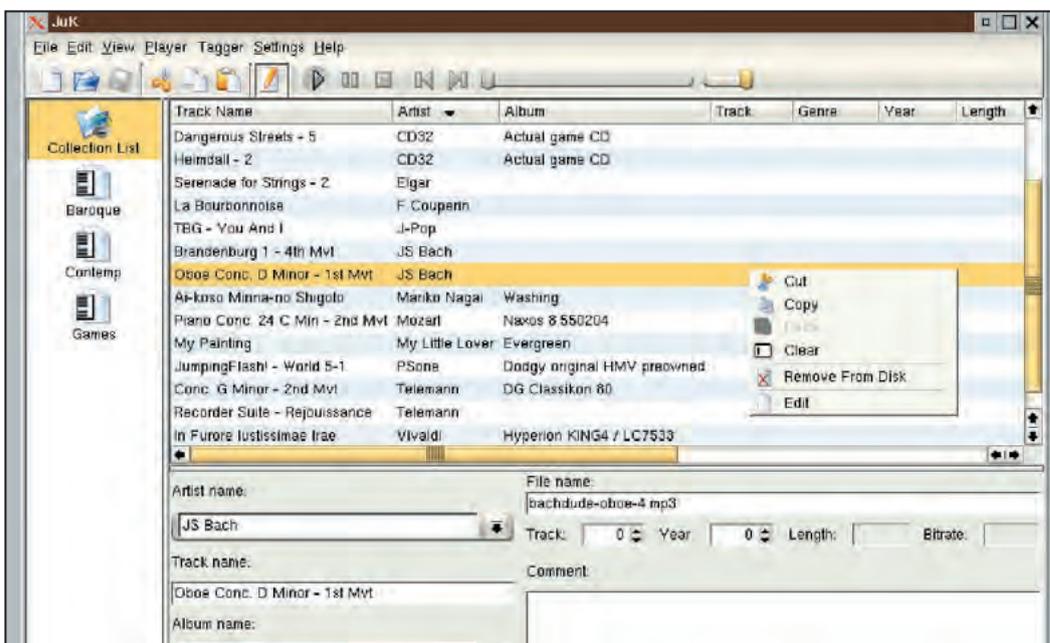
To build the latest *JuK*, you'll need a KDE 3.1 installation (including all the necessary development libraries and headers) with the *kdemultimedia* bits, and the *id3lib* library (required for the MP3 tag operations). Building should go hitch-free with those in

place, but Mandrake and SuSE RPMs are available from the project's site.

JuK's interface is built around a split window, with the playlists arranged down the left-hand side and individual songs on the right – the latter list can be sorted by name, artist, genre, year etc. Currently, *JuK* supports MP3s and the increasingly popular Ogg format, and .m3u lists can be loaded too (a life-saver with large, hand-ordered collections). By default it uses KDE's own *aRts* for output, although *GStreamer* is supported as well.

Songs can be dragged-and-dropped into other playlists, while tags for single or multiple song selections can be modified. There's little in the way of configuration options; the app is a breeze to use though, and there's already a fairly concise online manual in place. *JuK* stands well as a no-nonsense, simple music jukebox which fits in nicely with KDE, but a few features some regard as essential – like searching and skinning – aren't present in this current release.

Scott Wheeler, the main developer, is now concentrating his efforts on version 2.0 – it's due to include a UI redesign, configurable keybindings and other new features. *JuK* is definitely one to watch, then, and we're very much looking forward to seeing it in the next KDE release.



JuK's main window, here demonstrating the built-in tag editor.

LIVE MINI-DISTRO

LNX-BBC

■ **VERSION** Pre-2.0 2003-03-30 ■ **WEB** www.lnx-bbc.org

Almost every serious Linux user has, at some point, managed to make a box unbootable – either by removing the wrong package, botching a config file or incorrectly configuring a kernel. It's all part of learning the system, and that's when a rescue floppy or CD comes in very handy. Mini-distros like *muLinux* and *TomsRootBoot* have been popular in the past, and the likes of Knoppix have done well in demonstrating Linux to newcomers. *LNX-BBC* means *Linux Bootable Business Card* and, as the name suggests, it's a small live distro intended for use on those business-card CD-Rs.

At 50MB there's plenty of space to cram in a good amount of essential utilities, yet it should still be small enough for most micro-size CD-R discs. The *LNX-BBC* project has been quiet in recent months, but the developers are heading for a 2.0 release and the latest snapshots, as on test here, are looking good. You simply need to write it to disc with *cdrecord* (or an alternative graphical tool), pop it in the drive, set the BIOS to boot from CD, and go.

LNX-BBC runs as a live filesystem; essentially, this means that it runs entirely off the CD, so the OS and data currently on the machine won't be affected. All that's needed is a 386

with 16MB of RAM. Once booted, you're presented with a minimal but functional Linux system, with SSH for remote logins, *thttpd* for web serving, *VIM*, *Joe* (*jmacs* for *Emacs* fans), Perl, Python and more. There's even a basic framebuffer X installation with *HackedBox*, and it's all built around kernel 2.4.19 and *glibc* 2.2.5. Although it's not competing with the marvellous

Knoppix as a demo CD to sway Windows users, *LNX-BBC* is still an excellent choice for those needing a robust recovery CD, portable Linux kit or hard-drive-less server. There are a few areas which need cleaning up, but otherwise it's definitely worth writing a copy and keeping it safe for the emergency that's bound to occur one day...



The boot screen is about as exciting as it gets, graphically.

FIREWALL SETUP TOOL

Jay's Iptables Firewall

■ VERSION 0.9.95 ■ WEB www.wallaby.be/firewall/

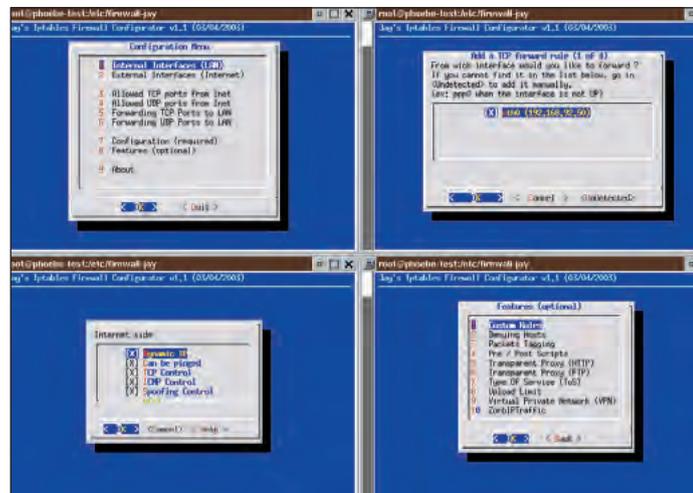
There's a big, bad world on the Internet, and keeping systems secure from crackers and other anti-social types is a never-ending chore. If you run a network, being able to restrict access to protect the boxes is essential; many distributions include basic firewalling tools for such configuration, but it can get tricky and programs like *Jay's Iptables Firewall* ease the setting-up process.

JIF is a text-based application (so it's ideal for running on headless boxes over SSH/telnet), and to use it you'll need a kernel 2.4 box with a recent *iptables* binary, along with Perl to run the script and the 'dialog' program for displaying the config system. The README lists various kernel options that need to be compiled in or built as modules to support *JIF*'s numerous features, but

the stock distro kernels we tested it with worked fine.

During installation, *JIF* adds *fw-jay* to the system's startup scripts, creates `/var/lib/firewall-jay` containing the main rules script, and the first run builds `firewall.config` in `/etc/firewall-jay`. Enter `firewall-config.pl -n` and a console-based configurator (similar to the Slack and Debian installers) pops up. This offers nine submenus, and it's a cinch to operate – for example, to forward a port (TCP or UDP), you're prompted for the interface, port number, and lastly host (and its port) of the other box.

The Configuration section assists in enabling NAT (for sharing an Internet connection), logging options and other extras, while the Features screen offers the ability to run scripts before and after starting up and



Main menu, port forwarding, Internet settings and extras screen.

defining a custom rules file etc. Other features include spoofing and SYN flooding control, TOS tuning and upload limiting.

The range is impressive, and apart from the odd slipup in translation (the

author is French), it's easy to work with. Networking gurus will naturally prefer to get their hands dirty with the command line, but for those who want to set up a basic firewall with minimum hassle, *JIF* is a great choice.

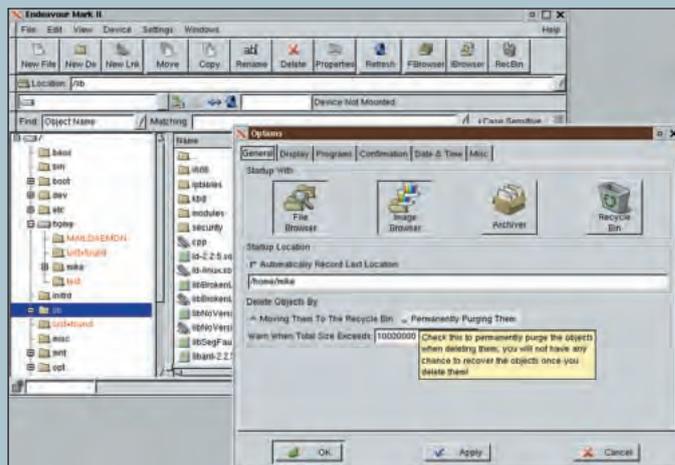
FILE MANAGER

Endeavour II

■ VERSION 2.3.0 ■ WEB <http://wolfpack.twu.net/Endeavour2/>

Hard drive sizes are growing at a stupendous rate and most every-day computer users have an enormous stash of files on their machines. Being able to

navigate and organise files via the shell prompt is useful knowledge, but fully-fledged graphical file managers are now more important than ever, with efforts to emulate Windows



Endeavour's chunky file manager window, with the options box in front.

Explorer, Norton Commander and even RiscOS currently being worked on. Quite why anyone should wish to parrot some of these managers in a 'faults-'n'-all' way rather than come up with a sleeker, more intuitive tree system is open to question; but maybe recent migrants from Windows will appreciate being able to rummage through their 80GB+ of accumulated jumble in a familiar fashion. Our roundup way back in issue 4 of Linux Format looked at File Managers, and an intriguing new entrant to the field is *Endeavour II*.

As the name suggests, this program follows on from the first *Endeavour* which was halted at version 1.13.0 for a rewrite. The latest RPM (as supplied on our coverdisc) is slightly dated, so it's better to compile from source to get the newest version – GTK+ 1.2.10 is required, and *Endeavour II* will make use of *lmlib* and optional accessory addons if installed.

Sporting an *Explorer*-like two-paneled display, with the directory tree on the left and detailed file list on the right, extra toolbars can be enabled for mounting disks and performing

searches. The unique icon set is plain and quiet, but there's a plethora of extra buttons which can be added to the toolbars. Initially, no applications are bound to file types, but fortunately a list (`eg /etc/mailcap`) can be imported.

Endeavour II is an impressively comprehensive file manager – a graphical *df* (disk space monitoring) tool is included, as is drag-and-drop and a history of recent operations performed. A supremely fast thumbnail-based image browser has been thrown in for good measure, too. The options box provides a pleasing array of options to tweak, while tooltips, status-bar hints and tiny help boxes all contribute to an easy-going application.

Reliability is top-notch; the program didn't fall over through any operation, search or configuration we tested, and it thoroughly speeds along. A little more general polish and work on the icons would be good to see, but if you're looking for a flexible, slimline, featureful and solid file manager with no superfluous graphical frills, take *Endeavour II* for a spin.

Controlling this beast is like driving a JCB in an ice-rink.



AIR RESCUE SIM
Search And Rescue

■ **VERSION** 0.7.21 ■ **WEB** <http://wolfpack.twu.net/SearchAndRescue/>

72.4 MPH
 Thr: 100%
 1007 AGL
 1008 RDC

Apparently, saving somebody's life is one of the most rewarding things in the world. Unfortunately, sipping beer in front of the computer is more comfortable than roaming seas and mountains in search of people with broken limbs or no arm-

bands. Thanks to the efforts of some talented programmers, though, we can now zoom around in a large helicopter, enjoying artificial adrenaline rushes as we airlift a 'victim' to safety.

SAR is satisfyingly simple to compile and install – providing you have OpenGL

(Mesa) on your system (and DRI is best for decent speed), it'll build without complaint. It'll also use the *libjsw* and *YIFF* libraries, for joystick and sound support respectively, but if they're not installed you simply need to pass the right flag to the configure script. *SAR*'s

main window is tiny on first run, but thankfully it's fully resizable through the window manager.

To get to grips with the controls, *SAR*'s data pack – as supplied on our coverdisc – provides six training missions; these typically involve flying to a location, carefully retrieving the stranded soul, and returning him to a nicer place. To make proceedings more interesting, the environment tries hard to hamper progress – fog, rocky shorelines and the dire visibility of night-time all come into play.

The 'Free Flight' option is far less intense. Simply choose an area (Desert, Los Angeles or the imaginatively named 'Test World'), an aircraft (ranging from a dinky nimble helicopter to a skin-tearingly fast Lockheed 1.7 Mach jet), and explore the scenery. Using PgUp and PgDn for throttle and the cursor keys for control is easy enough, and all the aircraft react convincingly.

SAR's worlds are enjoyable to cruise around in, but with some work it's possible to create new scenery, locations and even missions (stored as plain text files!), adding to the longevity. Despite the absence of guns, rockets or bad guys to waste, it challenges belief and manages to prove that helping people is actually surprisingly soothing.

PINBALL GAME

Emilia Pinball

■ **VERSION** 0.2.0 ■ **WEB** <http://pinball.sourceforge.net>

Pinball's ancestry is found in Pachinko-like machines of the early 1930s, and its instant popularity resulted in the Mayor of New York imposing a ban on it in 1948. Why? The game supposedly 'corrupted the morals of children and bred crime, rackets and gangsterism'. *Emilia Pinball*, then, should help continue the spread of 'juvenile delinquency', providing you have at least a 300MHz chip, OpenGL (Mesa) and SDL installed. The RPM on our coverdisc was built for RH 7.3 but should work elsewhere, and the source is also available for building.

Pinball sim designers are always faced with a difficult choice: does realism or fantasy make a better game? *Dragon's Revenge* and *Kirby's Pinball Land* opted for the latter, and included heaps of bizarre bonuses

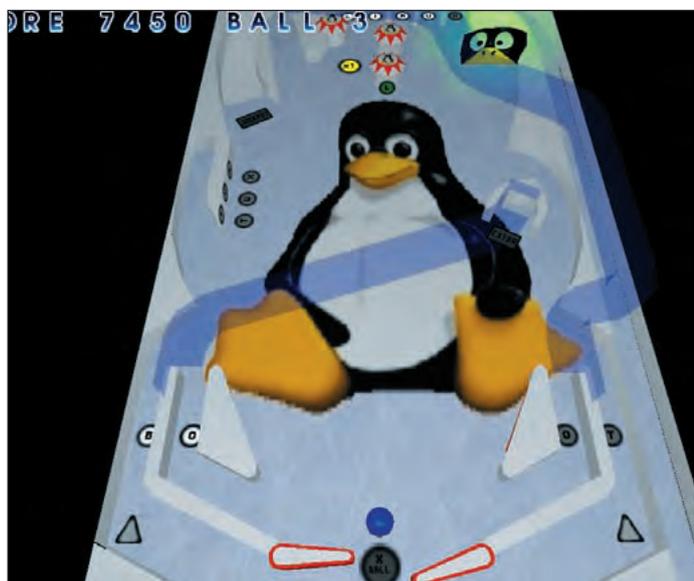
and sub-games to improve the game's lifespan. *Emilia*'s two supplied tables stick with the former, though – the standard pegs, buzzers, bumpers and chutes are all here, and there's even a tilt function to simulate abuse of the machine.

Emilia's controls are straightforward enough – the left and right shift keys control the flippers, space bar tilts, and enter is used to launch the ball. Tux (the Linux penguin mascot) is the basis of the first table's theme, with its giant image spread across the centre and bonuses for hitting all letters of 'BOOT' and 'LINUX' etc. Tux is evidently hungry though, so be careful when the animated Penguin head in the top-right opens up – he'll swallow the ball.

Professor, the other table, is less complete (and rather garish); its main

development over the original is the addition of an extra flipper further up. Still, the lack of bonuses and slightly awkward design results in Tux being far more playable, and although new tables can be created with a bit of effort, we'd like to see a few more

thrown in with the main bundle. *Emilia* is fun, silly and casual, with competent ball physics and fantastically fitting chilled-out background 'choons'. It won't eat up endless hours of free time, but it's good to have around as a quick desktop game. [LXF](#)



Egad! Tux's carefree smile won't stop my final ball disappearing...

LAPTOPS MEGATEST

Surrounded by seven top-of-the-range portable PCs buzzing away compiling kernels, Paul Hudson is in hardware heaven reviewing the current state of Linux laptops...

Linux on laptops has long had a very rocky history, mostly because laptops often use custom pieces of hardware that require specialist support. Sites have sprung up offering help on a case-by-case basis for laptops – what each unit contains and how easy it is to get parts

of Linux working. As time has gone by, more and more of the troubles with laptop Linux have disappeared thanks to better autodetection of hardware. Very often you're capable of installing Linux and having everything work first time.

But what if we use a distro without hardware detection? Are laptops

capable of running Linux without the help of programs like *kudzu*? Our hardware test platform is Debian 3.0, perhaps one of the least-forgiving distros available – can it work on today's laptops without hassle?

Leading laptop manufacturers sent us their top-range hardware so that we

could install Linux on them. If you want to know how well a particular laptop performs with Linux, read on. All prices include VAT, and are the manufacturers recommended retail price – you should shop around to get the best value. Also, each unit came with a built-in Ethernet port and also one (small) firewire port. >>



LaptopsMegatest

SONY VAIO PCG-GRS615SP

Sony's venerable VAIO series continues the march towards laptop supremacy – light, attractive, and always coming with top hardware.



SPECIFICATION

Price: £1650 (inc VAT)
CPU: 2GHz P4-M
RAM: 256MB
Screen: 15" XGA
HD: 40GB
Extras: Memory Stick, CD-RW/DVD, WLAN

Powered by a 2GHz P4 and 256MB RAM, the Sony offering comes in towards the lower half of our test group on specification alone, but on the other hand it performs admirably – this is a testament to the high-quality kit that is used internally. Sony products traditionally come in an attractive case, and our review model was no different. The keyboard is easy to type on, and the trackpad is the usual first class kit seen in the VAIO range. Added bonuses include a slot for Sony memory sticks, wireless LAN support, and also a CD-RW/DVD player to boot – not bad! The memory stick slot is, of course, only really useful if you have other stick devices around, like a PDA or digital camera.

The quality of the XGA screen is second-to-none, and looks considerably better than some others in this group test – excellent for users who want to use the VAIO in brightly lit conditions. The graphics are powered by an ATI Radeon 7500, so is suitable for some degree of gaming, however you may have a little trouble getting drivers for X.

Around the back we have three USB slots – very generous for a portable PC – and the usual assortment of external connectors. The PCMCIA slot is a little odd, though, as it has only one dust flap rather than the usual two, which, I suspect, might let dust inside the slot

– probably *not* a good thing, you'd agree. With regards to portability, there's room for two big batteries, so battery life shouldn't be an issue – however the unit isn't particularly light. The Toshiba, Acer, and Fujitsu machines are both lighter than this Sony, so if portability is a particular issue you may want to stick with just the one battery.

At the price, the performance offered by the VAIO is quite good – particularly if you consider that it doesn't have the fastest CPU in this test by quite a long way.

www.vaio.sony-europe.com

BENCHMARKS

hd	1.32
apache	1.57
mysql	1.66
compile	1.52
oggenc	2.4
Overall	1.69

VERDICT

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

Attractive, fast, and quite well-priced – a good buy, with a good pedigree.

LINUX FORMAT RATING

8/10

TOSHIBA SATELLITE PRO 6100

Doesn't have the most attractive design by a long shot, but its performance belies the specification it comes with – well worth a try.



SPECIFICATION

Price: £1350 (inc VAT)
CPU: 1.7GHz P4-M
RAM: 256MB
Screen: 15" XGA
HD: 30GB
Extras: WLAN, GeForce 4 graphics

First impressions often count the most, and our first impression of the Toshiba was none too good – the keyboard has been designed in a most peculiar manner, with the delete key coming next to the space bar, and the cursors being hard to hit without practice. Furthermore, the lack of a trackpad (instead we have a Thinkpad-esque 'nipple') and the confusing layout of mouse buttons mean that this laptop has actually got a learning curve in terms of all the knobs and switches!

However, it does have Nvidia-powered graphics, which means first-class graphics reliability, easy access to drivers, and the added assurance that games such as *Unreal Tournament 2003* will run with little fuss.

The unit itself is designed a little unusually, and although it's attractive enough, it's unlikely to win any design awards. However, one advantage the Satellite has is weight – it's the second lightest in our test, which makes for a very good level of portability. Add to that the low-power CPU and altogether the Satellite makes a great choice for users wanting long battery life. There are 2 USB ports around the back of the unit, plus wireless LAN built-in – a fairly good level of extras for a laptop.

Performance-wise, this laptop comes somewhere in the middle of the pack, which is excellent

considering it has some of the lowest specs in the test – again, this is another case of people not getting quite what they expected from raw figures! For the price, this is a fairly good buy – particularly if battery life is the most important thing in your criteria. However, when regarding price/performance, the Satellite loses out to several competitors, and probably won't be the first choice for power users.

At the very least, the Linux support for this laptop is impeccable – everything will "just work".

www.toshiba.co.uk

BENCHMARKS

hd	1.24
apache	1.57
mysql	1.36
compile	1.35
oggenc	1.85
Overall	1.48

VERDICT

Features	9/10
Performance	7/10
Ease of use	6/10
Value for money	10/10

Odd design, and a terrible keyboard. Needs work, but it's cheap at least.

LINUX FORMAT RATING

7/10

COMPAQ PRESARIO 2500

The second largest machine in the pack – a fast CPU, but also poor performance figures? Read on...



SPECIFICATION

Price: £1599 (inc VAT)

CPU: 2.8GHz P4

RAM: 512MB

Screen: 15" XGA

HD: 40GB

Extras: CD-RW/DVD, WLAN

The Compaq machine is amongst the largest laptops I have encountered in the past five years, however the size factor is minimised somewhat by the fact that Compaq put a lot of thought into making the space appear well-used – the design is very attractive, and the unit as a whole, despite its size, makes an excellent desktop replacement unit. Both the keyboard and mouse are of particularly high quality, which is always good to see – especially as the keyboard is full-size.

Internally, the unit contains a 2.8GHz P4 and 512MBs of RAM. The processor *didn't* appear to be a Mobile P4, unlike most of the others, which you would think would add a lot of extra performance at the expense of battery life. However, despite the particularly beefy internals – the most powerful in this group test – this laptop persistently dragged in our tests. This enigma eluded us through all the tests – for one reason or another, presumably due to particularly low-quality non-CPU internals, the Presario 2500 refused to perform anywhere near what one would expect of the CPU. This is a great shame, particularly as many people will purchase units based upon raw numbers quoted.

With regards to extras, the Presario comes bundled with a CD-RW/DVD drive, Radeon graphics, and a wireless LAN built-in. The screen is quite good, but not too clear from wide angles.

There are three USB ports – two on the back and one on the side. Also on the side are volume controls – however, it doesn't take away from the fact that this laptop is the among the heaviest and the largest of the group.

So, to conclude, it's big, it's unjustifiably slow, and the battery life is amongst the worst in the test. For the price it's not terrible, but I would be surprised if it featured high on anyone's list. If you're looking for a way to build up some muscle, pop this thing in your backpack – otherwise, you're probably better off with one of the other laptops.

www.hp.co.uk

BENCHMARKS

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

VERDICT

Features	8/10
Performance	4/10
Ease of use	8/10
Value for money	5/10

Performance you'll forget at a size and weight that you wish you could forget.

LINUX FORMAT RATING

6/10

ACER TRAVELMATE 535LC

Attractive and feature-packed, with a top-end CPU to boot. Can it beat off the rest?



SPECIFICATION

Price: £1585 (inc VAT)

CPU: 2.4GHz P4-M

RAM: 512MB

Screen: 15" XGA

HD: 40GB

Extras: CD-RW/DVD, floppy drive

The Acer definitely wins top marks for design – the unit has a curvy theme throughout that is reflected all the way down to the excellent keyboard and trackpad. On that note, both the keyboard and trackpad were amongst the best seen in this test. The trackpad has very sensibly designed buttons, and the keyboard has a natural key-press feel – although the Home and End keys require the Function key to be pressed, which is a minor pain at first, but you get used to it after a while.

On the peripheral side, the Travelmate contains both a CD-RW/DVD combo drive and a 31/2-inch floppy drive – something few other units in the test had. Internally, the Acer unit has a 2.4GHz P4 with 512MBs of RAM to power it, and it certainly performs fairly well in most of the tests. Around the back are the usual array of connectors for printers and external monitors, but also three USB ports – more than enough for most people, although some of the other laptops have four, just to make sure!

With regards to portability, the Acer laptop is amongst the lightest we had, and the battery life was surprisingly long (approaching 2 1/2 hours) given its very high specifications. An extra battery would weigh the unit down a little more, but would probably still leave it lighter than, for example, the Presario laptop,

and would also give a significantly longer battery life.

Overall, this is a very impressive laptop, and comes out top for value and performance. Acer are onto a winner with this unit, and I wouldn't be surprised if this unit came out top on most people's list. Of course, if you're not all that interested in raw performance, there's still a lot to be had here – great looks, solid Linux support, and unbeatable battery life.

www.acer.co.uk



BENCHMARKS

hd	1.08
apache	1.69
mysql	1.92
compile	1.75
oggenc	2.85
Overall	1.86

VERDICT

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

It's fast, well-designed, full of functionality, and portable enough for Speedy Gonzales.

LINUX FORMAT RATING

10/10

LaptopsMegatest

FUJITSU-SIEMENS LIFEBOOK E7010

Super-slimline, ultra-portable, with a surprising amount of power under the hood



SPECIFICATION

Price: £1679 (inc VAT)
CPU: 2.2GHz P4-M
RAM: 256MB
Screen: 15" XGA
HD: 40GB

The lifebook was the lightest laptop in the test, and has a very low-key profile – only light colours are used in the unit, with lots of chrome to give it a very polished and modern look. The keyboard and trackpad are both above average, although the trackpad could benefit from being less indented into the unit, as it feels quite small right now. Particularly, the keyboard does have a good collection of built-in functionality – adjust brightness, sound volume, screen stretching, etc. Why aren't these features standard on all laptops?

Around the back there's two USB ports plus IrDA and the other usual suspects. The graphics are powered by the same ATI Radeon 7500 seen in some of the other laptops, and performs quite well as long you don't mind the usual problems getting your ATI card to work. On the other hand, the screen is excellent and has a good range of brightness levels so you can knock it right down and save as much as power as possible.

Inside the box there's a 2.2GHz P4-M waiting and a quarter gig of RAM – both of these are obviously twinned with some great hardware internally, because the performance of this machine was outstanding, as can be seen in the test results. Scoring 2.62 on the Ogg encoding test, which is the test that is most CPU-dependent, shows they've obviously

managed to feed the 2.2GHz CPU with lots of data – well done, Fujitsu! Also, the MySQL benchmark was another clear case of the Lifebook tearing away from the pack: this really is a performance laptop, despite not having the fastest CPU in the test.

If you're looking for great performance and are willing to pay that little extra for it, the Fujitsu might well be the machine for you. Add to that the fact that the Lifebook is incredibly light and portable and this is an all-round great laptop that's definitely worth the extra money. www.fujitsu-siemens.co.uk

BENCHMARKS

hd	1.28
apache	1.57
mysql	1.81
compile	1.70
oggenc	2.62
Overall	1.80

VERDICT

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

Excellent performance but a price tag to make you think twice.

LINUX FORMAT RATING
 8/10

MESH EXPLORER GX PRO

The mean machine of this test – chunky and curvy, but will it perform?



SPECIFICATION

Price: £1585 (inc VAT)
CPU: 2.8GHz P4
RAM: 1024MB
Screen: 15" XGA
HD: 60GB

This was the highest spec machine we received, and it was also (co-incidentally?) the largest and heaviest machine we received. With such a high-power CPU and enough RAM to handle anyone's uses for the next five years or so, this is a machine built to last. There's also one of the new Radeon 9000 cards inside to boot, which is guaranteed to give superior graphics performance on all fronts.

The keyboard on the machine is wonderful – full-size, keys all in the right place, and also with lots of special function options on the F keys. There's also a CD-RW/DVD combo drive and a floppy drive built-in, so there's no need to switch peripherals around as with most of the others.

Around the back there are no fewer than *four* USB ports, firewire, plus a digital out sound socket. It's good to see Mesh have used their space wisely! However, on the left there's only one PCMCIA slot – this might prove annoying if you need to have a Bluetooth adaptor and something else (with a built-in network adaptor and a built-in modem, *is there much else?*) plugged in at the same time.

On the performance front, the Mesh computer leads the pack by a long way. A long, *long* way. This is a direct result of Mesh utilising the desktop P4 CPU (not the mobile version) to its proper use – unlike the

Compaq machine. To make it clear, this is a laptop that compiles the Linux kernel in just over two minutes – an amazing feat by anyone's judge. The Explorer leads the way in most of the tests, and is an excellent desktop replacement machine.

If you can handle the weight, this is the machine that will give you the most for your money. The price is particularly low, and I can imagine Mesh getting a lot of orders for it because people must believe the price is a typing error. At the price given, this is the best you'll find, in our opinion. www.meshplc.co.uk

BENCHMARKS

hd	1.00
apache	1.83
mysql	2.17
compile	2.30
oggenc	3.42
Overall	2.14

VERDICT

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

Performance to die for encased in a box that might kill you if you carry it too long

LINUX FORMAT RATING
 9/10

EVESHAM VOYAGER

Small, light, with good build quality – will it make the short list?



SPECIFICATION

Price: £1479 (inc VAT)

CPU: 2.0GHz P4-M

RAM: 512MB

Screen: 15" XGA

HD: 60GB

Evesham has a long history of producing well-rounded machines – it's no co-incidence, after all, that I run Evesham machines in my home! As such, I had high expectations of this

unit, and wasn't disappointed. The unit itself is an attractive black-and-silver job, with a slot-loading DVD/CD-RW combo drive, and four USB ports around the back – a great start! It's let

down, though, by only have one PCMCIA slot, however this might not be a problem for most people given that it has a network device built-in like the other laptops here. One very positive plus is that the Evesham network card device was already built into the stock Debian kernel – a very rare thing indeed. As such, installation was a snap.

The unit was fairly heavy given its size, although Evesham have managed to cram a lot into the space – there's three speakers, and the screen quality is excellent. The keyboard comes with a good selection of built-in function options to manipulate the screen and the speakers, although there's not really enough range on the screen brightness – it goes from "very bright" to "slightly brighter". I for one often prefer "very dim", and that's not an option here.

With regards to performance, the Evesham doesn't lead the pack by quite a margin – but then, it doesn't have the highest specification here either. For what it has, a 2GHz P4-M, the unit performs very well, and "felt"

the quickest out of all the machines, although clearly the benchmarks point to it being an average performer. The relatively slow CPU (when compared to some of the others here) will mean that the battery life should be quite good for most usages.

www.evesham.com

BENCHMARKS

hd	1.18
apache	0.81
mysql	1.63
compile	1.65
oggenc	2.23
Overall	1.50

VERDICT

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

A good all-round performer that would make anyone proud – and the price is right to boot.

LINUX FORMAT RATING

9/10

BENCHMARKS EXPLAINED

Our in-house benchmarks are: *bonnie* and *hdparm* to test hard drive performance ("hd"), 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 averaged to produce an overall score. The blue bar

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' on the right, the blue value appears less than the red.

BENCHMARKS

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

ABOVE: Example Benchmarks box.

Megatest results comparison

	SONY VAIO	TOSHIBA SATELLITE	COMPAQ PRESARIO	ACER TRAVELMATE	FUJITSU-SIEMENS LIFEBOOK	MESH EXPLORER	EVESHAM VOYAGER
Price	£1650	£1350	£1599	£1585	£1679	£1585	£1479
CPU	2GHz P4-M	1.7GHz P4-M	2.8GHz P4	2.4GHz P4-M	2.2GHz P4-M	2.8GHz P4	2.0GHz P4-M
RAM	256MB	256MB	512MB	512MB	256MB	1024MB	512MB
Hard drive	40GB	30GB	40GB	40GB	40GB	60GB	60GB
Extras	Memory stick, CD-RW/DVD, WLAN	GeForce 4, WLAN	CD-RW/DVD, WLAN	CD-RW/DVD, floppy	N/A-	N/A-	N/A
BENCHMARKS							
Hard Disk	1.32	1.24	0.83	1.08	1.28	1.0	1.18
MySQL	1.66	1.36	1.11	1.92	1.81	2.17	1.63
Apache	1.32	1.57	1.22	1.69	1.57	1.83	0.81
Compile	1.52	1.35	0.96	1.75	1.70	2.3	1.65
oggenc	2.4	1.85	1.71	2.85	2.62	3.42	2.23
overall	1.69	1.48	1.17	1.86	1.80	2.14	1.5
Score	8	7	6	10	8	9	8

LaptopsMegatest

CONCLUSION

Each and every laptop installed Debian perfectly first time – this is a testament to the increasing support for laptops available under Linux. If Debian works, pretty much *anything* will work – Linux is definitely very usable on today's systems, and you don't need to be an expert either. Some pieces of hardware, particularly WLAN and ATI video cards, will be distro-dependent – while Red Hat and Mandrake would have no problem with them, they need extra, non-trivial work to run successfully with Debian.

With our criteria being desktop replacement hardware, battery life wasn't all that important. As such, we didn't use *apm*, and several of our

laptops was entirely trouble-free, which shows Linux has come a long way in recent years to get common drivers into the distributions, and the results are very noticeable. Laptop help sites still have an important part to play, though – be sure to check out www.linux-laptop.net if you're having problems with your system, or, better yet, post your own experiences there to help others.

All the laptops reviewed, with the exclusion of the Presario, are of good quality and offer performance quite a way ahead of our yardstick 866MHz machine. However, both the VAIO and Fujitsu machines deserve particular mention owing to their superior build quality, these two should top your list

“Linux is definitely useable on today's laptops – you don't need to be an expert either; though hardware like WLAN and ATI video cards will be distro-dependent.”

tests (particularly the MySQL *Super-Smack*) would really drain battery life if it wasn't plugged into the mains!

That said, using Linux on these

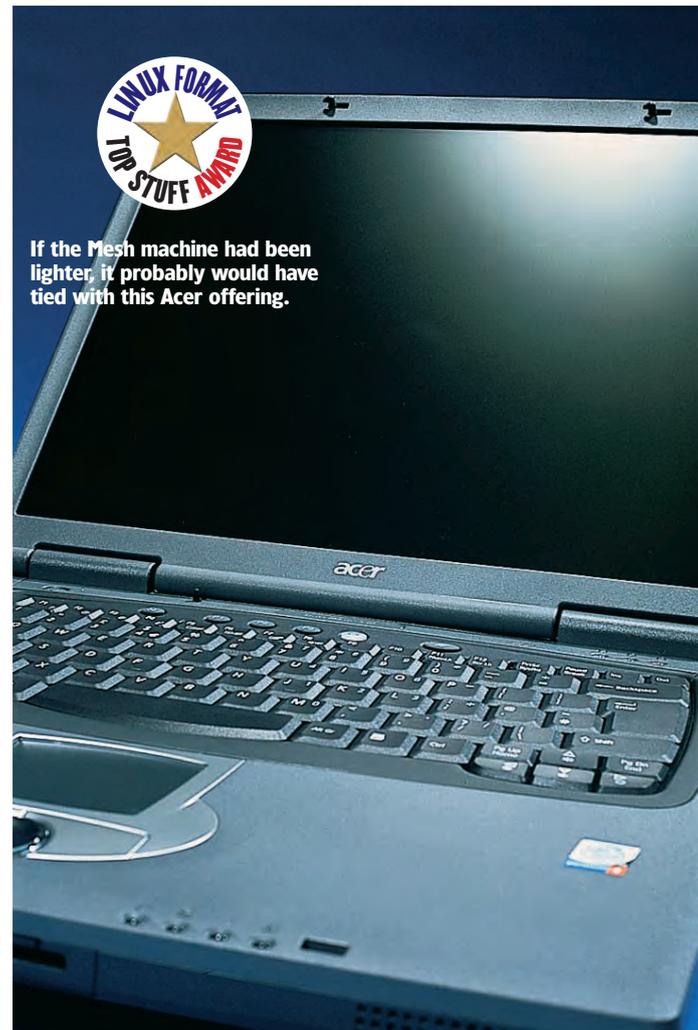
of options. The Lifebook, particularly, is very light, which will be a major boon for regular travellers. Both have excellent screens that work very well



Choose one or two batteries with Sony's VAIO.



Fujitsu is well worth a bit of extra cash.

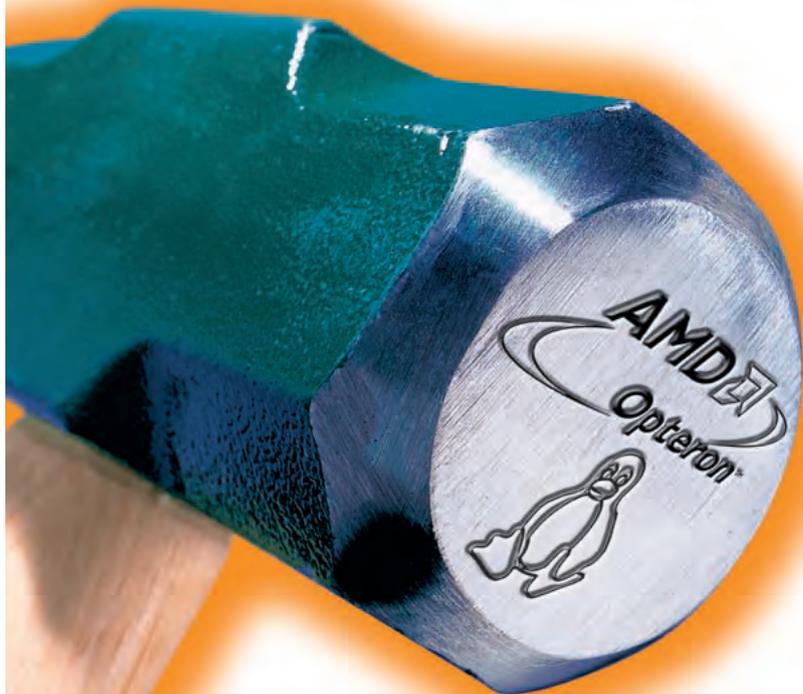


in low-light conditions, also. If you're looking for a sure winner that is guaranteed to please, Evesham should be your first port of call – the Voyager was an all-round solid laptop, that I feel would have led the pack if it had higher specifications. As it stands, the Evesham unit is less expensive than its rivals here, and that alone might be enough to recommend it to some!

The only reason we felt unable to give the Mesh laptop a Top Stuff award was because of its weight – a laptop, after all, should at least be relatively easy to carry short distances without risk of pulling a muscle. It is, however, a marvellous laptop – Mesh has excelled itself by combining power that would put most desktops

to shame at a price that makes me want to buy one right now.

Overall, the Acer comes top of the group, and not only because it has such superior performance for a small laptop. Design-wise, the laptop is a dream to use and is almost approaching the status of 'fashion accessory'! The performance is incredible for a mobile CPU, and shows that Acer must have put a lot of thought into its design beyond looks, and with three USB ports and a CD-RW, there's lots to be had for the money. So, congratulations to Acer – the TravelMate is a great laptop at a great price, and it thoroughly deserves the *Linux Format* Top Stuff award. [LXF](http://www.linuxformat.co.uk)



After much hype and anticipation, AMD is finally launching the first of its x86-64 CPUs. Previously code-named “Sledgehammer”, Opteron is now out in the wild. Paul Hudson shows us what it is, how it works, and how it aims to change the world...

AMD announced the x86-64 specifications all the way back in August 2000, and finally, in April 2003, the first x86-64 CPU was released. *Opteron*, previously known only as *Sledgehammer*, is the name for AMD's 8th-generation enterprise-class processor for servers and high-end workstations. Between the original announcement and the product release, x86-64 was renamed AMD64, and the two are interchangeable – although the official term is AMD64. The unique selling point of AMD64 is that it does what its name suggests: 64-bit operation combined with the same x86 technology that companies have relied on for so many years now. Not only does this allow adopters to preserve their investment in 32-bit technology, but it also allows them to ease the transition to the more advanced 64-bit model – one computer, powered by Opteron, can handle both 32-bit and 64-bit applications seamlessly.

If you've been watching the 64-bit market, you'll have seen both the Itanium and the Itanium 2 being

launched by Intel and HP since the AMD64 specification was first aired, and you might be forgiven for wondering when AMD will get its product out the door. Luckily for AMD, both the Itanium (*Merced*) and its successor the Itanium 2 (*McKinley*) – collectively known as the Itanium Processor Family, or IPF – have both received a cool reception from the industry, with the processor family even earning the title “Itanic” from several wry analysts. The problem with Itanium is that it wholly replaces x86-32 with a brand new architecture that bears little to no similarity to x86-32. While it can run in 32-bit compatibility mode, it requires special emulation to do so, and so performance suffers. Read the box *The Rise And Fall Of Itanium* on page 53 for more information.

The Opteron was officially launched on April 22nd with much fanfare at a special event in New York, which we attended (see box, *The Launch* on page 55). Now in the hands of consumers, it's time to prove its worth – will it have trouble like Itanium, or is Opteron the future?

Hammer Time!

cover feature



Why should companies switch to Opteron? Before you can truly understand the benefits of AMD64, it's important to first understand how we got to the current state of play.

x86's heritage

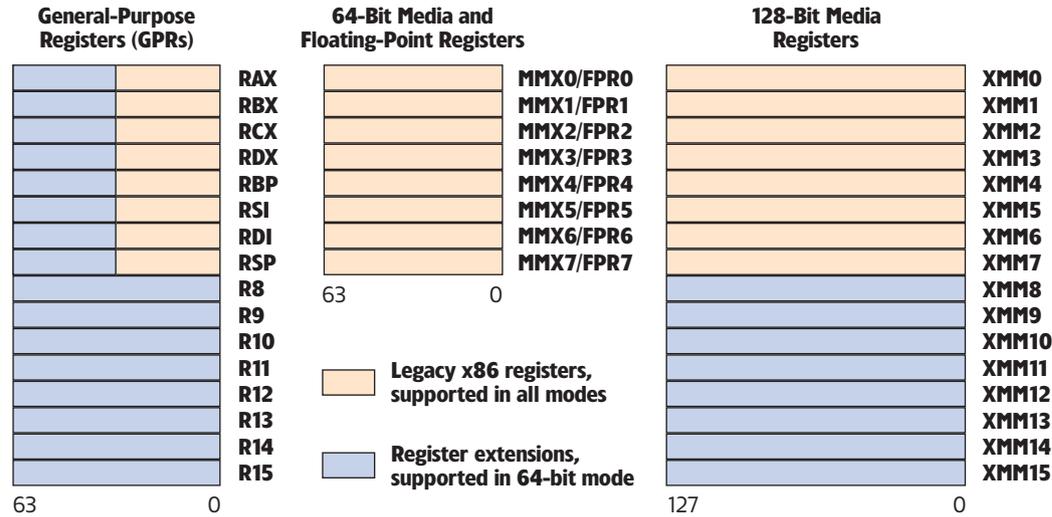
Unlike many other processor architectures, x86 has a very clearly defined history. Early on there was the 4004, then the 8008, then the 8080, and finally 8086. 8086 was Intel's first 16-bit CPU, capable of addressing up to 1MB of memory. The 286 followed (with a relatively uninteresting 80186 in between), introducing protected mode memory access, allowing lots more memory to be accessible to new programs without breaking existing ones. The 386 followed, introducing 32-bit flat addressing, which promised to do away with memory segmentation altogether. AMD and Cyrix started competing with Intel by manufacturing Intel-clone 386 CPUs.

The 486 was a key release, making several important changes to the architecture. Firstly, the 486 was the first to include level 1 cache on the chip, reducing the number of times main memory needed to be accessed. Furthermore, the later models had a ubiquitous maths co-processor (x87) built in, adding 80-bit floating point mathematics capabilities to the CPU which, at the time, made *Doom* run considerably faster!

The Pentium CPU, so named as to allow Intel to trademark the brand, made few changes to the x86 architecture. However, it did introduce MMX in January 1997, which was the first major change to the x86 instruction set in a decade and added 57 new instructions designed to improve multimedia operations. However, at this point, both AMD and Cyrix were actually creating better CPUs than Intel. For example, AMD's K5 CPU was capable of executing six out-of-order instructions at once, allowed register renaming, and also speculative execution – three things not seen in Intel's chip until the Pentium Pro was released.

The Pentium Pro was the first CPU to be specifically designed for 32-bit code. The 386, 486, and Pentium CPUs could *run* 32-bit code, but they were designed to run both 16-bit and 32-bit code in equal amounts. The Pentium Pro also included an integrated level 2 cache, which greatly increased both the speed and the cost of the CPU.

Owing to a mixture of its price and the fact that it performed 16-bit operations very poorly at a time when Windows 95 still relied heavily on 16-bit code, the Pentium Pro didn't sell too well for Intel, and it took the company 'til May 1997 to produce the Klamath CPU, known on release as the Pentium II. The PII corrected most



As seen above, the AMD64 architecture extends all GPRs to 64 bits and also increases the amount of GPRs to 16.

of the problems with the Pentium Pro, and was generally a faster and better CPU than the Pentium MMX. AMD's clone of the PII was first the K6, which again outperformed the PII clock-for-clock, and also later evolved into the K6-2. This new CPU included AMD's new 3DNow! extensions to the x86 architecture, which offered another boost for multimedia applications, and was much more popular than the MMX extensions.

With the Katmai CPU, released as Pentium III, Intel added 70 more multimedia instructions to x86, although left much of the CPU entirely unchanged from the Pentium II. This turned out to be a big mistake when AMD released the Athlon CPU, a much-enhanced version of its K6-3 CPU, and managed to trounce Intel's CPUs quite drastically. In order to turn around its falling fortunes, Intel worked hard to entirely redesign their x86 platform, and in the end released 1.4 and 1.5GHz versions of the Pentium 4. The Pentium 4 is, architecturally, light-years ahead of the Pentium III, however in order to take advantage of its new architecture, significant optimisation is required – optimisation that, at the time of launch, was just not available. As a result, the P4 took quite a bit of hammering from tech sites for the first six to nine months of its existence. However, the P4 does have two advantages in its favour: it can ramp up clock speeds very quickly and to very high levels, and it also came with SSE2: 144 new instructions for x86 to further enhance multimedia applications.

This is where the tale of 32-bit processing now ends. Intel's 64-bit Itanium CPU is not part of the x86 family, which left a gap open for another company to squeeze into if they were capable of producing an x86-compatible 64-bit CPU. In August 2000, this materialised in the form of AMD64, designed and created by Intel's leading x86-32 rival: AMD.

Hammer Time!

The Pentium 4 CPU, despite being much faster than the 386, still bears a great resemblance to its ancestor. For example, the main set of x86 instructions (that is, not the MMX or SSE/SSE2 [collectively called XMM] additions) remain the same. They both have the same number of registers, and they both use the same style of

maths co-processor. While other architectures have been more willing to break backwards compatibility, x86-32 has mostly seen a history of tweaks and speed hikes.

When Intel came to design its own 64-bit CPU, it threw the x86 architecture out and started afresh. Unfortunately that has stifled adoption of the Itanium architecture because most companies have very large investments in the x86 CPU architecture in both the hardware and also software that runs on it.

AMD took a different option, and decided to create a CPU *evolution* as opposed to Intel's *revolution*. The Hammer family of CPUs, code-named *Sledgehammer* for servers and *Clawhammer* for workstations, were designed around a new AMD64



“Linux Enterprise Server for AMD64 enables customers to combine the stability and security of Linux with the performance enhancements available only through the 64-bit architecture”

RICHARD SEIBT, CEO
SUSE LINUX AG



Operating Mode		Operating System Required	Application Recompile Required	Defaults		Register Extensions	Typical
				Address Size (bits)	Operand Size (bits)		GPR Width (bits)
Long Mode	64-Bit Mode	New 64-bit OS	yes	64	32	yes	64
	Compatibility Mode		no	32		no	32
				16	16		16
Legacy Mode	Protected Mode	Legacy 32-bit OS	no	32	32	no	32
	Virtual-8086 Mode			16			
				Real Mode	Legacy 16-bit OS		16

The Opteron's variety of modes support full backwards compatibility, but 64-bit Long mode is the preferable one.

Opteron

architecture, which retains complete compatibility with 32-bit code whilst adding the option of 64-bit computing. Moreover, it performs 64-bit and 32-bit operations natively – no emulation is required.

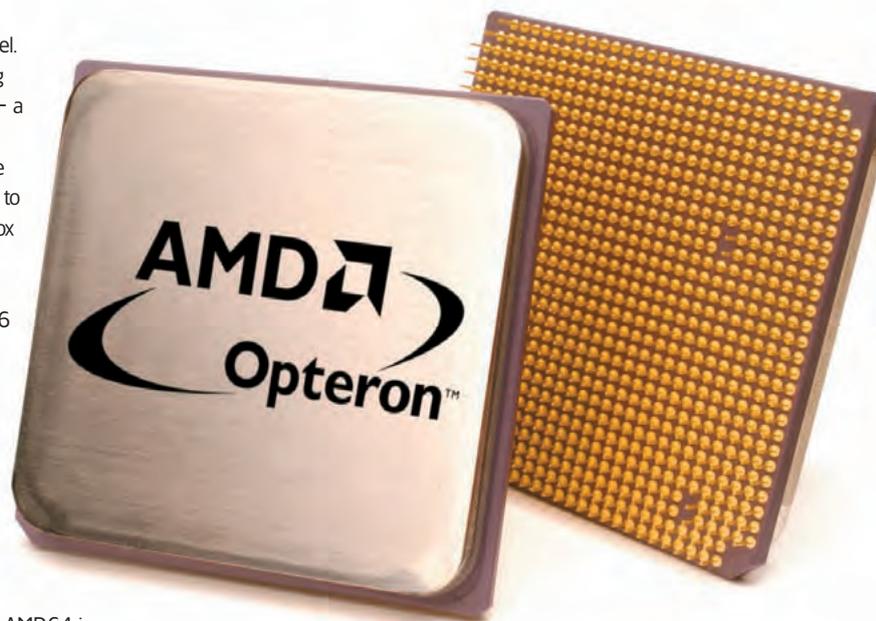
So, how does AMD64 handle both 32-bit code and 64-bit code? The architecture specification includes two modes of operation: “Long Mode”, and “Legacy Mode”. Legacy mode AMD64 CPUs operate as existing x86-32 chips do – they support 16-bit and 32-bit applications and also support 16-bit and 32-bit operating systems. Processors that implement AMD64, of which the Opteron is the first, boot into Legacy Mode by default, and therefore are entirely compatible with existing software – by default, the Opteron looks like a 32-bit CPU to software.

Long Mode, though, is where the magic happens, and can be split into two submodes: 64-bit Mode and Compatibility Mode. Both of these submodes require a 64-bit operating system to run; the differences lies in the applications: 64-bit Mode only supports 64-bit applications running on a 64-bit operating system, whereas Compatibility Mode supports 32-bit applications running on a 64-bit operating system.

Compatibility mode means that users can switch to a 64-bit operating system immediately whilst continuing to use their existing software – the operating system will perform better, which will yield some performance

boosts at the application level. However, it is of course Long Mode that is the ideal goal – a pure 64-bit processing environment. Apart from the inherent advantages offered to 64-bit programs (see the box *Why 64-bit Computing Is Needed*), AMD took the opportunity to revise the x86 platform to help bring it up-to-date a little. However, because of the need to retain backwards compatibility when running 32-bit programs, these new features are only available in 64-bit mode.

One of the most highly anticipated revision made in AMD64 is the addition of new registers. AMD has doubled the number of general purpose registers (GPRs) to sixteen, and also doubled the number of XMM registers to sixteen also. These 32 registers are, of course, lengthened to 64 bits each. Having more registers is key to improving performance, particularly when there were so few to begin with. Registers store small amounts of information directly in the CPU and provides instant access to that data – it doesn't need to be moved from cache or main memory. More registers mean more data can be held on the CPU, which means less shuffling around is required to get results of calculations. However, some are already saying that 16 GPRs aren't enough –



RISC architectures often have registers numbering in the hundreds, thereby allowing complete programming loop unrolling to take place inside the registers. For the time being, though, having just 16 registers is likely to make quite a drastic speed improvement by itself once programs are recompiled to use them – almost certainly a minimum of 10%.

Another key feature available exclusively in 64-bit Mode is instruction pointer-relative data addressing, which allows programs to address locations anywhere in the 64-bit address space relative to the instruction pointer.

These new features are only available in 64-bit Mode, where applications have specifically been compiled to make use of the functionality. However, fortunately it generally does just take a recompile to get the extra functionality offered by 64-bit Mode.

With AMD64, AMD took the opportunity to incorporate both the XMM extensions, SSE and SSE2. As mentioned already, there are twice as many XMM registers in the AMD64 specification as in Intel's own chips, which will mean a further increase in performance. Of course, the biggest XMM performance increase will be the XMM instructions themselves: thanks to Intel pushing so hard to have them adopted, these are now used extensively in multimedia programs such as *3ds max* and *Quake 3*. As a result, just by incorporating the XMM

instructions, the Opteron is going to be able to leap ahead in performance when compared to the Athlon chip.

When operating in 64-bit Mode, the Opteron switches to using a new flat memory segmentation model. Segmented memory was previously used as a method to allow operating systems to isolate programs from each other, thereby increasing reliability. However, because most modern operating systems do this all in software, there is a lot of wasted space in the x86-32 architecture. This has been stripped out for 64-bit Mode Opterons, which allows new 64-bit operating systems to have much simpler code to handle memory management.

So, not only does AMD64 allow much larger data to be worked on, but it also adds XMM support, cleans up a lot of very old legacy operations that no longer have a place in today's environment, and does so in such a way as to provide complete compatibility with existing 32-bit software.

AMD64 and Linux

Thanks to the hard efforts of the community, Linux has supported AMD64 for some time now. The new 64-bit kernel is based on the existing i386 port, and a lot of effort has been made to ensure that the new features of the Opteron are used to best possible advantage. SuSE and AMD have been working very closely to develop the new software, and

Learn More

Links and treeware

If you're interested in learning more about AMD64, AMD provide an excellent set of books under the banner, **AMD AMD64 Architecture**. If you're in the US or Canada, you can order a complete set of these books for delivery to your home. Readers outside of North America have to settle for the online version, available from the AMD site.

Though it hasn't been updated for some time, www.x86-64.org is still worth reading – there's quite a bit of info available there on how to write code for AMD64, and also how the effort to make AMD64-compatible versions of various free OSes is going.

AMD's website contains a great deal of info on Opteron – how it works, from where it's available, and such. It is



First-class reference for the IPF and kernel development in general.

of course a fairly one-sided opinion, and you'll find a better representation of Opteron by running a Google search.

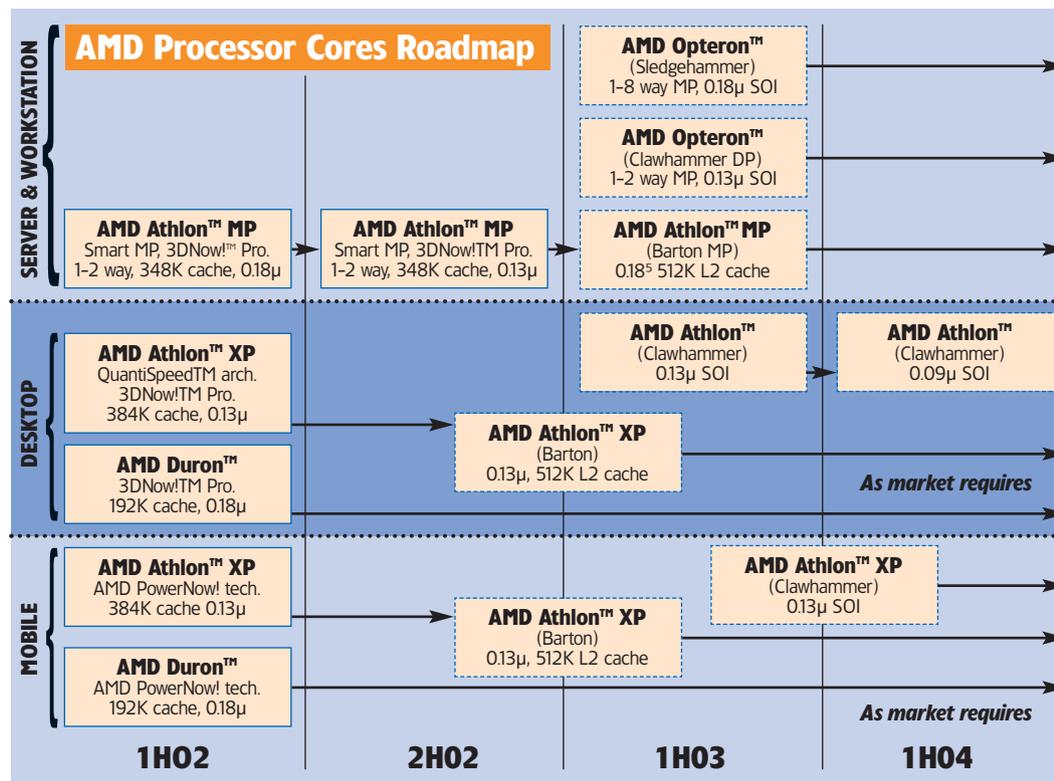
If you're interested in learning more about IA-64 and Itanium, the book *IA-64 Linux kernel* (Mosberger and Eranian, Prentice Hall ISBN 0-1306-1014-3), documents how the team at HP ported Linux to the IA-64 platform, and includes excellent information on the IPF and lots of insight into the Linux kernel itself. It's a bit pricey though.

naturally, it's released as GPL as part of the 2.5 kernel. Customers of UnitedLinux, such as companies who bought SuSE's excellent *Linux Enterprise Server 8* product, will be able to benefit from this code already, as the UnitedLinux kernel has many 2.5 backports, including AMD64 support. As such, SLES was the first server operating system available to fully support Opteron's architecture.

Red Hat's *Enterprise Server* product does not support AMD64 at the time of writing, and it seems they have plans only to integrate it into the next version of the software. MandrakeSoft, for its part, launched Mandrake Corporate Server 2.1 for Opteron on the day of the Opteron launch, although we've yet to have time to review it fully. Red Hat has promised us that a beta will be available by the time you read this.

AMD64 & other OSes

In late 2002, Microsoft delivered a development release of AMD64 Windows to several of its industry partners, and has recently announced that it is developing native 64-bit versions of its Windows XP and Windows Server 2003 operating



systems for the Opteron and Athlon 64 platforms. Specific launch dates are a little vague, with MS saying that beta releases are expected in the middle of 2003. It's important to remember that the support of

Microsoft and Windows will provide an important boost to the Hammer architecture, and could spur large-scale adoption of the CPU.

Both FreeBSD and NetBSD have work underway to support the AMD64

architecture, with NetBSD already having the port ready for public use.

So, with Linux, Windows, and two of the three major BSDs developing support for Opteron, it is quite likely that adoption will be swift. With SLES

The Rise and Fall of Itanium

Lack of backwards-compatibility

The original Itanium CPU was launched on May 29th 2001, although it lacked fanfare. Indeed, some contend that it was never really launched at all, because it wasn't used much, had little media coverage, and was replaced by its successor a little over a year later.

Both the original Itanium and the Itanium 2 form the Itanium Processor Family (IPF), and constitute an entirely new 64-bit CPU architecture, named IA-64. With IA-64, Intel partnered with HP to produce a new 64-bit system that would be the best CPU they could design, and this involved coming up with a new processing paradigm, known as Explicitly Parallel Instruction Computing (EPIC). Borne from the Very-Long Instruction Word (VLIW) paradigm, EPIC relies on compilers to handle instruction scheduling in order to minimise CPU complexity.

This "feature" is Itanium's biggest advantage and also its biggest disadvantage. Traditional x86-32 CPUs use peep-hole optimisation at run-time to make code execute faster – that is, as instructions are loaded into the CPU, it checks to see whether it can reorder

things and/or speculatively execute code to improve speed. However, it does so using a "peep-hole" – it has a very narrow view of the code in order to not slow things down too much and also to keep the amount of instruction scheduling-related silicon down. EPIC relies heavily on the fact that during the compilation stage, the compiler is able to have knowledge of all parts of a program at the same time – what variables it has, when they will be accessed, *how* they will be accessed, etc. As such, it is able to generate far more optimised code than CPU-level peep-hole optimisation can. EPIC allows compilers to specifically mark chunks of code that can be executed in parallel, which can provide a major speed boost.

However, the problem relies in the fact that, as you can imagine, Itanium relies heavily on a *highly-optimised* compiler – without one, Itanium CPUs will certainly not perform at their best. As such, some people have rechristened EPIC to stand for "Expect Perfectly Intuitive Compilers", because compilers must be very advanced to properly take advantage of Itanium's architecture.

Another drawback is that because Itanium is so different from x86-32, it requires an emulation layer. As discussed earlier, this emulation layer isn't fast by any means despite being implemented in hardware (although a new, software based emulation system has just been announced which is claimed to increase performance of emulated code). This means that although companies opting to switch to Itanium can retain backwards compatibility with their existing applications, they must accept very low performance when running emulated software – satisfactory performance can only be achieved by changing software along with the hardware. Furthermore, the compatibility that exists for x86-32 is only for applications: Itanium must run on a 64-bit operating system.

Programming for Itanium is quite a jump from programming for x86-32 – it has 128 integer registers and 128 floating-point registers, and also has a much more complicated assembly language. While there's no doubt that Itanium is certainly a great chip in its own right, it currently has no place in the x86-32 world. It's possible

that Madison, the expected third member of the IPF, may change this by increasing the performance of x86-32 code, however, there is one other option for Intel: Yamhill.

Yamhill, the supposed codename for a rumoured Intel project to produce an AMD64-compatible CPU, is right now unconfirmed and merely circulating by means of Chinese whispers. However, given that AMD allows licensing of AMD64, Intel would be crazy to not be developing an Opteron competitor just in case the technology is a big hit. Of course, that would leave Intel following in AMD's footsteps for the indefinite future – probably not a situation Intel want to be in. However, if companies move over to 64-bit for new computers – which is seeming increasingly likely as the business benefits become more apparent – Intel might lose out either way.

A lot rides on Madison: if Intel can bring x86-32 performance up to speed, then release a fast and optimised compiler for little cost to take advantage of the technologically clever EPIC system, Itanium might have life in it yet.

Opteron

« 8, SuSE has worked very hard to make sure companies wanting to switch to Opteron have a reliable, well-supported Linux platform available.

Can't touch this!

Hot topics right now are: "how fast can Opteron users expect their new computers to be", and "will it perform better in 32-bit mode than the IPF"? The second question is easier to answer, so we'll cover that first. Put simply, the answer is: almost certainly. Because Itanium shares no relation with the x86 architecture, it's difficult but not impossible to emulate a 32-bit CPU with software only. As a result, Itanium defines a special interface to translate x86-32 instructions into Itanium (IA-64) instructions that can be (and currently is) implemented in hardware. This translator, known as the Intel Value Engine, has the tricky job of converting x86-32 instructions into IA-64 instructions, which can then be executed. As you can imagine, this is quite a speed hit.

The Opteron, however, defaults to 32-bit mode by default, and becomes a particularly fast Athlon with SSE and SSE2 enabled for all intents and purposes. That is, it executes 32-bit instructions at the same speed or faster than a clock-matched Athlon.

With regards to the speed users will find their shiny new Opterons running at, it depends on how you look at the equation. Firstly, if you use applications that make use of SSE and



"I like to say they [Intel] are big, strong, and wrong... they can't stop us!"

JERRY SANDERS

Founder and chairman of the board, AMD

SSE2 (that's pretty much any 3D software/game), you'll see an immediate and considerable speed boost. Once software is specifically compiled for 64-bit mode, you'll see a further 10-15% speed improvement from the use of the new registers, although this could rise much higher for computationally-intensive code.

One factor that might muddy the issue is AMD's choice of model rating for the new Opteron. Each Opteron will be given a three-digit number to represent its speed and scalability. The first digit will either be a 1 for uniprocessor boxes, 2 for two-way systems, 4 for four-way, or 8 for eight-way. Digits two and three are a relative performance figure, and will start at 40. So, an Opteron 140 will be the slowest uniprocess Opteron machine, whereas an Opteron 141 will be faster than the 40 (faster by an unknown amount), and an 899 will presumably be the fastest eight-way Opteron machine available. Confused yet?

The reason, so we've been informed, that the second two digits started at 40 is purportedly so that people don't confuse the numbers as a three digit clock speed. Precisely how numbers such as 866 are don't look like a clock speed, we're not sure. However, this is the naming scheme that is to be used, and hopefully it won't confuse potential customers too much.

What's next?

The Opteron is the first chip to be built on the AMD64 architecture, but it almost certainly won't be the last. As Opteron is designed to be used in servers and workstations, AMD are working on a new version of its Athlon XP desktop CPU, dubbed Athlon 64/Clawhammer, that will also make use of the AMD64 architecture. This will allow AMD to push 64-bit computing to everyone. However, there are no plans as yet to discontinue production of the x86-32 Athlon XP or Duron chips, despite the

fact that a Clawhammer CPU would be able to emulate (and outperform) each of them.

There are also plans afoot to launch a Mobile Athlon 64 at the same time as the desktop Athlon 64, which will presumably contain much of the same power-saving functionality seen in the existing Mobile Athlon XP.

In the first half of 2004, AMD are slated to release 0.09 micron version of the Opteron (code-named *Athens*), the Athlon 64 (code-named *San Diego*), and the Mobile Athlon 64 (*Odessa*), which should allow them to continue increasing clock speeds.

Whether or not there will be an Opteron 2 out with a revamped architecture is as yet unknown. Intel are still making quite drastic changes to the Itanium, trying to get it right. With the Opteron, due to the fixed x86 heritage behind it, it's more likely that the interface will stay the same, with some parts being shuffled and tweaked internally to increase performance.

Why 64-bit computing is needed

But what's this I've heard about a diminished cache-hit rate?

We've had 32-bits for such a long time now, that it might seem a bit unnecessary to double things. 32-bits means 32 binary digits, so 32 1s and 0s are used to represent numbers. If we have a 32-bit unsigned integer (that is, it can only be positive), it can hold numbers from 0 to 4,294,967,295 (2^{32}). Now, because memory is allocated using bytes, it is also referenced using bytes – when a program wants to read from a byte, it simply looks up its number. As seen, numbers max out at 4 billion or so, which means that the maximum byte of memory that a 32-bit program can reference is number 4,294,967,295. This means that the range of memory that can be accessed is 4,294,967,296 bytes, which is divisible by 1024. So 4,294,967,296 bytes is 4,194,304 kilobytes, which is 4096

megabytes, which is 4 gigabytes. What this means is that the maximum amount of memory a 32-bit program can reference is 4GBs – sounds like a lot, no? Well, consider the topic of databases. Large companies buy very powerful servers designed to hold their entire database in RAM for the fastest access – in situations like this, 4GBs soon runs out.

Switching to 64-bit numbers squares the maximum value of integers, so the largest amount of memory that can be addressed becomes 18,446,744,073,709,551,616 bytes – 16 Exabytes! Of course, this is all virtual – there are other limitations generally involved when things are implemented in reality. With the Opteron, 40 of the 64 bits can be used, which gives 1 Terabyte (a thousand gigabytes) of memory. Furthermore, the

Opteron has a 48-bit virtual memory address space, allowing 256 Terabytes to be accessed.

So, lots more RAM is made available – but what else does switching to 64-bit help? One simple advantage is in the field of encryption, where very large numbers need to be processed quickly. 32-bit CPUs can perform 64-bit mathematics (such as long long integers in GCC), however the operation needs to be split into two for calculation, then recombined. 64-bit CPUs can natively perform 64-bit mathematics, and so avoid this speed hit.

As you can see, there's lots to be gained by switching to 64-bit – and that's before we consider maths-intensive code such as digital content creation! However, are there any downsides? Well, there is one: switching to 64-bit processing

does mean a slight decrease in the cache-hit rate. That is, because 64-bit data is twice the size of 32-bit data, less information can be stored in cache, which means it's less likely that requested data will be found in the cache (a *cache hit*). Digital Equipment Corporation found that this increase will cause a 5% performance penalty in programs that were similar apart from their bit size. To get around this, 64-bit CPUs support 32-bit data operands in order to allow programmers to keep their operands short when possible to minimise this penalty. AMD64 defaults to 32-bit data operands when working in 64-bit mode, and only uses the larger numbers when specifically instructed – this provides the best performance solution, and is easy for compilers to make use of.

Glossary

3DNOW! AMD-designed multimedia extensions to x86-32

ATHLON AMD's x86-32 CPU brand

ATHLON 64 AMD's proposed AMD64 desktop CPU brand

BIT Binary digit, a 0 or a 1

BYTE Eight binary digits

CISC Complex Instruction Set Computer, uses complex instructions that are often decoded to RISC instructions

CLAWHAMMER Codename for AMD's desktop 64-bit CPU, Athlon 64

CPU Central Processing Unit, forms the core of your PC's processing abilities

EPIC Explicitly Parallel Instruction Computing, allows compilers to specify which instructions can be executed in parallel. An extension of VLIW.

FPU Floating Point Unit, handles mathematics-specific capabilities of your PC. Usually built into the CPU

GIGABYTE 1024 Megabytes

GPR General Purpose Register, a register than can be used for general programming

IA-64 Intel's Itanium architecture

INSTRUCTION Simple binary digit that translates to one CPU operation when called

INSTRUCTION SCHEDULING the act of re-arranging instructions in a non-destructive manner to improve performance

IPF Itanium Processor Family, includes Itanium and Itanium 2

ITANIUM Intel's IA-64 implementation

KILOBYTE 1024 Bytes.

MADISON Codename for Itanium 3

MCKINLEY Codename for Itanium 2

MEGABYTE 1024 Kilobytes

MERCED Codename for Itanium

MMX Intel-designed multimedia extensions to x86-32

OPTERON AMD's server and workstation 64-CPU CPU

REGISTER A storage area directly on a CPU able to hold a small amount of data

RISC Reduced Instruction Set Computer, uses simpler instructions for faster execution at the expense of larger software

SLEDGEHAMMER Codename for AMD's Opteron

SSE Streaming SIMD Extensions. More Intel-designed multimedia extensions to x86-32

SSE2 Streaming SIMD Extensions 2. Even more Intel-designed multimedia extensions to x86-32

SIMD Single Instruction, Multiple Data. The technique of allowing one CPU instruction to act on multiple registers

TERABYTE 1024 Terabytes

VLIW Very-Long Instruction Word, method for allowing compilers better control over execution, while making CPUs simpler

X86 Common name for 386-compatible CPUs

X86-32 32-bit x86 CPU architecture

X86-64 Original name for AMD64 64-bit architecture

AMD64 64-bit x86 CPU architecture

X87 The mathematics architecture used in x86 PCs

XMM Internal name for SSE and SSE2 combined

Note: recent standardisation has changed the meaning of kilobytes, megabytes, etc. In order to preserve ease of reading, we have used the traditional usage of megabytes: 1024 kilobytes.

Conclusion

The Opteron is the biggest step forward AMD has ever made, and consequently also its biggest risk. However, having worked so hard these past years, AMD hasn't rushed to market with the Opteron – it has put a lot of thought into how to give customers the most value for money, and has also worked hard to recruit the best partners to help make AMD64 a success.

Will Opteron be a success? That's a hard question to answer, given companies' general dislike of changing what works. The Itanium fell into this trap by breaking native compatibility with existing popular hardware, but the Opteron doesn't do that – indeed, users can switch to Opteron immediately and the absolute worst-case scenario is that they'll continue to run their 32-bit systems and therefore only get a particularly powerful 32-bit CPU.

The one thing that might endanger the success of Opteron is that AMD has long been considered "almost as good

as Intel, but much cheaper". While this belief has been true at times in the past, it would likely damage Opteron sales if it were seen as a cheap version of Itanium, or "the poor man's 64-bit CPU". The server market carries with it a much higher margin level than the desktop market, which should in theory allow AMD to reach profitability quickly if they can keep its prices at a good level. However, if AMD somehow manage to sink back into its old persona of being billed as the "value option", this increase in revenue will certainly be jeopardised.

It's our opinion that AMD64 has a strong proposition for both the consumer and enterprise market. Long term popularity is likely to be dictated by the availability of software to take advantage of it, but that doesn't seem much of an obstacle in the Linux world at least. Time will tell, and the stakes may be high for AMD, but the Opteron and the AMD64 chips that will follow certainly seem to have a future. 

The AMD64 Launch

Paul Hudson reports from New York, Tuesday 22nd April.

I flew to NY to the Opteron launch event, and were mightily impressed – AMD has spared no expense in getting Opteron noticed by the media at large, and are eager to make sure that everyone who wants to know about Opteron can get lots of information.

Many VARs attended, including MSI, Nvidia, IBM, Oracle, Microsoft, and others. Crucially, SuSE, Red Hat, and Mandrake were all there, and more than willing to disclose lots of information about their plans for Opteron. As such, Opteron has a lot of support from launch – several of the speakers got up and declared that Opteron was pretty much the best thing that had happened to the x86 industry since, well, ever. A recurring theme was "It only took us X days to port to AMD64". For example, IBM ported DB/2 to Opteron/Linux in just two days, despite it having 10,000,000 lines of code – quite an achievement!

Perhaps the best quote from the trip was from Jerry Sanders himself, overheard at a private VAR dinner the night before Opteron launched: "With Opteron, AMD is going to change the world, and unlike Steve Jobs we're going to change the real world." With unsubstantiated rumours flying that AMD and Apple might have something in the works (OS X for Opteron?!), this seems like quite a risky comment to make!

Both Hector Ruiz (AMD CEO) and Jerry Sanders (Chairman of the Board), gave an excellent presentation making the situation quite clear: AMD64 is the way forward. Some laughter was heard during the media Q&A session when one reporter questioned whether there had been any "poisoning of the wells" from Intel or AMD. Humorous, yes, but less so considering that both myself and several other reporters there questioned company representatives at the launch, asking what Intel thought of them supporting AMD. The (scary!) response was more often than not that they had received a call from an Intel representative before they came to the Opteron launch asking the vaguely threatening question about whether they really wanted to go, and that some vendors had pulled out almost certainly as a result of the call. Co-incidence? Possibly – we'd like to think so. Eager to give Intel the chance to say it was all just a coincidence, we gave them a call "All we can say is that the reports are speculation and rumour" came the response. I'll leave you to make your own mind up.

Gene Kim, VP of Global Sales and Marketing, RackSaver said: "We're very

proud to announce that based on our Quadrix-64 platform, our 4-Way Opteron, we have reached the highest throughput score ever achieved for a 4-way system as measured by the TPC council ever. Not only that, but that score based on that same platform also has the lowest price and performance ratio, beating even Dell!"

At the launch, Opteron models 240, 242, and 244 were announced. Clearly heralding the start of mass confusion regarding this naming convention, AMD representatives from the Strategic Marketing department were unable to tell us precisely how fast these CPUs were in Gigahertz, although continued questioning from other AMD sources led to the answer 1.4, 1.6, and 1.8GHz respectively – not terrifically high speeds, but it leaves the Opteron a lot of space to grow.

On the subject of performance, early reports are already coming in of Opteron blowing competition out of the water. AMD's official results show model 444 Opterons are 37% faster than Xeon 2.8GHz running SPECint, despite the massive clock speed difference. Microsoft's own MMB benchmark results gave model 424 Opterons a score of 15,520 in real world tests, compared to just 13,200 for the next best (4 2GHz Xeon MP's). The same test had two model 244 Opterons trouncing four 1.6GHz Xeon MP's – wow!

I had the chance to talk to people from Red Hat, Mandrake, and SuSE – both were strongly behind the Opteron movement. Red Hat intend to support Opteron in all three of its Enterprise Linux platforms in version 3 of the family. Importantly, Red Hat said they will allow users to upgrade to an Opteron-enabled v3 system as long as the company's subscription is paid up. As to whether Red Hat will ship the x86-32 and AMD64 versions in the same box a la SLES 8, "That decision has not yet been taken". Further uncertainty arises from the fact that the current release of Red Hat's Enterprise Linux isn't supported for use on Opteron systems. Nick Carr, Red Hat's Enterprise OS Marketing Manager, informed us that while customer support will dictate its Opteron policy on existing releases, v3 of its Enterprise Linux platform will definitely support the chip.

SuSE didn't really have much to say, which was odd considering nearly all of the Linux machines there were running SLES8. Mandrakesoft gave us an early copy of its new Opteron Corporate Server 2.1, and we're looking forward to reviewing that in a future issue.

What on Earth is... THE HURD?

Richard Drummond answers your questions regarding the operating system kernel whose place Linux usurped.

»» What is this free operating system called the Hurd?

The Hurd itself isn't actually an OS – it's an operating system kernel, just like Linux is a kernel not an operating system. And, like Linux, the Hurd is designed to be a kernel for the GNU Operating System.

»» I don't get the distinction. What's a kernel and what's an OS?

Simply put, an operating system is the set of software which make a computer easy to use and easy to write software for. It controls and provides access to hardware resources, it furnishes the environment which allows you to run other programs such as application software, and it provides the interface with which the user interacts with the machine (such as the windowing systems favoured

by modern operating systems). The kernel is the central core of an operating system and provides the really essential services upon which the rest of the OS and other software depends – like process management, memory management, a file-system and so on.

»» OK. So what's the GNU Operating System?

The GNU OS is the free reimplementation of the Unix operating system that was begun by Richard Stallman and the Free Software Foundation (FSF) in 1983.

The GNU OS and the Linux kernel are the two principal components in every Linux distribution. For example, the *bash* shell, the *gcc* C compiler, and Emacs are all part of the GNU OS. When you use Red Hat or Mandrake or any other current Linux distribution, you are using the GNU OS – whether you

realise it or not. So great is the role of GNU software in forming a usable Linux-based operating system that the FSF maintain that Linux distributions should properly be called GNU/Linux. However, that's a debate for another day.

»» So the Hurd is meant to replace the Linux kernel?

You could say that, yes. But, actually, the Hurd predates the Linux kernel. It was always the intention of the FSF to create a kernel for their free Unix replacement, but for one reason and another very little progress was made early on. The Hurd project – as we know it today – wasn't begun until 1990 and a bootable Hurd kernel wasn't available until 1994.

The first public release of Linux was in 1991 – after the Hurd had been announced. Many people who started using the Linux kernel back then regarded it as stop-gap measure to get the GNU OS up and running until the Hurd was ready for release. Linux was considered a toy and thought not to be portable to



architectures other than the x86-based PC, a platform dismissed by hackers who grew up using Unix on 'real' computer hardware. Eventually, however, Linux matured and stole the limelight away from the Hurd. This is partly attributable to the ever-increasing popularity and power of the PC platform – but it is also due to the fact that developing the Hurd turned out to be a lot more difficult than anybody expected.

»» So, why bother with the Hurd now that we have Linux?

Why not? The Hurd project was in full swing by time the Linux kernel was considered mature, and it didn't make sense to the Hurd developers to abandon all that work. Besides, the Hurd is fundamentally different than the Linux kernel and from any other mainstream operating system kernel. It's a good idea to have two projects trying to solve the same problem using radically different approaches. The most important reason for the existence of the Hurd project, however, is that the developers enjoy working on it. It's fun.

»» What makes the Hurd so different from Linux?

It's a matter of architecture. The Linux kernel is a traditional monolithic design, whereas the Hurd is built on a micro-kernel architecture. This distinction is a technical one.

Typically, a kernel provides its services by executing software routines in a special mode of the CPU that it is running on. This kernel mode – or supervisor mode – allows the kernel to employ the special or privileged CPU instructions that it needs to perform such tasks as virtual memory management and process scheduling – which aren't available to software running in normal user mode.

The difference between a monolithic kernel and a micro-kernel is in how much of the kernel's services are provided by code running in kernel mode. A monolithic kernel such as Linux provides all of its services in kernel mode – including device drivers, network stacks (such as TCP/IP), and filesystems. A micro-kernel provides only the very

basic services such as process management, memory management and inter-process communication in kernel mode – everything else can be implemented in user mode.

»» So which of these is actually better: a monolithic kernel or a micro-kernel?

It depends who you ask! Micro-kernels were a hot topic in academic circles during the 1980s but fell out of favour in the 1990s (although they are now gaining in popularity again). Linus Torvalds is a famous example of somebody who has always been a vocal opponent of the micro-kernel approach.

The principal advantages of micro-kernels are generally considered to be flexibility and ease of development. In a micro-kernel much of the kernel code executes in user-mode just like any other software. Thus it can be developed, tested and debugged just like other software. It also makes the kernel more modular – with fewer interdependencies between kernel modules. This also makes development easier, and means that modifying the kernel is easier when requirements change. Thus a micro-kernel can be more flexible and more scalable. In contrast, in a monolithic kernel – where everything runs in the special kernel mode of the processor – testing and debugging can be very difficult. And because the kernel is one monolithic slab of code, interdependencies between different parts of the kernel can make modifying it difficult.

Ironically, the Linux and the Hurd kernels can both be seen as evidence against such arguments. The Hurd has proved to be notoriously difficult to develop and debug, whereas as you'll know, development on the Linux kernel takes place at a lightning pace. Moreover, the Linux kernel is actually very modular. Still, Linux can be difficult to modify, as anybody who participates in the Linux kernel mailing list can attest. Very often changes to one module or sub-system in the unstable testing branch of the kernel tree can cause wholesale breakage in other kernel components.

»» All very nice, but what difference is there to the end user between each architecture? Which one is faster?

The performance characteristics of micro-kernels versus monolithic kernels has always been a hotly debated issue. Traditionally, the argument has always been that monolithic kernels offer better raw performance, but that micro-kernels provide better interactive performance. You can see why if you consider the architecture of each.

All modern operating systems are multi-tasking – that is, they provide the illusion that more than one program can be run at one time, while in actual fact the processor is still only doing one thing at a time. This illusion is created by task switching or scheduling. Processor time is shared between all the tasks running on the machine, and each process in turn is allowed to run for a short period or time-slice before it is put to sleep (pre-empted) and another process scheduled and allowed to run.

How is this relevant? Well, typically, code running in kernel mode cannot be pre-empted. That means that when a program calls a kernel service that requires code to be executed in kernel mode then that program cannot be de-scheduled and another program run until that kernel service completes – no matter how long it takes. In a monolithic kernel this happens when any kernel service is used. Thus a monolithic kernel is bad news if you want tasks to get a bite at the processor at guaranteed, regular intervals – which is a requirement of real-time operating systems, and is essential to getting a good interactive response from an operating system. Micro-kernels don't have this problem since most kernel services run in user mode. On the flip side of the coin, if your interest is raw performance, then the monolithic kernel can usually offer better throughput of data for the very same reasons: code run in kernel mode can't be interrupted so it can be more efficient.

Because of the way they are designed, the determining feature of micro-kernel performance is how efficiently interprocess communication (IPC) operates. Early micro-kernels often gave poor



WhatOnEarthTheHurd

« performance due to inefficient IPC. This is one reason why Torvalds and others are so dismissive of micro-kernels. Modern second-generation micro-kernels offer much better performance.

»» Enough theory. Tell me more about the Hurd.

The Hurd is implemented as a set of user mode services hosted over a micro-kernel. Currently, the Hurd is built on the GNU Mach micro-kernel.

»» Wait a minute. Mach? I've heard of that. Isn't that used in Mac OS X?

Yes. GNU Mach is derived from the Mach micro-kernel originally developed at Carnegie Mellon University during pioneering research into micro-kernel based operating systems in the 1980s. A version of Mach is also the basis for Mac OS X.

In fact, the Mach micro-kernel has popped up in various operating system projects over the years – and seemed at one point to be set to dominate the field of OS development. It was the basis for NextSTEP, Steve Jobs' first attempt at Unix for the masses when he got the boot from Apple; Apple themselves sponsored a port of Linux to the Power Mac based on Mach called MkLinux; and famously there's now Mac OS X (itself derived from NextSTEP). It's not only Apple and ex-Apple employees that have been interested in Mach, though. There was also a traditional BSD Unix implementation based on Mach, called Lites; and IBM were to use Mach as the basis of their port of OS/2 to the PowerPC architecture (which never appeared). Funnily enough, thanks to Mac OS X, Mach is now the most widely-deployed Unix kernel.

Various parties have contributed to Mach development over the years. The Open Software Foundation (the OSF) dabbled for a while producing OSF Mach. Utah University picked up where CMU left off, and produced Mach 4 – upon which GNU Mach 1.2 is based – and OSKit Mach – which is the basis for GNU Mach 2.0. Meanwhile, MacOS X is hosted on OSF Mach 3.0. Did you get all that?

»» If there's been all these attempts at building Unix and Unix-like operating systems on top of Mach, why bother with Hurd? What's new?

Glad you asked. All the previous efforts to create production operating system based on Mach have opted for what is known as a 'single server' design. The user space portion of the kernel – which provides the remainder of the kernel services that Mach itself doesn't provide and supplies the application interface to kernel – is implemented as a single program; hence the term 'single server'.

In fact, most of these single-server implementations have been ports of a traditional, monolithic kernel to run on top of Mach. NextSTEP, Lites and MacOS X all use BSD Unix implemented as a Mach server, while MkLinux obviously uses a Mach-based port of the Linux kernel.

The Hurd, on the other hand, is rather different. It opts for a radical 'multiple server' design, where kernel services – such as user authentication, network stacks, filesystems, and so on – are implemented as a series of separate programs which run in user space.

(As an aside, this is where the name 'Hurd' is derived from. It's a play on words, a long-favourite diversion of hackers. 'Hurd' stands for 'Hird of Unix-Replacing Daemons' and then 'Hird' stands for 'Hurd of Interfaces Representing Depth'. The Hurd developers seem quite proud of the fact this is first time a software project has been named by a pair of mutually recursive acronyms.)

A single-server design fails to leverage the full flexibility that a micro-kernel core can potentially offer and suffers from many of the drawbacks of a traditional monolithic kernel: that is, the kernel is effectively still implemented as a single, monolithic chunk of code. The Hurd, however, takes full advantage of its micro-kernel basis to offer unparalleled flexibility. A user can choose to run which services they need, modify the behaviour of any of the services, or replace services with different implementations.

»» How do programs communicate with the Hurd's kernel services?

The Hurd provides all the usual features of a Unix kernel, but these are implemented in a novel way that makes it much more flexible for the user and much easier to extend. A core tenet of the Unix philosophy is that everything is a file: directories are files, devices are files, sockets and pipes are files. The Hurd is the same on the surface, since it is designed to be Unix-like, but the underlying mechanism is much different.

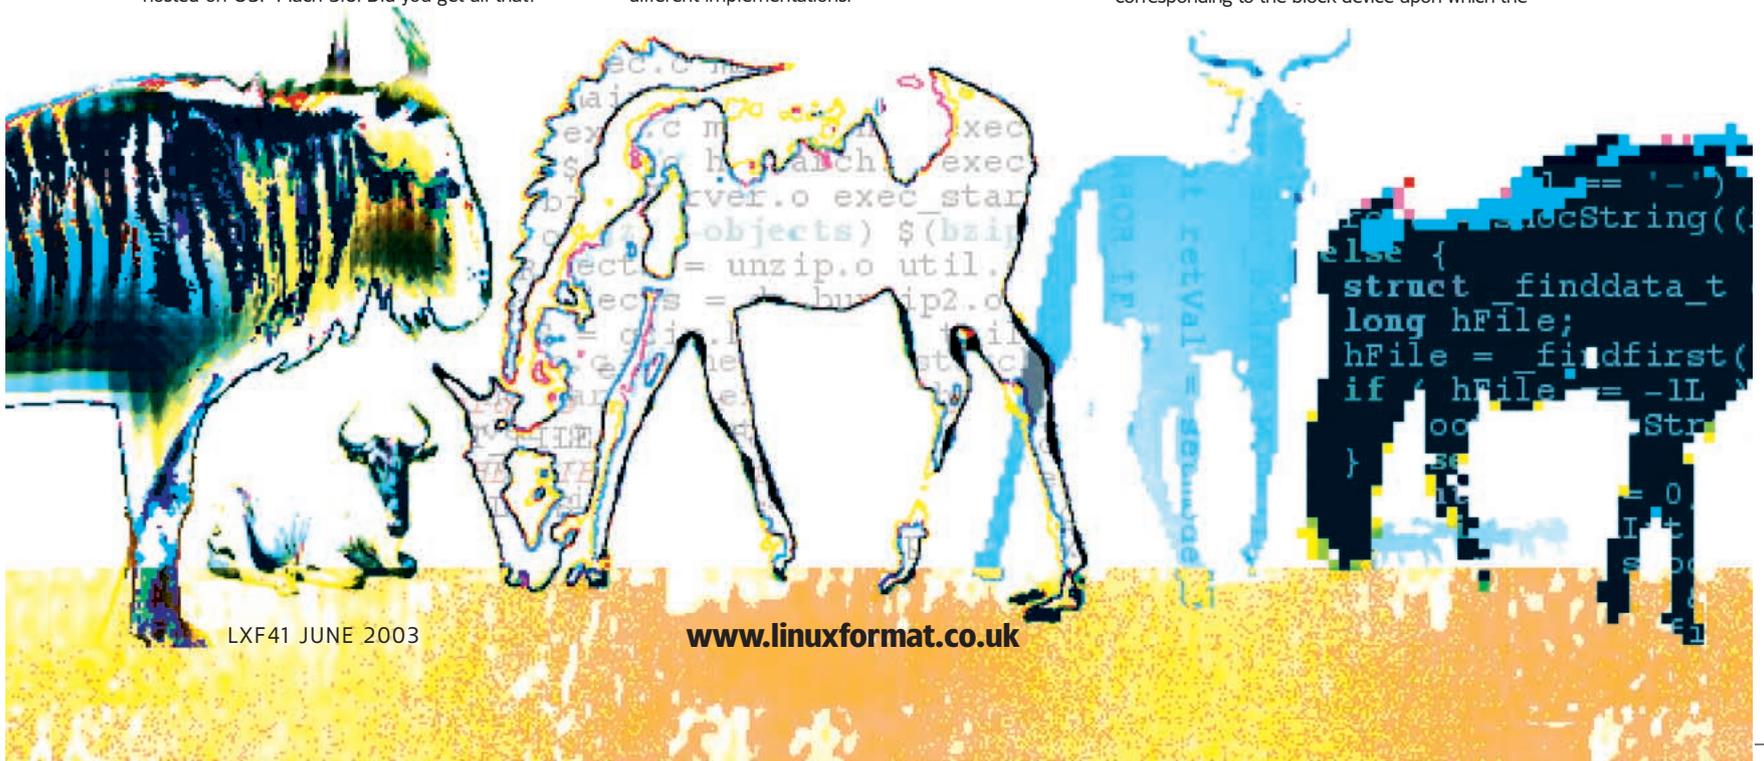
In the Hurd, a program uses kernel services – not by entering kernel mode and running the appropriate kernel routine – but by using the Mach micro-kernel's IPC system. In this scheme a program communicates with another program by sending a message to its message port. It is exactly the same when using kernel services since kernel services are just regular programs. But how do you know which port to send messages to? Well, in the Hurd, a server is identified by a file path relative to the root filesystem (or, more properly, the inode pointed to by that path). Thus, in the Hurd, while everything might be a file, every file identifies a message port on some server.

Inodes may represent regular files or more abstract services. Special Unix files such as named pipes are supplied by the fifo server; sockets by the pflocal server and symlinks by the symbolic link server.

A Hurd server which provides an interface through the filesystem in this manner is known as a translator. A translator is a regular program whose job it is to go between a file path and the program accessing that path and perform some kind of 'translation'. The translator is just one of the concepts that makes the Hurd so flexible.

»» What are some examples of translators in the Hurd?

A filesystem server which implements a regular filesystem is just one example of a translator. In this case the translator provides access to a device file corresponding to the block device upon which the



filesystem is stored. You mount a filesystem in the Hurd by associating a translator and a device with a particular file path in a manner somewhat similar to mounting a filesystem on a traditional Unix.

Another example of a translator is the pfinet server, which implements the TCP/IP protocol (and is accessed by default via the path /servers/socket/2). Bringing up a network connection in the Hurd requires that you associate the pfinet server with a particular network device such as an Ethernet card.

»» I still don't get it. What makes this translator concept so powerful?

Translators are especially powerful because the user requires no special privileges to use a translator – they just need permission to access the file path that the translator is being attached to. Also, translators are just regular user-space programs. Users are free to create, modify and use translators as they see fit. Compare this to a traditional Unix, where only the superuser can mount filesystems and where adding support for a new filesystem involves modifying the kernel, recompiling and rebooting. What's more the Hurd provides a set of shared libraries to promote code re-use and to make writing new translators very easy.

»» What filesystems does the Hurd support?

The Hurd currently includes translators which understand the Linux *ext2*, BSD's *ufs* and the ISO9660 filesystems. Support for the FAT filesystem and for NFS is also available. In addition, the *ftpfs* server provides a transparent, filesystem-like interface to FTP sites.

»» I'm sold. The Hurd sounds intriguing. How do I try it out?

The focal point for development on the Hurd is the Hurd pages on the FSF website (point your browser at www.gnu.org/software/hurd/). These pages contain a lot of useful information on the Hurd and

also links to where you can download the Hurd, get installation instructions and so on.

While you can build GNU Mach, the Hurd and most of the GNU OS from source code, most people wishing to try out the Hurd are not that masochistic. The GNU site hosts a binary GNU/Hurd distribution provided as a series of (currently four) CD images. Download these, burn them to CD-R and install away. (See [ftp://ftp.gnu.org/iso/](http://ftp.gnu.org/iso/).) Alternatively, you can try Debian's distribution of GNU/Hurd. This is built from the same source tree as the various Debian Linux ports and includes all the same great Debian infrastructure, such as the package update system APT. A large number of pre-compiled binary packages are included with Debian GNU/Hurd.

The Debian GNU/Hurd port is currently in the unstable phase of development, and, as such, official installation media have not been released. However, if you visit www.debian.org/ports/hurd/hurd-cd you can find a list of mirrors where you can download unofficial ISO images from which to install the Hurd.

»» What hardware platforms does the Hurd support?

The Hurd only works on an x86 PC at the moment. A port to the Alpha platform has also been started.

»» Will I have problems getting Hurd driver support for my hardware?

The Hurd lacks the maturity of the Linux kernel and the support of the big hardware vendors that Linux enjoys, so expect hardware support in the Hurd to be poorer. In particular, GNU Mach 1.2, the basis for current releases of the Hurd, make use of block and other drivers from the 2.0.x series Linux kernels – which have been superseded a long time ago. GNU Mach 2.0, the current development target for future Hurd releases, is built on OSKit Mach has a more modern and flexible driver architecture, so expect hardware support to improve.

»» What software can I run on the Hurd?

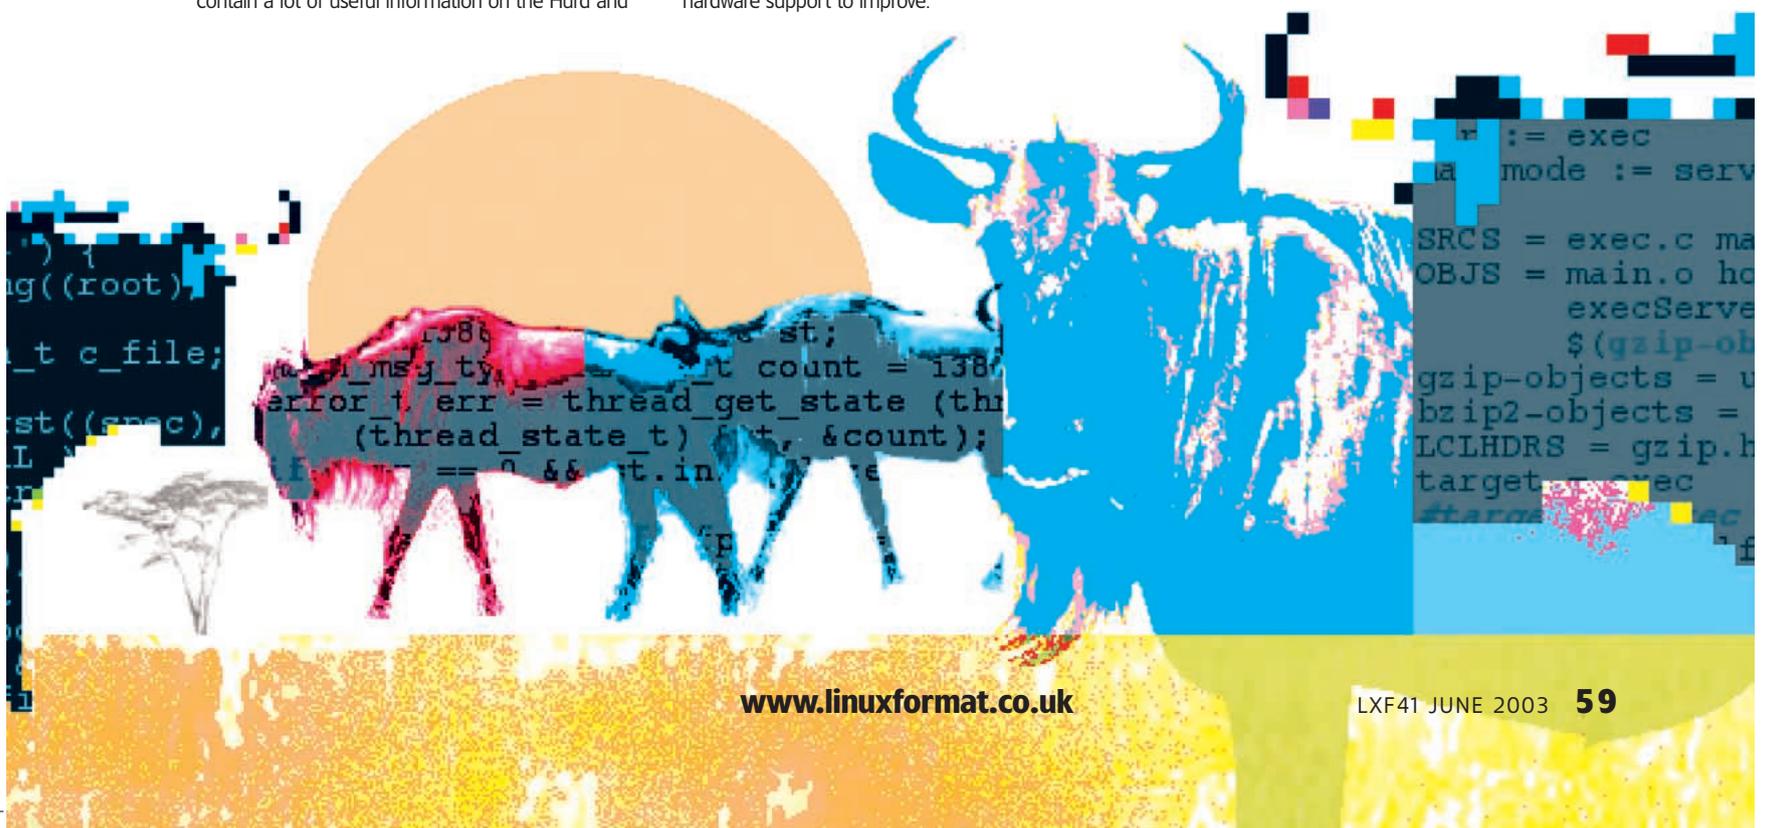
The Hurd is POSIX compatible, so if you have the source code, then most Unix software can be compiled and run on the Hurd. Most of the same software that will build and run on Linux will work on the Hurd, too. No binary compatibility is provided, so you cannot at the moment run Linux or BSD binaries on the Hurd – although such possibilities have been discussed.

In real terms, most of the free application software that makes up a modern Linux distro will work on the Hurd. Thus Xfree86 and the GNOME desktop will both run on the Hurd. Problematic software will be anything that use APIs that are specific to Linux – such as the /dev/fb or /dev/v4l interfaces – or those that make assumptions for Linux which are not valid on the Hurd. For example, most Unices including Linux assume that there is an upper limit on the length of file path allowed. This isn't the case with the Hurd, so not allocating files paths dynamically may cause a whole heap of problems.

»» When will the Hurd be finished?

How long is a piece of string? Seriously, a software project is only finished when nobody uses that software any more. The Hurd has a long way to go before it is ready for production use – a long way from that much sought-after stable '1.0' release. It's stable enough for everyday use, although somewhat lacking in features in some areas.

The Hurd project is looking for volunteers to contribute to all areas of Hurd development; so if you want to bring that '1.0' release nearer, lend a hand. If you think that Linux has become too mature and too dull, live on the cutting-edge of OS development and get the excitement back by choosing the Hurd. Visit www.gnu.org/software/hurd for more info. [LXF](#)



Tutorials

Our experts offer help and opinions on a whole host of Linux applications

Your guide to getting things done!

Whether you are just starting out in Linux, or an experienced veteran, there's always more to learn. Every issue of *Linux Format* is packed full of practical advice, and nowhere is it more concentrated than in our tutorials pages.

Here you'll find expert guides to all sorts of things, from Basic Linux usage to understanding and deploying network solutions, from simple script coding to the complexities of Perl regular expressions, Java server apps and more. We aim to bring a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact us, by post, by email (linuxformat@futurenet.co.uk) or log on to our website and post your suggestions in our special forums? (www.linuxformat.co.uk). Hope to hear from you soon!

Nick Veitch EDITOR

How code is represented

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap
between lines:
begin
  mniWordWrap.Checked := false;
end;
```

Usually, you'll find the code on our CD/DVD too.

THIS MONTH TEACH YOURSELF...

RPM Building

The mysteries of the spec file revealed – create your own binary or source packages with ease **p62**

Cable Modems

All you need for Ethernet connectivity is a card supported by the Linux kernel **p66**

The GIMP

Blend modes are available in all the painting tools, allowing non-destructive editing **p70**

PHP

If you want to make your PHP up to 326 times faster, why not replace chunks of it with pure compiled C? **p74**

Python

Control flow of Python programs, decisions, repeat actions and iterate over the contents of a container **p78**



Blender

Get more familiar with modelling by creating basic shapes in different perspectives using primitives and split views **p82**

Runtime Revolution

PART 2: Construct a basic image-viewer application and learn how to implement your own pull-down menus **p86**

TIP OF THE MONTH!

comm operates much like **diff**, except it's designed for more general, everyday use. When supplied with two sorted lists, **comm** will output three columns: column one contains all the items in the first list that are not in the second, column two contains all the items in the second list that are not in the first, and the last column contains items that are in both the lists. Here's an example: The file `paulh-fave-movies` contains *Aliens*, *X-Men*, *Blade*, *Star Wars*, and *SFW*, all on separate lines. The file `nickv-fave-movies` contains *Aliens*, *Blade*, *Pride and Prejudice*, *Sense and Sensibility*, and *SFW*.

comm

Comparing these two lists is simple:

```
comm paulh-fave-movies nickv-
fave-movies
```

The result is that **comm** outputs three columns, with column one containing *X-Men* and *Star Wars*, column two containing *Pride and Prejudice* and *Sense and Sensibility*, and column three containing *Aliens*, *Blade*, and *SFW*. **comm** can be further extended by calling it with a combination of **-1**, **-2**, and **-3** to make it not print a certain column. For example:

```
comm -13 paulh-fave-movies nickv-
fave-movies
```

would only print out favourite movies of

Nick's rather than Paul's. Furthermore, `comm -12 | wc -l`

would count how many movies were liked by both Paul and Nick.

The requirement to have files sorted might cause problems, however **comm** can read from a pipe if necessary by using a dash at the end to make it read from standard input. So this code:

```
sort paulh-fave-movies | comm -12
nickv-fave-movies -
```

would sort the list of Paul's favourite movies, compare it against the already sorted list of Nick's favourite movies, and output only movies they both liked. **LXF**

RPM PACKAGE CREATION

The Spec Files

Feel like going out tonight? Why bother when you could be sitting in, learning how to create funky RPM packages with **Warren Brown**...



The Red Hat Package Manager is fast becoming the package manager of choice for many distros such as Red Hat, SuSE and Mandrake, to name but three. RPM makes it easy to install, upgrade and query and erase packages. It can also check against its database to find out which package contains a certain file, whether certain files have been changed, and if you install the *rpmdb* package, you can have even more power with commands such as **rpm -q redhatrequires**. Programmers can also tap the power of RPM through libraries, making their software RPM aware.

All in all a very capable system, which is simple to use, extremely powerful and a great tool to have in your toolkit.

Why would you want to...

It's not compulsory to create packages, but it could make your life easier. Imagine catering for a helpdesk of about 150 staff, all running Red Hat. The system in use is developed in-house, with a new release about every 4 weeks – bug fixes, upgrades, etc. The developers don't care that you have to install the new system on 150 computers, they just know that the Perl modules attach to the C module, the C module attaches to the Python module, the Python module attaches to the Graphics modules... you get the idea? They've written a basic Makefile for the system, but it's sloppy, and things start to break when the new Perl module tries to call the C module that was written when Adam was a lad – not nice! So you decide to call upon the darker side of RPM and package the software into one incredibly sexy RPM which can be installed using cron jobs which run nightly checking for new versions of the software, and installing if necessary. They don't pay you enough for these bouts of genius! Pulling out your favourite *Teach Yourself RPMs* book, you start hacking away at a Spec file, then at the stroke of midnight all is complete! (Cue: insane laughter and thunderstorm) All you need to do is to put it on the ftp server and wait for the cron job to start, all the machines will be installed with the latest software plus any dependencies.

Non-sysadmin types can also use RPM to package software that hasn't been packaged yet. Just doing the usual configure, make and

make install (CMMI), can lead to future packages getting confused, not knowing that you had installed the latest library already. It's often better to package CMMI software that you want, and let RPM do the dirty work for you.

Finally, packaging your own software is useful when you want to install it using a kickstart disk. Just package the software, add it to the RPMS directory on your CD/NFS or FTP server, and update the kickstart file, then do an installation. Once the software is installed on your machines, it's relatively simple to upgrade the machine using *up2date* or *urpmi* – granted Red Hat don't have your updates, but you could install your own Open Source *up2date* server which would!

Spectacular SPEC files 101

Check if you have *rpm-build* installed. If you don't, it should be on your installation disks or from <http://rpmfind.net>. This contains *rpmbuild* that will be used to create RPMS. It also creates a directory structure under */usr/src* which has the following paths: **redhat/BUILD, redhat/RPMS, redhat/RPMS/athlon, redhat/RPMS/i386, redhat/RPMS/i486, redhat/RPMS/i586, redhat/RPMS/i686, redhat/RPMS/noarch, redhat/SOURCES, redhat/SPECS, redhat/SRPMS**

You may have some or all of these, depending on your system. Consider this to be the playpen/scratchpad of the RPM system, tread warily in its domain, who knows what forces you may unleash! (Sorry, too much sci-fi channel, and coffee I guess.)

SPEC File Format

Spec files reside in the *redhat/SPECS* files directory. They have structure and RPM expects certain things to be in a SPEC file, and will complain if they are not.

The main sections are Preamble, Prep, Build, Install, Files, and Clean. The first section, Preamble, contains information about the package that you want to build.

Summary:	Dr Evil's package	# Short note of what the package does
Name:	KillAustin	# Name of package – included in package name and package filename
Version:	1.0.0	# Version of software – included as above
Release:	1	# Release no - included as above
License:	GPL (programs), relaxed LGPL (libraries)	
Group:	Applications/Assassins	# How should it be grouped with other packages
Source:	ftp://ftp.evillair.com/pub/dr/KillAustin-%{version}/%{name}-%{version}.tar.gz	
		# Tells RPM what the file is called in SOURCES and where the source code can be found
Url:	http://www.evillair.com	# Points to the website.
Packager:	warren@work.net	# Who did the spec file
BuildRoot:	/var/tmp	# where the building of the package will

Some lesser-used RPM commands

rpm -qf /usr/bin/metacity	gives you the package that contains the file metacity
rpm -q --scripts bash_completion	displays pre and post installation scriptlets.
rpm -q --info bash_completion	displays header of your spec files and description
rpm -q --changelog bash_completion	displays changelog part of spec file
rpmbuild -ta packag.tar.gz	works if the package contains a spec file.

Installing *rpmdb* will give your RPM system even more power.
Commands such as:

rpm -q --redhatrequires sawfish	are now available to use!
rpm -q --redhatprovides /usr/bin/vim	

take place

That's a simple example of a Preamble, nothing too difficult there. Next comes the Prep Section. This section clears the area and gets things ready for the build. For example:-

%prep**rm -rf %{{buildroot}}/KillAustin****tar -zxvf \$RPM_SOURCE_DIR/KillAustin.1.0.0.tgz**

It consists mostly of Bash and Unix commands (get out your Bash scripting book, you may need it!). Those two lines are so common that the gurus who developed RPM made a macro and called it **%setup**. So we can rewrite the above as:

%prep**%setup**

The Build section can take a bit more out of you than the other two sections, sometimes it's easy, other times a bit more complicated. Again, it is basically just a Shell script. Let's go with easy for now:

%build**make**

Yep, that was it! Hang around though this is just a simple spec file example. Next we have **%install**, which is similar to make install when using Makefiles.

%install**make install**

The files section tells RPM what files you want in the package, we don't want our whole system to go into the RPM. So this section needs some careful thought! We need to create a file list of all the items that need to be included inside the RPM:

%files**%doc AUTHORS COPYING ChangeLog INSTALL NEWS README TODO****/usr/local/mybin/KillAustin_tui****/usr/X11R6/bin/KillAustin_gui****/etc/KillAustin.cfg**

Simple, but there's a gotcha, in that some of those directories may not exist, causing the RPM to fail. It may install on your box, because you have those directories.

The last section we'll deal with here is the **%clean** section. This allows RPM to clean up after the build.

%clean**rm -rf /tmp/mini_me**

Now on to building the package. The spec files should be in the redhat/SPECS directory. We use the following to build the package:

rpmbuild -ba <spec_filename>

That produces a lot of output to the screen, while rpmbuild does its job. End result should be:

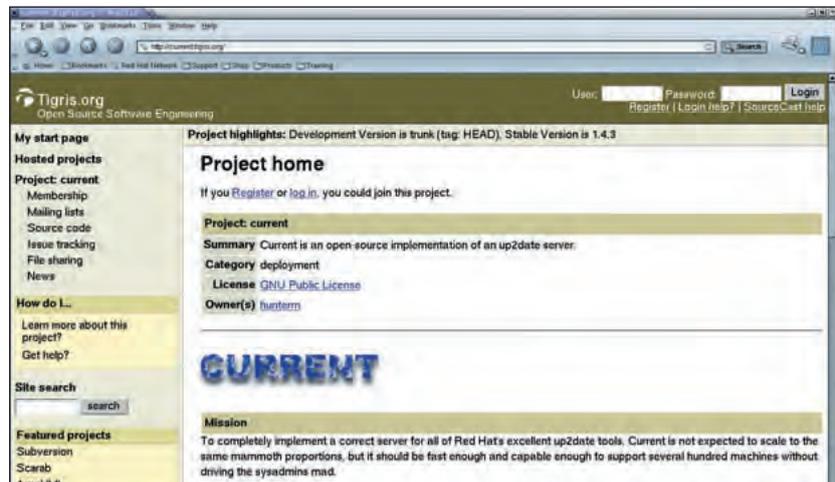
Wrote: /usr/src/redhat/SRPMs/KillAustin-1.0.0-1.i386.src.rpm**Wrote: /usr/src/redhat/RPMS/i386/KillAustin-1.0.0-1.i386.rpm**

Two files – an src rpm and a binary rpm.

Macros

Macros are tools that can make life easier. Macros such as **__mkdir** are just aliases for system commands, but they know where the command is on the host system. For instance, it could be in /usr/bin or /opt/bin/bin etc. We saw an example of a macro above called **%setup**, which can have arguments such as **-c** that tells **%setup** to create a directory, then change directory to it first before unpacking the archive. Useful for tar files that can mess up your current directory.

There's also a macro called **%makeinstall** that can be used instead of just **%make**. Be careful when using macros, however, because Red Hat's macros might not be the same as Mandrake's



which may not be the same as SuSE etc. Some useful macros are **%configure %makeinstall %setup**. The macro file is located in /usr/lib/rpm/macros.

My first RPM

We will need something to package. So I scoured the Internet and found a program called *bash_completion*. This program is just a Shell script, which needs copying to a directory of our choice. Download the tar.gz file from this location:

www.caliban.org/bash/index.shtml#completion. It needs to be put in the redhat/SOURCES directory for *rpmbuild* to use it. There is an RPM version but we are going to do the packaging this time, check out the original anyway to see what the packager did. **Note: Red Hat 7.3 and below may have to use rpm -ba instead of rpmbuild -ba**

To create a package, we need to be pretty clued up on how the original installed itself, so this means getting to grips with the README file, installation files and anything else we can get our hands on. In the case of *bash_completion*, there isn't much to read. Here is a summary of the ReadMe file:

INSTALLATION

```
-----
...

```

```
bash=${BASH_VERSION%.*}; bmajor=${bash%.*};
bminor=${bash#*.*}
if [ "$SPS1" ] && [ $bmajor -eq 2 ] && [ $bminor '>' 04 ] \
  && [ -f /etc/bash_completion ]; then # interactive shell
  # Source completion code
  . /etc/bash_completion
fi
unset bash bmajor bminor
```

This code checks that the version of bash that is parsing the code is later than 2.04 and, if so, sources the bash completion code.

If you are not putting the bash completion script in somewhere other than /etc/bash_completion, you should ensure that \$BASH_COMPLETION is set to point to it before you source it.

Version checking can be controlled by using **requires** in the **%prep** section. I'm going to move the script to /usr/bin, which means setting BASH_COMPLETION to /usr/bin. There's nothing

<http://current.tigris.org>
and www.nrh-up2date.com
are just two RPM sites
that can easily be found
using Google.



and production packages. Libraries should go in `/usr/lib`, binaries in `/usr/bin` and the include files into `/usr/include`.

Here's the spec file for `qDecoder`:

```
Summary: CGI Libraries
Name: qDecoder
Version: 7.0.1
Release: 2
Copyright: GPL
Source: %{name}-%{version}.tar.Z
Group: Application/Libraries
BuildRoot: %{_tmppath}/%{name}-%{version}
%description
CGI development library for C

%package devel
Summary: qDecoder Development libraries
Group: Development/Libraries
Requires: %{name} = %{version}

%description devel
qDecoder allows you to write cgi programs in C/C++

%package examples
Summary: qDecoder examples
Group: Development/Libraries
Requires: %{name} = %{version}

%description examples
qDecoder allows you to write cgi programs in C/C++

%prep
%setup -q

%build
%{_mkdir_p} %{buildroot}{%{_includedir},%{_libdir}}
./configure --prefix %{buildroot}/usr
%makeinstall

%install
pushd %{buildroot}%{_libdir}
%_rm -f libqDecoder.so
ln -s libqDecoder.so.2 libqDecoder.so
popd

%files
%defattr(-, root, root)
%doc AUTHORS CHANGES COPYING DISCLAIMER INSTALL
README REFERENCE
%{_libdir}/*.so*

%files devel
%defattr(-, root, root)
%{_includedir}*
%{_libdir}*.a

%files examples
%defattr(-, root, root)
%doc contrib examples

%clean
```

Not root enough to package

Packaging as a normal user

To create RPMs when not root, then do the following: Create the directory structure that RPM needs to build packages with:-

```
mkdir -p redhat/{BUILD,RPMS/{athlon,i386,i486,i586,i686,noarch},SOURCES,SPECS,SRPMS,tmp}
```

Then tell RPM where the new directory structure is. Create a file called `~/rpmmacros` and put the following into it:-

```
%_topdir /home/<USER>/redhat
```

```
%_tmppath /home/<USER>/redhat/tmp
```

```
%packager yourname@somewhere.com
```

where `<USER>` is the user creating the packages. In this case it would be `/home/warren/redhat`.

```
rm -rf %{buildroot}

%post
/sbin/ldconfig

%changelog
* Fri Sep 06 2002 Greg Wildman <greg@somewhere.com>
- added %clean for tidying up builds.
- added examples package for developers.

* Wed Sep 04 2002 Warren Brown <warren@somewhere.com>
- Version: 7.0.1.2
- Summary: Devel build included

* Wed Aug 23 2002 Warren Brown <warren@somewhere.com>
- Version: 7.0.1
- Summary: New build.
- New automated build.
```

The start of our spec file is pretty much the same as the previous one that we did. So nothing to discuss there.

Next we have the a new section labeled `%package devel`. It looks very much like the start of the file. Notice that it requires the `qDecoder` package to be installed, that's because we need to have the `qDecoder` libraries. The dev package should only contain the header files. The next section `%package examples` allows the user to install a package containing examples, this could have been put into the dev package, but at least it now gives you a choice. The `%prep` section is quite straightforward; the `-q` after the `%setup` stops the untarring from appearing on the screen.

`%Build` gives `qDecoder` a prefix of where to create initial directories, namely `%{buildroot}/usr/{%{_includedir},%{_libdir}}`. Then we use the `%makeinstall` macro. `%install` and `%files` which are reminiscent of what we did previously. Now look at the section `%files devel` and `%files example`. We are telling `rpm` that these packages only contain those files when it creates the `rpm` packages. The rest of the file is as we did before. Notice now, however, that two people worked on the RPM, so the changelog came in handy – Greg Wildman knows what I did, and vice versa! Now go ahead and build the RPMS with `rpmbuild`, then install with RPM.

When building RPMS, you should try to make them as generic as possible. Most of the spec file is script-based, so it shouldn't be too difficult to put in conditionals, for distribution type etc.

More Resources

Maximum RPM tells you all about how to use RPM and how to create packages for it. www.redhat.com/docs/books/max-rpm/ Also, there is no harm in looking at other people's spec files. Just download the `src rpm`. [LXF](#)

FASTER CONNECTION

Getting more from the Internet



Maurice R Kelly suggests: if you're tired of your slow dial-up service from your ISP, why not try a cable modem?

A recent government report cited an increase of more than 200% in the number of permanent Internet connections in the UK, which, for the most part, can be attributed to the explosion of broadband access. While the term "broadband" can really be extended to cover just about any high-speed, always-on connection to the Internet, in the language of the 'man in the street' it generally means one of two things – a cable modem or an ADSL connection. I'm going to be taking a look at the former, with the aim of passing on advice in order to help you get started with, or get more out of a cable modem.

So why get a cable modem?

There are two key factors to bear in mind if deciding to take the plunge into broadband – availability and speed. A cable modem is designed to be left on at all times, so that it has a permanent

connection to the Internet. This means that the availability of your Net connection is limited only by the time it takes to switch on the computer, or start a Web browser – ie you no longer have to wait for the dial-up procedure to complete.

So once you are connected the other factor – speed – comes into play. Cable modems connect to cable networks which are capable of high bandwidth allowing you to shove more data down the line, which means you can get the information more quickly – think of it in terms of a pipe with a larger diameter. Current speeds for cable modems range from 128KBps to 1MBps (although these figures are heading towards 150KBps and 2MBps). Bearing in mind that a dial-up modem has a maximum speed of 56KBps, it is obvious that there are significant performance benefits to be had.

To get a cable modem you need to live in an area which, to carry on the pipe analogy, is 'plumbed' for cable access – unfortunately this rules out a lot of non-metropolitan areas in the UK. The predominant cable companies are now ntl and Telewest, (cable modem access called ntlworld and Blueyonder respectively) in various locations around the UK. Have a look at the box marked "Essential Web-sites" for ntlworld and Blueyonder availability tools. If you can't get access in your area, you could contact the companies directly to enquire if future service is planned, or you could check for ADSL availability instead.

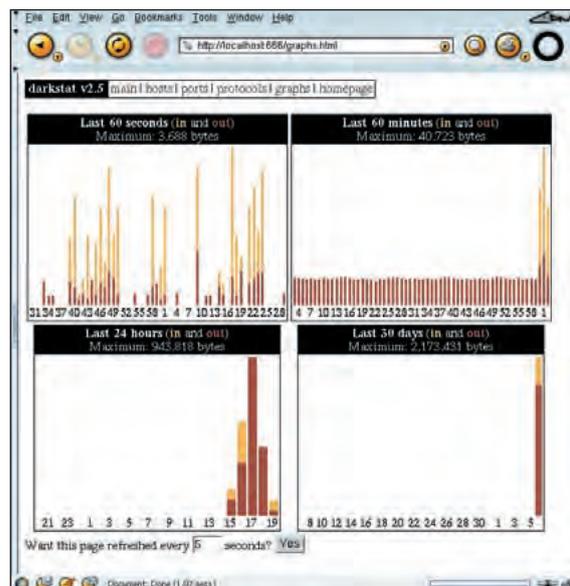
Getting hooked up

Before going out and ordering a cable modem, it is worth taking time to consider the hardware requirements. Any cable modems we have seen have had at least one of two means of connection to your computer – Ethernet and USB. Most recent computers come with some sort of USB option available, but if you want to use this with Linux then you're going to be in for a hard time. A look at the Linux-USB Device Overview at www.qbik.ch/usb/devices reveals that some cable modems do work with USB but support appears to be rather patchy – I couldn't find my own modem among those listed.

Considering that connecting to the cable modem via USB requires the use of the CPU for processing, it's preferable to use Ethernet whether you're running on Linux or not. All that's necessary to use a cable modem with Ethernet is a card supported by the Linux kernel. These can be picked up very cheaply nowadays, and although it is an inconvenience to install and configure, Ethernet is much better supported under Linux than USB.

Connecting your Linux box to your cable modem isn't rocket science, but can be troublesome for the inexperienced. If you've been using your cable modem with a Windows box it can make things a lot easier, as you can use the software supplied by the ISP in order to get registered. If you're attempting a fresh connection with Linux, then you'll have to perform some of the steps manually.

The key thing to remember is not to try to do everything at once. Get your Ethernet card configured to use DHCP first of all,



The *darkstat* connection monitoring tool is not as fully featured as *ntop* but aims to be less resource-hungry and more stable.

Linux? What Linux?

Something to bear in mind is that the cable companies do not seem to like Linux at all – support operatives are usually only trained to support Windows and MacOS (if even that.) There have been reports of tech support staff refusing to help if a caller mentions that they are using Linux, so it's probably best not to phone up

with any problems that may be OS-specific. If you are having problems with your service which you believe are nothing to do with your OS setup, then try connecting a Windows machine to the cable modem to identify the problem as being with the modem/network. You can also use this machine when talking to customer support.

which probably means installing the *dhcpcd* software. You could use other DHCP clients, but *dhcpcd* has had a better track record with the cable modem providers than *pump* and *dhclient*. With *dhcpcd* installed you need to set your Ethernet card to use DHCP – if your distro has one then start up your networking configuration tool, or alternatively on many distros you could use *linuxconf*. (Note: if you want to do it by hand, you could try editing `/etc/sysconfig/network-scripts/ifcfg-eth0`, and make the BOOTPROTO and IPADDR parameters read “dhcp” and “” [blank] respectively. Only modify these if you really can't find a configuration tool for your distro though!) Some configuration tools allow the user to specify which DHCP client to use, so set it to *dhcpcd* if it's possible to do so.

If you've already registered with your ISP using Windows, then you should be able to take down the cable modem interface and bring it back up again (as root) using:

```
ifdown eth0
```

```
ifup eth0
```

If you see no error messages, you should check that the process was completely successful by issuing **ifconfig** – ensure that the cable modem interface appears and has a sensible IP address. If you encounter no errors, then your connection should be running smoothly and you should be able to start using the modem to access the rest of the world. If you are having problems then check all your cable connections, the LEDs on the cable modem, and `/var/log/syslog` (which is where *dhcpcd* logs error messages by default.) You may wish to try invoking *dhcpcd* directly using **dhcpcd -d eth0** and monitor the logging. Check with **man dhcpcd** for more info on using it if the debug messages don't give you the information you are looking for.

First-time users of ntl's service will have to register themselves once the modem connection has been brought up for the first time. Navigate to <https://autoreg.autoregister.net/start.html>, and follow the process to completion (which should also provide you with ntlworld services such as email and Web space.) Once it has finished you should bring down the interface (remember **ifdown eth0**) and, following a power cycling of the modem, bring the interface back up again. The full expanse of the Internet should now be available to you!

For Telewest customers, the process is decidedly easier. After the engineer has installed the cable modem, you should phone up and register the MAC address of the network card connected to the cable modem; **ifconfig eth0 | grep HWaddr** should present this to you. If you have multiple computers or network cards then you can register up to five machines at <http://selfcare.blueyonder.co.uk> – only machines with registered NICs can be successfully connected to Blueyonder.

Feeling generous?

If you are living in a shared house, have a wired-up family, or if you just have multiple computers, you might want to think about connection sharing the cable modem. The best way to do this is get yourself an old PC, a couple of network cards and a cheap hub, and take a little bit of time to read up on the masquerading (IP-Masq) features of the Linux kernel. It is highly likely that your current kernel is capable of IP-Masq, but in the event that it isn't you should whip out *LXF39* and check the article on customising the kernel for details on how to compile a new one. You should then read the IP-Masq HOWTO for specific information about what kernel modules you need to enable.

You will need to define one of your Ethernet cards as internal and one as external – use your networking configuration tool to

Essential UK cable modem websites

ntl: service availability
<http://sales.ntl.com/>

Blueyonder service availability
<http://info.blueyonder.co.uk/publish/availability.html>

nthellworld - news and discussion
www.nthellworld.com/

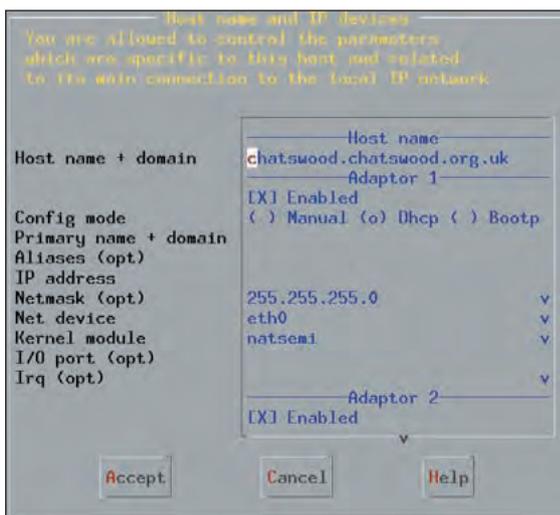
Linux firewalling information
www.netfilter.org/

Linux IP-Masquerade HOWTO
www.e-infomax.com/ipmasq/howto/c-html/index.html

Darkstat connection monitoring tool
<http://purl.org/net/darkstat>

Cable modem troubleshooting tips
<http://homepage.ntlworld.com/robin.d.h.walker/cmtips/>

www.by-users.co.uk/faqs/cablemodems/about/



If your distro doesn't come with a specific configuration tool, there's a good chance it provides *linuxconf*.

give the internal card a fixed IP address (from a private range such as 192.168.0.x) and set the external card to use DHCP to get its IP address from your ISP. Bring up your network interfaces and verify that you can access the external world from your gateway machine. If you can, then you're ready to use *iptables* to establish a basic firewall and provide masquerading services. Have a look at the script below for a sample firewall configuration (provides forwarding and drops incoming connections from the outside world):

```
#!/bin/bash
IPTABLES=/sbin/iptables
INT_INTERFACE=eth1
EXT_INTERFACE=eth0

# This is required to enable IP forwarding
echo 1 > /proc/sys/net/ipv4/ip_forward

# Set default policies
$IPTABLES -F INPUT DROP
$IPTABLES -F OUTPUT ACCEPT
$IPTABLES -F FORWARD DROP

# Clear out tables from previous configuration
for table in filter nat mangle
do
    $IPTABLES -t $table -F
    $IPTABLES -t $table -X
done
```



TutorialCableModems

Setting up your network card to connect it to a cable modem can be made easy if your distro provides an appropriate network configuration tools (such as this one in Mandrake 9.0).

```

<< # Handle forwarding
$IPTABLES -A FORWARD -m state --state
ESTABLISHED,RELATED -j ACCEPT
$IPTABLES -A FORWARD -m state --state NEW -i
$IINT_INTERFACE -j ACCEPT
$IPTABLES -A FORWARD -j DROP

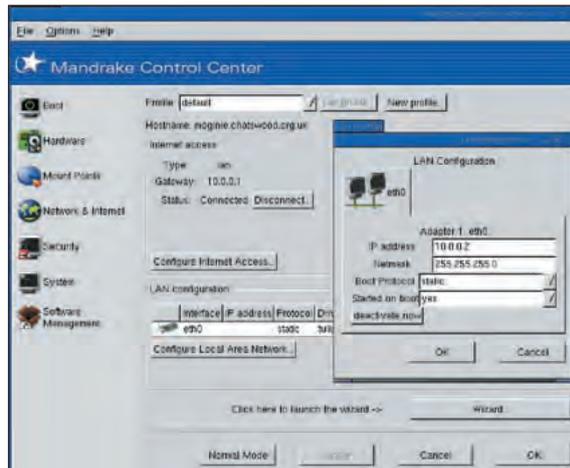
# Handle incoming
$IPTABLES -A INPUT -m state --state ESTABLISHED,RELATED
-j ACCEPT
$IPTABLES -A INPUT -m state --state NEW -i
$IINT_INTERFACE -j ACCEPT
$IPTABLES -A INPUT -i lo -j ACCEPT
$IPTABLES -A INPUT -j DROP

# Masquerading
$IPTABLES -t nat -A POSTROUTING -o $EXT_INTERFACE -j
MASQUERADE

```

Once you have the firewall up and running, you need to install network cards in your internal PCs and configure them for your network. They will need IP addresses, so you should assign each one an address from the private range that your gateway was assigned from. If you chose 192.168.0.x, then it is convention for the gateway to have the address 192.168.0.1, and other machines to have subsequently numbered addresses.

Once they are networked, you need to configure your internal PCs to use the masquerading box as their gateway to the Internet. If they are Linux machines you can set the gateway in most networking configuration tools, or by hand edit the file `/etc/sysconfig/network`, and update the GATEWAY parameter to correspond to the IP address of your internal Ethernet card on the masquerading machine. If you have Windows clients, then update the gateway in your TCP/IP settings dialog. Don't forget to update the DNS server addresses in your client machines to correspond to those of your cable ISP – you can pull them out of `/etc/resolv.conf` where `dhcpcd` should put them whenever you get your IP address.



Of course, you could always run a DHCP server on your gateway machine so that the client machines can get their IP address and DNS server configuration with less manual intervention from you. It requires a little extra work for you initially, but makes it a lot easier to add more machines to the network at a later stage. If you want to do so, install the `dhcp` software from ISC (www.isc.org/products/DHCP/ – although check for packages on your distribution CDs first) – it isn't very hard if you read the DHCP mini-HOWTO and the documentation available at www.isc.org.

Running your own servers

There are three very good reasons why people with dial-up connections don't often run server software on their own machines. Dial-up connections don't support a very high data transfer rate, the connections are not permanently up, and they tend to jump from one IP address to another every time the connection is re-established, which happens a lot when people have to pay the phone bill, and need to use the lines for making voice calls as well. In summary the problems are speed, availability and location (the first two should sound familiar).

Basic Firewalling

Always use protection

If you are even remotely thinking about connecting a Linux box to a cable modem without using some sort of firewall, then stop right now. Distributions are a lot more security conscious these days, and while you get more control over the list of services running from start-up you still need to put protection round those services that you do need to run for day to day usage.

Some distributions have firewalling tools and whilst these are useful and time-saving, it doesn't do any harm to roll your own firewall – then you really have control over what is happening. If you haven't already done so, make sure you have `iptables` installed – you may need to configure your kernel to support firewalling, but most distributions come with support in their stock kernels.

When deciding on your firewall, the key issue is to decide what sort of access you want to allow from the outside world – the safest thing to do is block all new incoming connections, while permitting established connections to continue to exist (ie a connection cannot be made to your machine from outside, but if you established the connection from inside, the

communication will continue.) It becomes more complicated if you want to access specific servers from outside your firewall, but we'll come to this later. The script below will create a basic firewall blocking all incoming connections: –

```

#!/bin/bash
IPTABLES=/sbin/iptables
$IPTABLES -F INPUT ACCEPT
$IPTABLES -F OUTPUT ACCEPT
$IPTABLES -F FORWARD DROP

# Clear out existing rules
for table in filter nat mangle
do
    $IPTABLES -t $table -F
    $IPTABLES -t $table -X
done

# Accept established connections, and those on the
loopback interface

```

```

$IPTABLES -A INPUT -m state --state ESTABLISHED,
RELATED -j ACCEPT
$IPTABLES -A INPUT -i lo -j ACCEPT

```

```

# If we haven't ACCEPTED yet, then REJECT.
$IPTABLES -A INPUT -j REJECT

```

If you do want to open certain ports for incoming connections (say for ssh, on TCP port 22,) then add a line such as the following before you get to the end of the INPUT chain (ie before the `$IPTABLES -A INPUT -j REJECT`):

```

$IPTABLES -A INPUT -m state --state NEW -p tcp --
dport 22 -j ACCEPT

```

Put the code into a script, and call it from one of your startup scripts (somewhere like `/etc/rc.local`) so that you are firewalled every time you reboot the machine. Firewalling can be a complex subject and it is worth reading into it.

Try the Firewall HOWTO for some background reading, as well as David Coulson's articles in *Linux Pro* (March and April 2003 editions.)

As far as speed and availability are concerned, broadband overcomes these problems to an extent as it provides upload speeds exceeding those possible with standard modems. So if someone connects to your server across a cable modem link, they'll get a reasonable download speed. Of course the speed is dependent on how much you pay for your service – for example, on ntlworld, the current 128k service has an upload speed of 64kbps, while the 600k and 1Mb services have upload speeds of 128kbps and 256kbps respectively. An upload speed of 128kbps should be more than enough to maintain a personal homepage (if it's not too graphics-intensive of course!)

Similarly, due to the always-on nature of cable modems, people can connect to your servers for as long as you leave your computer running. If you do plan to use your cable modem connection to run a server (whether it be Web, mail, *ssh* or whatever) you might consider acquiring an old machine to put those services on, which will leave your personal machine for your own work. It makes sense from an economic point of view (an old low-spec box with no monitor costs less to run than your power-hungry desktop beast!) and security-wise (there's no point in needlessly exposing your main machine with all your personal documents on it.)

The "location" problem with dial-up connections is that the IP address of the average dial-up connections changes from connection to connection. Cable modems tend to be on the network for long periods of time, and as a result, the IP address your computer receives is usually fairly static – it's not unusual to have the same IP for months at a time. As a result, your servers won't seem to move about as much as they would on a dial-up connection. Of course, it's still not perfect so you might want to consider dynamic DNS services (see the *Dynamic DNS Services* box on the right.)

From a security point of view, when running servers it is worthwhile remembering that software vulnerabilities exist and get exploited all the time. Make sure you are up-to-date and run only essential services. See the box on the left entitled *Basic Firewalling* for more details.

Don't go overboard

Broadband is great – you can download so much more than you ever could with a standard modem, and many people spend a lot of time downloading. Unfortunately there's a good chance that the truly unlimited services could be coming to an end as ntl recently announced download capping. At the time of writing, the ntl cap was a limit of 1GB per day (which can be broken three times a fortnight.) It sounds like a lot but if you plan on downloading ISOs of Linux distros, you could soon eat up your allocation.

Of course if you share a connection among a personal network you might want to keep an eye on your combined consumption through the use of a connection monitoring tool such as *ntop* or *darkstat* (one of our recent Hotpicks.) Once installed you instruct these programs to monitor the data throughput on a particular network interface. They both have built in Web servers which can display a wealth of information about total bandwidth usage and breakdowns of the actual usage by protocol and destinations.

To install *darkstat*, download the source, check that you have the library and development packages for *libpcap* installed, and do the usual build routine:-

```
./configure
make
```

Dynamic DNS Services

Domain IP management made easy!

Running your own server would be a lot simpler if you had an easily rememberable domain name – you could hand it to your friends once so you don't have to let them know when your IP address changes. Fortunately you can do just this with dynamic DNS services, like those provided by www.dyndns.org. They (and lots of other organisations) provide a number of free services to allow you simpler managing of a domain name for a machine whose IP address changes.

Signing up to the service is pretty simple, and there's plenty of help on hand for the minimal amount of effort required – you should decide whether you want to use the dynamic or static DNS service – they are pretty similar, except that the dynamic service propagates updates more frequently, and requires you to update the IP address at least once every 35 days. After sign-up it can take a few hours for your new URL (which can be yourname.dyndns.org, or one of 33 other domains) to become properly active. Once active you just have to sit back and let people visit, until the next time your IP address changes when you can log back into the DynDNS.org control panel and enter the new IP.

Of course, for the ultimate in no-fuss, you can run a piece of client software that will update the address without having to log into the Web-site. One such program is *ddclient*, which has its official home at <http://burry.ca:4141/ddclient/> – it's on a bit of a slow connection so you might want to Google for "ddclient" to find alternative locations. The tarball contains an executable called *ddclient* which is actually a Perl script. It should be copied into */sbin*, and a configuration file (*ddclient.conf*) should be created in */etc*. A basic config file can be found below for the static DNS service – change **use static=no** for the dynamic service.

```
daemon=300
syslog=yes
mail=root
pid=/var/run/ddclient.pid
```

```
protocol=dyndns2
server=members.dyndns.org

login=your_dyndns_login_name
password=your_passwd
```

```
static=yes \
server=members.dyndns.org, \
protocol=dyndns2 \
your_dyndns_host.dyndns.org
```

Fortunately the author includes a number of sample scripts and a reasonable README file in the tarball. You can perform updates manually (or scheduled in a cron job,) through a *dhcpcd.exe* script, or by running *ddclient* in daemon mode (where it monitors your IP address for changes.) Most people using cable modems will be using the *dhcpcd* DHCP client so it is easier to update your address using the *dhcpcd.exe* script. Don't use the sample provided by *ddclient* – instead insert the listing below into the sample *dhcpcd.exe* file which you should find in */etc/dhcpc*. Insert it into the case statement for the situation when a new IP address is detected. There should be a complete *dhcpcd.exe* file on the CD-ROM.

```
new) logger -s -p local0.info -t dhcpcd.exe
"interface ${INTERFACE} has been configured
with new IP=${IPADDR}"
# ==== Put your code for the case interface
has been brought up with new IP address here
```

```
case "${INTERFACE}" in
eth0)
logger -t dhcpcd IP address changed to
${IPADDR}
ddclient -daemon=0 -syslog -use=ip -
ip=${IPADDR} >/dev/null 2>&1
;;
*) ;;
esac

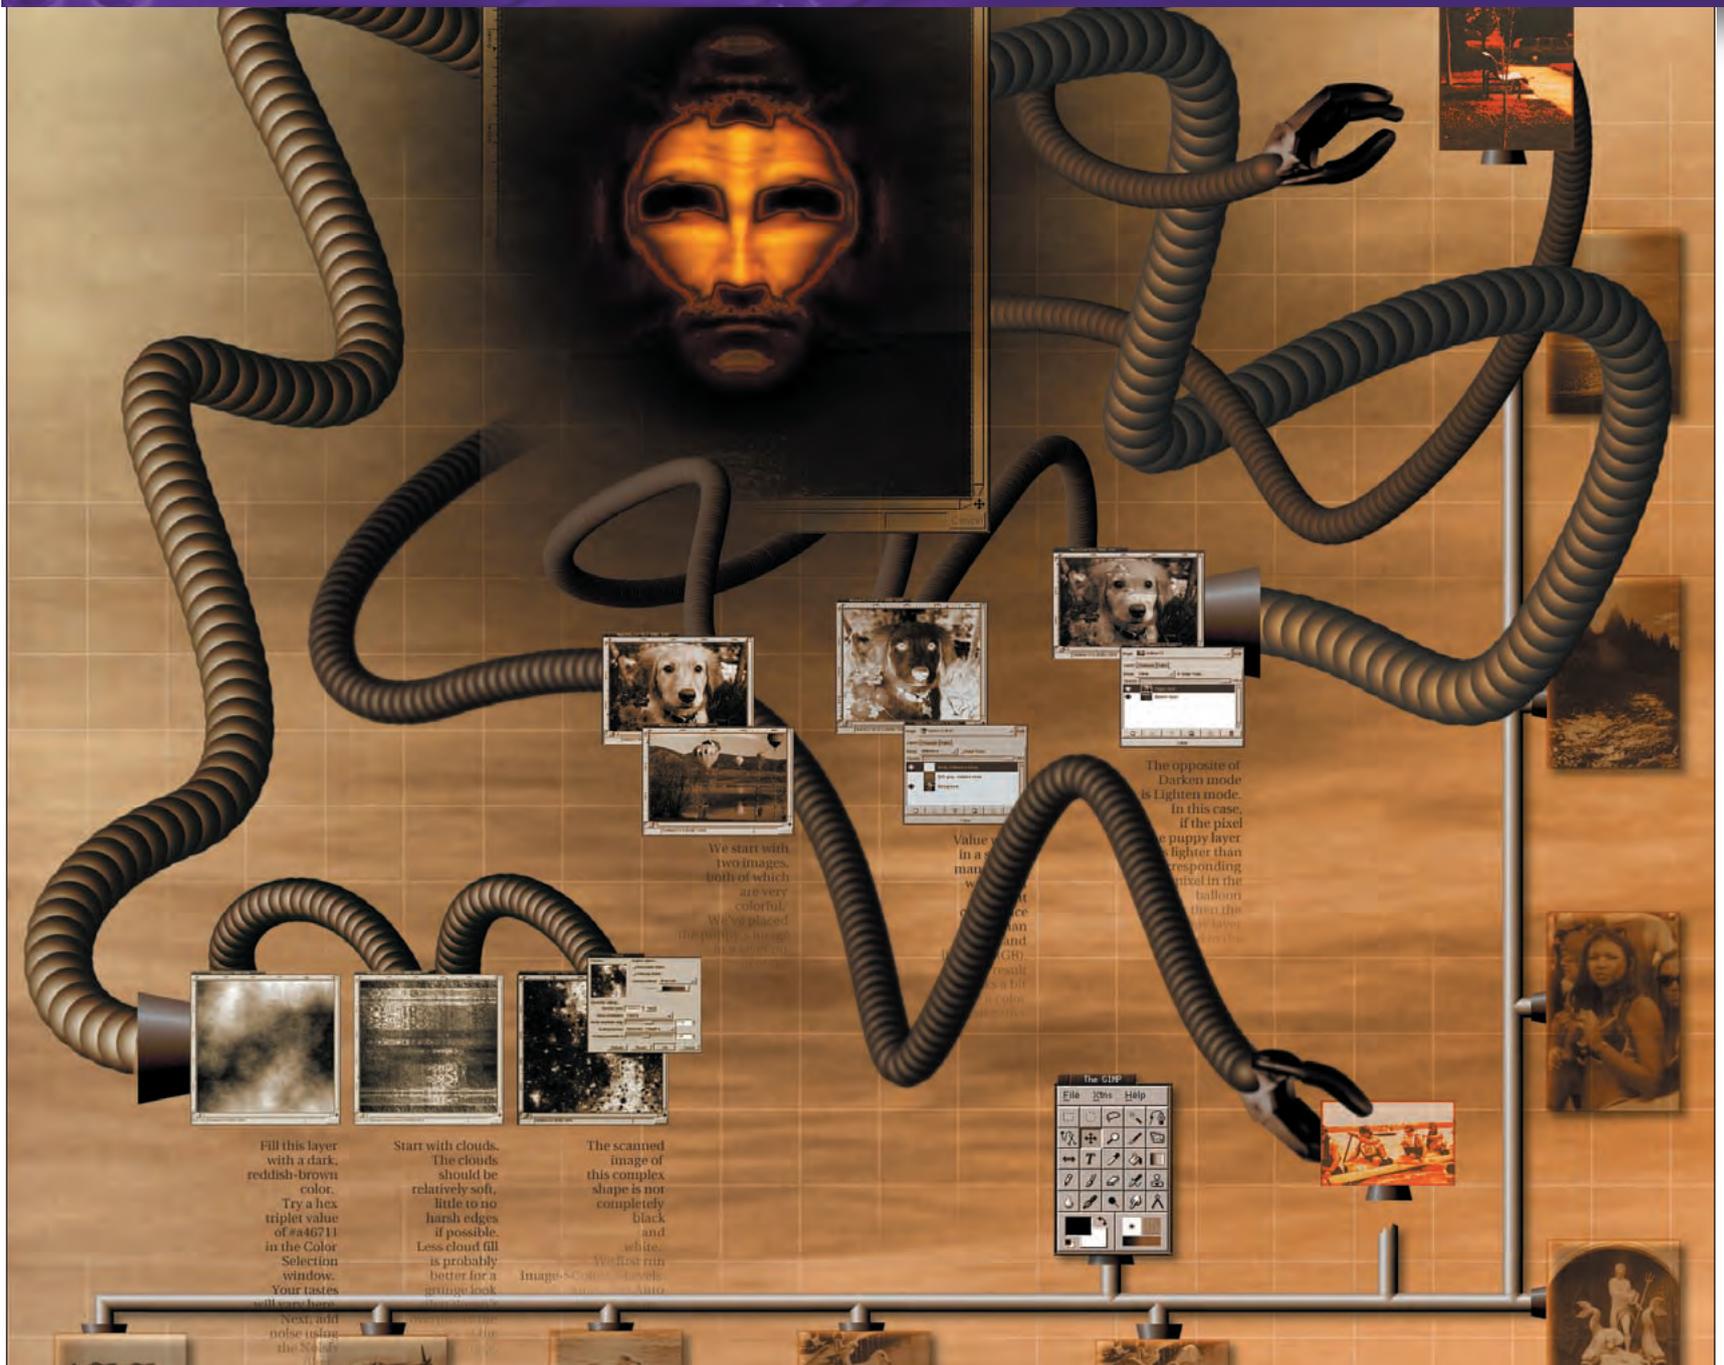
# ==== End
;;
```

su -c "make install"

Once installed you can simply run it from the command line (as root) and point your Web browser to <http://localhost:666/> in order to check the statistics it generates. To run *darkstat* in the background you should execute it using **nohup darkstat &** (log messages will go to a file called 'nohup.out'.) For a masquerading machine specify the external interface using **-i eth0** and let it know about your internal network using **-I 10.0.0.0/255.255.255.0** Check out the full range of *darkstat* options that are available using **darkstat -h**.

In parting...

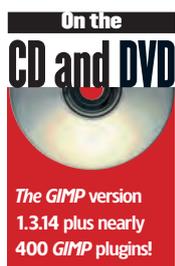
We hope that this article has been a benefit to the intended audience – those thinking about getting a cable modem, those who had one but were unsure about using it with Linux, and those who were already using it but wanted to get a little bit more from it. Just remember to keep it secure and enjoy the opportunities that broadband brings you! [LXF](#)



NON-DESTRUCTIVE EDITING

Blend modes

PART 4 Blend modes can modify colour content with simple brush strokes. **Michael J Hammel** shows us how...



Modifying an image by changing the pixels might seem the only way to create effects with *The GIMP*, but such changes are destructive in nature, making it hard – if not impossible – to recover the original image info should you change your mind. A better solution is to create effects without changing the pixels – what is called non-destructive editing – through the use of Blend Modes.

When painting with a paintbrush, the pixel data provided by the brush stroke normally replaces the pixel data over which the stroke

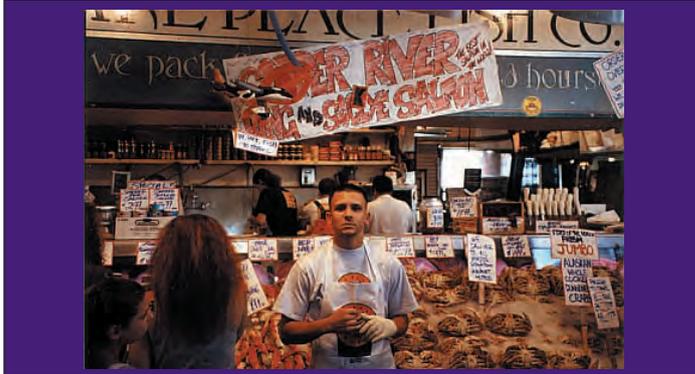
is applied. Blend modes allow brush strokes and other image data to be combined in ways beyond the standard replace option. Blend modes are available in all the painting tools (Paintbrush, Gradient Tool, Bucket Fill, etc) as well as in the Layer and Channels window. In the case of the Layer blend modes, changes can be applied in a non-destructive manner.

A blend mode is a method describing how a pixel will change when composited with another pixel. In the case of layers, blend modes define how pixels in a higher level layer will be combined with pixels in the next layer down. One use of Layer blend Modes is to produce a colour negative from an image by placing a white layer on top of the image and setting the white layer's blend mode to Difference. The colour negative can then be desaturated and used to create a complex selection or layer mask. The tutorials here show all the different blend modes available in *The GIMP* and how they can affect image data. While reading these tutorials, keep in mind that not every image works well with every blend mode. Experimentation is key to understanding our topic this month.

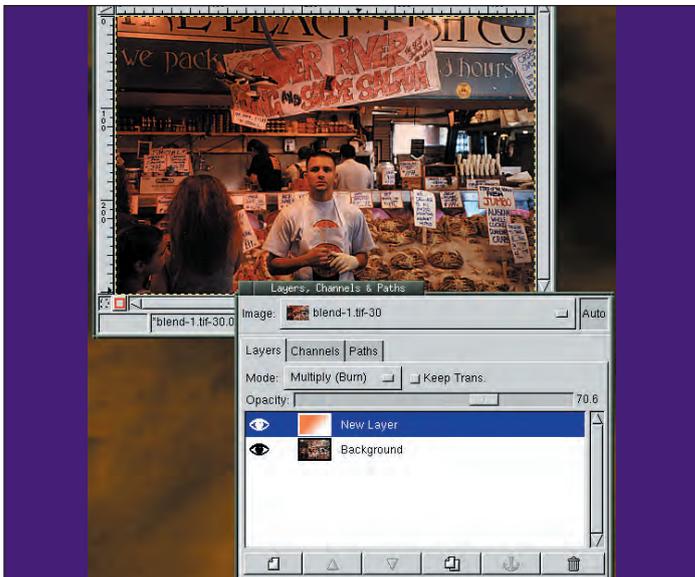
MULTIPLY, DIVIDE, DISSOLVE

Blend modes can be used to lighten dark areas of an image, add colour to lifeless images, and produce unusual affects, all without actually modifying the

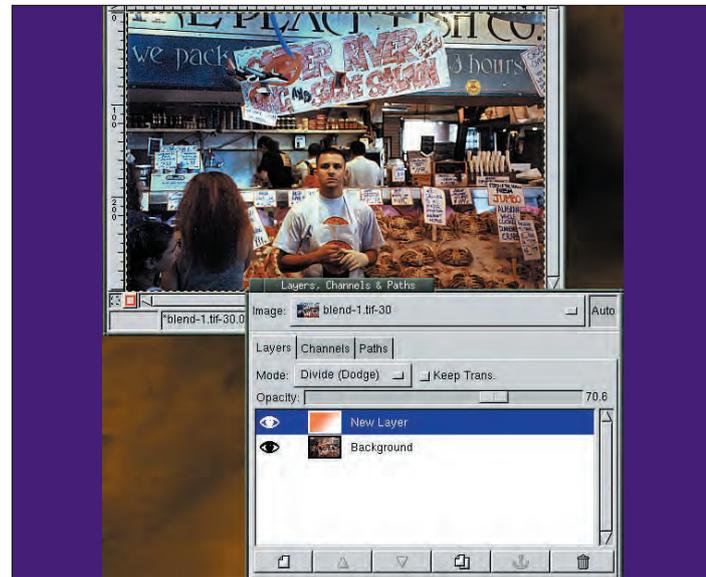
image. In the following set of images we'll see three blend modes change the mood of an image with very little effort on our part.



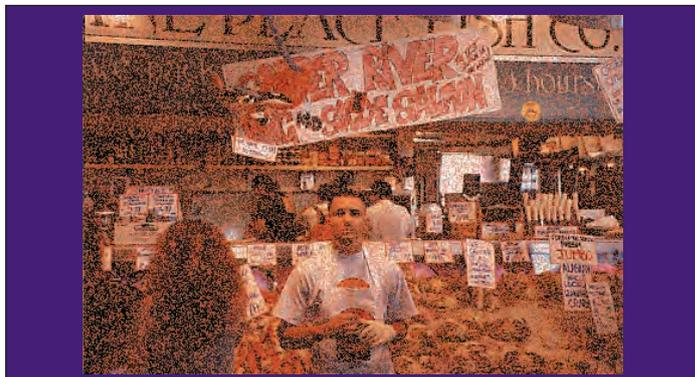
1 We start with our original image. The picture shows a popular market in Seattle lit by low-power bulbs and plenty of reflected outdoor sunshine. There is a definite warm feeling from the red and orange colours in the centre and bottom half of the image, but we can change this to be much more dramatic using blend modes.



2 Adding to the warmth is easy using the Multiply blend mode. We add a new layer above the original image. Change the foreground color to a reddish-orange and the background to white. Apply a linear gradient (using the Gradient Tool from the Toolbox) starting in the upper left corner of the new layer that flows to the bottom right. Now change this new layers blend mode to Multiply (Burn). The result is that the hanging light in the upper left corner appears to be orange and casts an orange light on that side of the market. We've reduced the opacity of the new layer a bit to reduce the harshness of this effect. Multiply mode always makes the result darker or leaves it unchanged. Multiply a black layer and you get black. Multiplying a white layer will not change the visible image.



3 Alternatively, we can change the mood to be brighter by removing some of the warm colours. To do this we simply change the new layers blend mode from Multiply (Burn) to Divide (Dodge). Now the light in the upper right is brighter and casts a stronger blue coloured light on the market. A Dodge blend with white has no effect.



4 A less dramatic effect can be generated by selecting all of the background layer (CTRL-A) and doing a Bucket Fill using the Dissolve blend mode. Double click on the Bucket Fill tool in the Toolbox to open it's Tool Options window. Set the opacity to about 50% for this image and the Mode to Dissolve (we're using the Tool blend mode here, not the Layer blend mode). This is one method of adding grain to the image. A Burn blend with black has no effect.



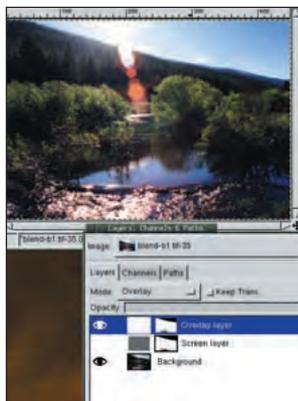
« SCREEN, OVERLAY, BEHIND

The mode you use on any given image to bring out detail will depend on the image itself. Each mode works well with different colours. In the previous tutorial you saw how a relatively well lit image could be enhanced with Multiply and

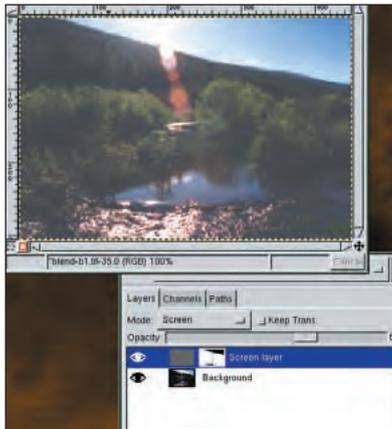
Divide. Underexposed images are common with lower-end digital cameras that don't have many manually user-definable functions. In a poorly lit or underexposed image you might try Screen or Overlay, as we see below.



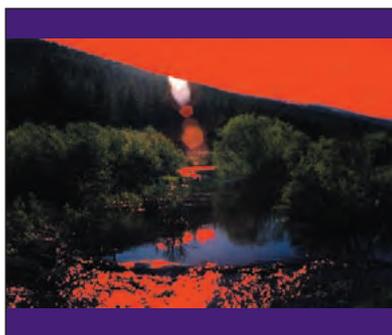
1 Our original image, taken outside and facing the sun. While not a bad shot, dark areas are underexposed. We need to bring out detail in the foreground trees.



3 Add another layer, filled with white. Add a layer mask. Click on the gray layer's mask to make it active and copy it. Click on white layer's mask and paste. Anchor floating layer into the new layer's mask. Now change the white layer's blend mode to Overlay. This time we get a better result.



2 Add a layer and fill it with a dark gray. Add a layer mask to this layer. Make the original layer the active layer, copy it (CTRL-C) and then make the layer mask in the new layer active by clicking on it. Now paste the copy (CTRL-V) and anchor it to the layer mask. Using the Brightness/Contrast filter (Image>Colors>Brightness-Contrast...), increase the contrast in the mask so it becomes a black and white (no gray) image. Invert the mask. We've just isolated the dark areas. Now change the layer blend mode for the new layer to Screen. The effect is that the dark areas are all lightened, even with the opacity reduced to lessen the effect. Screen normally has an effect of projecting two photographic images over each other so that a little of each shows through. A fully white screen will produce a white image, while a fully black screen produces all black.



4 In step 3, white overlay makes light areas lighter. Black overlay makes dark areas darker. A 50% gray overlay has no effect. To see what parts of the layer mask are not affected, duplicate the original layer, add a layer mask and copy the mask from one of the other layers into it. Then we apply the layer mask (press and hold the right mouse button on the layer name, then select Apply Layer Mask). Select the entire layer (CTRL-A). Now select a red foreground color (any color will do but choose one that will stand out) and do a Bucket Fill in this layer. Set the Bucket Fill blend mode (in its Tool Options window) to Behind. This fills in only the transparent pixels in the layer with the fill color.

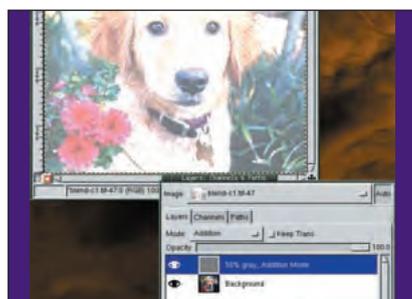
ADDITION, DIFFERENCE, SUBTRACTION

There are multiple blend modes that do very similar things, such as the Screen and Addition modes. Brighter areas in a Screen layer will brighten an image more than the darker areas. With Addition, each pixel is added directly to the

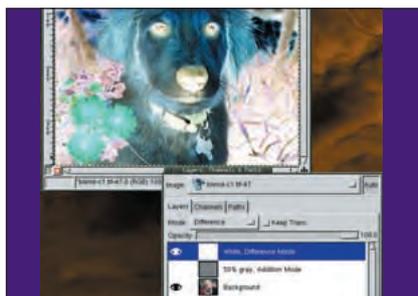
lower layers without regard for how bright or dark that pixel might be. Here we look at Addition, Difference and Subtraction. These tools can be used to create colour negatives of images.



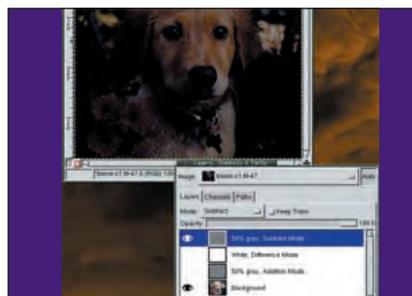
1 The original image is colourful and fairly well balanced with respect to brightness. While we don't really need to modify this, we can easily show how Addition, Difference and Subtraction will affect it.



2 Adding a layer of 50% gray and setting the layer blend mode to Addition we again see a washed out appearance of the original image. Addition adds each pixel of the current layer to the composited pixels of the layers below it, with the maximum result being white.



3 Adding a layer of white and setting the layer blend mode to Difference we get a colour negative of our image. The Difference mode subtracts the current layers pixels from the composited pixels of the layers below. The absolute value of the result becomes the new composited pixel.



4 Duplicating the 50% gray layer and changing the mode to Subtract we get a darkened image. If we set the white layer to subtract, our image is completely turned to black. Subtract works like Difference except if the resulting composited pixel is negative, the pixel is made black (a zero value).

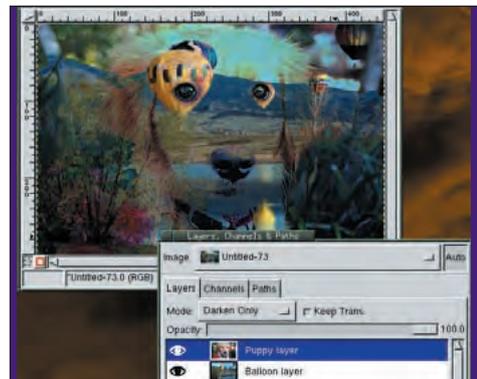
DARKEN, LIGHTEN, VALUE

The Darken and Lighten blend modes work opposite of each other to merge a layer with the layers below it. A ghosting effect will occur with the Value blend

mode, especially if the layer being blended is desaturated (Image>Colors>Desaturate).



1 We start with two images, both of which are very colourful. We've placed the puppy's image in a layer on top of the balloon image for this example.



2 Setting the puppy layer to Darken mode produces a soft ghosting effect. In this mode, if a pixel in the balloon layer is darker than the corresponding pixel in the puppy layer then the composited image uses the balloon pixel. If the puppy pixel is darker then the puppy pixel is used.

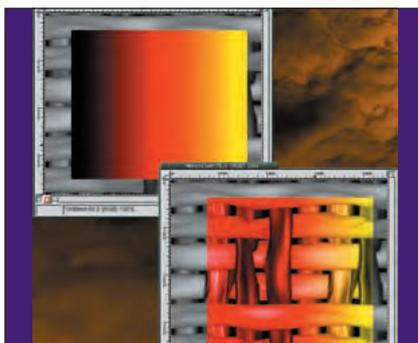


3 The opposite of Darken mode is Lighten mode. In this case, if the pixel in the puppy layer is lighter than the corresponding pixel in the balloon layer then the puppy layer is used in the composite. The result is another ghosted image, though this time a bit lighter overall.

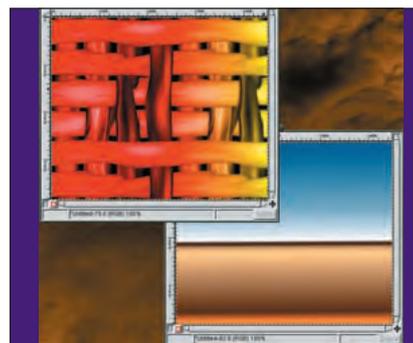


4 Value works in a similar manner but works in a different color space (HSV) than darken and lighten (RGB). The result looks a bit like a color negative blended with the balloon image. There are many color spaces images can be processed in though most GIMP users will work directly in RGB and leave color space conversions to filters to handle automatically.

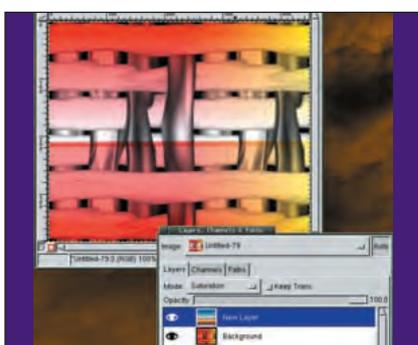
COLOR, HUE, SATURATION



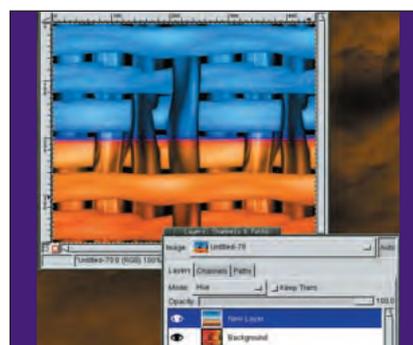
1 Blending with the Color mode is the fastest way to colourise a desaturated image without actually modifying the image itself. The lightness of the lower layer defines the structure of the composite image while the color (a combination of the Hue and Saturation) in the blend layer defines the color.



2 Working with Saturation and Hue modes is a little more difficult to explain. Both work in the HSV color space instead of RGB (though this is done invisibly, without the user knowing it). We'll use two coloured patterns to demonstrate these last two modes.



3 In Saturation mode we can see the image structure (the weaving) of the original layer is kept intact, along with some of the colouring. What changes is the intensity of the colours. Where the Saturation layer is dark, the weaving is dark. Where the Saturation layer is light, the weaving is light.



4 In Hue mode the color intensities (and Hues) follow the blend layer while the image structure remains mostly with the weaving layer. This is a simplification of how Hue and Saturation work and it will take lots of experimentation before you find the true usefulness of these two blend modes. [LXF](#)

SPEEDIER SCRIPTS

Practical PHP Programming



What's faster than PHP code? Surely nothing! Paul Hudson shows you how to make your scripts run 326 times faster!

Everyone knows that PHP is faster than a speeding ticket, but can it be made to go faster? C programmers have for many years trumpeted the fact that their language is extremely fast and therefore capable of handling performance-critical tasks. However, very often you'll find that when C programmers *really* need performance, they use inline assembly code.

Open up your Linux kernel source (you *do* have the kernel source to hand, right?), and pick what you consider to be a CPU-intensive operation. I chose `arch/i386/lib/mmx.c`, the code that handles MMX/3dNow! instructions on compatible chips. Inside this file you'll see lots of quite complicated C code, but also extended instances of assembly code wherever speed is optimal. In fact, if you change directory to the root of the Linux kernel tree, try this command:

```
grep "__asm__" ./-r | wc -l
```

That searches all the files in the Linux source distribution for instances of assembly, and counts the number of files that match. In 2.5.65, the number of files in the kernel source that use assembly once or more is the rather ominous number of 666! So, C programmers using assembly is quite a widespread thing.

PHP programmers, although blessed with a naturally fast language, can also use a lower-level language for speed-critical operations – although in our case, C is next in the food chain. While it's possible to use assembly code from PHP (through C, as C programmers do), there's more than enough speed improvement just switching to C, so that's what we will be covering here.

Please note that, within the extent of the space available, this is a no-holds-barred article – prior knowledge of C is required, knowledge of assembly would be good, and *very* good knowledge of PHP is mandatory. Furthermore, in order to provide the most detailed description of how things work, this tutorial has been split into two parts. I hope you will agree it's worth it!

The C Perspective

PHP itself is written in C, as are Flex and Bison (see this month's *Linux Pro*), the lexer and parser that PHP uses internally. The process of executing PHP code works by matching various parts of code against pre-defined lists of acceptable grammar. For example:

```
T_IF T_LEFTBRACK T_CONDITION T_RIGHTBRACK
T_LEFTBRACE T_STATEMENT T_RIGHTBRACE
```

In that piece of pseudo-grammar, T stands for "Type". It will match a statement that starts with **if**, then an opening bracket, followed by any boolean condition, followed by a close bracket,

then an opening brace **{**, a statement, then a closing brace **}**. Sound familiar? PHP uses the same sort of rules – although on a much more complicated level – to parse your code.

PHP has hundreds of such rules, and, when it matches them, it calls appropriate internal C functions to handle the statement. For example, when PHP matches the following rule (this is direct from the PHP source code):

```
T_DOLLAR_OPEN_CURLY_BRACES T_STRING_VARNAME '['
expr ']'
```

The *Zend Engine* will, amongst other things, call `fetch_array_begin(&$$, &$2, &$4 TSRMLS_CC)`, which uses items 2 and 4 of the rule (`T_STRING_VARNAME` and `expr`) to read and return an array item. So, as you should be able to guess, that particular handler is for accessing arrays inside strings, eg: `foo['bar']`

Because the code to execute your script is just compiled C code, it means that no matter how fast your PHP code is, it still has to be interpreted then executed as normal C code. PHP is not compiled to native machine code at any point, so there is never any chance of it out-performing C, or generally even coming close to the performance of C.

Faster and faster

So, the way to make your PHP code faster is to replace chunks of it with pure, compiled C. In PHP, this can be done in three ways: writing your own module, editing the PHP source code, or editing the *Zend Engine*.

Writing a module for PHP is the accepted way to add functionality, and there are many modules available in PHP to do all sorts of tasks. However, modules are the slowest way to add functionality, particularly if calls to `dl()` are required to dynamically load the module each time a script needs it.

Writing functions directly into the PHP source code is faster than using modules, but only really possible if you're working on your own server. Finally, writing functions directly into the *Zend Engine* provides the biggest performance boost, but basically confines your script to your own machine – not many would be willing to patch their *Zend Engine* code to try out your code! There is actually a surprising boost for shifting code into the *Zend Engine* – when Andrei Zmievski converted `strlen()` into a *ZE* statement as opposed to a function, he reported a 25% speed boost.

With such a big gain to be offered, you're probably thinking *everything* should be put directly into the *Zend Engine*. However, it's important to realise that there's a big trade-off between speed

Flex and Bison

For more information about Flex and Bison see the first part of our new *LINUX PRO* tutorial series starting on page 11.

and manageability, and generally modules come out top because they operate more than fast enough for most needs.

C vs PHP

To give you an idea of quite how much faster C is compared to PHP, I wrote a very simple C extension and compared it with its PHP equivalent. Here's the PHP script:

```
<?php
$start = time();

for ($count = 1; $count < 1000000; ++$count) {

    $j = 0;
    for ($i = 0; $i < 999; ++$i) {
        $j += $i;
    }
}

echo "PHP time: ", time() - $start, " (number: $j)\n";

$start = time();

for ($count = 1; $count < 1000000; ++$count) {
    $result = lxf_hardwork();
}

echo "C time: ", time() - $start, " (number: $result)\n";
?>
```

`lxf_hardwork()` is the module function I've written in C. Don't worry about how to create and install modules yet – we'll get to that later. For the time being, here's the source code to the `lxf_hardwork()` function:

```
PHP_FUNCTION(lxf_hardwork)
{
    int i = 0;
    int j = 0;

    for (i = 0; i < 999; ++i) {
        j += i;
    }

    RETURN_LONG(j);
}
```

`PHP_FUNCTION` and `RETURN_LONG` are both C macros to avoid lots of complicated code in source files, and they can be ignored for now. The rest of the code simply performs exactly the same thing as the PHP code, just in C – as you can see, the two are very similar linguistically.

Executing the PHP script first runs through two loops adding up numbers, then runs through another loop and calls our C function. This could have been optimised further by putting the outer loop into the C code also, but leaving it inside PHP allowed me to tweak the number of iterations without a recompile.

When the script is run, it outputs how long both PHP and C took to execute the loops. If you're not sitting down, I suggest you grab onto something before reading on!

PHP took a total of total of 1,956 seconds to run through the loops. The C code, in comparison, took just *five seconds* to do exactly the same. Of course, when you consider the loop is only 999,000,000 iterations and that this is an 800MHz PIII able

Special Warning

PHP manual – read with care!

It's not often you hear me say this, but the PHP manual is *not* the best place to check for info on writing extensions. The reason for this is because the information in the online manual is an edited version of one of

the chapters from *Web Application Development with PHP 4.0* (Ratschiller & Gerken, New Riders ISBN 0-7357-0997-1). Although *WAD* is a good book in itself, it's quite old – much of what is in there just

doesn't apply any more. The online version available in the PHP manual has a number of edits to bring the work up to speed, but the end result is that some information is correct and some is not – read with caution!

therefore to do 800,000,000 operations a second, five seconds sounds quite a lot. However the loop in `lxf_hardwork()` function compiles down to the following assembly:

```
.L319:
    addl %eax,%edx
    incl %eax
    cmpl $998,%eax
    jle .L319
```

From the label `L319`, add `i` to `j`, increment `i` by one, compare it against `998`, and if it's less than or equal to `998`, re-do the loop. So, there's actually four instructions in there, one of which is a jump, which is a branch instruction and therefore incurs more of a speed hit than the others. So, albeit somewhat simplified, I hope you can see that five seconds really isn't all that much – it's as fast as the computer could go!

In the example code above, we saw a 326x speed improvement when switching to C. Naturally the example is hardly from a real-world piece of code, but suffice to say that converting to C is likely to give a huge performance boost no matter what you choose to do with it.

Before we begin

If you're still reading, you're hopefully all set to write your own PHP extension. Extension writing in PHP is actually fairly easy, because the PHP team have put a lot of work into making the process as streamlined and foolproof as possible. Furthermore, as you'll discover, the *Zend Engine* is a remarkable piece of software that really takes much of the hard work away from programmers. You will need to have the PHP source code on your system.

For the purpose of this tutorial, we'll be creating an extension for PHP that handles tar files. To do this, our extension will use the *libtar* library created by Mark D Roth, available from <http://freshmeat.net/projects/libtar/>. *libtar* is available under the BSD licence, so we're free to use it for our needs. You'll need to have the *libtar* development files on your system.

Just to make sure we're all reading from the same songsheet, I want to briefly discuss the tar format. TAR (short for Tape ARchive) was designed to handle tape backups, but has been in general use for quite some time. Put simply, a tar file is a concatenation of files that are not compressed. Using tar, many files become one file, which can then be compressed using gzip or bzip2. Tar files by themselves are uncompressed, and approximately equal in size to the sum of the files it holds.

First steps

To get you started with a module, PHP includes `ext_skel`, which creates the skeleton of an extension. To run `ext_skel`, go into the `ext` directory of the PHP source code, then type:

```
./ext_skel --extname=tar
```

`ext_skel` creates for you a `tar.h` file and a `tar.c` file to contain our




```

PHP_ADD_LIBRARY_WITH_PATH($LIBNAME, $STAR_DIR/lib,
TAR_SHARED_LIBADD)
AC_DEFINE(HAVE_TARLIB,1,[
])
AC_MSG_ERROR([wrong tar lib version or lib not found])
])
-L$STAR_DIR/lib -lm -ldl

PHP_SUBST(TAR_SHARED_LIBADD)

PHP_NEW_EXTENSION(tar, $ext_shared)
fi

```

Near the top you can see the **PHP_ARG_WITH**(tar, for tar support, line. Other important lines are:

```
SEARCH_FOR="/include/libtar.h"
```

This locates the header file required for *libtar*, which is *libtar.h*. Also, these two lines are crucial:

```
LIBNAME=tar
LIBSYMBOL=tar_open
```

LIBNAME is used as part of the GCC compile line. In this case, **-ltar** is used. **LIBSYMBOL** should be set to a symbol contained in the **LIBNAME** library. **tar_open()** is a function contained in *libtar*, so that's what I've used for **LIBSYMBOL**. If you're wondering why this is important, configure actually writes out a short C program that calls the **LIBSYMBOL** function, then tries to compile and link that program against **LIBNAME** using GCC. If the compilation succeeds error free, it means the *libtar.so* exists and it contains the reference we're looking for, which means it's a legit copy of *libtar* for and not, for example, a file that is "Lopsided Igloo Bureau for Tuning All Radios". In other words, these three crucial lines all make sure the system is capable of compiling our extension.

Configure, compile, install, and test

Now we're done with *config.m4* – **cd** back to the PHP source directory and type **./buildconf**. This generates the configure script for PHP, and will include our new tar extension if all is well.

To make sure *buildconf* succeeded, type **./configure --help** and look for the line **--with-tar**. If the *m4* file was good, you should see the line somewhere in there, and also **Include tar support** in the column next to it. On my screen, the **Include tar support** column is one character off to the left compared to the others. If you recall, the default *m4* file had a line in there saying **Make sure that the comment is aligned** – this is what that comment was referring to. If your comment is out of alignment add or remove spaces in the *config.m4* file (line five, if you've used the above *m4* file) to correct it.

The next step is to run:

```
./configure --with-tar
```

You may want to add other PHP extensions to your configure line if you use them, but the above is enough to test our new extension.

As the output from configure flies by, you should see the following three lines somewhere in there:

```
checking for tar support... yes
checking for tar files in default path... found in /usr
checking for tar_open in -ltar... yes
```

The first line signals **yes** if **--with-tar** was specified on the command line. If **--with-tar** was used, configure checks for the location of the header file we specified (*libtar.h*), and outputs where it found it, which is line two. The final line is our library

check, and makes sure that the **tar_open** symbol is in *libtar.so*. If any of these tests fail, configure will stop with a warning, and you can read *config.log* to see where the problem is.

Once configure is complete, type **make** to compile PHP and the tar extension. **make** is likely to take quite some time, depending on the speed of your computer.

Once **make** has finished, **cd** into */sapi/cgi* – this is where the **PHP CGI SAPI** is placed once built, pending installation. Type **./php -m** to have PHP output a list of modules available – you should see **tar** in there, probably between **standard** and **tokenizer**. If so, you're successfully compiled your first PHP module!

To perform a slightly better test, **cd** into the PHP source directory and run these commands:

```
su
make install
exit
php -f ext/tar/tar.php
```

tar.php was created by *ext_skel* and calls the function **confirm_tar_compiled()**, which is a default function defined in *tar.h* and *tar.c* that that simply confirms the module compiled correctly. So, if your tar module works fine, you should see the message "Congratulations! You have successfully modified *ext/tar/config.m4*. Module tar is now compiled into PHP!" **LXF**

Make your mark

Brainstorms 'R' Us

Would you like to get your name in the mag and learn about stuff you're most interested in?

We're looking out for ideas for new *Linux Format* Practical PHP tutorials, and where better to look than to you, the reader? If, while reading past issues of *Practical PHP*, you've thought "I wish they'd covered XYZ in more depth...", or "I really want to know how to use...", then now's the time to get your voice heard!

Send an email to paul.hudson@futurenet.co.uk with your ideas – all the good suggestions that you send in will be covered in future issues. So far, the topics we have covered in some depth include MySQL, XML, CLI, GUIs, media generation, templates, and more. If you're short of ideas, you're certainly welcome to write in with comments about prior issues – we're always looking to improve the overall quality of tutorials.

--with-tar is there, although the description is a character out to the left.

NEXT MONTH

Having had a special eight page PHP tutorial in last month's issue, it just wasn't possible to run another long tutorial this time – at least not without renaming the mag *PHP Format*! So, this tutorial will be continued next month.

At this point, you've got a working extension to PHP – although it doesn't do much. Next month we'll look at how to use *libtar* in the extension by writing a function **tar_list()**. If you want to create your own extension in the future, simply repeat the steps covered in this issue – next tutorial will be *libtar*-specific.

LET YOUR LOVE FLOW

A walk in the park with Python



Patrick O'Brien continues to chaperone as the prenuptial Python action heats up with Shakespeare in the park.

Last month we talked about containers – structures designed to hold one or more Python objects. In particular, we looked at strings as containers of characters, lists and tuples as ordered containers of arbitrary Python objects, and dictionaries as unordered containers of key/value pairs. We ended with a container full of love letters.

This month we'll build on what we've learned so far, but our focus will be on action. In this tutorial you'll learn to control the flow of your Python programs: how to make decisions, how to repeat an action a number of times, and how to iterate over the contents of a container, making use of each item in that container. And to see how these concepts relate to the real world, I'll show you how all these techniques have been used in some recent and exciting developments involving the *wxPython* GUI toolkit.

The good thing is that I've got plenty of new information for you; the bad thing is that I hardly know where to begin. Should I start with the changes to the *wxPython* project, which are new and exciting and might motivate you to learn Python, or should I start with the fundamentals, which will help you understand the real-world examples? I do so want you to love Python. Where should I begin; how should I decide?

Is your love true?

Unlike my simple dilemma, decisions of the heart are often the hardest to make. You need to be a good judge of character, and be in touch with your own true feelings. The great Bard himself had to say on the matter:

"This above all: to thine own self be true, And it must follow, as the night the day, Thou canst not then be false to any man."

Making decisions in Python is, thankfully, more straightforward. All you need is an understanding of Python's notion of truth values and simple Boolean logic. "To be, or not to be, that is the question." The answer is this – I'll start with the fundamentals and leave matters of the heart to you. If your love is true, you'll read on.

Notice how my last sentence, a decision point for you, dear reader, began with the word "if". Simple decision points in Python are also tied to the **if** keyword. When you begin a statement with **if**, you are telling the Python interpreter that it should evaluate the expression following the **if** keyword, and if the expression is true, execute the block of code that follows. (This is also called 'branching' or 'conditional execution'.) Let's look at a simple example, using Python's interactive mode:

```
>>> love = True
>>> if love:
```

```
...     print "Roses are red;"
...     print "Violets are blue;"
...     print "When it comes to Python,"
...     print "my love is true."
...
Roses are red;
Violets are blue.
When it comes to Python,
my love is true.
>>>
```

Here's what happened in our example. We created a boolean object, named **love**, by assigning it a value of **True**. "True" and "False" are built-in constants representing, well, true and false. (Remember that Python is case-sensitive, and notice that the constants have an initial uppercase letter, followed by lower-case letters.) Since our love was **True**, the four print statements in the code block were executed. (In Python, code is grouped into blocks through the use of indentation – code indented to the same level, usually a multiple of four spaces, is considered to be part of the same block.)

If not, what then?

We can also have a compound **if** statement, where we evaluate more than one expression until we find one that evaluates to true (using **elif** clauses), and where we execute a final option if all previous evaluations were false (using a final **else** clause). Here is an example that evaluates the contents of a string:

```
>>> love = 'Python'
>>> if love == 'Python':
...     print "Shall I compare thee to a summer's day?"
... elif love == 'Visual Basic':
...     print "Out, damned spot. out, I say!"
... else:
...     print "Get thee to a nunnery!"
...
Shall I compare thee to a summer's day?
>>>
```

When Python evaluates an expression, it evaluates any non-zero number or non-empty container as true; it evaluates empty containers, the number Zero, and the special object 'None' as false:

```
>>> if "":
...     print "Not empty"
... else:
...     print "Empty"
...
Empty
>>> if "Something":
...     print "Not empty"
```

```

... else:
...     print "Empty"
...
Not empty
>>> l = []
>>> if l:
...     print l
... else:
...     print "Empty list"
...
Empty list
>>> l = [1, 2, 3]
>>> if l:
...     print l
... else:
...     print "Empty list"
...
[1, 2, 3]
>>>

```

Let me count the ways

The **if** statement is one way to control the flow of execution in a Python program, the **while** loop is another. Where the **if** statement lets us make decisions based on the evaluation of a condition, a while statement allows us to repeatedly execute a section of code until a condition is no longer true. Here is an illustrative example:

```

>>> count = -1
>>> while count < 5:
...     count += 1
...     if count == 0:
...         print "How do I love thee? Let me count the ways."
...     else:
...         print "I love you", count, "ways."
...
How do I love thee? Let me count the ways.
I love you 1 ways.
I love you 2 ways.
I love you 3 ways.
I love you 4 ways.
I love you 5 ways.
>>>

```

That's not exactly poetic, is it? Let's see if we can't make our love counter a little more refined, to keep Miss Browning from groaning in her grave:

```

>>> numbers = {1: 'one', 2: 'two', 3: 'three', 4: 'four', 5: 'five'}
>>> count = -1
>>> while count < 5:
...     count += 1
...     if count:
...         number = numbers[count]
...     if count == 0:
...         print "How do I love thee? Let me count the ways."
...     elif count == 1:
...         print "I love you", number, "way."
...     else:
...         print "I love you", number, "ways."
...
How do I love thee? Let me count the ways.
I love you one way.
I love you two ways.

```

```

I love you three ways.
I love you four ways.
I love you five ways.
>>>

```

In our second version we made use of a dictionary to translate from numeric values to their English-language equivalents. We also added another **if** statement handle the first time through the loop, when count is equal to zero, which we didn't define in our **numbers** dictionary. We topped it off by treating a count of one as a singular noun.

I did it for love

The last type of flow control that we will discuss is the **for** statement, which allows us to iterate over any sequence, including the contents of a container object, such as a string, list, or tuple.

Using our **numbers** dictionary from the previous example, we can extract all the key/value pairs using the following code:

```

>>> for key, value in numbers.items():
...     print "I love you", value, "ways."
...
I love you one ways.
I love you two ways.
I love you three ways.
I love you four ways.
I love you five ways.
>>>

```

In order to have an iterable sequence, we've used the dictionary's **items()** method to return a list of tuple pairs:

```

>>> numbers.items()
[(1, 'one'), (2, 'two'), (3, 'three'), (4, 'four'), (5, 'five')]
>>>

```

The **for** statement allows us to bind variable names to the items in the iterable sequence. In this case each item is a tuple containing the key and value, leading us to pick the ever-so-clever variable names, **key** and **value**. We covered dictionaries, tuples, and lists in great detail last month, so I won't dwell on them here. But I will point out one potential pitfall with the previous example.

Remember that dictionaries are inherently unordered. We happened to get lucky when Python maintained the same order in which we created our original **numbers** dictionary, but unfortunately that coincidence won't always be the case. Since you often want to work with a dictionary in a particular order, the following example overpage illustrates a better approach. In this case, we'll count down, rather than up:



To be, or not to be

That's the boolean question...

Believe it or not, up until last year Python did not have a Boolean data type, nor did it recognise **True** or **False** in any predefined way. How could this be? It's pretty simple, actually. Python's notion of truth values is broader than the simple Booleans **True** and **False**. In fact, it's usually more helpful to think of Python's truth evaluation as looking for "something" vs "nothing,"

where something evaluates to true and nothing evaluates to false. That should help you understand why empty containers, such as empty strings, lists, tuples, and dictionaries, evaluate to false.

The details surrounding the introduction of a genuine boolean type into the Python language are documented in PEP 285: www.python.org/peps/pep-0285.html

PEP stands for Python Enhancement Proposal, a formal mechanism for suggesting and debating significant enhancements or changes to the Python language. All PEPs are available for your entertainment and enlightenment at the official Python website at the following address: www.python.org/peps/

```

<< >>> numlist = numbers.keys()
>>> numlist.sort()
>>> numlist.reverse()
>>> numlist
[5, 4, 3, 2, 1]
>>> for number in numlist:
...     print "I used to love you only", numbers[number], "ways."
...
I used to love you only five ways.
I used to love you only four ways.
I used to love you only three ways.
I used to love you only two ways.
I used to love you only one ways.
>>>

```

While I could provide more information about each of the control flow statements we've just discussed, I'd rather move on to some more advanced examples. If you'd like more details, the Python tutorial, written by Guido van Rossum, also has an introduction to the control flow statements covered in this issue: www.python.org/doc/current/tut/tut.html

The real world

In previous articles I've mentioned my *PyCrust* project, which is an interactive, graphical Python shell, written using the *wxPython* GUI toolkit. There have been some interesting developments with *wxPython* and *PyCrust*, and since they relate to what you've learned about Python so far, I thought I'd share some of the details as an example of how you can put Python to use to solve real problems. To understand these details, we need to talk a bit about namespaces.

What's in a name(space)?

William Shakespeare claims "That which we call a rose by any other name would smell as sweet." And while he may be right in principle, we still have a need to name things, especially in our computer programs. The real issue for us is this: how does the Python interpreter know about "that which we call a rose?" By that I mean, when we define our own variables, and then refer to them later, how does Python know to which object the variable name is bound?

In other words, how does Python provide the context in which your code executes? The answer is a **namespace**. Python uses namespaces to keep track of names that are part of its built-in capabilities, as well as names that you define by creating variables. And it does this using dictionaries.

So in one sense, a namespace is just a dictionary. Remember dictionaries from last issue? A Python dictionary is a structure that contains mappings between keys and values. If you know the key, you can retrieve the value, in the same way that you can look up the definition for a word in a paper dictionary. Here is a simple example of creating and accessing a Python dictionary:

```

>>> d = {'x': 42, 'y': 'I love Python'}
>>> d['x']
42
>>> d['y']
'I love Python'
>>> d.keys()
['y', 'x']
>>> d.values()
['I love Python', 42]
>>> d.items()

```

```

[('y', 'I love Python'), ('x', 42)]
>>> d
{'y': 'I love Python', 'x': 42}
>>>

```

Since a dictionary is a mapping between a key and a value, and a variable is simply a binding (or mapping) of a name to an object, it should be apparent that a Python dictionary is a perfect structure in which to implement a namespace. You can see this in action by running Python interactively, either from the command line or using a shell, such as *PyCrust*.

When you run Python interactively, Python creates a namespace that is accessible using the built-in function **globals()**. This function returns the **__dict__** dictionary associated with a module or an interactive session. Here is an example of accessing the namespace of an interactive session:

```

[pobrien@localhost pobrien]$ python
Python 2.2.2 (#2, Feb 5 2003, 10:40:08)
[GCC 3.2.1 (Mandrake Linux 9.1 3.2.1-5mdk)] on linux-i386
Type "help", "copyright", "credits" or "license" for more
information.
>>> globals()
{'__builtins__': <module '__builtin__' (built-in)>, '__name__':
 '__main__', '__doc__': None}
>>> x = 42
>>> y = 'I love Python'
>>> globals()
{'__builtins__': <module '__builtin__' (built-in)>, '__name__':
 '__main__', 'y': 'I love Python', '__doc__': None, 'x': 42}
>>> namespace = globals()
>>> namespace['x']
42
>>> namespace['y']
'I love Python'
>>> namespace['z'] = 'Ziggy'
>>> globals()
{'y': 'I love Python', 'x': 42, '__builtins__': <module '__builtin__'
 (built-in)>, '__name__': '__main__', 'z': 'Ziggy', 'namespace':
 {...}, '__doc__': None}
>>> z
'Ziggy'
>>>

```

Notice how the dictionary returned by the **globals()** function changed as we created the **x** and **y** variables. And by storing the result of the **globals()** function call in a local variable (that we called **namespace**), we were even able to add to the **namespace** dictionary, creating a **z** string variable by adding it directly to the **namespace**.

If you are new to Python, you may have to reread this section a few times before it makes complete sense. But even if it doesn't, you should at least get a sense of the power that Python gives you to control and manipulate your programming environment. That power is what allowed us to remove a blemish from the *wxPython* project.

wxPython gets a facelift

If you've installed *wxPython* and looked at some of the example code, or the demo program, you may have noticed that all the *wxPython* objects have a **wx** prefix, like *wxFrame*, *wxWindow*, and *wxButton*. That's because *wxPython* came about at a time when many similar libraries recommended that you make use of them using the following import statement:

```
from wxPython.wx import *
```

We haven't really discussed modules and importing (that's our topic for the next issue), but what the 'import' statement does is import everything from the **wx** module (which resides in the *wxPython* package) into the namespace of the module doing the importing. To safeguard against conflicts between names defined in the namespace of the module doing the importing and names defined in the **wx** module, the **wx** objects were all given a prefix of **wx**.

Suffice it to say that while **import *** is a convenient shortcut, it is a potentially dangerous shortcut whose use is now frowned upon. Instead, the preferred method of importing looks like this:

```
from wxPython import wx
```

Importing the module means that you must include the module name as a prefix (**wx**) when you want to reference an object in that module. The result is two **wx** prefixes, and code that looked like this:

```
from wxPython import wx
```

```
class MyApp(wx.wxApp):
```

As you can see, the **wx** prefix on the objects contained in the **wx** module is now redundant, and having to type so many **wxs** is a pain. In fact, it was a pain that started to annoy me so much that I decided to do something about it. The solution I developed creates a new **wx** package and dynamically populates its namespace with objects from the old **wx** module, minus their **wx** prefixes.

The new wx package

Here is the magical code that you need in order to create the new **wx** package on the fly:

```
"""wx package
Provides a way to drop the wx prefix from wxPython objects."""
__author__ = "Patrick K. O'Brien <pobrien@orbtech.com>"
__cvsid__ = "$Id: __init__.py,v 1.1.2.1 2003/04/04 17:51:27 RD
Exp $"
__revision__ = "$Revision: 1.1.2.1 $"[11:-2]

from wxPython import wx

import types

d_new = globals()
d_old = wx.__dict__

for old, obj in d_old.items():
    if type(obj) is types.ModuleType or old.startswith('_'):
        # Skip modules and private names.
        continue
    new = old
    if old.startswith('EVT_'):
        # Leave name unmodified; add to the new wx
        namespace.
        d_new[new] = obj
    elif old.startswith('wxEVT_'):
        # Leave name unmodified; add to the new wx
        namespace.
        d_new[new] = obj
    else:
```

```
if old.startswith('wx'):
    # Remove the 'wx' prefix.
    new = old[2:]
    # Add to the new wx package namespace.
    d_new[new] = obj
```

```
del d_new
```

```
del d_old
```

```
del new
```

```
del obj
```

```
del old
```

```
del types
```

```
del wx
```

Let's look at the basic steps taken in this code. The name of this module is `__init__.py`, which is a special name that turns an ordinary directory into a Python package, which allows you to create a hierarchy of modules. The name of the directory that will contain this `__init__.py` file is the **wx** directory. The next version of *wxPython* (which will already be available by the time this tutorial appears in print) will create the **wx** directory and install this code. Doing so will allow an import statement that looks like this:

```
import wx
```

When you **import wx** this way, the code in this `__init__.py` file gets executed. The first thing it does is import the old *wxPython.wx* module (the old **wx** is a single module, not a package). The namespace of the old **wx** module is `wx.__dict__`; the **namespace** of the new package is retrieved using the `globals()` function. Both **namespace** dictionaries are stored in local variables where they can be manipulated.

The **for** loop then iterates over all the objects in the old **namespace**, skipping modules and private objects, leaving event-related object names unchanged (those starting with **EVT_** and **wxEVT_**), and trimming the **wx** prefix off the rest of the names. The objects from the old namespace are then added to the new namespace using the modified names and ordinary dictionary syntax:

```
d_new[new] = obj
```

The last few lines are simply housekeeping, deleting local variables that don't belong in the new **wx** namespace.

The result is a new package, with a new namespace containing clean, prefix-less names. And all of this is achieved without touching any of the existing code, ensuring backward compatibility for all existing *wxPython* programs, while allowing new code to use the new syntax immediately. The only price is a slight delay at startup as the new **namespace** is created.

What about PyCrust?

I did say that there were interesting developments with *PyCrust* as well, didn't I? Unfortunately, I've run out of space. So you'll simply have to wait until the next issue to find out! But I will leave you with a few hints. If you've used *PyCrust*, you know that it is an interactive Python shell, which means that it provides features, such as autocompletion and calltips and code coloring. And it provides these in a completely dynamic fashion – the autocompletion lists are built on-the-fly, and the calltips reflect the code as it was just defined. Wouldn't it be nice if those dynamic features weren't limited to the shell environment, but were also available when you were editing Python code. Sounds like a tasty proposition, doesn't it? That would be like having your pie and eating it too... [LXF](#)

NEXT MONTH

We look at Python's packaging mechanisms: modules, packages, search paths, module loading and reloading, and modules as singletons.

3D ART

Modelling with Blender

PART 2 Jono Bacon sketches the basic methods involved in modelling to enable you to create basic shapes.



ILLUSTRATION BY PHILIP - WWW.BLENDER3D.ORG

Last month we took a look at the magical world of *Blender*, installed it and had a look at exactly what it can achieve. To finish off, we took a glimpse of some of its toolbars and what they do. This month our goal is to get more familiar with *Blender* and to take a look at some basic modelling. By the end of this month's tutorial you will be able to create basic shapes in different perspectives.

Before you get onto the juicy stuff though, if you're a newcomer to the world of 3D modelling, have a read-through of the *Primer* boxout on the left.

Last month we had a brief look at splitting the interface and adding a shape. Split views allow us to have a different perspective on the 3D scene that we are shaping. It is important to remember that these views work together in tandem to enable you to see all sides of a 3D scene and ensure that you model it effectively.

3D modelling

A primer

3D modelling is the act of creating something in 3D, and originally would have been undertaken exactly in a manner as the name suggests by building a real representation of the desired finished article, either life-size or in miniature. Nowadays though, modelling is typically done using a 3D modelling app such as *Blender*, and the modelling process involves these steps:

- Model
- Lighting
- Texture
- Render
- Output

These processes are constructive and each builds upon the previous one with the Output stage being the final image/animation/sequence. This is how the process works:

MODEL

First we will create a number of **meshes**. These are wireframe shapes that comprise the entity that you are creating in 3D. In this stage all of

the detail of the entity is created using various tools within *Blender*.

LIGHTING

By default the 3D scene is in complete darkness, so the next stage is to add suitable lighting. This is not a simple case of just switching on a light, but there are different lamps and effects that can be used.

TEXTURE

The next stage is to give these bland objects life by applying textures and materials to them. It is this stage when the object really becomes 'real' and the quality of the textures and materials has a huge impact on how effective your 3D models are.

OUTPUT

In this final stage your 3D scene is output to a file format such as TIFF or JPG or an animation is output to a video file such as AVI. This stage is when your final work can be viewed without being first loaded into *Blender*.

Modelling

When we begin modelling, we use a variety of **primitives** – simple shapes such as cubes, cylinders and cones that can be used to build together into a more complex shape. It is a good idea before we start to just look at objects near you and try and visualise them as these primitives; you will then be able to develop the skill of seeing objects as groups of primitives; a skill that every modeler should have.

Although we can group shapes together, we get a finer grained control of our model by adjusting the separate points on the shape; points that are referred to as **vertices**. You can see an examples of vertices in **Fig 1** which shows two cubes, with the one on right showing its vertices.

Getting started with primitives

We will begin by looking at different primitive shapes and how we can view them. You will first need a brand spanking new Blender project that can be started by selecting File>New; clicking on Erase All from the popup window if required. The default view is the top view, and you can see our default plane. We will first get rid of this by pressing the Delete key and clicking on the popup window. In this

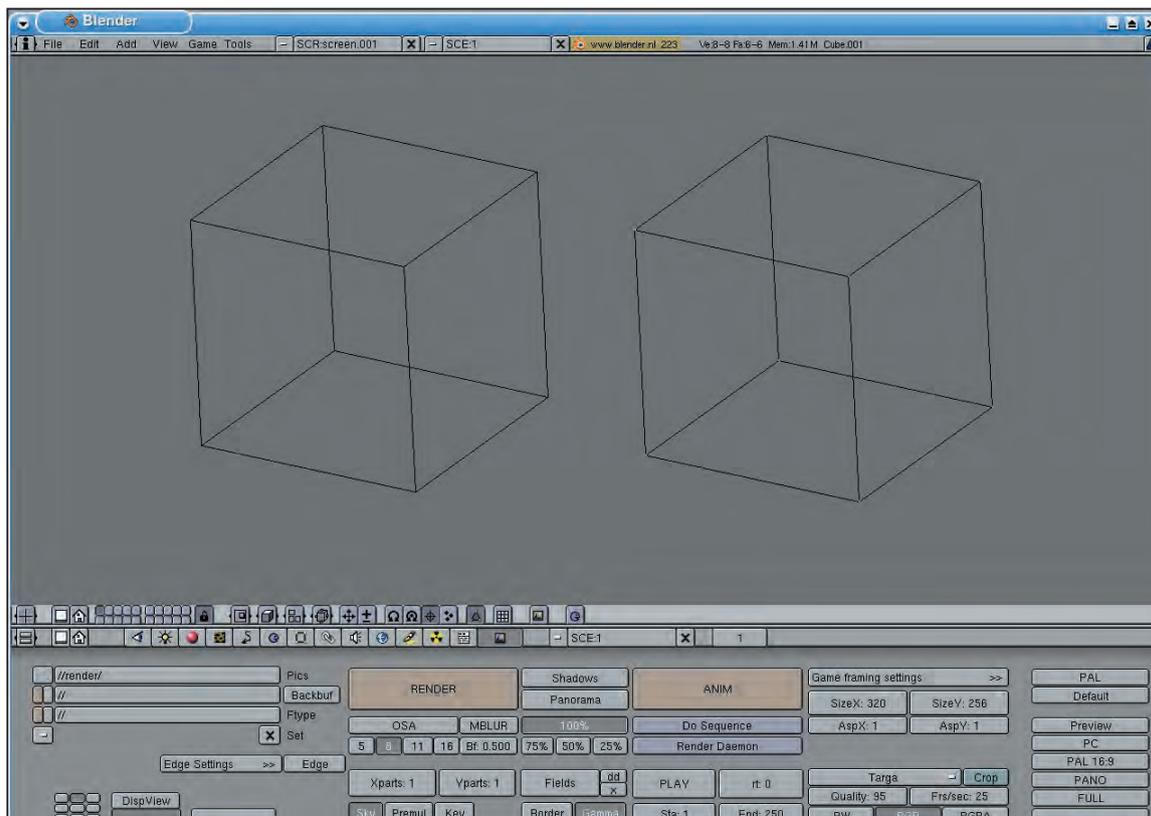


Fig 1: Vertices are the points that make up a shape.

view we can also see the top of the camera (the triangle at the bottom of the screen – we'll discuss the camera more later).

Right, let's now add a cube. To do this we hit Space to bring up the toolbox and select Add>Mesh and then cube. You will now see a square with a yellow vertices at each corner. You may be wondering why you can only see a square, but remember that you are in top view, so you can only see the top of the cube; hence the square.

We will now make use of the *Blender's* multiple views and look at another perspective. To do this we need to split the current view. We looked at this last month, but to recap – bring the mouse pointer down to the border of the topmost toolbar and you will see the cursor change to a resize pointer. You can now click the right mouse button and select Split Area from the popup window. You will then see the re-sizable divider and pressing the left mouse button will then split the view. Split the screen in half by clicking the movable divider to the middle of the window.

We now have two views and we can set a different perspective for each. On the right view you can click on the little button with a **T** in a yellow box to change the perspective. Click on the button and select the button with the **S** on it. This is the side view, and you will see the camera on the left pointing towards the cube. We will now add a third view, this with the camera in it. Move the mouse over to the right of the screen and bring it left to the divider to split it. Notice how you approached the divider from the right – this dictates which side is split. Now put the new divider halfway down the pane now creating a large area on the left and two smaller areas on the right.

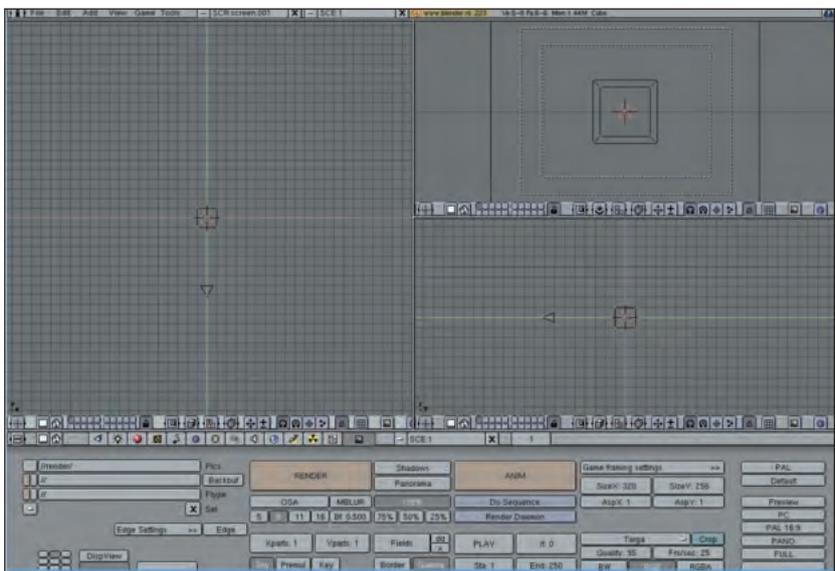
To select the camera view we need to use the Perspective Mode button; this has a 3D cube on it. Click on the Perspective Mode button below the top-right window and select the top button that has a ball and a curved line on it. This is the camera button and you will see the view change to the camera lens. Your screen should now look similar to that of **Fig 2**.

If you take a look at the camera view, you will see the front view of the cube and the back edge of the cube slightly smaller due to the perspective. It is important to remember that at the moment we are in wireframe mode, where every object is transparent.

Adjusting primitives

Now we have our cube clearly shown in our camera view, we will play with it a little and move it about. This was discussed briefly last month and will involve us going in and out of edit mode. First exit edit mode and the cube will turn pink. We can resize the cube by pressing the **S** key and moving the mouse until we have the desired size; we then use the left mouse button to select the new size.

Fig 2: A top view (left), camera (top-right) and side view (bottom-right).



TutorialBlender

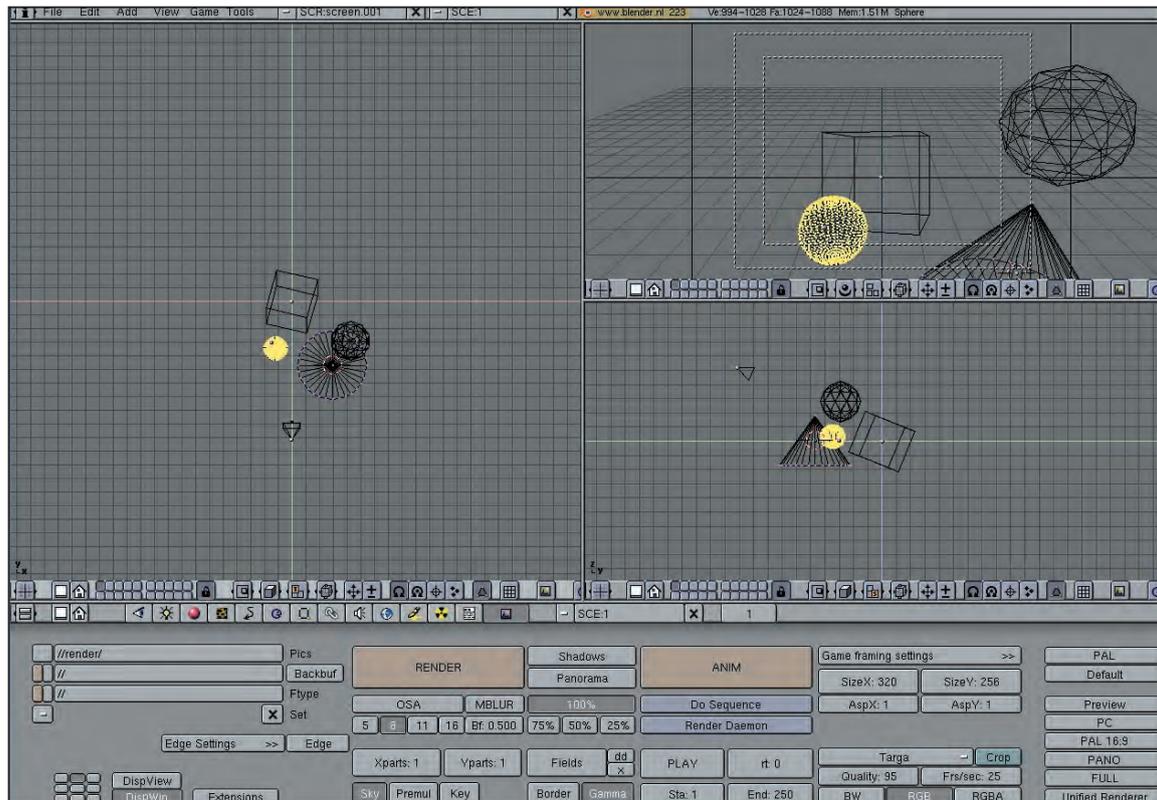


Fig3: Lots of primitives in our 3D scene.

◀◀ Notice when you resized the cube that while you were resizing it, each view is updated. This is an important feature to remember – *Blender* is a real-time modeler in that it shows the correct perspective of your object when you are modelling it.

We will now use another facility for general shape management; rotation. This is used in a similar manner and started with the **R** key. Have a play with using the rotate function in each of the different views and you will see how you can control the position of the shape to a large degree.

Using the camera

The camera is an important part of *Blender* and is essentially the eyes through which your 3D scene is viewed. The camera can be rotated and moved as well using the techniques we have covered and as such can be used to get different views on the scene. We can test this now by using the side view to move the camera up

(right-click the camera to select it and use the **G** key to move it). You can then rotate the camera to look down on the cube. You will notice when you move the camera how the camera view dramatically updates the scene. Experiment to get a feel for it.

Adding more primitives

We will now add some more shapes to see how they look in our scene. First move the 3D cursor to a new location in the top view and then use the toolbox again by pressing space to add other meshes such as a UVSphere, Icosphere, Pyramid and Cylinder. You can then use the rotation and scaling facilities to lay our objects out in the scene.

When adding the primitives to the scene, sometimes you may end up having more than one object group together when you move a shape. You can select one shape by going into edit mode by hitting Tab and then pressing **B** (selection box) and selecting the vertices of the shape by drawing a box around them – they will then go yellow to show they are selected.

Let there be light!

So far we have created a number of shapes in our 3D scene and we have our camera currently pointing at the scene. At the moment all of this has been done within the *Blender* environment and we have not yet generated any output (such as a picture). Unfortunately our 3D scene is currently wallowing in the dark and is currently unlit for us to see anything. This can be verified by hitting the Render key (F12) and you will see a window pop up with a big black area inside.

Adding light to the scene is a simple operation luckily, and can be done with the toolbox by selecting Add>Lamp. It is often best to add lamps in the top view and you will see the lamp appear in the scene. When you have added the lamp you can move it around to get a good position. You will see as you move the lamp above the shapes, there are dotted lines from lamp downwards. This is the direction that the lamp is pointing in.

Blender Resources

Get assistance from the community

There are a number of resources on the web for *Blender*, and since it was open-sourced recently, the future is incredibly bright for this powerful application. Here are some of the best resources available for *Blender*.

Blender 3D www.blender3d.org/

This is a showcase site for *Blender* and gives a taste of what *Blender* can do. This site also acts as a centralised and official point for *Blender* downloads and documentation.

Blender Community www.blender.org/

This is a well stocked *Blender* community site complete with news, forums and other facilities.

This site also houses a number of third-party projects such as the network renderer, different plug-ins and is a focal point for main *Blender* development. If you want to know the latest about *Blender*, this is the site for you.

Elysiun www.elysiun.com/

This is a computer art site with a heavy community of *Blender* users. The main benefit of this site is its popular forum and its wide range of tutorials and documentation on doing a variety of things with *Blender*. This site is a good resource for asking questions about *Blender*; there is a thriving community of people who can no doubt help.

When you render the scene, you will see that the light will hit some of the objects but not all faces of the objects. It seems that although there is light in the scene, there is not any particular direction of the light. This is indeed an issue, and we will now play with some lighting techniques to make the most out of our scene.

First, ensure that the lamp is currently selected (by right-clicking on it), and then select the lights button (one of the main buttons with a yellow emitting light bulb on it). You will see that a number of buttons will appear in the window area below the toolbar to configure the light.

There are a number of important buttons in this area, and the main ones are the Lamp, Spot, Sun and Hemi buttons. These are the different types of light available. Probably the most useful light for our current scene is the spotlight. Select the spot button and you will see that the Lamp button is unselected (thus changing our light from a lamp to a spot). You will also notice that the light in the 3D windows will have a wider projection coming out of the light. This area shows the extent of how far the spot light will reach.

Now we have a spot light, move the light so it is at an angle (and will hence get the low faces) and ensure that the circular size of the projection area (shown in the top view) covers all of the objects. You should see something similar to that of **Fig 5**.

There are a few settings that it is important that you are aware of within the light settings. These settings can dramatically affect the operation of your light. The first button is Energy slider. This slider affects how bright and powerful your light is, and the further right the slider, the stronger it is. The SpotSi slider is also useful and sets the size of the projected area that the light emits. This can be used as a strategy for only lighting the relevant parts of the scene that you choose. The Spotbi slider is used to adjust the level of edge softening that is on the outside of the projected beam of the spot light; useful for creating variations of light.

It is important to remember that most scenes will use more than one light and often use a number of lights. Try adding some more

Blender and the Game Engine

Collision detection code not open-sourced

One point of discussion that has been ripe on the Internet has been why the *Blender* game engine was revoked in the 2.26 release of *Blender*. The game engine is a powerful real-time engine for creating 3D games that are fully interactive and require little programming if any (programming is in Python where applicable).

When *Blender* was open-sourced, there was a licensing issue with the collision detection library used by *Blender* and it had to be replaced with a

freely distributable one that took time to develop. Since then the original Collision Detection library has been relicensed and can now enable the Game Engine to be brought back into the release chain sooner than thought, and has had a number of improvements added. The game engine is not something we are going to cover, so you are free to use 2.26 with it removed and it will not impact your ability to use this tutorial series. If you do want the engine though, you can use 2.25.

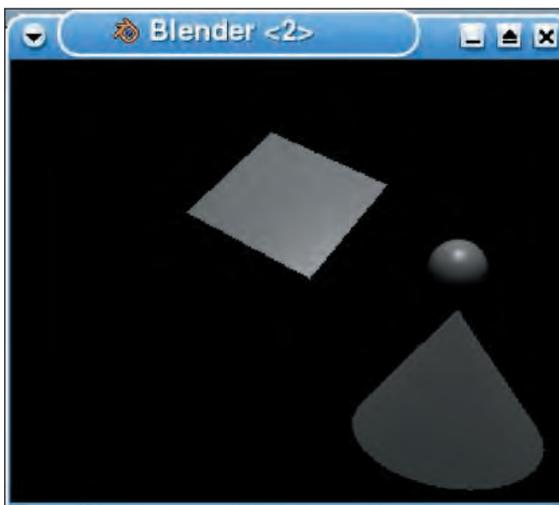
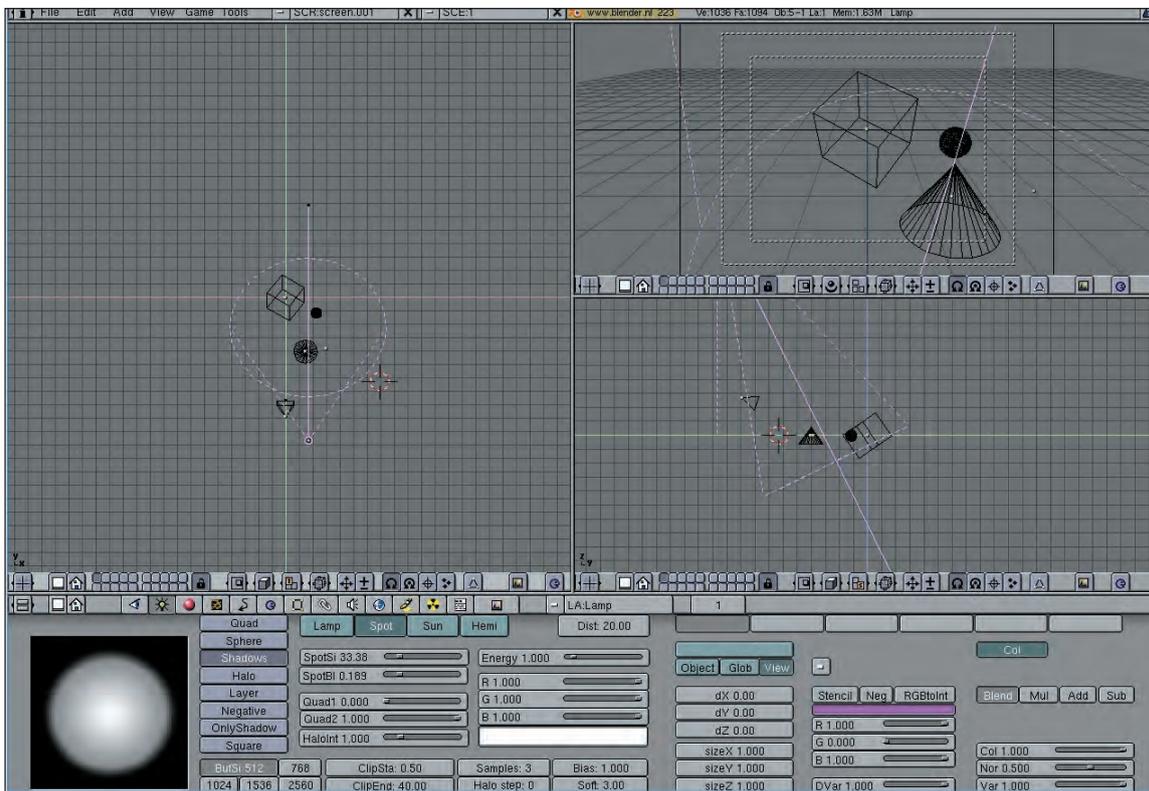


Fig 4: A rendered scene with a lamp.

spotlights, positioning them at different angles and moving the camera to different angles. Each time you change the scene, remember to render it by hitting F12. [LXF](#)



NEXT MONTH

We will be looking at some more advanced modelling techniques and better lighting. Each month we will be building on our skills and creating more and more powerful scenes.

Fig 5. Using lights within Blender.

RAPID DEVELOPMENT ENVIRONMENT

Runtime Revolution

PART 2 Get to grips with menus and images in Richard Drummond's guide to RR.

Last issue we looked at how to build a simple app in *Runtime Revolution*. in the evaluation version of this easy-to-use development suite on the coverdisc of that issue. With the basics under your belt, you are now ready to tackle something less trivial than the "Hello World" program we created.

This issue we are going to construct an image-viewer application. While this is still very simple stuff, it will give you the chance to learn how to implement pull-down menus in *RR* and how *RR* handles images – as well as picking up some more of *RR*'s scripting language, Transcript, along the way.

So here we go then. First, create a new application by selecting New Mainstack from *RR*'s File menu. It's a good idea to give objects memorable names, so, before you go any further, change the name of the newly created stack to 'ImageViewer' in the properties dialog and change its card's name to 'ImageCard'.

On the menu

Our image-viewer is going to be menu-driven, so next we want to add some pull-down menus. The *RR* tool palette has a selection of widgets for constructing menus, but a much simpler method is to use the built-in Menu Manger. This tool automates the process of adding pull-down menus to a stack – and takes care of their size and position on screen.

Open the Menu Manger by selecting it from the Tools menu. When the Menu Manager window opens, hit its New... button to create your menu. Another dialog should pop up now containing a default set of menus – shown in a list – which will be added to the stack. These default menus are File, Edit and Help. For our purposes we don't need an edit menu, so delete that from the list. While you're about it, change the menu bar's name to something more meaningful, such as 'MainMenu'. Now hit the 'OK' button. The Menu Manager window should look something like Figure 1.

Back in the Menu Manager, the list on the left shows the names of the menus that make up this menu bar, while the list on the right shows the menu items in the currently selected menu. Select the 'Help' menu in the left list, and you'll see the items '&Help' and '&About' appear on the right. (The ampersand is a place-holder and, when it precedes a letter in a menu or menu item title, that letter will be underlined when drawn on the screen; *RR* calls this a 'Mnemonic', and it is done to mark the keyboard shortcut associated with that item. The actual shortcut will be shown as suffix to the item name, separated by an oblique.) Then delete '&Help' because we don't need that item.

Next we are going to write an event handler for the 'Help' menu. With the '&Help' menu still selected in the list, click the 'Edit Script' button below, and the script editor will open.

When a menu item is selected in an executing *RR* application, that item's menu object is sent a 'menuPick' event with the name of the menu item selected as a parameter. Thus, to respond to the user selecting a menu item, we need to write a handler which accepts 'menuPick' events. We'll do this now for the 'Help' menu.

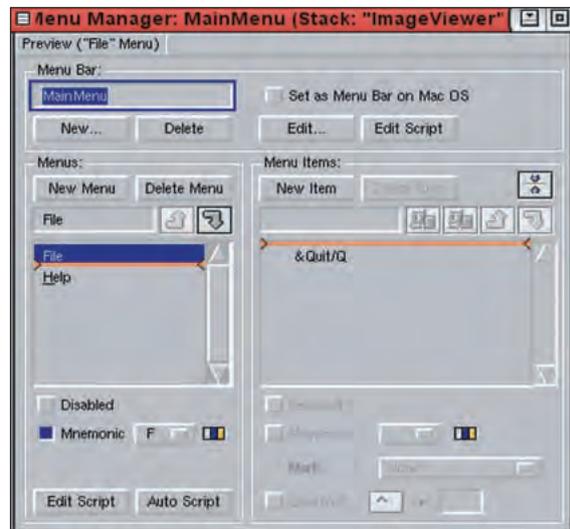


Figure 1: *Runtime Revolution*'s Menu Manger vastly simplifies the task of adding menus to an application.

In the script editor enter

```
on menuPick
answer "LXF ImageViewer"
end menuPick
```

This handler is dead easy: since there's only one menu item, we don't need to check which menu item has been selected. If we're in this handler, then the user selected the 'About' option from the 'Help' menu. Our response is to open a dialog displaying some information about the program. This is what the **answer** command above does. Look up this command in the Transcript dictionary for more information.

Next, choose the 'File' menu in the left list. We need to add an item to this menu, so click the 'New item' button above the list on the right. An entry called 'Item 2' will appear below the existing item '&Quit/Q'. In the text field above the list, change the item's name to '&Open/O'. (This is the menu entry which will open an image for display.) We want this to appear first in the menu, so click the up arrow above the list of menu items to move it up in the list. Next click the button in the top right, the one that looks like a pair of facing vertical arrows with a line between. This will add a menu divider after the currently selected item – in this case, between the open and quit items.

Image is everything

Before we can add a handler for the 'File' menu we need to do some more work. We need an image for the 'Open' option to operate on! We'll take care of this now.

Find the window that represents the application's main stack: it'll be called 'ImageViewer' and it will contain the menu bar that we have just been creating. To add an image object, select the Image tool from the tools palette (it's in the bottom left-hand corner) and

Image formats

What can RR understand?

The range of image formats that *RR* understands actually depends somewhat on the platform that you are running it on. All platforms support a core set of image formats including GIF, PNG, JPEG, BMP and XPM. In addition, when running on UNIX you can use EPS images (as long as you have the Display PostScript library, *libdps*, installed) and on MacOS (all versions, including MacOS X) you can use PICT images.

drag out a box in the stack window to create a new image object and define its position and dimensions. You want the image to fill the whole of the window below the menu bar. Open the properties dialog for this new object and give it the name 'MainImage'.

To open and display an image file in this image object we need to set its 'filename' property to the image file's file path or URL. For our app to work, we want the 'Open' option in the 'File' menu to open a file dialog to request a filename from the user and cause that file to be loaded into the image 'MainImage'. How's that done?

Well, there are several ways we could tackle this. To keep things nice and modular, we'll implement a custom event handler in the card 'ImageCard' which will respond to 'loadImage' events and set the image's 'filename' property accordingly. The menu then simply has to ask the user for a file and send that file path to the card as the parameter of a 'loadImage' message.

First the custom handler. Select the card 'ImageCard' in *RR*'s Application Overview window and click the 'Script' button above to write a handler for the card. In the Script Editor, enter

```
on loadImage imageFile
  set the filename of image mainImage to imageFile
end loadImage
```

The **on** control structure begins an event handler, in this case for our custom event 'loadImage'. This event also takes a parameter which we make available here to our handler with the identifier 'imageFile'. The body of the handler sets the 'filename' property of our image to the value of the parameter we received with the 'loadImage' message, and thus causes the file to be loaded and displayed.

To hook this up to the menu, go back to the Menu Manger and select the 'File' menu in the list on the left. Here's the script we want to attach to that menu object:

```
on menuPick pWhich
  switch pWhich
  case "Open"
    answer file "Select an image to display"
    call "loadImage it" of card "ImageCard"
  break
  case "Quit"
    close defaultStack
  break
end switch
end menuPick
```

We could enter all of this manually, but to make less work, you can use the MenuManger's Auto Script function. This will generate a skeleton handler for the selected menu – including all of its menu items. The above code was generated this way – all I had to do was fill in the bits between the **case** and **break** statements for each menu item. So, go ahead and click the 'Auto Script' button and then the 'Edit Script' button and fill in the code as above. (See **Figure 2**.)

This time our event handler takes a parameter **pWhich** which will contain the name of the menu item selected. The **switch**

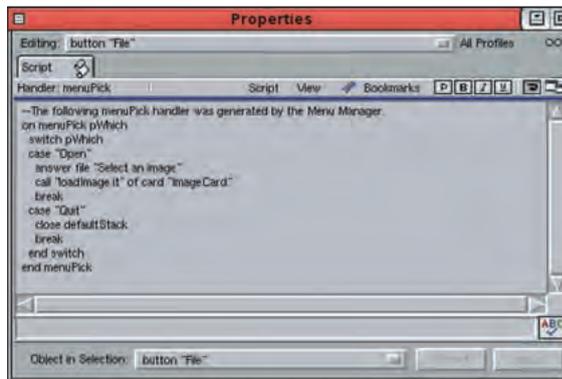


Figure 2: **Transcript** makes writing an event handler to respond to menu selections very easy.

control structure lets us examine the parameter **pWhich** and execute a different set of statements depending on its value. Each option in a **switch** statement – each case – is terminated by a **break** statement. (The switch/case construct will be familiar to users of the C language.) For example, when the value of **pWhich** is 'Quit', the statement **close defaultStack** will be executed (which causes our application to quit).

When the 'Open' item is selected, we first ask the user for an image file to load with the **answer file** statement and then we cause that image to be loaded by passing the file path to the card 'MainCard' in a 'loadImage' message. How this is done is worth examining further.

The file path returned by **answer file** isn't explicitly identified. However, it is returned in an implicit variable – the variable referred to by **it**. The keyword **it** identifies a special local variable which is used with a variety of Transcript commands to obtain a result. Here **it** contains the file path of the image that the user selected.

We pass the 'loadImage' message to our 'MainCard' object with the **call** command. This is fairly straightforward and invokes the named handler of the specified object. Here we also pass a message parameter, the value of the variable **it**.

You can now test the ImageViewer application by entering *RR*'s browse mode – either by selecting the browse tool on the palette or from *RR*'s Tools menu.

If you test out the application for long, you'll probably discover a number of problems with it. For instance, the image object into which the image file is loaded will have its dimensions set to the size of the image being loaded; nasty things happen when this is larger than the dimensions of your stack! Also, when you cancel the file dialog, the image currently loaded will disappear. Finally, the image file is persistently attached to the application. Load up an image, close the application, then re-open the application and the same image is still displayed. This is probably not the desired behaviour.

Until next time, I'll leave these and other problems with the image-viewer as an exercise for you. [LXF](#)

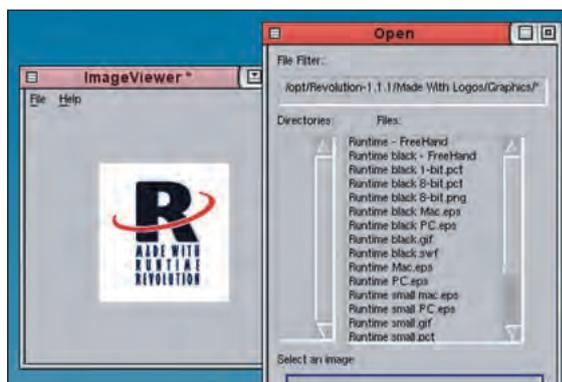


Figure 3: Our 'finished' application in action.

NEXT MONTH

In the third and final part of this short series, we'll tidy up some of the loose threads and add some final polish to your app.

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



Compiler concerns

Q I am running Red Hat 8.0 on a KT400 (Gigabyte GA-7VAXP) with an Athlon 1800. I am trying to upgrade the kernel to 2.4.20 but run into problems when building the modules. I copied the tarball from a LXF cover disk to /usr/src then:

```
bunzip2 linux-2.4.20.tar.bz2
```

```
tar xvf linux-2.4.20.tar
```

```
cd linux-2.4.20 cp /boot/config-2.4.18-14 .config
```

```
make oldconfig
```

```
make dep
```

```
make clean
```

```
make bzImage
```

```
make modules
```

The last command exits with the following error:

```
ld -m elf_i386 -r -o scsi_mod.o
scsi.o hosts.o scsi_ioctl.o constants.o
scsicam.o scsi_proc.o scsi_error.o
scsi_obsolete.o scsi_queue.o
scsi_lib.o scsi_merge.o scsi_dma.o
scsi_scan.o scsi_syms.o
```

```
ln -sf sim710.scr fake7.c
```

```
gcc -E -D KERNEL -
l/usr/src/linux-2.4.20/include -
traditional
```

```
-DCHIP=710 fake7.c | grep -v '^#' |
perl -s script_asm.pl -ncr7x0_family
```

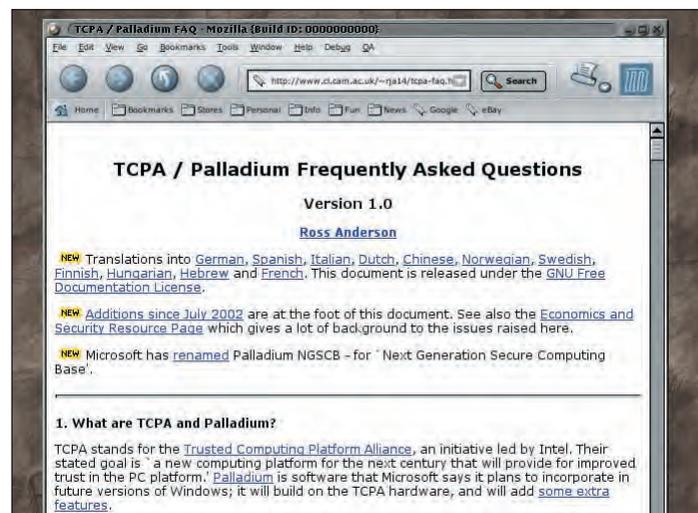
```
script_asm.pl : Illegal combination of
registers in line 72 : MOVE
CTEST7 & 0xef TO CTEST7
```

Either source and destination registers must be the same, or either source or destination register must be SFBR.

```
make[2]: *** [sim710_d.h] Error 255
```

```
make[2]: Leaving directory
`/usr/src/linux-2.4.20/drivers/scsi'
```

```
make[1]: *** [modsubdir_scsi] Error 2
```



The TCPA restricts what software can run on a specific system, and will probably stop Linux from working, particularly for those that dual boot.

```
make[1]: Leaving directory
`/usr/src/linux-2.4.20/drivers'
```

```
make: *** [ mod_drivers] Error 2
```

I've checked that I've got all the requirements listed in Documentation/Changes, and as a proving exercise, I carried out the same procedure on the Red Hat source (linux-2.4.18-14) which completed without problems. Any help would be much appreciated.

Barry Morrison, via email

A This error is due to a problem with the compiler version you are using. Specifically, compiling Linux kernels with *gcc 3.1* is likely to cause problems. You may wish to switch to a 2.95 build of *gcc*, to eliminate compatibility issues with the Linux kernel.

Of course, depending upon the specific *gcc* compiler you are using, upgrading it to a more recent version may also solve the problems. Usually the Linux kernel throws up a few compile problems not seen with other code, so it can often take a while for a specific build of *gcc* to be 'validated'

for use with the Linux kernel. This can result in modifications to the Linux kernel or *gcc*, so until everything is situated with *gcc 3* then *gcc 2.95* may result in a more productive outcome to your compile.

TCPA terror!

Q Following the information I have read on the subject of TCPA / Palladium, my understanding seems to be that their operation will depend upon the chip as well as the OS.

Question 1: Could Linux be installed and run correctly with these chips?

Question 2: Shouldn't we take advantage to purchase chips now instead of when this system will be installed, if we want to use Linux?

Marc Heerbrant, via email

A The whole TCPA 'initiative' may very well prevent Linux from running on Palladium CPUs at all. A TCPA system contains a device known as a 'Fritz' chip, which is similar in some cases to a BIOS, except that all devices need to be 'trusted' by it, as well as the operating system. Chances are, unless Linux is

certified for use with TCPA, it's not going to authenticate with the Fritz and the box won't boot up.

Assuming one can get Linux to boot on a TCPA box through fair means or foul, control then passes to the operating system, which may or may not pay any attention to the trustworthiness of certain applications.

TCPA is a long way from becoming mainstream, so it's far too early to go out and buy yourself a lifetime's supply of Intel CPUs. It's also worth remembering that TCPA is lead by Intel, and Linux runs on many other architectures other than Intel ia32 or isa64, so if at the end of the day people are not able to run Linux on their Intel boxes, then an alternative architecture may become far more popular with Linux users.

Good work practise

Q Hi, I've a few bash backup scripts which I've used on more flavours of GNU/Linux than there are mindless varieties of coffee in one of those shlock-sloppy coffee chains. A particular favourite is one that tars and zipped all the important directories,

backs these up to a zip drive, creates an image using *mkisofs*, then burns. I even used to set this as a cron job every Friday morning, but kept forgetting to stick in a CD-R disk when I booted up. But, I can't get the scripts running because of something perverse that SuSE has done to the way *cdrecord* works. All the relevant hardware belongs to users and I have even dabbled with setting the owner as myself, *cdrecord* has been similarly treated, but the only user that can burn is root, no matter what I do, such as creating a burner's group, then assign my normal user as being a member.

On every other flavour of GNU/Linux that I have used from Debian to Mandrake, I've never had this problem once I'd tweaked a few permissions on burning devices and software – this does not work with SuSE 8. Any ideas? My scripts may not be anything award-winning, but my view is: why waste time (especially with GUI front-ends) when you can script all that you need to do with a couple of lines, invoke the script and get on

```

david@macha:~ (pts/5)
mail@hacha:~$ telnet mail2.eircom.net 25
Trying 159.134.198.135...
Connected to mail1.eircom.net.
Escape character is '^]'.
220 mail02.svc.cra.dublin.eircom.net ESMTP
EHLO mail.davidcoulson.net
250 mail02.svc.cra.dublin.eircom.net
250 PIPELINING
250 8BITMIME
250 ETRN
250 XEXDATA
AUTH PLAIN
502 unimplemented (#5.5.1)

```

SMTP authentication is performed using a variety of methods, although the server may only support a few of them.

with everything else? I can only assume that SuSE have set things up rather idiosyncratically, if not to say somewhat perversely.

Any ideas, or should I just add a few lines of script to login root for the burning, then log out again? It's certainly easy to do this, but a little cumbersome if I'm trying to implement good working practise. Any advice would be very welcome. *Daithidh MacEochaidh, via email*

A This problem with *cdrecord* is not specific to SuSE, although many distributions have a work around for it. *cdrecord* requires root access to execute, so most distributions set the *cdrecord* binary to be SUID, so that independent of the individual who executes the program, it does everything as the root user. This can be a security problem, as if the program has an exploit, then it is possible to root on

A QUICK REFERENCE TO: Zebra

Having a large network quickly becomes complex to administrate, particularly when it is further subdivided into subnets for various groups or individuals. Each subnet, having its own IP block, requires a separate route on gateways and, of course, it doesn't take long to have hundreds of routing entries on multiple gateways. If one of these routes is missing or incorrect, then the particular section of the network which is referring to would become unreachable and effectively disconnected from the rest of the network.

To avoid user error, there are a number of protocols commonly used to distribute routing information over a network. The most common protocol is RIP, which is a very simple routing protocol. Any route on a router is broadcast to all other routers on the network, so there is no need to set up specific routes on

gateways as they will gather all the routing information from other gateways on the network.

RIP is a distance vector based protocol, so each route out of a router is just as good as any other. RIP decides the best route to use based upon the number of hops between the router and that network. Unfortunately this can mean that you may end up with a network with five hops to a network over 100Mbit, or two hops over a 56K modem.

An alternative to RIP is Open Shortest Path First, or OSPF. Each route has a metric assigned to it giving it a 'cost', which may be available bandwidth or actual financial cost of bandwidth over that link. A popular use of OSPF is for a redundant network, so if one link fails another takes over.

While these are standard protocols on most routers, Linux requires a service to be running in order to make use of these.



Zebra handles all the major routing protocols used by IPv4 and IPv6 networks.

Originally gated was used on Linux systems, however that is now a commercial product and the source code is no longer distributed. Instead, one has to use *Zebra*, which can handle RIP, OSPF and BGP for both IPv4 and IPv6 networks. Installing and configuring *Zebra* for RIP is fairly straight forward, although using OSPF or BGP requires knowledge

of how routers interact and how the protocol works. The main advantage of *Zebra* is that it has a configuration front end very simple to IOS on Cisco boxes, so anyone familiar with using IOS to configure their routers can easily switch to *Zebra* without having to learn most of it all over again.

Docs and downloads for *Zebra* can be found at www.zebra.org.

FREQUENTLY ASKED QUESTIONS PHP

FAQ I'VE APACHE INSTALLED. HOW DO I SET PHP UP WITH THIS?

PHP can either be compiled into *Apache*, which requires the *Apache* source code and a complete rebuild of the webserver, or it can be compiled as a loadable module, or DSO. Compiling it as a DSO allows PHP to be upgraded without having to recompile *Apache*, as you can simply replace the PHP module with an updated version.

Generally when compiling a DSO one would use the *Apache Extension System*, or APXS. To compile PHP using APXS, one simply runs `./configure` with the `--with-apxs=/usr/local/apache/bin/apxs`. To compile PHP into *Apache*, one needs to have the *Apache* source code available and run `./configure` with `--with-apache=/path/to/apache_1.3.26`

FAQ I DON'T HAVE ROOT ON THE WEB SERVER. AM I ABLE TO USE PHP?

PHP can be used as a CGI, rather than as an *Apache* module. You simply need to insert our PHP script into your `cgi-bin` directory and add the full path to the php binary as the first line:

```
#!/home/david/bin/php -q
```

PHP is compiled as a stand alone binary if `--with-apxs` and `--with-apache` are not included in the `./configure` command. Of course, one would generally want to compile

the binary on the same system that will be providing CGI capabilities so that libraries are available, although PHP could be compiled statically if necessary.

FAQ HOW DO I EMBED PHP CODE WITHIN MY HTML PAGES?

Before being able to put PHP code in HTML pages it is important to setup *Apache* to process certain documents as PHP scripts. This requires an 'AddType' directive in `httpd.conf`:

```
AddType application/x-httpd-php .php3
```

```
AddType application/x-httpd-php .php
```

```
AddType application/x-httpd-php .phtml
```

```
AddType application/x-httpd-php-source .phps
```

This will have `.php3`, `.php` and `.phtml` parsed out as containing PHP code. `.phps` scripts will be formatted and colour coded so that it is easy to read the PHP code embedded within the HTML.

We can test the PHP installation with a simple `.php` file containing

```
<?php
phpinfo();
?>
```

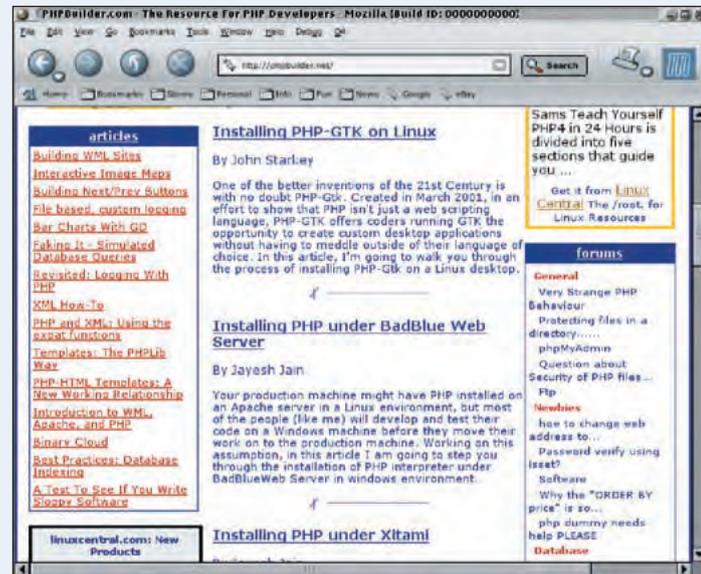
If PHP is working properly then a very fetching information page will appear, giving information about the PHP installation. If the information page fails to appear, checking the source code produced in the browser

authentication type PLAIN

What has to be set where to overcome this problem? Nothing special had to be set up when configuring WinXP's *Outlook Express*, so I presume my ISP's mail server isn't asking for anything unusual.

Within *Evolution*, under the Sendmail Tag in the editing of mail preferences mode, I have tried both checking and unchecking the authentication box – unfortunately to no effect. The mail server, using *OE*, does not require a password for sending mail.

I use *wvdial* to set up the connection, and have no references to authentication in `wvdial.conf`.



PHPBuilder.net is a great resource for PHP programmers, with example code and tutorials.

window should show the unprocessed PHP code, meaning that the web server does not know what to do with it. Either PHP is unavailable, or the above `AddType` entries are missing.

FAQ CAN I CONNECT TO A DATABASE WITH PHP?

PHP supports connectivity to a variety of databases, from `gdbm` and other flat file databases, to comprehensive database servers, including MySQL, PostgreSQL and Oracle. Access to MSSQL and *MS Access* is also possible through the use of the PHP ODBC capabilities.

FAQ HOW DO I QUERY MYSQL FROM A WEBSITE?

Querying a MySQL database is done with a number of different PHP calls. The first is to connect to the actual database, with `mysql_connect`, then to perform the query and process the results:

```
$dbh=@mysql_connect("hostname",
"user","password");
@mysql_select_db("database",$dbh);
$result=mysql_query("SELECT id
FROM mytable "
"WHERE name like \"%bob%\" "
"AND act='1'");
```

the box through it. You should be able to solve all your problems by simply doing the following:

```
# chown root.root `which cdrecord`
# chmod +s `which cdrecord`
```

SMTP? Sorted!

I have Red Hat 8.0 installed and for the most have successfully configured Internet access with *Evolution* and *Mozilla*. Incoming mail works fine, and so does *Mozilla* for Web access.

However, when trying to send email with *Evolution1.0.8-10* I get the following message:

```
SMTP server mail2.eircom.net
does not support the requested
```

Firstly this problem has nothing to do with your PPP connection with *wvdial*. The authentication which the message is referring to is the authentication your mail client is performing against the SMTP server. SMTP supports authentication, to permit you to relay messages through the server then on to the Internet. This can be done using a variety of methods, including LOGIN, CRAM-MD5 and in your case PLAIN. PLAIN text authentication simply sends the username and password over the connection in plain text. LOGIN isn't much better, since it simply runs the username and password through a base64 encoder.

```
if(@mysql_num_rows($result)==0) {
    print "<b>No rows</b>";
    exit;
}
print "<h1>Results</h1>\n";
while($row=@mysql_fetch_array(
    $result)) {
    print $row[id]."<br>\n";
}
```

FAQ CAN I WRITE STAND-ALONE SCRIPTS WITH PHP THAT DON'T NEED APACHE?

As with running PHP scripts as CGIs, we can run them as regular command-line scripts through a PHP binary. When running PHP scripts on the command line we can make use of other capabilities, such as process control, allowing us to handle signals and fork new processes.

```
<?php
$pid = pcntl_fork();
if ($pid == -1) {
    die("could not fork");
} else if ($pid) {
    // we are the parent
} else {
    // we are the child
}
?>
```

For UI integration, GTK+ libraries are available for use with PHP, allowing graphical applications to be written in PHP and have them executed from the command line.

For more on PHP, read *LXF's* regular PHP tutorial series and read *PHP And MySQL Web Development* that's reviewed on page 34 this issue.

CRAM-MD5 is the most secure authentication method over a non-SSL connection, as the actual password is never sent over the connection, as a hash is built against a token sent by the server. This stops what are referred to as 'replay attacks' where someone captures your traffic and just dumps it back to the server in an attempt to authenticate against it.

As the server you're using does not support PLAIN method authentication then you will need to configure your client to use a different authentication method over the connection. Of course, depending upon the actual mail server, it may not support authentication at all. The

exactly requirements depends upon the configuration of the server.

We would suggest that you check with the administrators of the mail server to find out how to appropriately configure your mail client. Looking at the response we get when we sent an **EHLO** command to the mail server, it does not support authentication:

```
250-mail04.svc.cra.dublin.eircom.net
```

```
250-PIPELINING
```

```
250-8BITMIME
```

```
250-ETRN
```

```
250 XEXDATA
```

If the mail server belongs to your ISP, then you probably don't need to authenticate against it, assuming you're mailing from their IP block.

Print with Samba

Q Please could you help me with the following problem: My home network is

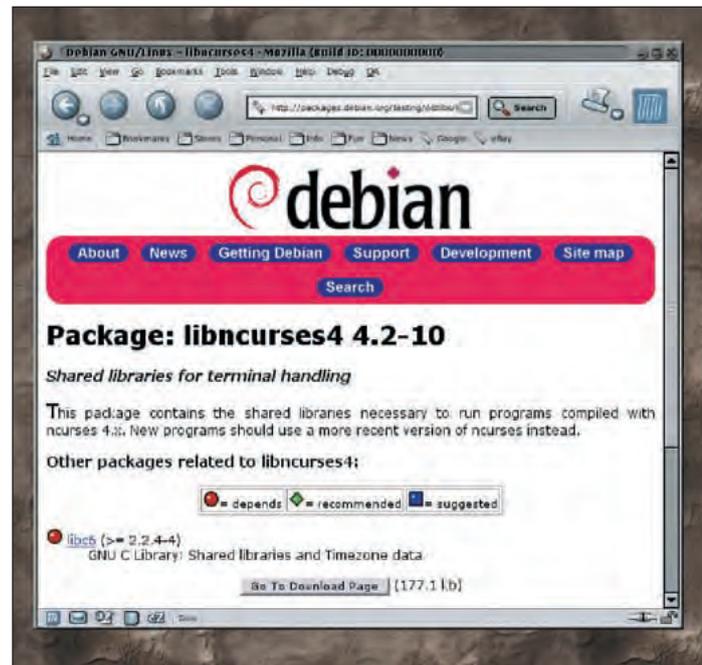
currently in a state of flux, moving from a Windows 2000 peer-to-peer workgroup to Mandrake Linux 9.0. My printers are connected to machines still running Windows 2000, and they are shared and accessible from other Windows 2000 machines. I have CUPS installed on the Mandrake Linux 9 machines, but cannot work out how to configure them so that they can print to the printers shared by Windows 2000. Please could you illustrate what should be configured on a Linux PC to access a printer shared by Windows 2000 (or, I guess, Windows NT or XP). Thank you in advance.

Kevin Lawton, via email

A Printing to a Windows box from Linux is done using the *smbspool* utility, which is part of the *Samba* suite. This allows print jobs to be sent to a Windows system over your network.

The first step to solve your problem is to symlink the full path to *smbspool* to */usr/lib/cups/backend/smb*, so that CUPS knows where to find *smbspool* to print with. You can then go to the CUPS administrative interface and add a printer, giving the device in the format **smb://server/printer** including any required authentication information in the standard URI format

```
smb://user:pass@server/printer
```



libncurses is used by console applications, but older binaries compiled against earlier versions may not work.

You can find more information on this at <http://mumford1.dyndns.org/~bs7452/linuxhelp/cups.html>

Debian depends

Q The magazine gets better and better, keep up the good work!

I have been trying to get *WordPerfect Office 2000* onto Debian 3.0, but get a dependency problem. My query concerns *libncurses4*; now, as you are probably aware, Debian has *libncurses5*, which I assume is not backwards-compatible? Can I therefore put *libncurses4* into Debian 3.0? If so how, and where can I get *libncurses4*? I do not seem to get anywhere searching the Internet for *libncurses4*.

John Bywater, via email

A Looking through the information that is provided at <http://packages.debian.org/> there is indeed a *libncurses4* package available for backwards-compatibility with older programs which are not linked to *libncurses4*.

However, this is only for the 'testing' distribution of Debian, rather than 3.0, which is the stable tree. You can install this side by side with *libncurses5*, by simply doing:

```
apt-get install libncurses4
```

As it is a major version number change, it is not improbable that there

are binary incompatibilities between *ncurses* versions 4 and 5. As there is a separate version of *ncurses4* available for Debian, then there are likely to be incompatibility issues.

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: lxformat@futurenet.co.uk

Answers



Hijack security

Q Thanks for providing the GNU/Linux community with such a great reference and resource. My question deals with what is going to be a huge networking issue as the community grows and technology becomes more available (and cheaper).

I run a Linksys wireless network that connects three Windows computers and my GNU/Linux box, which uses an Orinoco Gold card. I am running a complete install of Red Hat 8.0. I have a Symantec (Norton) Firewall and have antivirus software on all of my Windows boxes. After I installed and registered Red Hat, I decided to use my 60-day telephone support because I couldn't find any information on security on a wireless network that has another OS's firewall and A/V software running on it.

I called Red Hat and spoke with a tech. After explaining what the problem was with my setup and then asking about a recommended security software app I could use to prevent attacks (which I've been told are frequent), he stated "You don't need any firewall or antivirus software; just ensure that you keep your kernel updated."

I am worried that it sounds all right, but was he correct? So, can you tell me what I need to do to protect my wireless system that runs Windows Me and XP and Linux? I don't want to change what I have, just add what I need to my Linux box.

Any help is greatly appreciated!
Fred Gibbs, Bedford NH, USA

A There are two different security aspects you need to take into account in your environment. First, the Internet is rather an insecure place and there are several ways an intruder will try to gain access. You also need to bear in mind that somebody within range of your wireless network could potentially

'hijack' a connection to your network. I'll cover both aspects briefly.

You neglected to mention how you connect to the Internet. I'll assume that you have some sort of router/gateway in place already. If this device assigns you a non-routable IP (an address such as 192.168.x.x or 10.x.x.x, amongst others) then you already have a certain amount of security as no attacker from the Internet can access these PCs directly. However, this protection is only as good as the device providing the network address translation. If you have a real Internet IP address, then you really need some sort of firewall. In theory, if you make sure that only the absolutely required network services are running and that every package, which can listen on the network, is up to date you should be safe. But in practice, most people don't update packages immediately after errata are released; even if you did, it would still not protect you from a Denial of Service (DoS) attack.

So a firewall is always a good idea, in particular as *iptables* comes with most modern distributions. There are several good *iptables* scripts available on the Internet so sniff around on Google and you shouldn't be disappointed. With regard to wireless

security you should make sure that you change the default SSID of your access point. Anybody with a sniffer will still be able to pick up your network though. Turning on WEP and setting a 128 bit encryption key should prevent your traffic being sniffed. As an added security setting I would highly recommend turning on MAC address filtering if your access point supports it. MAC filtering will allow you to set which MAC addresses your access point will talk to. You will need to add the MAC addresses of all the wireless adapters on your network. If a MAC is not listed then the Access Point will not communicate with it.

New possibilities

Q After converting my office machine to Linux, I want to be able to mount a directory onto my machine at home via the Internet. Both boxes are running Debian Woody. I also have a firewall running Smoothwall protecting my home machine. Is it within Linux's capabilities to mount a directory like this in a way that is both secure and reliable? As being able to do so would open up a whole new window of networking possibilities!

Luke Woods, via email

A It is possible to connect two networks (or two computers) together securely using the IPSec protocol. This has been implemented into Linux under the guise of FreeS/WAN. Basically a secure tunnel between the two networks is established. FreeS/WAN provides security by being able to tell

if a packet has been modified en route between the two networks. This would pose a problem on your network as *Smoothwall* probably does masquerading (NAT) on your home PC's IP address. Although it is possible to run FreeS/WAN through a NAT'ed network it's tricky and not recommended. You would probably find it easier to load FreeS/WAN onto your firewall directly. I'm not certain if *Smoothwall* supports FreeS/WAN so you may be better off loading a pure Linux gateway, which can do both NAT and IPsec. That way you can connect to your home gateway from work and browse your home PC's Samba shares that way. I'd highly recommend taking a look at www.freeswan.org as they have excellent documentation on their site.

Red Hat RPMs

Q OK, there's gotta be an obvious answer that I'm missing. How can I tell which x86 optimisation an installed RPM has? For example, for *glibc* on Red Hat 7.2, I could be using either:

```
glibc-2.2.4-32.i686.rpm
glibc-2.2.4-32.i386.rpm
```

`rpm -qi` gives me all kinds of info including the source rpm used, but does not tell me if I installed the right one for my CPU. And since i386 works all ix86, how can I tell?
François Caen

A You can supply a custom query format to rpm in the command line. For example:

```
$ rpm -q --qf '%{NAME}-
%{VERSION}-%{RELEASE}-
%{ARCH}'n' glibc
```

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



WIN A SHARP ZAURUS PDA FROM RACKSPACE MANAGED HOSTING

www.rackspace.co.uk

Every month, the best question related to Systems Administration that a LXF reader sends in wins a Sharp Zaurus SL-5500 running Linux.

With industry support from the likes of IBM, SAP, Oracle, and Sybase, the Zaurus features a QWERTY keyboard behind a sliding cover. The 65,536 colour screen is 240x320 resolution. An Intel® StrongARM® 206MHz processor and 64MB memory makes the Zaurus ideal for downloading software or playing MP3 and MPEG1. The docking station that fits into its USB connector, allowing users to synch with their desktop,

schedules, address book and ToDo lists. The email app supports POP3 or IMAP email servers. Write and send emails and browse the Internet with a compatible modem card, or using the infra-red port and a mobile phone; or any replied, forwarded, or composed emails are transferred to your outbox the next time you synch. SD and CF slots allow upgrades and addition of software.

Email all your sysadmin questions and problems to Hans Huberland: sysadminqa@rackspace.co.uk Every published query on these system admin pages wins a Rackspace T-shirt.

It's only 74mm wide by 138mm long and a svelte 18mm thick – so the Sharp Zaurus really is Linux in your pocket!



★ Star Question – PDA winner!

This issue's lucky winner is **Robin Brands** – your new Zaurus PDA will be with you shortly!

I recently completed a project to set up several diskless PCs on a mainly Unix-based LAN. The PCs boot Slackware Linux from a Sun Ultra 5 using Etherboot. This all works well, and has given us some very cheap X terminals by recycling our redundant Office PCs, but I now want to use these PCs to access a non-X Windows Host via Telnet. My question is, how can I use one of the virtual terminals to automatically telnet into a Host on the network without a logon to Linux (ie telnet as the shell, not *bash* or *csH*)? I want the

end user to select, say <CTRL><ALT> <F3>, and telnet straight into the remote HOST. I've tried editing the */etc/inittab* to have the following line

```
c3:34:respawn:/bin/telnet
192.168.9.240
```

Needless to say, this did not have the desired effect, I still got a Linux login prompt. It's probably worth me pointing out that I am aware of potential security problems with what I am attempting to do, but these are not an issue with this particular LAN. The network has no external

connections, and very few users, all of whom can be fully trusted not to tamper in areas not concerning them.

Robin Brands, via email

The configuration you have is nearly correct. *init* does not call *bash* directly. *init* calls *mingetty*, *agetty*, or another *getty* then calls *login*. *login* checks */etc/passwd* and gets your shell from there. The format of entries in *inittab* is as follows:

```
name:runlevel:type:command
```

Although the name of your action is *c3* this does not tell it to run on *tty3*. You can make it run on

any *tty* you choose by redirecting STDIN, STDOUT and STDERR for a particular binary there. The first field (name) is simply a four character identifier. Your *inittab* entry should read:

```
c3:34:respawn:/bin/telnet
192.168.9.240 >/dev/tty3 2>&1
</dev/tty3
```

Bear in mind that if the machine is booted it will keep respawning telnet connections. If the telnet daemon on the other server has a short timeout and logs these timeouts it could quickly create very large log files.

should return `glibc-2.2.5-42-i686` or `glibc-2.2.5-42-i386`.

There are many more fields you can grab from your RPM database. For more info see the Maximum RPM book online at www.redhat.com/docs/books/max-rpm/

Flamebait?

Q Due to the many problems with *Sendmail*, I have finally decided to switch to something else. I know of three alternatives; *qmail*, *exim* and *postfix*. I was wondering which would be best recommended? I wish to have *procmail* integration,

and *Maildir* support would be a definite bonus as my current email files take a long time to go through, although I believe that this could be done with *procmail*(?). Ease of installation would also be a deciding factor, however it isn't 100% necessary.

Dave Wickham, via email

A This is a very difficult and controversial question. Almost as dangerous as "which is the best Linux distribution?" *Qmail*, *Exim* and *Postfix* all support *procmail* integration and the *Maildir* format. All three claim to be easy to use (easier than *sendmail*) but your mileage may vary. I personally find *sendmail* with

m4 very easy to configure but that's because I've been using it for some time. *Qmail* has some very good virtual hosting plugins available and both *Postfix* and *Exim* are very popular with people moving from *sendmail* due to similarities in how they operate. It's also worth noting that only *Exim* is licensed under pure GPL. *Postfix* and *Qmail* both have other more restrictive licenses although they are still free.

Teen angst

Q I work for a company where I am required to enforce content filtering, as there are loads of teenagers under the age of 16 whose curiosity causes me

headaches. What content filtering software would you recommend for quality and value for money?

Bradley Robinson, via email

A Hi Bradley, check out LXF39's DVD. It contains a distro called *Censornet*. If you don't have the DVD you can download it from www.censornet.com. The software performs exactly what you are looking for (and then some) by acting as a gateway to the Internet. You can set permissions and do logging per computer for *www*, *irc*, and much more. The company also offers a blacklist update service, which keeps an up-to-date list of offensive sites for approximately £100 per year.

Answers

MDK booting bother

Q I have installed Mandrake 9.0 from the December LXF disc, I had a problem in that it would not boot from the DVD.

I managed to install it with a boot floppy made with rawwrite.

I have tried to upgrade KDE using the LXF 38 DVD, and have encountered the following problem, I have used the following:

Opened a root terminal on the desktop, then

```
cd /mnt/cdrom/Desktop/KDE3.1/
Mandrake to enter the Mandrake directory.
```

```
rpm -Fhv *.rpm i18n/*.rpm
```

I then get a message:

Failed dependencies and a long list lib files required for programs the sort of file required are libaudio.so.2 . libmad.so.0, libstdc++.so.5, libxft.so.2, libfontconfig.so.1, lib.so.6glibc.2.3, libstdc++.so.5 glibc_3.2.2, Mandrake mime, KDroot warning, kdbase-service menu etc.

I have also tried `rpm -Uvh *.rpm i18n/*.rpm` this gives the same result but includes a list of i18n files. This is the first time I have tried installing from the command line. Any ideas?

Ken Weeks, via email

A It looks like your Mandrake installation is missing a great deal of RPMs. These will be available from your installation DVD and can be installed with the `rpm -Uvh <file>.rpm` command. As boring as it sounds, you're going to have to go through each dependency in turn and install the specific required RPM to fulfill that dependency. If you prefer you can use the Mandrake Control Center to install them for you, although you will still have to pick each individual RPM out in turn.

XINE DVD playback

Q Firstly, when we download a source file what is the command to make source file into rpms ?

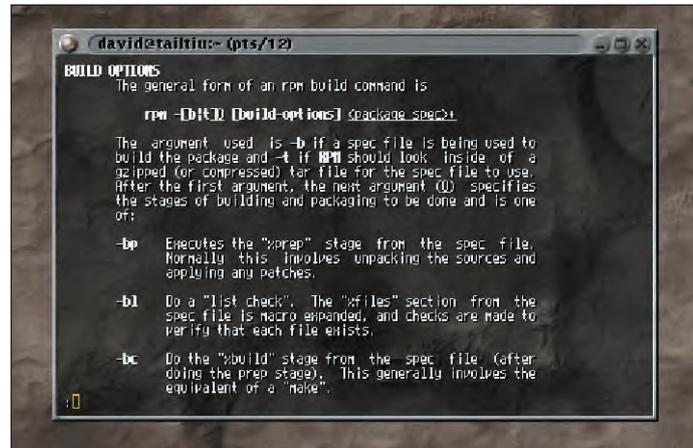
When I download `fvwm`, I do the following:

```
tar vxfj fvwm.bz2
```

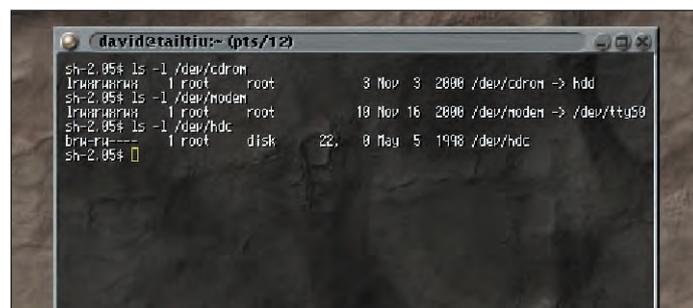
```
./config
```

```
make
```

```
make install
```



Building a RPM from the source tree usually takes a number of steps, but `bb` does it all for you.



`/dev` entries are either device nodes, with a major/minor number set, or a symlink to another device.

`make clean`

but the files are everywhere inside `/usr/bin` on the computer. What happens if I want to install a newer version of `fvwm` and some file are new? This way of installing is just too messy.

I understand that there is a `/rpm` directory (after `untar` it) with all the files required to make a rpm but I am just not sure what the command is to do so.

Secondly, when I was reading your tutorial on `devfs` (March 2000) I went to see my `/dev` and couldn't find folder for my dvd (`/dev/dvd`) but only `/dev/cdrom` subsequently I can't watch DVD on XINE. I went to XINE FAQ and it says that I can playback DVD by doing `ln (I don't know if a flag is needed) /dev/dvd /dev/hdc` (`hdc` is where my dvd is). I thought that Mandrake 9 will automatically detect that it's a DVD-ROM instead of CD-ROM. according to the tutorial in March 2002 I can just do:

```
mknod /dev/dvd
```

```
ln /dev/dvd /dev/hdc
```

and that's it, or should I do the

registration as described in the tutorial? If so can you give more detail please?

This might sound a little dumb, but I was just wondering if `ppp` can just create `dvd` in `/dev`, then how does the kernel know that `/dev/dvd` is DVD but not other stuff? For example, what if you put:

```
ln /dev/something /dev/hdc
```

Is kernel that smart to know that something is dvd? I am just not understanding the `devfs` concept of making everything as a file.

z2226477, Australia

A Taking your first question, you can build a RPM from the source code with the `rpm -bb` command. There are numerous stages of an RPM build which you can step through in turn with `rpm`, but `'-bb'` will build everything and produce the RPM for you. This should be done from the top level directory of the source code, which will contain a `*.spec` file, which is the build instructions for the RPM. If you want to build from source without an RPM, it's usually best not to install it into `/usr`, as this is where packages are installed. Use `/usr/local/<package>/` instead so each page is in its own directory and can easily be removed if

you don't want it anymore. Then simply add `/usr/local/<package>/bin` to the `$PATH` variable.

As for your DVD/CD question.

`/dev/cdrom` will be a symlink to your actual device, such as `/dev/hdc`. This link will have been created during the installation, and it makes no difference to anything if it's a CD-ROM or a DVD drive actually attached to `/dev/hdc`.

Your instructions for creating `/dev/dvd` are slightly skewed, as you only need to do

```
ln -sf /dev/hdc /dev/dvd
```

The kernel doesn't know anything about `/dev/dvd`, as anything which accesses `/dev/dvd` will really access `/dev/hdc` - `/dev/dvd` is simply a convenient alternative name for the device. If you want to create something in `/dev` then you either need to make the device with `mknod`, which is passed a major/minor number which defines the device, along with an entry to define it as either a block device or a character device. Something like a serial port is a character device and a hard drive or CD-ROM drive is a block device. The major/minor is what is used by the kernel to figure out what the device actually is and how to access it.

Generally you will not want to use `mknod` to create devices unless you know what you're doing, but you may want to use `ln` to create links to specific devices with more useful names, such as `/dev/mouse`, `/dev/modem`, `/dev/zip` and so forth.

If you're using `devfs`, note that when you reboot, any modifications you make within the `/dev` directory will disappear. You will want to look at `/etc/devfs`, as there are entries in there which allow you to automatically create the symlink when the device is registered the kernel. **LXF**

Posting to the forum

The LXF online community

Got a technical question? Other LXF readers may be able to help!

The forums at www.linuxformat.co.uk have a section dedicated to technical queries, hardware, programming languages and general help. As well as being able to call on `lxfadmin` (when there's no deadline!) and the ever-present 'anonymous', the forums are also frequented by Linux heroes like **Jeremy, Nelz, Fingers99, Rhakios, Erin** and many others brimming with knowledge and experience of using Linux in a wide variety of situations.

missed one?

LINUX FORMAT BACK ISSUES

Every month *Linux Format* brings you the best tutorials, the essential reviews and the latest news. But if you've missed out on a must-read feature or a vital programme from our expertly compiled CDs and DVDs, order your back issue NOW! And remember, you need never miss an issue of your favourite Linux mag, if you subscribe to *Linux Format* (see overleaf for more details).



May 2003

Product code:
LXFB0040(cd)
LXFD0040(dvd)

DVD HIGHLIGHTS:
Vector Linux, Dropline
GNOME, Lopster, jEdit,
QtParted, QuickRipDVD,
BaitAndSwitchHoneyPot,
AdiosBootCD, zTappy,
MondoRescue

MAGAZINE FEATURING:
SCO v IBM \$1 billion lawsuit,
Blender tutorial, LXF Awards
2002 results, Linux on Apple
Mac, Gigabyte SRI13 server,
EFI to replace BIOS, VariCAD
& IEMS mailserver reviews

CDs HIGHLIGHTS:
Mozilla 1.3, Opera
update, IMP server,
Kino, IceWM, ECLiPT
Roaster, VariCAD



April 2003

Product code:
LXFB0039(cd)
LXFD0039(dvd)

DVD HIGHLIGHTS:
CensorNet, K3b, VIPS,
FrozenBubble, Linux
Open Media Box, IDMS
Linux, WhiteDune, OBM,
AstaroSecurityLinux

MAGAZINE FEATURING:
Customising Linux, Evesham
Linux PCs, Open publishing
with WIKI, Linux PIMs and
PDAs, SME Server review,
File Managers Roundup,
Serial Networking, GnuCash

CDs HIGHLIGHTS:
SME Server, TightVNC,
GNOME 2.2, VOCP,
CoyoteLinux, UT2003
update, Blender, Qt,
BootEverywhereLinux



March 2003

Product code:
LXFB0038(cd)
LXFD0038(dvd)

DVD HIGHLIGHTS:
Python, The GIMP,
Lycoris, Yellowdog,
Knoppix, Cygwin,
DansGuardian, Senken,
AbiWord, X Northern
Captain, Eclipse

MAGAZINE FEATURING:
You, Copyright & The Law,
X apps on Mac OS with Fink,
SVG versus Flash, Audio
Editors roundup, Python
tutorial, GIMP secrets, SuSE
Office Desktop examined

CDs HIGHLIGHTS:
KDE 3.1, UNIX Tools For
Windows, Konstruct,
phpMyAdmin, X Net
Strength, AutoUpdate



February 2003

Product code:
LXFB0037(cd)
LXFD0037(dvd)

DVD HIGHLIGHTS:
IPCop, Postnuke,
Lllypond, Squid, Allegro,
Cocoon, Fandango,
gPhoto, Xrmap, Ted,
AdvancedPHPdebugger,
Flightgear updates

MAGAZINE FEATURING:
Lindows on test, Serial ATA,
Window Managers roundup,
Arkeia review, Xft2, build your
own Linux backup box, Java
compiling, Jool Kwartz review,
New Larger Linux Pro mag

CDs HIGHLIGHTS:
Yoper, KOffice 1.2.1,
Parted, DOM ToolTip,
Guarddog, GL-117,
PythonImagingLibrary



January 2003

Product code:
LXFB0036(cd)
LXFD0036(dvd)

DVD HIGHLIGHTS:
Evolution PIM/calendar/
email client, Freevo
personal video recorder,
OpenZaurus, PilotLink

MAGAZINE FEATURING:
Optimisation tips and tricks,
Netbox Cubit review, build
your own desktop, Xandros,
Cubit, CORBA, Realsoft 3D,
Instant Messenger roundup

CDs HIGHLIGHTS:
Racer, Freeduc, Linux
From Scratch, GNOME
Meeting, Opera 6.1, Film
GIMP, phpOpenTracker,
Tuxpaint, Snort



Christmas 2002

Product code:
LXFB0035(cd)
LXFD0035(dvd)

DVD HIGHLIGHTS:
Mandrake 9.0 - 3 CD
download edition, jEdit,
BastilleLinux, Phoenix,
VideoLAN, KDE 3.0.4

MAGAZINE FEATURING:
Huge C++ IDE RoundUp,
LXF Awards nominations,
gadgets for Xmas, User
Mode Linux, compiling code,
process accounting

CDs HIGHLIGHTS:
3CDs, including 2CD
Mandrake 9.0, SpamX,
Stellarium, Krename,
ImagePress, Nogger,
DigitalMusicCentre

Other back issues also available, please call for further information, giving the issue number and/or date required

To order, call the back issues hotline on **0870 4448645**

or **+44 1458 271178** for overseas orders

Remember to quote the issue code shown above and have your credit card/switch card details ready.

Prices (inc post and packing): CD – UK£5.99,

US & Eur £6.99, rest of world £7.99

DVD – UK £6.99, US&Eur £7.99, rest of world £8.99.

LINUX FORMAT

www.linuxformat.co.uk

SUBSCRIBE AND

ON A YEAR'S SUPPLY OF LINUX FORMAT

That's 3 issues free!



Title _____ Initials _____ Surname _____

Address _____

Postcode _____ Country _____

Email address _____

Daytime telephone number _____



Media with passion
Future Publishing Ltd, Cary Court,
Somerton, Somerset, TA11 6BR



Instruction to your Bank or Building Society to pay Direct Debits

Originator's Identification Number

7 6 8 1 9 5

Please send this form to address shown

1. Name and full postal address of your Bank or Building Society branch

To: The Manager

Bank/Building Society

Address

Postcode

2. Name(s) of account holder(s)

3. Branch sort code

(from the top right-hand corner of your cheque)

4. Bank or Building Society account number

5. Instruction to your Bank or Building Society

Please pay Future Publishing Direct Debits from the account detailed on this instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Future Publishing and if so, details will be passed electronically to my bank or building society.

Signature(s)

Date

Ref No (Office use only)

Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

You may photocopy this blank form and then complete it. You cannot fax or e-mail Direct Debit forms.

QUOTE ORDER No. **LXFP41**

LINUX THE UK'S BEST-SELLING LINUX MAGAZINE

FORMAT

Please choose your method of payment **1** or **2**

1. Direct Debit (UK only – double CD Edition) – I wish to pay £27.45 every 6 months

or Direct Debit (UK only – DVD Edition) – I wish to pay £29.95 every 6 months

Please complete the form on the left

I understand that I will receive 13 issues during the next 12 months.

2. Cheque/credit card 13 issues

	CD	DVD
UK	<input type="checkbox"/> £54.90	<input type="checkbox"/> £59.90
Europe (including Eire)	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
North America	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
Rest of the World	<input type="checkbox"/> £91.90	<input type="checkbox"/> £96.90

Cheque (payable to Linux Format. Sterling cheques drawn on a UK account.)

Visa

MasterCard

American Express

Switch

Issue No/Valid date _____

Card No _____ Expires _____

Signature _____ Date _____

UK READERS RETURN THIS COUPON TO:

Linux Format, Subscriptions, Future Publishing Ltd, Freepost BS4900, Somerton, Somerset, TA11 6BR.

OVERSEAS READERS RETURN (POSTAGE PAYABLE) TO:

Linux Format, Subscriptions, Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6TB, UK

WEB: www.linuxformat.co.uk then click on **Subscribe**

Please tick this box if you would prefer not to receive information on other offers.

ORDER HOTLINE ☎ 0870 444 8645

OR FROM OVERSEAS ON +44 870 444 8645

OFFER ENDS: 19.06.03

SAVE £24.47!*

– THE LEADING LINUX MAGAZINE



13 x

LINUX

HAMMER TIME!

FORMAT

AMD'S LATEST, MOST AMBITIOUS PROCESSORS EVER – RUNNING LINUX!

The inside track on AMD's revolutionary Opteron CPUs – will 64-bit x86 computing live up to the hype? PAGE 50

LINUX PRO

e-vote

JACOBOX

JEE PhD developer lowered to death?

HOWTO:

Build your...

LAPTOP MEGATRENDS

= SAVING OF £24.47*

- **THE ESSENTIAL** Linux read for all Linux users – can you afford to miss out?
- **EVERY ISSUE** comes with the ideal supplement for IT professionals – *Linux Pro* covers trends, technology, security issues, industry comment and thought-provoking case studies
- **GUARANTEED DELIVERY** direct to your door so you never miss an issue!

RISK-FREE GUARANTEE!

You can cancel your subscription at any time and receive a full refund on all unmailed issues, no questions asked

SPREAD THE COST OF SUBSCRIBING

For under £15.00* every 3 months you can pay by convenient Direct Debit. Simply complete the form overleaf and return it back to us today.

Call now on 0870 444 8645 to make sure you receive the next issue

* Prices and savings based on DVD edition.

www.linuxformat.co.uk

LXF41 JUNE 2003 **97**



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 CD/DVD being physically damaged we'll send you a new, working version within 28 days. Send your defective disc – complete with your name, address, and a description of the fault – to:

**Linux Format, Future Publishing Disc Department, 3B
Athena Avenue, Elgin Industrial Estate, Swindon, SN2 8HF.**

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?

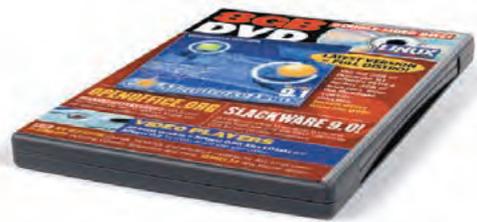
You do not have to use Linux to burn the ISO to a disc. All the Linux-specific bits are already built into the image file. Programs like *cdrecord* simply dump it to the disk. If you don't have a CD-writer, find someone who does have one, and a DVD drive, and use the CD burning software on their computer. It can be Windows, MacOS, AmigaOS whatever.

No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions it is also possible to mount the images and do a network install, or even a local install from another disk partition. The methods often vary between distributions, so check on the distro vendors website for more information. [LXF](#)

Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format DVD*. Some of the information here also applies to MDK CD users.

Now you have installed Mandrake 9.1 (see the CD pages for details), here are a few pointers to using it.

Programs are accessed via the menu button at the bottom left of the screen. Feel free to play around. Provided you are logged in as a normal user and not root, you can't do any damage by experimenting.

MDK9.1 – Logging in

Some people have complained that Mandrake 9.1 no longer allows you to login as root. This is not quite true, although the default setup provides no root login option. There are two ways to change this. You can change the default display manager to KDM

or GDM in the hardware section of *Mandrake Control Center*. Alternatively, you can enable root logins with the Mandrake display manager in the System>Login Manager>Users section of *KDE Control Centre*. Untick root in the list of hidden users and it will appear next time you log in. Having said all that, there really is no need for you to log in to the GUI as root, and there plenty of good reasons not to do this.

MDK9.1 – Hardware

Configuration of hardware, and a lot more besides, is done from the *Mandrake Control Center*. Start it from the configuration section of the menu, you will need to give the root password to run it. From here you can configure hardware such as printers, scanners and monitors, change your display settings, add and update software (see below for more on this)

MDK9.1 – Appearance

If you are running the default window manager, KDE, you can change just about every aspect of its behaviour from the KDE Control Center, accessed via the configuration menu. There are a huge number of options here, so that you can make your desktop fit in with your way of working, rather than you being forced to change to suit the desktop. Most sections have a button to reset back to the defaults, so you can undo the results of any experiments if you don't like them. You can also change settings for individual components with the right mouse button, click the desktop to change it's setting, a window title bar to change the window settings and so on.

MDK9.1 – Updating

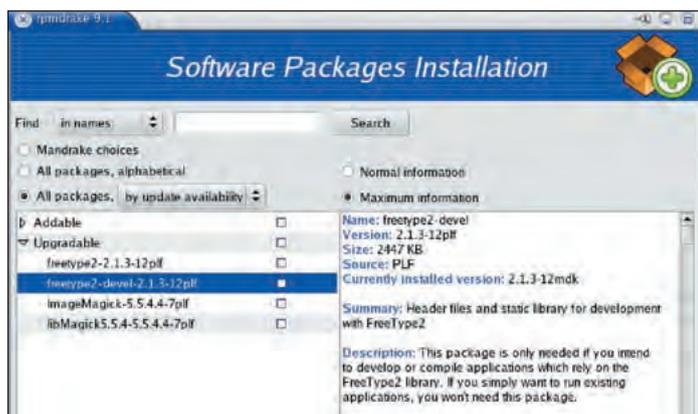
While you will have installed a lot of programs with Mandrake, sooner or later you will want to add others. Mandrake uses a system called URPM to manage software packages. This uses the



The *Mandrake Control Center* is where you configure hardware or networking and install new software, plus a lot more besides.



Change your desktop's behaviour and appearance in *KDE Control Centre*.



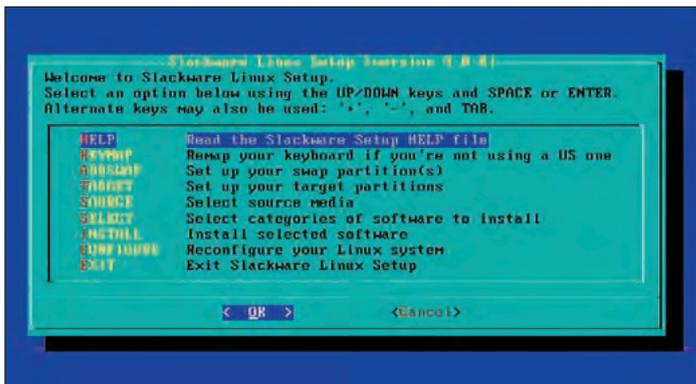
MDK's Software Manager, also known as *rpmdrake*, easily installs software from a variety of sources, taking care of dependencies automatically.



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.



This is the main installation menu, all tasks can be started from here.

command line tools *urpmi*, *urpmq* and *urpmf* and a graphical front end called *rpm-drake*. After installation, the only packages it knows about are those on your installation CDs or DVD, and an updates server if you chose the online update option.

If you installed from the CD issue of the magazine, you should add one of the Mandrake mirrors to the software manager. This will give you access to the programs on the DVD but missing for the CDs. You should also add an updates server if you did not check for updates during installation, this will give you instant access to security updates for any programs you have installed. To do either of these, select 'Software Management' from the Mandrake Control Center and choose 'Software Sources Manager'.

DVD users have an extra 2GB of programs from Mandrake's contrib collection. Add these to the software manager by typing the following command in a terminal window

```
urpmi.addmedia LXF Contrib
removable://mnt/cdrom/contrib/RPMS
with hdlist.cz
```

You can give the source any name you like, but **LXF Contrib** avoids confusion if you later add one of the online contrib sources.

There are other package repositories you may want to add. The Penguin Liberation Front Lair contains packages that couldn't be included in Mandrake for legal reasons, such as DeCSS for playing encrypted DVDs. There is an easy way to add as many of these repositories as you want in one go. Point your browser at <http://plf.zarb.org/~nanardon/index.php> and select the mirrors you want to add. Press Send and you will receive a page containing the correct *urpmi* commands to add them. Paste these into a terminal window and you can

add several repositories without typing a single command.

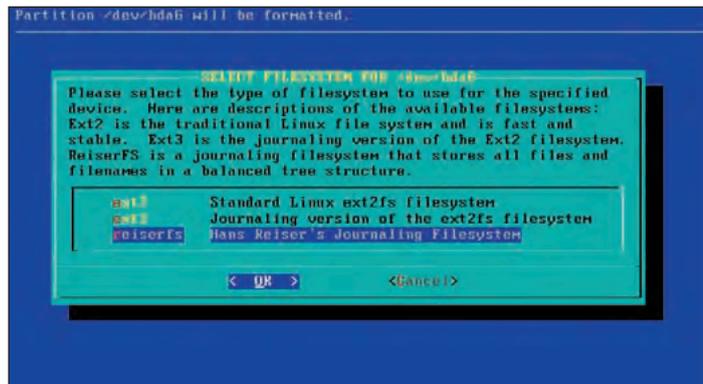
You need to be root to use *urpmi*, but you don't need to login in to KDE as normal, using the icon in the task bar, type **su** and give the root password when prompted. You are now able to run program as root in that terminal. When you have finished, press Ctrl-D to return to a normal user session.

urpmi is normally used for installing packages from a suitable repository, like a Mandrake mirror or PLF, but you can also use it to install from individual files, by giving the full path to the file. In this case, it works like using *rpm* with the upgrade option (*rpm -Uhv*) but also takes care of downloading and installing any dependencies that may be needed.

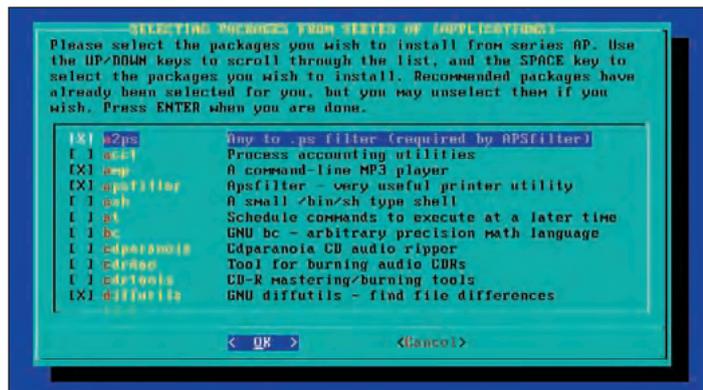
DISTROS SLACKWARE 9.0

Slackware is one of the oldest Linux distributions, but the latest release is certainly up-to-date. Here is a quick guide to getting started with Slackware. Boot from the DVD and log in. The first task is to add some suitable partitions to your hard disk. See the box for information on choosing partitions, then type **cfdisk** at the prompt to begin partitioning. *Cfdisk* is fairly logical in use, press **H** if you get stuck. When you have set up your partitions, it is safest to reboot to ensure the kernel has up to date information, then type *setup*.

This starts the installation program, with a menu listing the various stages. As you complete each one, you get the option to proceed to the next, so you only need to use this menu at the very start and end of the process. We will run through the main points of the installation, if you are unsure of anything, check the help section. After choosing your keymap, the first



Slackware gives a choice of filesystems. While ReiserFS is probably the fastest, on the whole ext3 is considered the safest choice.



Choosing individual packages to install reduces bloat in your system, or you can let the installer take care of it to save time.

main task is to set up your partitions. you should have created these partitions beforehand, this section is only for allocating which partition is used for what. At the very least, you should be setting up a swap and root (/) partition here.

If you think the installation process looks familiar, you would be right. Vector Linux, on last month's disc, is based on Slackware and has a similar installer, but here is where the differences show. While Vector just installed a preset collection of programs, Slackware gives you a choice. The choices initially involve choosing from a number of software groups, such as desktops, networking and games. They are all selected by default, you may need to unselect some if space is an issue. Now you have the choice of installing everything from the selected groups, choosing specific packages from the menus or a 'newbie' install that stops and prompts you before installing each package. While this may be very informative, it is very time consuming, and not particularly friendly to the newbie. This option is more suited to those who know exactly what they want and need precise control over the installation process.

When package installation is complete, it is time to choose a suitable kernel for your system. If the disc booted with no problems, the default kernel is already working for you, so choose *cdrom* from this window then */cdrom/kernels/bare.i/bzImage* from the next, these options should be highlighted by default.

Then create a boot floppy, and put it in a safe place. If you somehow manage to muck up your boot settings, this will enable you to get into your Slackware system.

Hardware configuration is next, first the modem, then the USB hotplug system, followed by installing the bootloader. Without this it will be impossible to boot your new Linux system. The simple option is usually the best choice. You then have a choice of three places to install LILO. The first is only suitable if you already have another bootloader installed in your hard disk's master boot record (MBR). This option leaves that bootloader untouched, you need to add an entry to it to call this one. Unless you already have a custom bootloader set up, don't use this option. MBR installs LILO on the hard disk's MBR, replacing the Windows bootloader if you have one, >>

CoverdiscDVD

but it adds a Windows entry to LILO's boot menu. The Floppy option does exactly what it says, and a boot floppy is always a good idea. This also means you can try out Slackware without altering your current boot setup, you can always add it to the MBR later.

If you are connected to a LAN or the Internet via ethernet, you can configure it next. If your network or ISP uses DHCP, you only need to give your machine a name. Otherwise, you will need to give an IP addresses for this computer, plus addresses of the DNS server(s) and any Internet gateway.

The installation ends with setting the password for the root user. As soon as possible after installation, you should create a normal user account for yourself, saving the root account for essential system administration

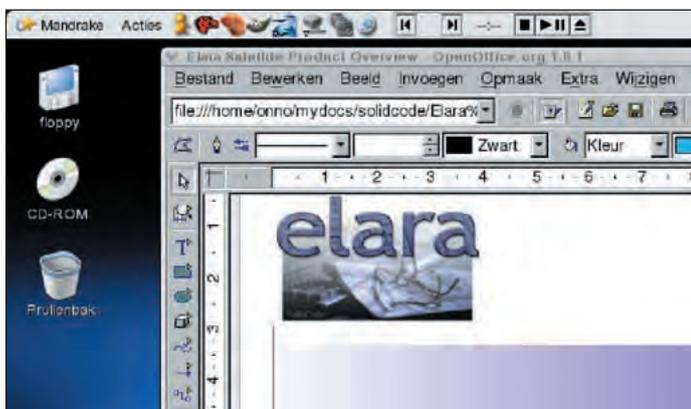
tasks. Select Exit from the main menu and then press Ctrl-Alt-Delete to reboot, or type **reboot** at the command line prompt.

OFFICE OPENOFFICE.ORG

OpenOffice.org 1.0.3 is the latest stable version, offering a number of small improvements and bugfixes. For those that prefer life on the bleeding edge, we also have the beta version of the upcoming *OpenOffice.org 1.1*.

Installation of the two versions is the same. It should be possible to install both 1.0.3 and 1.1 beta if you use a different prefix setting for each, although we haven't tried this. Installation is in two stages, the first part is done as root

```
cd /tmp
tar xzf /mnt/cdrom/Office/
```



Two versions of OpenOffice.org, the latest stable and the 1.1 beta for those of you wanting to see the future.

```
OpenOffice.org/OOo 1.0.3 LinuxInte
l_install.tar.gz
```

or the following for the beta:

```
tar xzf /mnt/cdrom/Office/
OpenOffice.org/OOo 1.1beta LinuxIn
tel_install.tar.gz
```

```
cd install
```

DVD SIDE 2 AT A GLANCE

Desktop

BackgroundBuddy
Bookcase
GKrellF@H
GKrellMailwatch
GKrellMMS
GnomeSystemMonitor
Heyu
KDE
KExchange
Konstruk
Sawfish
TakeAJoint
TuxCommander
VirtualNetworkComputing
Xbindkeys

Automatic wallpaper changer
Personal book collection manager
GKrellM plugin to monitor the Folding@Home client
GKrellM plugin which monitors mailboxes
Plugin for controlling XMMS from within GKrellM
GNOME process viewer and system monitor
Controls powerline CM11 remote control interface
Bug and security update to KDE 3.1
Converter for over 150 currencies
Install KDE with KOffice, KDevelop and Quanta Plus
An extensible window manager
Makes sharing directories under Linux easy
GTK file manager
Cross-platform Virtual Network Computing
An events grabbing program for X windows

Development

CommandLineCrypt
DOTCONFppLibrary
Gfime
GNULibrary
Graphics3D
GtkHTMLLibrary
ReginaRexxInterpreter

A command line interface to the crypt() function
A dotconf-like configuration file parser written in C++
MIME utility library
The C library used in the GNU system
A high-performance 3D graphics codebase
An HTML rendering/editing library
Portable implementation of the ANSI Standard for Rexx

Distros

DamnSmallLinux
PeanutLinux
Slackware
Warewulf

A 50 MB live distro with a functional desktop
A preconfigured distribution for Linux newbies
The Slackware distribution
A distributed Linux distribution

Games

AdvancedStrategicCommand
Freedroid

Turn-based strategy game
Diablo clone with the Tux as hero and the MS as evil power

Graphics

3DLissajousFigures
AcidRip
DVB-MPEGtools
Gmencoder
Gtranscode
KPlayer
MPlayer
OpenDP-500
Oxine
Vobcopy
Xine

A program that displays 3D Lissajous figures
GTK Perl DVD encoding frontend
Tools for manipulating MPEG transport streams
GNOME 2 frontend to mencoder
GUI frontend for transcode
KDE media player based on MPlayer
Movie player for Linux
Java server for the KISS DP-500 DVD player
Purely OSD-based xine frontend
Copies DVD .vob files to hardisk
Unix video player

Help

LDP

A complete mirror of the Linux Documentation Project

Internet

DailyStrips
ElmMEplus
Evolution
Host
Mozilla
Mutt
Netcourt
PlugggerForMozilla
POPsneaker

Automatically download your favourite online comics
Enhanced Elm with MIME, POP, IMAP and PGP support
GNOME mail client and PIM
Powerful command-line DNS query and testing tool
Alpha release of the new Mozilla version
Small but very powerful text-based mail client
Display tool for PPP logging and statistics
A modified version of the Pluggger plugin for Mozilla
Mailfilter for POP3 email accounts

Mobile

GPSMan
GTKpod

Graphical manager of Garmin GPS data
Platform-independent GUI for Apple's iPod, using GTK2

Office

ManhattanVirtualClassroom
OpenOffice.org
phpGroupWare

Web-based virtual classroom system
An Open Source version of StarOffice
A Web-based software suite

Server

Courier
Exim
FKEditor
Kpum
mod_ssl
mod_variety
PHP-Nuke
PythonDesktopServer
PythonWebObjects
YetAnotherLDAPAdmin

ESMTP/IMAP/POP3/Webmail server
Message Transfer Agent for Unix systems
The text editor for Internet
KDE pureftpd user manager
Apache interface to OpenSSL
Serve a random file from a requested directory
PHP Interactive Web Portal System
A Weblog and news aggregation system
Embed Python code in HTML
A general purpose Web-based LDAP administration GUI

Sound

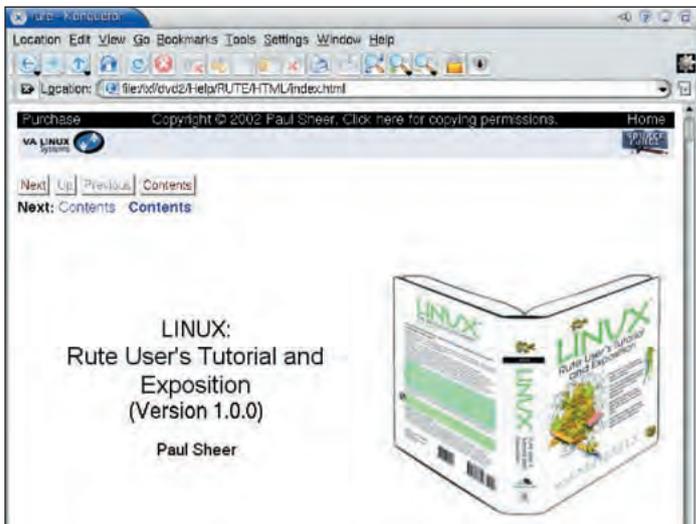
BenMP3
Disc-Cover
MP3Station
MusE
Snd
SoX
Wmusic
Xcplay

An MP2, MP3, OGG, CD, WAV, and modules player
Generate covers for audio cds non-interactively using cddb
Plain text tool set for managing MP3 playback in a car
Linux MIDI editor & sequencer
Sound editor for X/Motif
Sound Processing Tool
Remote-control DockApp for xmms
An nurses interface to XMMS

System

Autospec
Backup2l
ClamAntiVirus
GNOMEDebianPackageManager
JVerify
NSASecurityEnhancedLinux
OpenSSH
OpenUPsd
SambaTNG
SlackPkg
Swaret
Xeasyconfig

Program which creates RPM spec files from a tarball
Lightweight tar-based backup/restore tool
An anti-virus utility for Unix
A graphical manager for Debian packages
A system binary verifier
Linux kernel security enhancements and supporting utilities
Port of OpenBSD's free SSH release to Linux
A UPS monitoring daemon for some Belkin UPSes
SMB fileserver that concentrates on NT RPC and PDC
An automated package-management tool for Slackware Linux
Update, install, remove and search Slackware packages
A PPC-only XFree86 configuration helper



The new Help directory contains some extremely informative resources, such as RUTE and the Linux Documentation Project.

```
./install --prefix=/opt
cd ..
rm -fr install
```

The user setup is done for each user, so go back to your normal user login and type

```
cd /opt/OpenOffice.org/1.0
# or wherever you installed to in the
previous stage
./setup
```

SECTION HELP

This is a new section of the cover discs that will grow. It is intended to be a one-stop reference source. We have moved the mirror of the Linux Documentation Project from its hiding place in the Essentials directory into here. This contains a wealth of reference and tutorial information. This month's Help directory also includes the excellent RUTE. Hardly a day goes past on the *Linux Format* forums without someone being recommended

to read RUTE, so here it is. RUTE is a Linux manual, covering theory as well as practical applications and tutorials.

DESKTOP KDE

No sooner do we put a new version of KDE on the DVD than they release another one! We skipped 3.1.1, as it came so soon after 3.1, but now there's 3.1.1a. These are security and bugfix releases, they don't add any significant new features to KDE. We have only included the source code this time, so you'll have to compile it to install. This is not a particularly difficult task, although it can take quite a while to unpack, compile and install each package. To ease the task, we have included *Konstruk*t on the DVD. Alternatively, you can unpack each archive and do the usual **./configure && make && make install** routine. If you do this, you should install kdelibs first, followed by kbase. The rest of the packages are optional



Guess what? Yet another update to the KDE desktop – 3.1.1a.

Creating CD ISO images from the DVDs

The nitty-gritty of installing

We have decided to continue to use *Jigdo* to create CD ISO images from bootable DVDs. While some people would prefer us to put CD ISO images onto the DVD, this is inconvenient for the majority of DVD users, who would prefer to install direct from the DVD without having to create intermediate discs and then change discs during the installation. If you wish to install Mandrake or Slackware on a computer without a DVD drive, we have made the process of creating ISO images for this, whether from Linux or Windows, as easy as possible.

Jigdo did cause a few problems for some people when we first used it, but we have refined the system and it is no longer necessary to call *jigdo* directly, everything is handled for you by a script. It no longer expects any software to be installed on your computer, everything runs directly from the DVD. The instructions that follow assume that your DVD drive is mounted at /mnt/cdrom for Linux or D: for Windows users. Make the appropriate changes if your setup is different.

For those already running Linux

Open a terminal window and type

```
sh /mnt/cdrom/Distros/Mandrake/mkiso
```

or

```
sh /mnt/cdrom/Distros/Slackware/mkiso
```

The ISO images will be created in your home directory and compared with the originals' MD5 checksums. If you want to create the images in another directory, give it as an argument, such as

```
sh /mnt/cdrom/Distros/Mandrake/mkiso /home/me/iso-images/
```

Creating the ISO with Windows

To build the ISO image when running Windows, open an MSDOS shell and type

```
D:\
```

```
cd Distros/Mandrake
```

or

```
cd Distros/Slackware
```

```
winnkiso
```

The ISO image will be created in C:.

Once you have created the ISO images, they can be burned to CD in the usual way. See the Essential Disc Info page for some information on this. There are detailed instructions on burning ISO images to disc, using various Linux and Windows programs, at <http://www.linuxiso.org/viewdoc.php/howtoburn.html>.

The Slackware script creates the standard single-CD installation disc, whereas the DVD contains more. To create a disc with a different package selection, see the readme file in the isolinux directory.

Permission denied

If you get a "Permission denied" error message when trying to run *mkiso*, it means your DVD has been mounted with the *noexec* option. This is the default when set up to allow normal users to mount devices, as a security measure. The quickest way to fix this is to type, as root:

```
mount /mnt/cdrom/ -o exec,remount
```

and may be installed in any order. Make sure you read the documentation for *Konstruk*t before using it. In particular, you need to change FILE_SITES in the configuration file to point to the KDE directory on the DVD, otherwise *Konstruk*t will try to download the files from the Internet.

DISTROS PEANUTLINUX

Mandrake and Slackware are both well-established distributions, providing a comprehensive selection of packages. In addition to the mainstream distributions like these, Red Hat, Debian and SuSE, there are many smaller distributions.

Smaller both in size of the distribution and the organisation behind them. Peanut Linux is one such distribution. Installation is a breeze, thanks to the way it is largely pre-configured.

It is rather like last month's Vector Linux in this respect, but Peanut uses the RPM system for package management, so it is able to use the thousands of existing binary packages to install new software. The ISO image is only 340MB, barely half a CD, but it installs a highly functional system from there. Clearly, you don't get heavyweight programs like *OpenOffice.org* and KDE, but we have those on the DVD already if you need them. [LXF](#)

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 DONCASTER

URL www.doncaster.lug.org.uk
Contact Andy Smith

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.sderby.lug.org.uk
Contact Dominic Knight

36 BELFAST (BLUG)

URL www.belfastlinux.cx
Email russell@belfastlinux.org

37 WILTSHIRE

URL www.wiltshire.lug.org.uk
Contact Jason Rudgard

38 SOUTH LONDON

URL www.sl.lug.org.uk
Email edo@perceptiondm.com

39 CHESHIRE

URL www.sc.lug.org.uk
Contact Anthony Prime – enquiry@sc.lug.org.uk

40 NORTH WALES

URL www.northwales.lug.org.uk
Contact Andy Hutchings A-Wing deltaone@virgin.net

REVISED
DETAILS

41 MIDLANDS

URL <http://midlandslug.port5.com/>
Contact Pete Thompson

42 CUMBRIA

URL www.cumbria.lug.org.uk
Contact Jamie Dainton

43 DORSET

URL www.dorset.lug.org.uk
Contact John Robinson

44 SHROPSHIRE

URL www.shropshire.lug.org.uk
Email shropshire@lug.org.uk

45 SOUTH WEST

URL www.southwest.lug.org.uk
Email southwest@lug.org.uk

46 SOUTH WALES

URL www.swlug.org.uk

47 NORTH LONDON

URL www.kemputing.net/lug/anlug-aims.html
Email jason@voyagercomputers.co.uk

48 MALVERN

URL www.malvern.lug.org.uk
Contact Greg Wright

49 HUDDERSFIELD

URL www.hud.lug.org.uk
Contact Dave Naylor – knocker@caramboo.com

50 NOTTINGHAM

URL www.nottingham.lug.org.uk

51 ST ALBANS & LUTON

URL www.lust.lug.org.uk
Contact Michael Culverhouse – mike@easily.co.uk

52 WREXHAM

Contact Paul Kersey-Smith
Email paul@pkls.fsnet.co.uk

53 PRESTON & LANCS

URL www.preston.lug.org.uk
Contact Phil Robinson

54 DERRY

URL www.derry.lug.org.uk

55 ISLE OF WIGHT

URL www.iow.lug.org.uk
Contact David Groom – info@iow.lug.org.uk

56 SCARBOROUGH

URL www.scarborough.lug.org.uk

57 BLACKBURN

Email matt@consultmatt.co.uk

58 YORK

URL www.york.lug.org.uk

59 LINCS

URL www.lincs.lug.org.uk

**60 HULL**URL www.hull.lug.org.uk**61 WALTON-ON-THAMES**Contact William Mutch
Email rael@freeuk.com**62 GLOUCS & COTSWOLDS**URL www.gloucs.lug.org.uk**63 WEST OF SCOTLAND**URL www.wos.lug.org.uk**64 SOUTH STAFFORDSHIRE**URL www.staffs.lug.org.uk**65 MANSFIELD**URL www.mansfield.lug.org.uk**66 BORDERS**URL www.linux.bordernet.co.uk**67 BIRMINGHAM**URL www.sb.lug.org.uk**68 COVENTRY**Email info@coventry.lug.org.uk**69 NEWARK & LINCOLN**URL www.newlinc.lug.org.uk**70 BEDFORDSHIRE**URL www.beds.lug.org.uk**71 LINCOLN**URL www.lincoln.lug.org.uk**72 LOUGHBOROUGH**URL www.loughborough.lug.org.uk**73 EXETER UNIVERSITY**Email 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 paulfjohnson@ukonline.co.uk**81 BRIXTON**URL www.communitytechnology.org.uk/~linuxhome**82 ST.ANDREWS, FIFE**URL www.standrews.lug.org.uk
Email stuart@nx14.com**83 NUNEATON**URL www.nuneaton.lug.org.uk**84 ISLE OF MAN**URL www.iom.lug.org.uk
Email helix@manx.net**85 AYLESBURY**URL www.aylesbury.lug.org.uk
Email 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 ORMSKIRK**Email 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@ntword.com**92 MENAI**URL www.menai.lug.org.uk**93 ABERDEEN**Email mark@arricc.net**94 SHETLAND**URL www.shetland.lug.org.uk
Email c_s_s_butler@yahoo.com**OTHER GROUPS:****B-LUX**URL www.b-lux.org/
Email blux-admin@lineone.net**FUENGIROLA GNU/LUG**URL www.fuengirola.lug.org.uk/
Email fuengirola@lug.org.uk**SCHOOLS**URL www.schools.lug.org.uk/**TR LINUX – TURKISH LINUX**URL www.trlinux.lug.org.uk/
Email c_s_s_butler@yahoo.com**YOUNG LINUX**URL www.young.lug.org.uk/

LUG OF THE MONTH

The Lord's Abode, Johannesburg, South Africa

Andrew Gargan writes:

The IT training Centre at *The Lord's Abode* has been going since March 2001. We started with 12 donated PCs (old Pentiums). We managed to get five working, and our own three machines. Originally they ran Win95 with one Red Hat 6.2 box doing file, print, web and general network services. All courses offered are free of charge – the project is aimed at those in our community who have the ability but not the opportunity to

study. Our main drive is basic computer literacy, but we also offer advanced training geared towards programming and web development.

Having gained some experience working with Linux, we switched all our lab over to Red Hat 8.0. Since we started we constructed a new lab with carpets and air conditioning and have bought 14 new workstations, having donated the old machines to the most hardworking students. We're preparing to submit our application

for accreditation with the ISETT-SETA and are confident that we will meet all their requirements.

Unfortunately we have no Net connection for our centre since the monthly costs are too high (we are self-funded) but we are hoping to get one when funds are available. For info and offers of support, contact Andrew Gargan avrin17@iname.com or Terence van Zyl tvanzyl@scientist.com



Worldwide Linux User Groups

Free Software users across the globe

Africa**EGYPT**

URL www.linux-egypt.org

GAUTENG, SOUTH AFRICA

URL www.glug.org.za

Email glugmin@revolution.org.za

Australia**ADELAIDE**

URL www.linuxsa.org.au

Email mtippet@anu.edu.au

ALICE SPRINGS

URL www.aslug.org.au

MELBOURNE, VICTORIA

URL www.luv.asn.au

Contact luv-committee@luv.asn.au

PERTH

URL <http://plug.linux.org.au/>

SYDNEY

URL www.slug.org.au

Europe**AUVERGNE**

URL www.linux-arverne.org

Email Cyril.Hansen@wanadoo.fr

COSTA DEL SOL (English speaking)

URL www.fuengirola.lug.org.uk

DENMARK

AIssund www.alslug.dk

Esbjerg www.eslug.dk

Fyns www.flug.dk

Midt-og Vestjylland www.mvjlug.dk

Nordjylland www.njlug.dk

Skåne Sjælland www.sslug.dk

Trekantsområdet www.tlug.dk

Vest-fyn www.haarby-net.dk/vflug

Århus www.aalug.dk

EIRE

URL www.linux.ie

Email root@linux.ie

URL www.dilu.org

Contact glossary@dilu.org

URL <http://midlands.linux.ie>

Contact midlands@linux.ie

Middle East**ISRAEL**

URL www.iglu.org.il/IGLU/

Contact webmaster@iglu.org.il

PALESTINE

URL www.lugps.org

Email isam@planet.edu

Asia**HONG KONG (multilingual)**

URL www.linux.org.hk

SINGAPORE – SLUG

URL www.lugs.org.sg

SRI LANKA

URL www.lklug.pdn.ac.lk

MYANMAR (formerly BURMA)

URL www.myanmarlug.org

Email afyde@balug.org

PAKISTAN

URL www.linuxpakistan.net

Email tux@clug.org

HYDERABAD, SINDH, INDUS VALLEY

URL www.geocities.com/slug_pk/

KASHMIR

Coming soon!

China

BEIJING (GB encoding, but mostly written in Chinese)

URL <http://mud.263.net.cn/~linux>

CHINESE LINUX USER GROUP

URL www.linux.org.cn

NANJING

URL <http://jllib.jlonline.com/njlug>

India**LINUX INDIA**

URL <http://linux-india.org>

ALIGARH LUG

URL <http://linux.amupost.com>

BOMBAY

URL www.ilug-bom.org.in

CHANDIGARH

URL www.geocities.com/vipinb

CHENNAI AND MADRAS

URL www.chennaiug.org/

CYBERABAD (CLUG)

URL <http://seeknew.freesevers.com/clug/>

DELHI

URL www.linux-delhi.org

KOLKATA

URL www.ilug-cal.org

MADURI

URL <http://linuxmadurai.tripod.com>

NORTHERN INDIA LINUX

URL <http://groups.yahoo.com/group/lug-northindia>

Spreading the word

Linux advocacy isn't just about great coding – you need good sales and presentational skills too, says **Jono Bacon**.

So far we have been discussing how we create some advocacy documentation. In the next few months we will now be looking at methods of getting Linux noticed by people. This month we will focus on some general issues for preparation, and we will then look at different cases over the next few issues.

When you are going out to formally promote and advocate Linux, it is important that you represent it in a both honest and well-prepared fashion. Linux is dependent in many respects on its community spirit, and while some members of the community code, write documentation etc, we still need people to advocate and promote the platform.

It is important that you first prepare your documentation and supporting information. Ensure your websites are online and working and that your paper documents are in good order.

When this is ready you just need to create the contacts. Getting contacts is dependent on who you

are targeting. The first place to look for contacts is on the Internet. You may find people looking for help or advice. It may be wise to contact them via email to maybe organise a meeting. Ensure that you do not spam people; contact should be warranted from either previous discussion or another reference.

When you have got your contact, get in touch and be professional and polite. It is always could to profess as helpful a stance as possible. The tone of your contact is dependent on the context; you will be more formal in a business setting for example. If you get a response to your contact you can then arrange to meet or discuss Linux further. If you don't get a response don't be disappointed; many people are simply too busy to answer generic emails and you should just write it off. Next month we will look at our first case study – schools. Schools are organisations with lots of users, little money and high IT requirements... ideal Linux candidates. Stay tuned!

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

EDITORIAL

Editor Nick Veitch nick.veitch@futurenet.co.uk
Reviews Editor Paul Hudson paul.hudson@futurenet.co.uk
Production Editor Matt Nailon matthewnailon@futurenet.co.uk
Art Editor Julian Jefferson julian.jefferson@futurenet.co.uk

Editorial Contributors

Jono Bacon, Neil Bothwick, Andy Channelle, David Coulson, Richard Drummond, Hoyt Duff, Michael J Hammel, Maurice R Kelly, Biagio Lucini, Patrick O'Brien, Mike Saunders.

COVER CD PRODUCTION

CD Editor Neil Bothwick

ART CONTRIBUTORS

Art assistance Martin Mulchinock
Photography Simon Lees, Gavin Roberts, James Wilson, Powerstock, Photodisc, Digital Vision
Illustration Paul Bateman, Chris Winn, Shane Collinge

ADVERTISING SALES

Key account manager George Gill george.gill@futurenet.co.uk
Business development manager Damian Hughes
Senior Sales Executives Giles Crosthwaite-Scott, Harry Maltby
Classified Sales Executive Stephen Hall

MARKETING AND PROMOTIONS

Marketing Manager Mike Hawkins
Subscriptions Executive Marie Spicer

PRODUCTION

Production Co-ordinator Diane Ross
Group Production Manager Clare Tovey
Advertising Production Co-ordinator Jo Crosby

MANAGEMENT

Publisher Kelley Corten
Group Publisher Dave Taylor
Publishing Director John Weir
Managing Director Colin Morrison

DISTRIBUTION AND CIRCULATION

Circulation Manager Jamie Malley
jamie.malley@futurenet.co.uk

Distributed by Seymour Distribution, 86 Newman Street, London W1T 3EX **Tel** +44 (0)207 3968000

Overseas Distribution by Future Publishing Ltd. **Tel** 01225 442244.

Overseas Licences

International Licensing Director Simon Wear
simon.wear@futurenet.co.uk **Tel** +44(0)1225 822798
Fax +44 (0)1225 788105 **Cell phone** +44(0)777 5641493

Contact Details

Linux Format, 30 Monmouth Street, Bath BA1 2BW
Tel +44 (0)1225 442244 **Email** linuxformat@futurenet.co.uk

Subscriptions and Mail Order

Phone +44 (0)1458 271178. See page 96
Email linuxformat.subs@futurenet.co.uk

Copyright No part of this publication may be reproduced without written permission from our publisher. We assume all letters sent – by email, fax or post – are for publication unless otherwise stated, and reserve the right to edit contributions. All contributions to Linux Format are submitted and accepted on the basis of non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing. Linux Format recognises all copyrights in this issue. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight. We cannot be held responsible for mistakes or misprints.

All CD-ROM demos and reader submissions are supplied to us on the assumption they can be incorporated into a future covermounted CD-ROM, unless expressly stated to the contrary.

Disclaimer All tips in this magazine are used at your own risk. We accept no liability for any loss of data or damage to your computer, peripherals or software through the use of any tips or advice.

Printed in the UK by Midway Clark (Holt) and TPL

© Future Publishing Ltd 2003

LINUX is a trademark of Linus Torvalds, GNU/Linux is abbreviated to Linux throughout for brevity. All other trademarks are the property of their respective owners

Future Publishing Ltd. is part of The Future Network plc.

The Future Network produces carefully targeted specialist magazines and websites for groups of people who share a passion. We aim to satisfy their passion by creating titles that offer superb value for money, trustworthy information, multiple ways to save time and money, and are a pleasure to read or visit. Today we publish more than 80 magazines and over 30 magazine websites and networks from offices in five countries. The company also licenses 32 of its titles resulting in over 60 local editions in a further 23 countries.

The Future Network plc is a public company quoted on the London Stock Exchange (symbol: FNET).

Non-executive Chairman Roger Parry

Chief Executive Greg Ingham

Chief Operating Officer & Managing Director, UK Colin Morrison

Group Finance Director John Bowman

Tel +44 (0)1225 442244

BATH • LONDON • MILAN •
 NEW YORK • PARIS •
 SAN FRANCISCO • WROCLAW
www.thefuturenetwork.plc.uk



Media with passion

NEXT MONTH

Issue 42 on sale Thursday 19 June

LET'S TALK ABOUT

The Xserver is one of the most common and uncontested parts of any GUI Linux system. It has been in development for years, and is a primary requirement for many installations and most software. But that doesn't mean it's perfect – far from it.

Find out how the history of X has left us with a system that works, but can be terribly inefficient. A system that is in continual development, but only by a relatively small core of programmers. With news of possible forks, we talk to key developers about the future of X, and what it will mean for virtually all Linux users.

GAMETASTIC!

Winex3 is released and alongside a detailed review, we'll be giving you a brief guide to some of the great games you can run with it, plus comments on future development from its creators



ON TEST

Next issue there will be a bumper crop of new servers from NEC, xinit and more. Plus we'll be getting the lowdown on Intel's Vtune software, trying out Red Hat's new Advanced Server distro and catching up with other great software like the newly released MySQL 4. Don't miss out.



PLUS LINUXPRO

24 pages of real-world Linux news, views and case studies for IT professionals

DON'T MISS YOUR COPY OF LINUX FORMAT!
 Get it delivered to your door every month – subscribe on page 96

The exact contents of future issues are subject to change

LINUX PRO

FROM THE MAKERS OF LINUX FORMAT

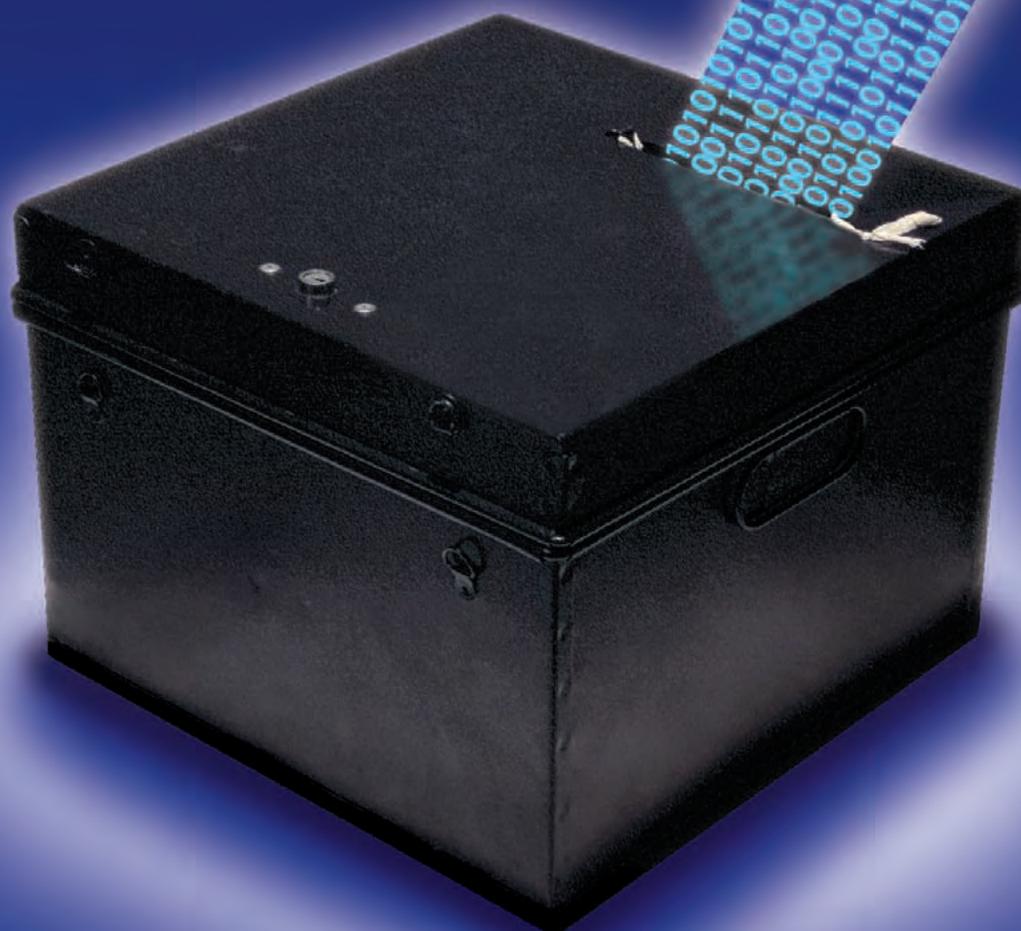
JUNE 2003

WIRED DEMOCRACY

e-vote



Electronic voting systems are designed to bring citizens back to the ballot box, but the security implications are immense. Could Linux and Free Software be the answer? **PAGE 4**



PLUS

NAS and SAN

Which storage technology will save your business? Find out the pros and cons

Embedded developments

Developing cross-platform applications the easy way – QTopia

.NET for Linux

Mono and DotGNU – bringing .NET technologies to the Linux platform

Thin Clients

Diskless terminals are back in demand and make good sense

Welcome

Twenty-four pages of real-world Linux for IT professionals

Can you imagine implementing a computer system, accessible to the general public, on which it's secure enough to bet the future of the country? Those proposing and implementing the various e-voting schemes run in local elections face just such a challenge, and the debate will continue for some years to come, we're sure. This issue we'll take a look at some of the security challenges encountered by such a venture, and examine how Free Software – a close ally of free speech – and Linux are getting in on the act as a part of the democratic process.

.NET has been a somewhat controversial technology since its birth. Many harbour suspicions about the motives of its creators, and see it merely as a way of derailing Java. But whatever the politics, it is a real technology that real developers are using today – shouldn't Linux developers be able to use it too? Fortunately there is work ongoing to bring .NET to Linux, enabling both Linux developers to create .NET code, and Linux users to run .NET applications. Find out more on page 6.

In our storage section this month we have an overview of NAS and SAN technologies. Through experience, it seems that many users opt for one or the other, without really getting the complete story on what the strengths and weaknesses are, when it might be more appropriate to use one or other, or even both. If you are considering your storage needs for the future, it's essential reference material.

Finally, we also have a feature and case study on using 'thin clients': the idea dates way back to the early days of computing, but it's never gone away, and in recent times is having a gathering renaissance. Understandable really when you look at the benefits in terms of cost savings, reliability, security... Find out more on page 20.

And, as always, if you have any comments on anything you have read here, or want to make suggestions for future issues, feel free to contact me at the address below.

Nick Veitch Editor
nick.veitch@futurenet.co.uk



“Whatever the politics, .NET is a real technology that real developers are using today – shouldn't Linux developers be able to use it too?”

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

Elsbeth Wales, Andy Channelle

Photography

Digital Vision, Powerstock

ADVERTISING

Key Account Manager

George Gill george.gill@futurenet.co.uk

Business Development Manager

Damian Hughes

damian.hughes@futurenet.co.uk

Senior Sales Executives

Giles Crosthwaite-Scott, Harry Maltby

Classified Sales Executive

Stephen Hall

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

JUNE 2003

CONTENTS

E-VOTING

A democratic revolution? **p4**

.NET FOR LINUX

Mono and DotGNU explored **p6**

QTOPIA

Trolltech's embedded toolkit **p8**

WRITING COMPILERS

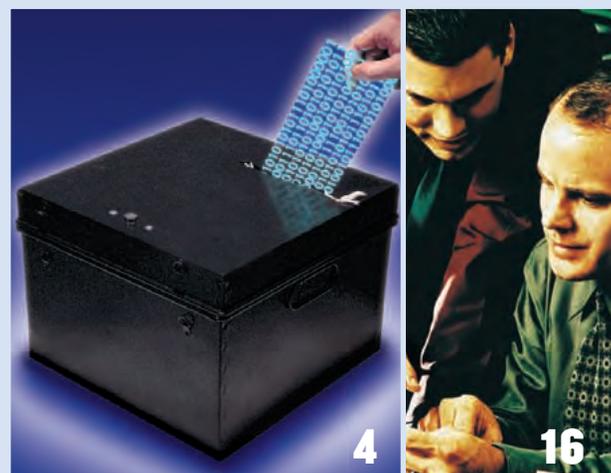
Flex and Bison explained **p11**

SAN OR NAS? OR BOTH!

Which storage technology? **p16**

THIN CLIENT COMPUTING

Get rid of desktop PC boxes **p20**



LINUXPRO 3

WIRED DEMOCRACY

e-vote

With election turnout spiralling downwards in most parts of the world, authorities are attempting to 're-engage citizens with democracy' by introducing simpler, more convenient ways of voting. **ANDY CHANNELLE** investigates what part, if any, Linux will play in the brave new world of e-democracy.



Democracy doesn't have the highest approval rating in 2003. Whether it is the hanging chads of Florida that swept George W Bush into power, the banana republic 'one candidate one vote' method or the sub-50 per cent turnout that characterises the elections of most western-style 'democracies,' voting is in need of a shake-up. One of the more popular methods put forward to improve the image of 'the vote' is electronic voting, or e-democracy. This, we should be clear, is not the same as Internet or online voting (see below) which received a lot of attention initially but is regarded by some as being about as secure as a chocolate padlock; Bruce Schneier, author of *Practical Cryptography* says an online voting system with sufficient safeguards would be "the first secure networked application in the history of computers". The more popular view of e-democracy is simply a replacement for the pencil and paper, place your **X** in the box approach which the UK has used since William Gladstone introduced the secret ballot in 1872. Instead of wielding a pencil, voters would make their choice known by selecting an option from a keypad.

As well as potentially being more cost-effective than printing millions of voting slips, computerised voting would also dramatically increase the speed and accuracy of vote counting, and while this may spoil the tension of election night, it would, it is suggested, lead to fairer and more efficient elections. On the flip-side of the coin, there are concerns about the security and stability of electronic systems and about authentication: is the person exercising their franchise who they say they are? Of course, these issues exist with the traditional methods of vote casting, but as 'it has always been done like that' most people don't even think about it.

Vote Linux!

There have been a number of trials of e-democracy and, after the problems experienced in the 2000 US Presidential race, they are set to expand creating a multi-billion dollar industry. The British government dipped its toe in the water in the local elections of 2001, using e-voting alongside postal votes, and expanded efforts in May 2003. One of the biggest trials took place in the Australian Capital Territories (ACT) elections of 2001 using a system built by Software Improvements Pty. The Electronic Voting and Counting System (eVACS) is an Open Source project running on commodity hardware – important for reducing costs – and a Debian base. Beyond core tasks of registering and counting votes, the system also incorporates an audio system for visually impaired voters, offers election info in 12 languages and uses a one-time barcode issue system for authentication.

The most visual element of the eVACS system is the hardware and this is divided into three disparate parts which, thanks to the very limited remit, can be built from basic (ie cheap) hardware. At the user interface are the booth machines which minimally require a Pentium 200MMX processor, 32MB RAM, video card (capable of driving a 19" or 21" monitor at 1280x1024/75Mhz for visually impaired voters) and sound card. Additionally the OS needs to be booted via a network or CD-ROM, thus removing the need for local storage which is regarded as a security risk. Beyond the commodity hardware, each machine requires a barcode reader, a keypad and a horizontal mount for the screen to

ensure secrecy. Each of the polling machines is connected to a central 'ballot box' machine via a standard ethernet cable.

In the backroom there are a number of data entry workstations which are used to enter paper votes in the traditional manual fashion and the electronic votes from the ballot box machine via a Zip disc. The final stage is the counting machine which uses a standard Hare-Clarke scrutiny – voters select a number of candidates by preference – to decide the outcome of the election. Hare-Clarke voting usually entails a long wait, sometimes days, for the result to be known, and this is where the real benefits of this system are found. The counting machines also double up on administrative tasks when voting begins such as creating the barcodes for authentication. Electoral Commissioner Phil Green was very pleased with the system. "A total of 16,559 votes were cast electronically. The electronic voting system proved reliable and secure, and combined with computerised counting, it delivered an election count that was close to 100% accurate."

The big issue here is the same as with paper votes: authentication. Rebecca Murcuri, academic and leading critic of fully electronic voting systems, claims that authentication and the lack of an 'independent, voter-verified audit trail' is the key problem. "The vendors and certifying authorities have taken a 'trust us' stance, claiming that the machines are "fail-safe" and that the internal record and tally constitutes an accurate reflection of the ballots cast on the machine"

In fact, she says machines have failed in real-world use – with choices displayed that were not selected and votes recorded wrongly. She also claims some of the machines enter a lock-down mode when the polls are closed which

"From the problems we had trying to get any info at all, it appears that e-voting will be as transparent as a big black metal box."

makes it impossible to check that votes had been cast properly for each candidate or issue.

However, these criticisms seem to be aimed mainly at vendors of proprietary systems who "have tied the hands of election officials and independent examiners by protecting their systems under restrictive trade-secret agreements, making it a felony to inspect the operation of the machines without a comprehensive court order." This of course doesn't apply to something like eVACS which is available for scrutiny by anyone. The eVACS software suite is fully GPLd and is available from www.elections.act.gov.au/evacs.tar.gz.

Or not, or maybe...

A recent MORI survey suggested that over half the UK population would be 'more likely' to vote in elections if they could do so by SMS, digital television, Internet and supermarket kiosks, and trials last year saw an increase of eight per cent in voter turnout where these alternatives were available. The big problem is that remote voting throws down a torrent of security problems far beyond the identity issues outlined above. The government's own consultation paper devotes much of its space to the concerns raised by the likes

GNU-FREE ABANDONED

Some say that if voting actually changed anything, it'd be illegal...

GNU-Free was a concerted attempt to create a secure remote/online voting system based on Open Source principles that grew out of a joint research project into the effects of the information revolution in activism and political participation at the University of Warwick. But, recently lead developer Jason Kitcat has had a change of heart on the validity of online voting and has switched his efforts to advocating the responsible – and free – provision of governmental services online. After a 'monumental' amount of work, Kitcat became increasingly disillusioned with the way technology can be used as a 'smokescreen behind which fraudulent plans can be hatched!

"While all of these ideas appear exciting and modern they have many risks. They are expensive (though in the long-term a cost saving may be found) and they won't improve turnout as some claim," he said. "The reasons people don't vote are to do with them feeling their vote doesn't make a difference – convenience is very low down the list."

From one of its high-profile advocates Kitcat, swayed by the arguments of security experts, has become a vocal critic. "Basically I'm opposed to the use of electronic technology in voting, as it will enable more potential fraud and problems than it is worth. And what is worse much of the fraud will be undetectable."

of Jason Kitcat (see box above) and Rebecca Murcuri, concluding that any online solution must be "at least as secure as traditional methods." The list of potential problems include the introduction of malware, trojans, viruses, lack of secrecy and the possibility of coercion and equipment failure. And that's without even going into identity and identity theft.

Still remote voting is seen as inevitable and there are a number of competing companies hoping to grab a slice of this lucrative pie. What is interesting about the UK position is that recently the government changed its rules for tenders following responses inspired by its "In the service of democracy" consultation document. Section 17 of the subsequent document (www.edemocracy.gov.uk/downloads/your_response_report.pdf) says: "As a result [of the consultation] some requirements were modified. In particular a requirement to support "publicly verifiable code" was introduced to strengthen the auditability of the system." This seems unambiguous but, as ever, it has the whiff of a fudge. Certainly the Athena Consortium, which finally won the contract to manage the online voting trials on May 1st 2003 are sure that 'publicly verifiable' means something entirely different. As you may expect with so much taxpayers money and their reputations on the line, the Consortium management are "extremely paranoid" about details of their system getting out. Organisation head John Ellis told us that he wouldn't divulge any information under any circumstances before, during or after the local elections. He said that whether the group chose "Solaris, Windows, OpenBSD or some other 'NIX" the potential for hackers to compromise the whole process was a danger they took very seriously.

As we go to press, the first significant test of online voting will be getting underway in a number of councils across the UK; and while this may bring about a revolution in democracy, from the problems that we had trying to persuade any of the relevant bodies to give us any information on the subject, it appears that the systems in which we place our hard-won franchise will be as transparent as a big black metal box. ■

NOTE

The office of the UK government's e-Envoy has yet to respond to any of *Linux Format's* questions about the transparency and accountability of defiantly closed-source systems.

Linux.NET

As many people know, Microsoft came onto the Internet bandwagon quite late into the game. However, once they restructured their business plans to take note of the WWW, they have made it their business to have a finger in every pie that was available. It's no accident that Hotmail is the world's most popular webmail system and MSN is the world's most popular website! The latest initiative from Redmond has been christened .NET, and is much more than just a marketing ploy.

What is .NET?

.NET isn't just the name Microsoft are giving to many of their products, so although *Visual Studio.NET* is out, and *Windows Server.NET* almost made it to release (it's now been renamed to *Windows Server 2003*, "to avoid confusion"), .NET is a technology, not a product in itself.

.NET is a set of technologies designed to take advantage of the Internet. As such, it works over the Web using Web Services, it is heterogeneous in that programs written for .NET compile down to an intermediate language not native machine code, and also the .NET interpreter system is designed to catch common flaws exploited over the Internet, such as buffer overruns.

Perhaps the most important part of .NET is its heterogeneous nature – it can convert a variety of source languages (such as C++, Java, and Microsoft's new language, C#) into intermediate code (known as CIL, Common Intermediate Language), which is then hot-spot compiled into native code on the execution platform. As a result, .NET is capable of running the same executables on a variety of platforms – wherever the interpreter is available.

If you're a Java convert, this will all sound like old news to you – after all, Java does much the same. Indeed, reading through C# code and the same program in Java will show quite how similar the two systems can be. .NET, though, was *designed* to accept multiple languages (unlike Java, which wasn't, but can anyway) and so makes its class libraries available to all languages. What this means is that a programmer who uses language X because it has libraries Y no longer need worry, as each language is cross-compatible. In fact, this cross-language compatibility extends to inheritance: you can design a class in C# then inherit from that class using Eiffel#, and you can even debug across the two languages simultaneously.

.NET includes a large collection of class libraries that allow easy access to XML, databases, GUIs, and such. A complete .NET implementation includes replicated class libraries, a compiler to CIL, and of course an CIL runtime interpreter. Collectively, these are known as the .NET Framework. Generally, C# is chosen as the basic language to compile, because it was designed to implement all parts

Microsoft is pushing its .NET principle hard – but how does it affect Linux users? PAUL HUDSON investigates...

of .NET – the language can be thought of as a means rather than as an end.

Linux and .NET

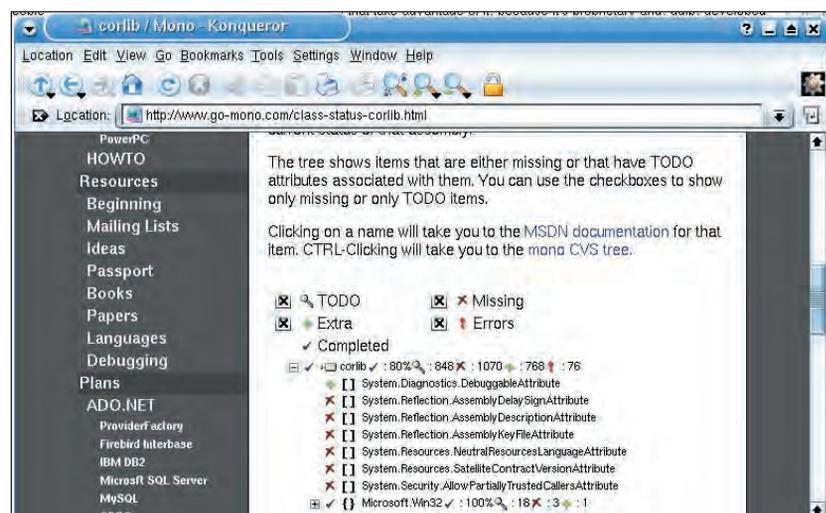
So, .NET is a technology that might be useful for Linux users. After all, if applications written for Windows are cleanly compiled down to CIL, they can be run on Linux natively. Two groups of developers recognise this possibility, so there are two free software projects currently being developed that provide .NET compatibility: DotGNU and Mono. The Mono project is being led by Miguel de Icaza, of GNOME fame, and co-founder of Ximian, which has lent some extra weight to Mono development. It's important to note that Mono and DotGNU aren't competing with each other – officially, DotGNU is designed to enable decentralised services and authentication, whereas Mono is there to provide the compiler, class libraries, and runtime.

Once Mono is ready for primetime, most Windows .NET applications should run fine on Linux. Other than the obvious result of "great, lots more programs for Linux are available, so people might be more likely to switch", there's also the side-effect that it would make programming for Windows and Linux exactly the same – a programmer trained to write C# on Windows could write exactly the same code on a Linux box and have it work exactly the same way. So, adding quality .NET support to Linux increases not only the number of applications that are available, but also the number of people capable of writing new applications.

Why use .NET?

Using a .NET-compatible language today allows you to take advantage of its cross-platform nature – as long as you don't use vendor-specific class libraries in your code, it should be portable to Windows, Solaris, etc, while still retaining all the

Want to know what Mono supports as of now? The Mono site tells you which of the class libraries are supported and what's missing in a very straightforward manner.



functionality – including the GUI. Also, you can benefit from .NET's cross-language nature, allowing you to distribute work over a group of people, each of whom are able to use the language they feel most comfortable with – at the end of the day, the compiled code will work together seamlessly.

One advantage to using .NET is the option to use *managed code*. This is where you place a lot of your programming responsibility into the hands of the .NET interpreter. The interpreter will perform automatic garbage collection, and also detect attempted buffer overflows, and halt them in their tracks. The obvious disadvantage to using managed code is that it takes control out of the hands of programmers – I for one like to have complete control over my code, and consider it a little sloppy to rely on garbage collection to clean up after yourself.

For those programmers willing to take on the .NET challenge, there's certainly a great deal of support – not the least of which is the set of class libraries that the .NET Framework includes. These class libraries allow easy access to particularly complicated and regularly used aspects of web programming. The most useful libraries are Windows Forms and XML, which give .NET-compatible languages much of the same functionality as their Java versions.

One particularly important thing I wish to make clear is that, in my opinion, you shouldn't feel 'dirty' using .NET. All too often I hear people saying that we shouldn't use .NET or technologies that take advantage of it, because it's proprietary and, gulp, developed by MS. It is my fear that this is representative of an increasing feeling of "not invented here" among Linux developers (particularly those who spell the Redmond company 'Micro\$oft') – that is, if software is developed by a big company, it should be shunned. Perhaps these people forget the roots of GNU and Linux, which were a group of people wanting to take existing proprietary UNIX and re-implement it as free software – precisely what DotGNU and Mono are doing for .NET. .NET is a technology that represents a big leap forward for programming languages: try to judge it on its individual merits without worrying about who invented it.

Be careful using .NET

It's important to remember that though its functionality is great, there are potential problems with .NET programming. Firstly, like Java, .NET is designed to be cross-platform, however, unlike Java, it is very easy to tie .NET programs to a specific platform – especially Windows. The cross-platform problem has several parts: OS differences, Windows-specific features, and native calls.

Firstly, OS differences comes down to specific things such as reading files. For example, whereas `c:\program files\foobar` might exist on Windows, it will not exist on Linux. Similarly, you may need to watch case sensitivity now and then, as this is often an issue, particularly with files. Windows-specific features include the Windows Forms GUI toolkit, the Windows registry, etc. Mono has compatibility for Windows Forms built-in, however it's only currently working solidly on Windows – Linux/OS X support is shaky. The problem here is that Windows Forms weren't designed for cross-platform compatibility, but instead they were designed to allow developers to switch to .NET whilst still retaining as

YOU CAN QUOTE ME ON THAT...

What people have to say about Mono and DotGNU

"With Mono and DotGNU, we hope to provide good alternatives to components of .NET, ones that will respect your freedom, and your privacy. You will be able to use the facilities of Mono and DotGNU either with, or without, the Internet, and using servers of your choice."

RICHARD STALLMAN, GNU project founder

"I want to be able to deliver four times as many free software applications with the

same resources, and I believe that this is achievable with these new technologies."

MIGUEL DE ICAZA, founder of GNOME and Mono lead developer

".NET is truly multi-language and the role of the framework is to provide a reasonable target to which all current languages can map"

BERTRAND MEYERM, creator of the Eiffel language



Installing Mono on your machine is usually painless. Debian users should go to www.debianplanet.org/mono/ and follow the instructions.

much of traditional Windows GUI programming as possible.

The final problem is native calls using Platform Invocation Services, or *pinvoke*. This allows a .NET-managed program to call a unmanaged program, such as a Windows DLL. Naturally this means leaving the cross-platform safety of .NET, so *pinvoke* isn't guaranteed to work, even if the shared library is available. This works in the same way as Java's Native Interface (JNI), and can be a great benefit as well as a great pain in the backside!

Conclusion

If MS has its way, .NET is certainly here to stay. Even if MS decides not to support it in the future, it's still a great technology and one that Linux users should be happy to take hold of and use. People regularly cite "Ah, but MS might introduce patents!" and the like as reasoning not to use .NET, but really – you have nothing to fear. Even if MS do introduce patents into .NET v2 (or whatever they call it), it won't stop Mono developers working with what they have already, and it certainly won't stop Mono developers creating their own implementation of proprietary code – look how well the *Samba* people have done.

Give .NET a try, and attempt to forget that it's a MS idea. We think you'll be pleasantly surprised! ■

Qtopia

gets bigger in its own small way



Qt is Trolltech's basic multi-platform application framework product offering developers a C++ tool kit for building multi-platform GUI and application development, allowing them to write applications which run natively not only on Linux/UNIX, but also on Windows and Mac OS X, as well as embedded Linux.

Based on Qt is Qtopia, the application platform for embedded Linux which has recently received a comprehensive update in functionality. It runs on PDAs, Communicators and a hybrid devices and is the platform chosen for Sharp's Zaurus PDA, which made it the first major vendor to ship a device with embedded Linux.

Qt/embedded is another product based upon Qt and is the application development framework and windowing system for embedded Linux.

Trolltech targets its products at developers who use the Qt multi-platform framework, corporate IT professionals who can build applications using the software, the mobile device manufacturer and of course, end-users, who can take advantage of the fact that applications built in Qt and Qtopia run and synchronise across multiple desktops and devices.

How Qt came about

Trolltech was formed in 1994 by Eirik Eng, now the president, with the idea of devising their own UNIX-based development tools after having decided that others on the market were inadequate. Being a new company they were also watching costs closely and turned to Linux because they could not afford to pay for expensive UNIX licences, explained Eng. Marketing funds were limited too so they decided to use the Internet as a distribution channel.

Trolltech operated a dual licensing model offering an open source version for development of Open Source software and a commercial version the development of commercial software. In order for that model to work, both products had to be exactly the same so the open source version is exactly the same as the commercial version sold to commercial companies on the UNIX and the Linux platforms.

The company experienced slow but steady growth until it secured external funding in 1999 and again in 2000 as a result of some work carried out for Linux distributor, Caldera systems (now SCO), which caught the eye of investors. Subsequently the pace of growth has quickened substantially and Trolltech also managed to emerge unscathed from the ensuing dotcom bust

Trolltech's major product is QT the core of which is a C++ library that is made available for different platforms such as

As the tributes from developers here show, Qt is increasingly finding favour with Linux for small form devices. ELSPETH WALES illustrates how it's also starting to be embraced by big-name vendors such as Sharp and IBM.



Eirik Eng, president and founder of Trolltech has used a dual licensing model since 1994 that is now common right across the Linux world.

Unix, or Windows and Apple Macintosh. In addition, there is a special version for Linux which is called QT embedded. This began life as a graphical user interface library and has grown to be a general application framework containing all the components that are needed to create professional applications, including everything from string handling, database productivity, inter-process communication, essentially everything needed to develop software..

Qt's multiple platform support

Qt supports multiple platforms such as Microsoft Windows 95/98/Me, NT4, 2000 and XP as well as several UNIX variants including obviously Linux, Sun Solaris, Hewlett-Packard's HP-UX, IRIX, and IBM's AIX, and the Apple Mac OS X. Porting it to different platforms requires a simple recompile, there are no emulation layers and no virtual machine, leaving developers with fast, elegant, native code to work with.

Qt 3.1, the latest version, is also object oriented, as easy to use and as powerful as earlier versions but out of significant functionality and introduces new library classes. These classes are fully featured and designed to reduce developer workload, providing consistent interfaces to speed learning, claims Trolltech.

The Qt library classes, or widgets, also known as controls in Windows terminology, provide standard GUI functionality. Qt includes a rich set of these widgets such as buttons or sliders or a way to access the database, to name but a few. Some of the latest introductions include alternatives for inter-object communication, called Signals and Slots, replacing the old and unsafe callback technique. It also features a conventional Events model for handling mouse clicks and key-presses as well as multi-platform GUI Applications which can use all the user interface functionality that modern applications call for, including menus, context menus, drag-and-drop, dockable toolbars, and balloon help.

Other key features listed by Trolltech include intuitive naming conventions and a consistent programming approach in order to simplify coding, the inclusion of Qt Designer which is a tool for graphically designing interfaces that supports Qt's powerful Layouts capability as well as absolute positioning. This can be used to create GUI designs or entire applications using an in-built C++ code editor.

Trolltech touts the short learning curve that Qt demands as being one of its key strengths because developers only have to learn one API to write applications that run almost anywhere. In addition to the rich set of standard widgets Qt

developers can also write their own custom controls. The framework also encapsulates four different platform specific APIs and the APIs for file handling, networking, process handling, and database access also make for a shorter learning curve, claim the company, as does support for Motif and ActiveX integration.

Qtopia and Qt/embedded

On top of this Qt library Trolltech has built Qtopia, having recognised in the middle of 1999 that Linux was going to have a great future for embedded systems, said Eng. "We saw that it was modular, it was easy to scale and, of course, the price was impossible to beat. We also saw that there was a big fragmentation in this market, there were a lot of people running proprietary systems and we saw that these were going to be exchanged for something else," he explained.

But on entering the market the company found that it was very different to its earlier market. Previously its customers were developers but from 2000 hardware vendors started to become customers and were not interested in building all the necessary software they needed around the library.

As a result, Trolltech recognised that the company needed a total platform, complete with simple applications, launcher, Window system, prompting the development of Qtopia. "So Qtopia is a complete platform much in the same way as a pocket PC needs a platform for hardware devices. So it really contains everything needed by a hardware vendor to create a PDA or handheld type of device," continued Eng.

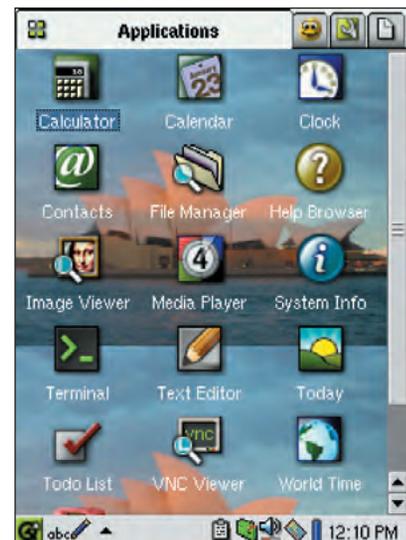
Aron Kozak, product marketing manager, explained that in many ways the Qt product family are the same in that all are based on Qt and Qtopia was developed with Qt/embedded so uses the same API in both areas. The biggest difference between these two is that Qtopia contains a few additional libraries such as a PIM database and Qtopia styles all packaged in an SDK, which is designed specifically to work with Qtopia, not Qt/embedded which Kozak described as being a much broader product.

According to Stacey Quandt, analyst for Linux and Open Source with research firm Giga Information, the key features of Qt/embedded and the Qtopia development toolset is a single embedded application source tree eliminates the need for multiple source trees, so with only a recompile



Above: **Unlike the current crop of 3G mobile phones, the Zaurus has no problem handling MPEGS and audio.**

Above right: **An easily customisable launcher makes any PDA running a Qt-based app equally appealing to both personal and business users.**



developers can have the option of portability to multiple embedded devices such as PDAs, set-top boxes, computer games, and smart phones, while applications developed using the Qt framework can run natively on Windows, Linux/Unix, Mac OS X.

From a development point of view, Kozak said he believed that the main advantages offered by Qtopia are ease of application development and portability to the desktop.

"Qt supports MS Windows 95/98/Me/2K, NT4 and XP, and UNIX variants like Linux, Solaris, HP-UX, IRIX, AIX and Mac OS X"

"Because Qt was developed to be a powerful desktop development tool, it has a lot of high-level features that allow a developer to insert functionality easily and quickly without having to write a whole bunch of low-level code. Now the same API is available on a PDA. What we have done is we have made it as small as possible so it runs properly on a PDA, but it is still just as powerful," he said.

Quandt added that developers would use Qt instead of other embedded development options such as Microsoft CE and the Palm OS, "Because with Microsoft CE and Palm OS, portability is restricted to specific platforms and this limits the flexibility of developers to target other platforms."

In her view, the advantage of using Qtopia is that it includes PIMs and utilities which simplify the ability of developing apps with variations in look-and-feel for mobile devices.

Developers

"Trolltech has been very successful in attracting developers, and there are already over 1500 applications, embedded applications, developed with Qtopia. Recent partnerships with Metrowerks, a part of Motorola, and IBM will only fuel further use by developers," commented Quandt.

Developers who have worked with Qt and Qtopia are generally enthusiastic about them. One professional developer, Tim Wentford, said that the best features of Qt/embedded (the toolkit which Qtopia runs on) are that it "is very portable and is compatible with

Below left: **The straightforward Zaurus desktop appeals to a broad range of users.**

Below: **Anti-aliased fonts and full-colour graphics make the GUI clear and attractive.**



EMBEDDED DEVELOPMENT

the Qt toolkit for desktop machines – I've had apps running under Windows, Linux and on the Zaurus (and the iPaq after installing *opie* [the Open Palmtop Integrated Environment]) with very few source code changes."

"Qt's slot/signal mechanism enables widgets to be more self-contained/independent, which means that you can, for example, reorganise the layout of a widget or even the way it works and it doesn't impact on the main app at all (or vice versa)," he added. Wentford said that it is also good at tidying up after itself in that widgets "usually have parents which can delete their children" meaning that you normally only have to worry about the topmost widget of a hierarchy, but he's not overly keen on this feature as it can encourage sloppiness.

Advantages of Qt

He felt that the biggest advantage that Qt has over other embedded operating systems such as Microsoft Windows CE and the Palm OS, is its portability but pointed out others too.

"The Palm environment has a very restricted memory model and is very 'stylised' in that you need to program specifically for the palm and special skills are required which are not portable to other platforms.

"CE isn't quite so bad, but you can tell that you are working on a C++ API which has been glued over a C API. Sometimes it's a struggle fit it into a C++ framework (I'm a C++ programmer) and you often end up with functions being made unnecessarily static and adding an extra argument which should have been the parent object. Qtopia is very obviously object oriented from the ground up and has a very clean API. It's very easy to extend by subclassing," said Wentford.

Some of the first features to be offered by the very earliest PDAs were world time and calculator functions, and of course Qt continues this legacy.



He has been working with Qt for about 2 years and develops professionally for mobile devices. Most of his work is on the server side but he has done a moderate amount of development on CE and Symbian clients, and contributed to Palm and Java clients. One project using Qt was an eBook reader called *OpieReader* (was *QtReader*) in order to read eBooks on the Sharp Zaurus, which runs under Qtopia.

Another developer, George Wright, said he felt that the main advantage Qt has over other embedded development environments is that it is based on the same source code that is used on the desktop side of many programs – Qt.

"Qt is a truly cross-platform toolkit which is available for Windows, UNIX and Mac OS X, and so if a program simply needs porting and it was originally written in Qt/C++, then a bit of code modification and a recompile is all that is needed. In the Palm OS world or EPOC, you would generally have to rewrite the interface again from scratch," Wright explained.

Other appealing features, he said, are support for Java applications, using various Java runtime engines, such as Insignia's Jeode which ships with the Sharp Zaurus as well as the existence of third party developments which give support for vital libraries like SDL, which is a standard on Linux for 2D graphics (used in many games). He described it as being a very programmer friendly programming framework making programming much simpler and neater than it would be in, for example, the MFC used by Microsoft.

Both above developers, however, highlighted areas which they felt that Trolltech could improve upon, many of which Kozak said have been addressed in the recent updated release of Qtopia 1.6. These updates include an overhaul of the document handling system, so that applications installed on external media without having to modify the internal filing system or RAM. Kozak said that a Qtopia PIM API has been added, which allows unified access to PIM information. The Qtopia PIM API will aid the developer in the forthcoming transition from the current XML format used to store data in the Qtopia 1.x series to the SQL database in the 2.x series.

Support for TrueType fonts and a font plug-in framework have also been added in version 1.6 and a large update on docs was also included in this latest version. Essentially, this is designed to be a technically more stable version.

Trolltech is working on other developments too which will broaden the appeal of its product suite, as Kozak explained: "We have almost completed the development of a Qtopia Windows SDK, we just need to ensure it works in various IDEs, to find an appropriate license, price etc. We hope to have it released this summer!" ■

HARDWARE TAKE-UP

Sharp, IBM... Who's next?

Sharp was the first major hardware vendor to adopt Qtopia, using it in a range of devices, including handhelds, PDAs and mobile devices.

In March the platform received a major boost when IBM signed a deal to licence Qtopia for use in a PDA design kit which will include the blueprints and software necessary to create a variety of Linux-based PDAs, from inexpensive consumer devices to more complex business models that have the capability to access corporate resources such as email systems.

Michael Karasick, Director of Embedded Development, IBM Pervasive Computing said that the rationale in licensing Qtopia was that IBM is focused on working across a variety of different platforms and that for this 'openness' is the key. It wants to develop open platforms and its work with companies such as TrollTech is part of the greater IBM Pervasive Computing strategy.

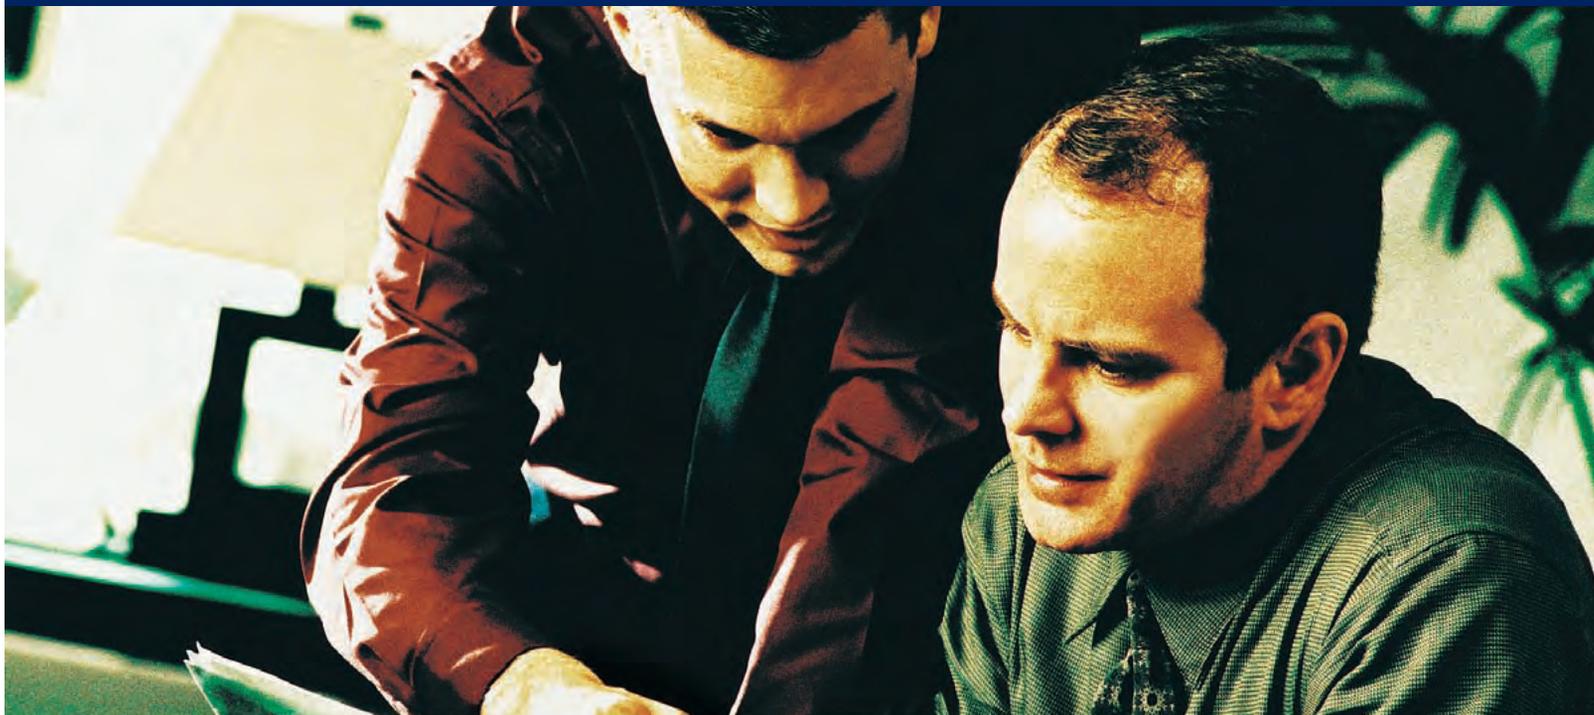


Kozak said that there are several other high-profile hardware vendors who are engaged in Qtopia projects, but are not yet ready to go public.

In spite of the platform's growing popularity it does have a long way to go before it can seriously take any market share from Microsoft or Palm, the two dominant players in the space.

Giga's Quandt is optimistic, however, that Metrowerks, a Motorola subsidiary combined with IBM's support, will broaden the use of Qt/embedded to more microprocessor architectures and hence its wider uptake.

STORAGE



NAS, SAN – neither or both?

Addressing your storage needs in the past, seemed a simple matter of checking the space left on your servers attached disks and buying in some more when things started to get full. The sheer volume of data created in the average organisation these days – with database-driven websites, groupware projects, and more and more users generating more and more data (in larger and larger files thanks to ‘advances’ in apps software) –

NICK VEITCH weighs up the options for advanced storage management.

mean that you soon come to a stage where it really is necessary to adopt a proper storage management strategy.

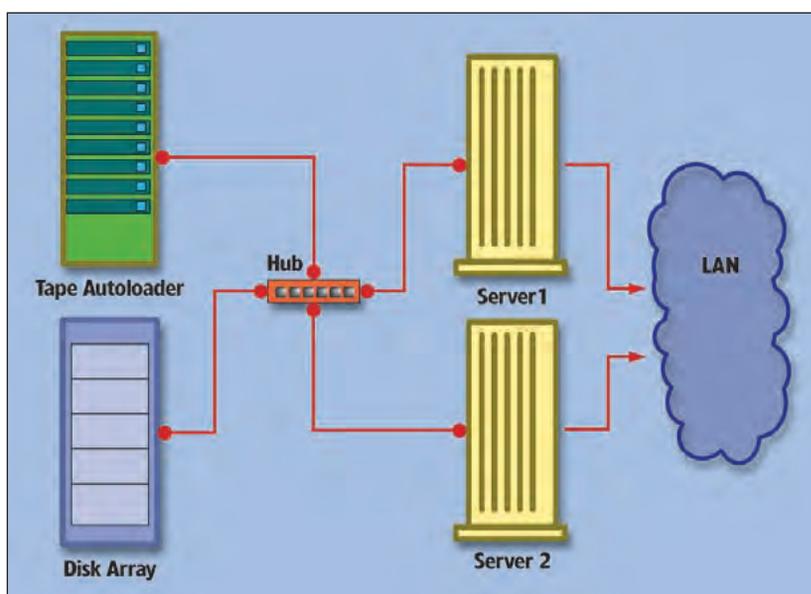
Two technologies have emerged in recent years to solve the problems of storage management and availability – Network attached Storage and Storage Area Networks. These palindromic terms seem to describe similar-sounding things, but as we see in this overview, they are quite different in terms of their technologies and the problems they solve.

WHY SAN?

To boil down the SAN concept to its lowest common denominator (because the actual definition seems to depend on who you are talking to), a SAN is a dedicated storage network infrastructure based on one of the available serial SCSI protocols. The protocols are where the confusion sets in. SANs are commonly constructed using Fibre Channel or more recently iSCSI. The connect fabric allows devices to access connected storage at a raw device level at very high speeds. This is one important difference between SANs and NAS storage, which are based on a shared networking filesystem instead, and is the key to what SANs are best at.

SANs were primarily developed to solve the backup headache. As servers grew larger and larger, therefore so did the backups. So, the practice of backing up to a locally attached tape simply wasn't viable any more. Even if you had the tape capacity (which usually meant migrating from a simple tape unit to a much more expensive autoloader), many people found that the backup cycle took so long that it impaired the efficiency of the server to do what it should be doing. In a significant number of cases, backups couldn't be completed in the desired time at all.

A simple SAN setup providing storage and backup for two servers.



FIBRE CHANNEL IN A NUTSHELL

It's not exclusive to fibreoptic cable, despite what the name may suggest

FIBRE CHANNEL IS MORE THAN JUST A different way to connect up devices. There are some important features of fibre channel that make it ideally suited to SAN. For simple connections and small SAN deployments, it isn't really necessary to know too much about the technology, but an overview of its features and limitations is handy.

Fibre channel is really a layer of protocols

used to transfer data to and from remote devices. In some ways it is analogous to IP, and in fact, you can route IP traffic over fibre channel. The real difference is the way that fibre channel interfaces to the host and devices.

A request to retrieve data would result in the data being retrieved physically from the device, passed on to the SCSI map of the fibre channel.

Data is sent in frames across the physical

network. At the other end the frames are decoded and recombined. As far as the initiating process is concerned, it just issues SCSI commands as it would if were connected directly to the device.

Fibre channel itself, while commonly used on fibreoptic cable, can equally be used on copper cables. The latter are cheaper and perfectly fine for shorter distances (up to 30m).

SANs and iSCSI/Fibre Channel could address these problems, and save money by consolidating backup functions. With a SAN, multiple servers could be connected across a dedicated serial SCSI network to an array of storage devices, whether these were disks, tape drives or autoloaders. With access to these devices at a raw level, complicated backups could be achieved, and thanks to the remote storage, backups and other files can be read by any other server that are connected to the system.

Although the infrastructure of a SAN is expensive, especially when it comes to adding fibre channel hubs and switches, by replacing multiple backup devices, and adding extra flexibility, it isn't all one-sided. The material costs may be high, but as well as the benefits of a more flexible backup infrastructure, it should free up other resources.

Typical uses for SAN

There are several key areas where SAN offers improved performance, flexibility or features over alternatives.

LARGE DATABASES When running a high-performance database, especially a very large one, it's hard to ignore

SAN. With database size on the increase its easy to see how the backup performance of SAN is almost vital. In simple performance terms, with money no object, SAN will deliver better numbers too.

LARGE NUMBERS OF FILES – Because SAN gives raw disk access, it is possible to use filesystem image based backup solutions, instead of reading the individual files. This isn't possible with a NAS solution, which will take significantly longer to backup the same amount of data.

VENDOR LIMITATIONS – Particularly an issue with databases, there are many software solutions that just don't support NAS.

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

WHY NAS?

The growth of the Network Attached Storage sector has been quite remarkable. There are now countless dedicated NAS solutions, and plenty of associated software to make the best use of them available on many platforms.

Essentially, NAS makes use of well-understood and highly developed technologies such as IP networking and networked filesystems. Isn't a NAS then just an NFS RAID array with an Ethernet card? When such devices were introduced, this probably wasn't far from the truth, but as NAS products have matured they do now offer a number of features which make NAS more viable and certainly easier to implement.

If you already have a fileserver, the basics of NAS are pretty easy to understand. Shared storage on networked file servers usually comes in one of four flavours. The most popular are SMB/CIFS (see boxout /LOCATION/) on Windows based networks and NFS on Unix based systems. Some people still use the Novell filesystem, and we



SANs require new dedicated infrastructure hardware such as this Vixel 7100 switch.

shouldn't ignore Apple's AFP.

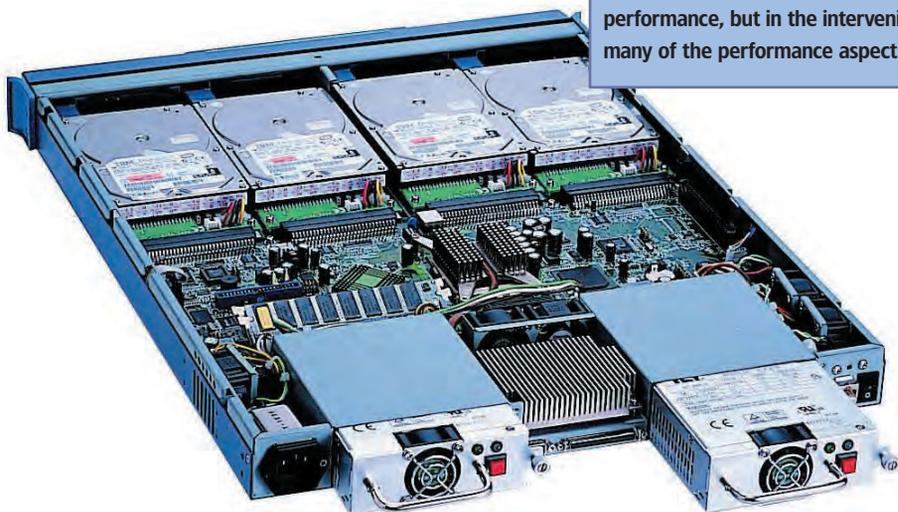
Modern NAS devices can support most of these protocols (which to be honest can be a pain to set up and administer by hand) and also introduce new features to optimise their performance, and critically, to implement failsafe features such as safe caching and better locking procedures. Because a NAS filer is designed from the ground up to be a fileserver, there are further

STORAGE

optimisations that can be made, typically by reducing the number of interrupts and memory accesses required for the data to route the required data from the disk. In essence, this is an optimised kernel that does nothing but support the required protocols. Typically this can be reduced to such a small size it can be permanently resident in hardware. Some NAS devices actually use several processors, each dedicated to a specific task.

Further hardware improvements can include a large amount of memory being available for caching often used files. Typically this is implemented in non-volatile RAM to prevent data loss should the system crash.

In some cases, manufacturers have even taken the step of creating their own specifically tuned filesystems to further improve performance. Coupled with



Gigabit Ethernet, NAS is no longer a slouch when it comes to access speeds, though obviously it does lag behind high-end SAN implementations.

Usually NAS filers are capable of supporting a few other protocols too, such as http for static web pages, FTP access and repository services.

Typical uses for NAS

Familiarity is a good reason for using NAS. You probably already have much of the infrastructure in the way of cabling and hardware that would be needed to deploy it. Even high end NAS devices are relatively cheap – you can get a 320GB Gigabit Ethernet NAS device for less than the cost of fitting Fibre channel adapters to a couple of servers and your storage device.

● DATA CONSOLIDATION

Sticking all shared file resources onto a NAS is probably the thing you'll find filers doing most often. That's primarily

CIFS IN A NUTSHELL

Common Internet File System

CIFS IS JUST MICROSOFT'S NEW NAME for the SMB-based network filesystem. SMB was originally conceived by IBM in the early eighties, but it's main proponent have been Microsoft, who have included support in every version of Windows.

Obviously the protocol has developed over the years. Originally it was designed for simplicity with multiple users accessing the same files. Much of this specialisation was at the expense of performance, but in the intervening years many of the performance aspects have

been addressed. For Linux and other Unix systems, NFS has historically been preferable, because it institutes the 'proper' permissions and integrates better in the Unix environment.

The truth is though that most IT environments these days are mixed, and therefore a combination of Windows and Linux servers or clients need access. Thanks to *Samba*, an implementation of SMB/CIFS for Unix systems, it is possible to use SMB shared folders for both, although there can be some limitations.

what they were designed for. The management benefits are pretty clear – a single shared NAS device is easier to backup and configure than having multiple shared user directories on dozens of different servers.

● DATABASES

Higher-end NAS devices are perfectly feasible for storing database information, with benefits in transparency, ease of management, scalability and availability. Snapshot backup features reduce the need for complicated bespoke backup systems, though backup speeds will typically be slower than SAN

● WEB APPLICATIONS

Web apps and groupware projects are also suitable for NAS, as they are much easier to scale to meet requirements – ideal for short projects too.

Simple rackmount NAS filers like this RareSystems 4100 offer a lot of features for the price.



There are also options for seriously large scale storage, like this IBM 300 NAS.

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

Which is best?

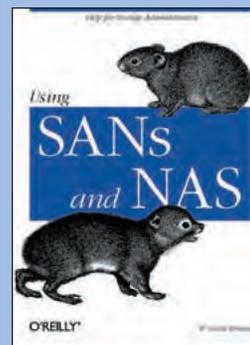
So should you choose to expand or update your

infrastructure with SAN or NAS? There is a valid case for both, and the answer to the question largely relies on what it is you intend to achieve, and what is important to you.

From a cost point of view, NAS has many advantages, but it would be no use saving all that money if a NAS-based solution doesn't achieve your goals – if you need a really high performance solution, for the moment SAN may be the answer. The other thing to bear in mind is that technology is changing, and faster LAN speeds, eg with 10GB ethernet will make NAS even more attractive.

Larger organisations tend to find that a mixture of both is the solution. Though it may initially appear to be an unnecessary increase in the IT management workload, it allows both technologies to be played to their strengths. ■

FURTHER READING



Using SANs And NAS by W.Curtis Preston (O'Reilly, ISBN:0596-00153-3), covers the technologies and how to implement them, and has practical management advice, network topology and backup strategy. Essential reading.

BRIEF GUIDE TO RAID

Mirroring and striping

Almost all NAS and SAN devices will use RAID, either as a redundancy protection or a performance booster. The term RAID means Redundant Array of Independent Disks, but there are several different ways to manage RAID storage, referred to as levels. (and more confusingly, some can be implemented with others). The basic principle of a RAID is to group together the physical storage medium into one virtual pool of data with some added error protection.

In broad terms there are two techniques involved – mirroring and striping. Mirroring is simply duplicating data on to more than one physical drive. Thus, two 30GB SCSI devices provide 30GB of storage. If something catastrophic were to happen to a particular sector of a disk, or even a whole drive, the data could be read from the working disk (this is where the redundancy comes in).

Striping is a technique used to improve performance. Striping is simply storing different chunks of data on different drives. Alternate data packets are read or written to different physical devices, almost halving the time it takes to read or write data. Pure striping is known as RAID 0, although, since there is no redundancy in this system, it isn't really RAID at all. Another restriction here is that it's quite important to use identical drives if possible, otherwise one mechanism will tend

to be faster than the other which will negate much of the speed increase.

Although you may come across RAID 0 being used on its own, it is more commonly used in combination with the basic mirroring of RAID 1. Though this in itself can be confusing – because there are two ways of doing it – striping and then mirroring or vice versa. Here is a rundown of all the available RAID modes.

● RAID 0

This is, as explained above, simply striping the disks in the array for extra performance. Although sometimes used on its own, it is most often used with RAID 1 mirroring.

● RAID 1

This is the first true level of RAID. Data is mirrored onto a duplicate set of disks. The disadvantage here is that all the data needs to be written to two drives. As both sets of disks contain the same data, reads can be performed on either disk, so the disks can share read requests. Reading performance is therefore usually much faster than on a single disk.

● RAID 0+1

This is a mirrored pair made from two sets of disks striped with RAID 0. This is created by first striping a set of disks and then mirroring them. This provides some protection from data

loss (you will still lose data if the two physical disks it is mirrored on suffer a failure). The main drawback here is in recovery times should a failure occur.

● RAID 1+0

This type of RAID is created by creating mirrored pairs in the first instance and then striping them with RAID 0. The main advantage to this over the previous method is that recovery times are much faster should an error occur.

● RAID 2

Not often used in practice, this uses a disk layout with a parity checking ECC code.

● RAID 3

Synchronised reads and writes optimise performance for large blocks of data, but it's only really useful for single-stream tasks – where one process is storing or reading a large amount of data.

● RAID 4

Uses a form of interleaved striping for parity data which is stored on a single parity volume. The advantage of RAID 4 over RAID 5 is that disks can be added to the array without recomputing the parity information.

● RAID 5

Pretty much identical to RAID 4, except in this case the parity data itself is distributed across all the drives in the array, removing a potential bottleneck.

flex and bison Compiler writing

Writing your own programming language may at first sound like a daunting task – after all, it surely must be more difficult than programming in an already-existing language? The answer might surprise you – UNIX has had tools to enable easy compiler creation for years now, in the form of *lex* (available as far back as 7th Edition UNIX in 1979) and *yacc* (“Yet Another Compiler Compiler”, 10th Edition UNIX). Naturally, there are now open-source equivalents of both of these, *flex* (the Fast lexer, available from www.gnu.org/directory/flex.html) and *bison* (a cheap pun on *yacc/yak*, available from www.gnu.org/directory/bison.html).

In order to design our own programming language, there are several requirements. Firstly, solid knowledge of C++ programming. Yes, this is the second tutorial by me this month that requires C programming; no, *LXF* has no plans as yet to run a C programming tutorial – write to us if you disagree! I specifically say that C++ is important as

If you’re itching to try something new with your system, why not write your own programming language? PAUL HUDSON shows us how easy it is, with the help of *flex* and *bison*.

opposed to C, specifically because it’s possible that we may need to use the STL to make life easier for you.

Secondly, you’ll need to have *flex* and *bison* installed. These can be downloaded from GNU directly, although it’s more likely you’ll already have them installed by default.

Thirdly, *bison* grammar is similar Backus-Naur Form (BNF), but I will be explaining grammar as we go, so don’t worry if you no little or nothing of it.

If you have all of the above, you’re ready to begin. Writing your own language can be surprisingly fun when you consider all the possibilities it entails. For the purpose of this tutorial, we’ll be creating the language SKYLang. SKY, if you’re interested, is the inverted ROT13 of “LXF”.

What is a compiler?

A compiler is more complicated than you may first think, so it may be best for us to think through any



FLEX AND BISON

◀ existing definition we have to make sure we're thinking on the same lines.

A compiler is a program that translates human-readable source code written in a high-level language (such as C++, Java, or, ahem, SKYLang) into machine-readable code. In essence, it converts programs from one language to another. Compilation can be broken down into various steps, the most basic of which is "Input, Analysis, Synthesis, Output". Steps 1 and 4 are simply the reading of the input language and the outputting of the target language – the real work goes on in steps 2 and 3.

The analysis stage can be broken down into three sections: lexical analysis, syntax analysis, and semantic analysis. Lexical analysis is the process of literally scanning through each character of the source input, trying to decide what the text is. For example, given this code:

```
myvar = 10 + 10;
```

a C compiler might see the following in its lexical analysis stage:

```
VARIABLE ASSIGN EQUALS NUMBER PLUS NUMBER
```

Apart from whitespace, which is ignored in C, each part of the source code is represented by an identifier, or *lexical token*. It is this group of lexical tokens that the lexical analyser (usually called the *lexer*) passes to the syntax analyser (usually called the *parser*).

The function of the parser is to extract grammatical structure from a group of lexical tokens. A parser contains a built-in grammar of rules that it uses to match syntactically valid input. That is, the parser steps through each token (often called a *lexeme*) passed to it by the lexer to decide which rule to apply to the instructions.

The parser literally jumps through various multiple case statements to decide which rule to apply, shifting more and more lexical tokens onto its internal stack until it finds a rule match. Consider the following rules:

- A number is one or more characters of 0-9.
- A variable is one character beginning with A-Z or a-z, then one or more numbers as described above.
- An expression is one number plus another number.
- A statement is a variable being assigned to an expression.

“A compiler translates human-readable source code in a high-level language (like C++ or Java) into machine readable code”

For example, when matching our C code above, a C parser would read something like this:

```
VARIABLE
```

```
VARIABLE ASSIGN EQUALS
```

```
VARIABLE ASSIGN EQUALS NUMBER
```

```
VARIABLE ASSIGN EQUALS NUMBER PLUS
```

```
VARIABLE ASSIGN EQUALS NUMBER PLUS NUMBER
```

The first four lines do not match any of our rules, so internally the parser will shift the token onto its stack and go onto the next token. However, line five contains two rules that are matched. Firstly, **NUMBER PLUS NUMBER** is an *expression* as defined above, so *bison* will execute any

THE AIM OF THIS SERIES

IN THIS SERIES WE HOPE to be able to teach you how to write your own compiler for our basic language, SKYLang. At first, we will be covering designing a lexer and parser using flex and bison to compile and interpret source code. Later on, given sufficient support for the idea, we may go further to discuss how to produce semantically valid, optimised intermediate code. Much later on, if we get emails from readers saying it would be of interest, we may alter the compiler so that it outputs code in x86 assembly – drop us a line if you're interested. Anyhow, we don't want get ahead of ourselves – for now, we're concentrating on flex and bison.

action we define for the rule NUMBER PLUS NUMBER, then convert those three tokens into one, EXPRESSION. This is called *reducing* a rule. As a result, we have another rule match, **VARIABLE ASSIGN EQUALS EXPRESSION** is a *statement* as defined above, so *bison* will execute any code we define for that rule. Parse complete!

As you can see, parsing is also fairly simple. However there are two key problems you will find when using *bison*: firstly, it has a very specific rule-set, and cannot parse ambiguous grammar, which is common sense, really. Secondly, *bison* has a one-token lookahead, which means it is able to look beyond the current token by one place to help decide what the current token is. Don't worry too much about this latter problem – you're not likely to encounter it unless you have more complicated grammar than we have room to discuss here!

Continuing on with how a compiler works, the semantic analyser comes after the lexemes have been discovered and after rules have been matched, and its job is to decide whether the code is semantically correct – that is, that the code actually means something valid. For example:

```
int i = "bar"
```

is a syntactically valid C line, but semantically incorrect – an integer holds a number, not a character string. Semantic analysis, as you can imagine, is tricky!

So far, we've covered the *analysis* stage – now onto the *synthesis* stage of compiler theory. If you're running out of patience with theory, don't worry – you can skip to the next heading, because the synthesis stage isn't required knowledge for all users!

The synthesis stage, as with the analysis stage, can be split into smaller parts: in this case, these are intermediate code generation; then optimisation. Intermediate code generation is exactly what it says on the tin: compilers generate an intermediate form of the source code in an internal structure. This might seem a little pointless, but it serves several purposes: firstly, everything up to this point has been machine-independent, whereas the optimisation phase of compilation and also the output phase are very much machine-dependent, so generation of intermediate code allows an easy split between the two. Secondly, compilers can be split into a front-end and a back-end, with a shared intermediate code structure, which allows compilation to be retargeted to other platforms more easily.

Optimisation is simply the process of repeatedly taking existing operations and replacing them with better operations that require smaller code or faster code. Common machine-independent optimisation techniques include constant folding, strength reduction, sub-expression elimination, and the like, but there is a lot that can be done at the machine-dependent level, for example making use of special processor instructions offered by MMX and SSE. Optimization works by looking for known patterns of slow code, and replacing them with code that is known to be better.

Finally, we reach the Output stage of compilation, where the optimised intermediate code is converted into our target language and outputted. Sometimes the code is outputted to a file, as seen in GCC where the target language is assembly. Other times, the code is outputted

direct to an interpreter, where it is executed directly, as seen in PHP. Interpreters merely go through each operation (“op code”) that is given to them and execute the appropriate instruction.

So, to summarise, a compiler takes source code as input, and converts it into tokens that in turn form rules. When a parser has matched its rules, it creates intermediate code which can then be optimised and converted to the target language. The use of intermediate code allows multiple languages to be compiled to one language, which can then be converted to native code for multiple platforms.

What are flex and bison?

flex and *bison* are two tools that work together wonderfully, but do not have to be used together. *flex* uses regular expression to match character strings to lexemes, which are then passed to *bison* to match against its rules. In order to use the two of them, you write source files in a special format, which are converted by *flex* and *bison* into portable C code. The grammar of the files they accept to create a lexer and parser are somewhat similar, and not too complicated at all.

flex

When given an input file of regular expressions and C code, *flex* generates our lexical analyser. *flex* doesn't require *bison* to operate – we can write a lexer that operates entirely standalone. The *flex* grammar is remarkably simple – here's a simple lex program:

```
%%
.\n      { ECHO; }
%%
```

Save that as `simple.l`, then run `lex simple.l` and then

`gcc lex.yy.c -ll` to compile it. All being well, you should have an `a.out` in your working directory. When you execute it, you'll find it works like the command `cat` with no parameters – any input you type in with a return at the end is repeated back to you. Press Ctrl-D to quit. The percent signs serve to separate the grammar into various sections – we'll come to that more soon. Line two is the key line, and contains two parts. Firstly, we have `.\n`, which is a regular expression. Secondly, we have `{ ECHO; }` which is the action **flex** should take if it matches the regular expression. `ECHO` is a macro which prints the matched text straight to output. The regular expression translates as:

```
.- match any character except new-line (\n)
| - OR
\n - new-line.
```

Therefore, our regular expression is “match any character except new-line OR any new-lines”, which is “match any character at all”. So, our program reads in every character, and echoes it out.

Now modify it to the following:

```
%%
[0-9]+  { printf("Number!"); }
.\n     { ECHO; }
%%
```

flex it, *gcc* it, then run it – this time, typing any number will output **Number!** rather than the actual number you typed. Example interaction might be:

```
test
test
number
number
99
Number!
```

SKYLANG

All aboard *Linux Format's* very own programming language!

IN ORDER TO CREATE A LANGUAGE, WE MUST first define what its capabilities are. There are lots of things one can do with languages, and as such they range from fairly simple (BASIC) to very complicated (C++, although, when you think about it, shouldn't that be ++C?).

Naturally we're space-limited here, but here's what I intend to program in for SKYLang:

- untyped variables
- built-in functions
- user functions
- loops

These four should be enough to allow a fairly complicated language, although you will probably want to add more functionality later.

Now, for the important part: definitions of the language.

A variable must begin with a dollar and a letter from the range A-Z and a-z or an underscore (`_`). It can then consist of one or more letters from A-Z, a-z, 0-9 or an underscore. So, `$foo`, `$f00`, and `$_foo99` are

valid variable names, whereas `$99`, `$4foo`, and `$!""_3` are not. A variable can be an integer, a floating point number, or a character string.

Users should be able to use variables interchangeably, so an integer can be read into a string without the need for type conversion.

A function consists of the reserved word “function” followed by the function name, an opening bracket, a comma-separated list of arguments, then a closing bracket. Functions may not take a variable number of arguments or optional arguments. The code contents of functions will be enclosed within braces. Valid functions will look like: **function myfunc(\$var1, \$var2) { SKYLang Code Here }**

A loop consists of the reserved word “for” followed by a variable, then an equals sign and an initial value for that variable 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 by 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 }**

Each statement should end with a semicolon. `=` will be used to assign variables, with `==` being used to test variables. `""` (a full stop/period) will be used as the concatenation operation to join strings together. Function calls that have no parameters still require brackets, eg: **myfunc();**

Strings are double-quoted. Variables inside strings will not be expanded. Double-quotes inside strings need to be escaped with a backslash, eg: **This \"is\" a test**

Make sense? If not, carry on reading – it will be explained more over the coming pages and issues. Owing to page constraints, the above definition may alter slightly as time goes by, but we'll do my best to stick to the above. Things like variable expansion inside strings might be possible if we're lucky space-wise!

FLEX AND BISON



```
number99
numberNumber!
```

```
99foo1092
```

```
Number!fooNumber!
```

flex matches its regular expressions in order, and executes only one action for each match – that's why the ECHO action isn't executed when a number is matched, because the number action calls **printf()** and stops there.

flex can be extended by using the space before the first %% to include raw C code. Adding a variable **int nc**, for example, we could count the number of times a number was entered. We'll come back to that in a moment – but first, *bison*.

bison

As you can see running our simple *flex* programs, *flex* by itself can be useful, however it is limited in that it returns lexemes singly rather than being able to match groups. This is where *bison* comes in – it allows us to define actions to take when certain groups of lexemes come in.

To create a *bison* parser, *bison* needs a BNF-like grammar to decide its rules. This grammar defines all the different tokens we accept, and our *flex* lexer needs to be rewritten to take advantage of these tokens. A *bison* grammar looks much the same as a *flex* grammar, as you'll see. Save this new *flex* grammar as `variables.l`:

```
%{
#include "variables.tab.h"
}%
%%
\[A-Za-z_[A-Za-z0-9_]* { yyval = strdup(yytext); return
T_VARIABLE; }
[0-9]+ { yyval = atoi(yytext); return
T_NUMBER; }
= { return T_ASSIGNEQUALS; }
.\n { ECHO; }
%%
```

In this new *flex* source file, the top part is now used to import C code, which needs to be enclosed in `%{` and `%}`, and is copied directly into the final lexer code for GCC to handle. We now have a new regular expression to handle variables – as you can see, a variable is currently defined as we planned for SKYLang: it starts with a dollar followed by a character or underscore, then is followed by any number of characters, numbers, or underscores. Note that the dollar sign needs to be prefixed by a backslash in order for it to be properly differentiated from the `$` regular expression symbol, which means “match end of line”. (see the box, *Regexps* above right for more information)

We also match equals signs and keep them handy.

Note that each time we match something, we return a special value – the Type (**T_**) of lexeme we matched. This type, defined in our *bison* grammar, is used in *bison* rules to decide what action to take. Also, we use two special variables, **yyval** and **yytext**, to hold information about the lexeme read. **yytext** is set by *flex* to be the full lexeme matched, eg “1024” for a number, or “\$ _foo” for a variable. However, **yytext** is not available inside *bison*, whereas **yyval** is, so we need to copy **yytext** into **yyval** either using a straight assign (for numbers, and we use **atoi** to

REGEXPS

A miniature arsenal of symbols to get *flex* working magic

```
foo match the string “foo”
^foo match “foo” at the
start of a line
foo$ match “foo” at the
end of a line
^foo$ match “foo” when it's
alone on a line
[Ff]oo match “foo” or “Foo”
[A-Z] match any upper-case
letter
[a-z] match any lower-case
letter
[A-Za-z] match any letter
[A-Za-z0-9] any letter or
number
[A-Z]+ one or more letters
[A-Z]* zero or more letters
[A-Z]{1,3} match 1 to 3
upper-case characters
-[0-9]+ match a number
with optional leading minus
[^0-9] any non-numeric
character (^ is “not” inside
square brackets)
[^0-9A-Za-z] any symbol
(not a number or a letter)
```

convert the string **yytext** to the number **yyval**) or the **strdup()** function for character strings. As `=` is always `=`, we don't bother putting anything into **yyval** – we just return **T_ASSIGNEQUALS**.

Here's the *bison* grammar to handle the above. Save it as `variables.y`:

```
%{
#include <stdio.h>
}
void yyerror(const char* s) {
printf(stderr, "%s\n", s);
}
main() {
printf("Beginning to parse...\n");
yyparse();
printf("Done parsing!\n");
}
}
%}
%token T_VARIABLE T_ASSIGNEQUALS T_NUMBER
%%
statements:
|
statements statement
;
statement: T_VARIABLE T_ASSIGNEQUALS T_NUMBER {
printf("Setting %s to %d\n", $1, $3);
}
%%
```

Most of the complex stuff there is done in the C section at the top. Firstly, we include **stdio.h** so we can write to `stdout` when we have errors. Then we define a function, **yyerror**, to handle any errors caught whilst parsing. *bison* requires that we define **yyerror** – it will not compile easily without it. The **main()** function simply calls *bison*'s default **yyparse** function to get parsing going. Note that you can pretty much ignore for now these two functions, as they won't change.

The important parts of the *bison* file come after the `%}`. First, we have **%token**, which is a *bison* declaration that defines which lexical tokens we are to expect from our lexer. After the `%%` and before the next `%%` is where we define our *bison* grammar rules. Each rule is defined as such: name of rule, colon, what the rule matches. Note that matches are often combined together using an `|` (OR symbol) which will match multiple things.

The first rule in the *bison* grammar, `statements`, will match either nothing (there is nothing before the OR symbol), that is, the end of the program, or “statements statement”. “statements statement” is recursive – the rule will match itself and the “statement” rule, which is defined later. This recursion allows *bison* to keep on reading statements by continually expecting more and more.

The statement rule is made up of tokens defined in the earlier **%token** statement. Here, a statement is defined as being a variable followed by an equals sign followed by a number, and if *bison* finds a match it executes the code inside the braces. The code for our statement rule simply

calls the `printf` function, however note that it passes to it as parameters `$1` and `$3`, which as you can see from the first parameter to `printf` are a `char* (%s)` and a number (`%d`). These `$` “variables” are replaced at runtime by *bison* with the corresponding part of our rule – so `$1` matches `T_VARIABLE`, and `$3` matches `T_NUMBER`. The reason `$1` is set to be a character string is because we sent `yyval` through from the lexer using `strdup`, which duplicated the string in `yytext`, whereas `$3`, our number, was sent through converted to a number using `atoi`.

Now it's time to build an executable of this. Because executable building is something you'll need to do quite a bit, we're going to use a makefile. As this tutorial isn't about how to write makefiles, I'm just going to give you the source code – you may want to read up on it by yourself if you're interested in learning how it works.

```
FLEX = flex
BISON = bison -d

CC = gcc

binary : variables.tab.o lex.yy.o
$(CC) variables.tab.o lex.yy.o -l

lex.yy.o : lex.yy.c variables.tab.h

lex.yy.c : variables.l
$(FLEX) variables.l

variables.tab.o: variables.tab.c variables.tab.h

variables.tab.c variables.tab.h: variables.y
$(BISON) variables.y

clean:
rm -f *.o
rm -f *.c
rm -f *.h
rm -f a.out
```

Save that as Makefile in the same directory as the other files. Note that the indents are tabs, not spaces – if you see errors like “Missing separator” while you're running `make`, check you're using tabs rather than spaces.

Now type `make clean`, `make binary`, then `ls` from the shell – you should see our `a.out` file ready to run. When it's executed, we should be able to have a slightly more intelligent “conversation” this time! Here's an example:

```
Beginning to parse...
```

```
$foo = 20
Setting $foo to 20

$bar = 40
Setting $bar to 40

$wombat = -10
-Setting $wombat to 10

$bestlang = skylang
skylang
```

```
help
help
10
Setting $bestlang to 10

10
syntax error
Done parsing!
```

That example shows up two things: firstly, what has been programmed already (setting variables to numbers) works fine. However, note how attempting to set `$wombat` to `-10` outputted `-Setting $wombat to 10`. The reason we see that is because our definition of a number doesn't include the possibility of a `+` or a `-` before it, so the minus sign slips through to our catch-all `.|n` rule, and is echoed out. However, the `$wombat` part is parsed successfully as a

“Note that indents are tabs, not spaces. If you see errors like ‘Missing separator’ running *make*, check what you're using.”

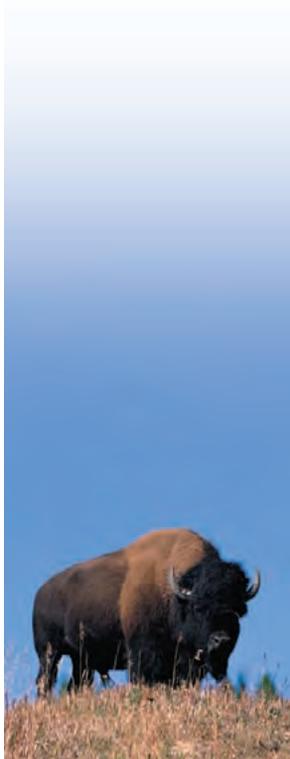
variable assignment, so it follows on. Also illustrated is my attempt to set a variable to a raw string (that is, an unquoted string). Unquoted or not, SKYLang does not yet currently have the capability to handle string variables, so it remembers that we tried to set a variable, and echoes out what it doesn't recognise – `skylang`. Again, typing `help` doesn't match any rule but our catch-all, so it's echoed out also. However, typing `10` completes the earlier rule started with `$bestlang =`, and so `$bestlang` is nominally set to `10`.

Finally, just to show what happens when the system is entirely unable to cope with input, I typed `10` at the end, to which we get a syntax error and automatic program termination. The reason this errored out when typing `help` and `skylang` earlier didn't is because this returns a token to *bison*, which can then syntax check it, whereas our catch-all returns nothing to *bison* – as far as *bison* knows, the words `skylang` and `help` didn't even get typed.

Conclusion

In this, the first part of our short exposition into compiler writing, you should be able to take away quite a bit of new knowledge. Firstly, I hope you now know quite clearly how a compiler works, what *flex* and *bison* are, and how they all work together to allow us to create our own language. Secondly, you should now be able to see that compiler writing really isn't that hard at all – one can rely on pre-existing tools to take away much of the hard work. Finally, hopefully you're encouraged that in approximately two pages of text we've managed to create a language that recognises variable assignment – there's lots more to come, and we'll be able to cram in much more now that the compiler theory is mostly out of the way!

As you can see, SKYLang is far from perfect. We'll be adding to it as time goes by, but in the meantime you're encouraged to play around with it yourself – try some new regular expression in the *flex* file, for example, to handle numbers that are both positive and negative. ■



THIN CLIENTS

RETURN OF THE Thin Client



Back in the days when computers were expensive and desktop computers were an expensive business tool exclusively for high-powered execs, most serious computing was done on thin clients. Those old enough will probably remember the green screens and tinny keyboards of the dumb terminals that populated university computer rooms and large corporate departments. But the PC heralded the end of those devices. Nobody wanted to use a slow, unappealing terminal when they could use a PC instead.

Centralised computing is back on the IT agenda, as NICK VEITCH discovers.

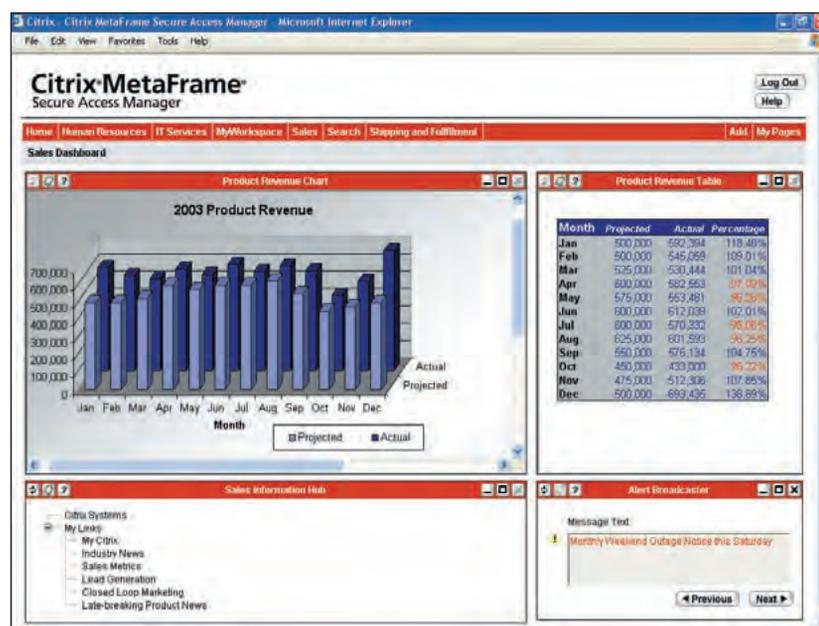
This may have been great from a user point of view, but it also caused problems, both for the users and for the people managing the systems in large organisations. Most of these problems stem from users finding ever more inventive and catastrophic ways of causing the PC to stop working – whether that be from installing software or just accidental mishaps like coffee spills, deleting system files or trying to reconfigure settings they don't understand. IT departments have never been busier, especially since many employers short-sightedly provide the barest minimum of computer training in an effort to keep costs down.

But there is an argument that this isn't the way IT should be run. The computers should be managed completely by the IT department. The users should just, well, use them to get their work done.

Graphical clients

The return of the thin client is mainly thanks to the ability of a server to deliver a graphical desktop over ethernet. There is more than one way to do this, including proprietary solutions such as Citrix, or the good old-fashioned free Xfree86 X11 server, used already on almost every Linux system.

Bandwidth can be a problem with X11 – running large numbers of clients can quickly swamp even a fast ethernet connection, although Xfree86 does have extensions for compression, or the connection can be handled through compressed SSH, with the added benefit of a little more security. VNC is a much more lightweight option – although it does still take up some space on the client and may be a little more effort to set up.



The Citrix server software can serve MS applications to a Linux based thin client.

PROS AND CONS OF THIN CLIENTS

Pros

- Usually Cheaper to implement
- Reduced IT management overhead
- More effective use of resources
- Scalable
- Various remote admin options

Cons

- Not suitable for some high performance tasks
- Limited usability of external devices
- High server/network demand
- Can introduce a single point of failure – eg if the network/server goes down

Hardware

It is possible to repurpose desktop PCs for thin client/terminal use, but for most people, the cost benefits in terms of support and longevity tend to lead to a dedicated thin client solution. One of the reasons for this is that a thin client solution is much easier to implement across a single client type – you don't really want different client hardware with different graphics cards, devices and what-have-you – they make management that much harder.

Considerations

Of course, thin clients aren't for everyone. Common pros and cons are listed opposite, but need more explanation. The most obvious disadvantage is the increased strain put on the network environment. Even tasks that aren't computationally intensive, such as word processing or web browsing, can actually be intensive from a network point-of-view. This, coupled with the cheap hardware often used, makes them unsuitable for some applications, particularly those involving intensive graphics.

In many cases, the cost benefits outweigh these considerations though. On the support and management side, everything can be done on the server – there's no need to have IT staff trailing round buildings trying to locate errant desktops anymore. Plus there is a more efficient use of storage space and processing power – imagine if you could group together all the disk space on every desktop PC in your organisation. That's effectively what happens with a thin client setup.

Not only can storage be managed more effectively, but also allows proper backups and usually dynamically scalable storage solutions. Depending on the server solution, this too can be scaled easily to meet with demand.

“The thin client solution is easy to implement across a single client type.”

FURTHER RESOURCES

LTSP – the Linux Terminal Server Project is a system for using standard or stripped down desktops as diskless clients running X.
www.ltsp.org

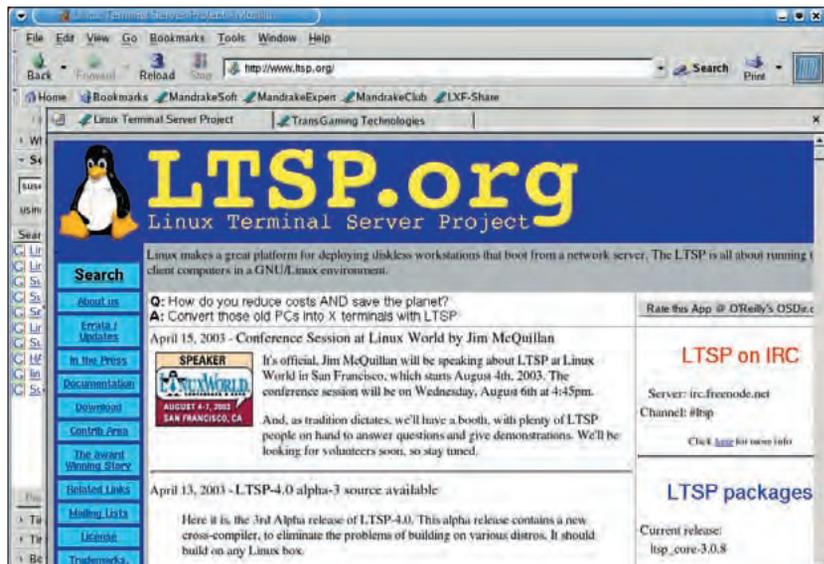
NEOWARE – One of the largest manufacturers of thin client devices, including the Linux based Cacio device featured in the case study on the next page.
www.neoware.com

CITRIX SYSTEMS – Provides software and services for servers to run thin clients from. The best-known solution is the Metaframe range of software.
www.citrix.com

NIC – Linux-based thin clients.
www.thenicstore.com/nic

THINGUIN – more info on Etherboot and various other thin client/terminal projects.
www.thinguin.org

LIN:WARE – customised Linux client software for turning PCs into thin clients, supports the ICA protocol used by Citrix. Also supply thin clients.
www.linware.com



The LTSP project has been going for a few years now, and offers software and a great deal of advice on thin terminals for Linux.



VNC is much more lightweight than serving X11 directly across a network, and is cross-platform too.



Etherboot is a project for creating network booting ROMs for Linux.

« Putting it into practice

Linux Pro caught up with John Webber from the University of Northumbria, who have an aggressive scheme to replace desktop PCs with Linux-based Neoware Capio devices, served from a Citrix server. How well does the thin client approach work in practice?

LXF: Hi, can you briefly explain how you came to be running a thin client solutions?

JOHN WEBBER: The university had some developments that thin clients were very appropriate for. They were specialist applications, databases. That's where we started with thin client. Then we discovered we were using the thin client as a desktop environment for administration.

When our open access areas needed an upgrade, we looked at implementing a desktop on Citrix, and then what sort of clients we could hang off that. Basically, from the point of view of manageability and updates that were happening on a regular basis in the open access areas, we started testing different terminals.

We finally decided on the Neoware Capio as that seemed to give us everything we needed – a nice interface, a very manageable device, and something we could attach a floppy disk drive to.

LXF: A floppy drive?

JW: The students still want to keep stuff on floppy disk. Why I'm not sure, because we give them about 60MB of centralised disk space, but they still like floppies.

We tested 50 Capios in an open access area and got some feedback from the students. A lot of the comments seemed to suggest that they found them no different to using a normal desktop PC. From our point of view that was excellent. From there we moved to 120 devices.

LXF: Are all the clients on the same site?

JW: No. We built a new resource centre in Carlisle, and rather than put in PCs, we installed more thin clients there. They are connected across JANET, the academic network, to servers in Newcastle.

LXF: And that has worked OK?

JW: From our point of view, the servers are based in Newcastle, so if anything goes wrong we can fix it here without having to go across to Carlisle.

LXF: A tremendous saving from the management and time-saving points of view.

JW: From the management point of view it's great. Also

the Capios have a facility to shadow directly to the device – taking remote control of the system. With other systems we can only do this when a user is logged in, but with the Capios we can shadow the actual device, so we can actually help people to login, which is useful.

LXF: What system did this replace?

JW: Our priorities within the university are the teaching labs – which tend to get the newest equipment. The open access areas were equipped on a sort of cascade level, where the older desktops got used; but now they have been replaced by Capios.

LXF: And how long have you been implementing thin clients at the university?

JW: I've been working on Citrix for about 2 years. In that time we have replaced two open access areas with thin clients. There are still running on PCs which we have, but as they die, they are being replaced with thin clients.

LXF: Why the Linux option for the client?

JW: The underlying operating system wasn't a major concern for us as long as the management tools were good. Having said that, we found the Linux systems a lot easier to lock down.

LXF: ...and cheaper

JW: Yes, there is a cost implication too.

LXF: You mentioned feedback. Were any users upset they couldn't do things they used to?

JW: No, not at all, because when they were on PCs they were locked down anyway. So they could still do everything they used to. They are mainly used for email, word processing, web browsing and a few other specific things like MiniTab. That's about it. Specialist teaching stuff is done in the labs, and there we have a tendency to use new PCs, one or two years old at the most.

Not much bigger than an external modem, there's no wonder that some users thought that they were being sold short when their clunky PC boxes were replaced!



LXF: Do you think that might change or are the thin clients only suitable for the open access stuff?

JW: Well, if we take it a stage further, we are now changing admin staff over to thin clients. A lot of them didn't want them to start with – because they perceived that we were taking their PCS away from them. I could take you to a couple of places where they wouldn't go back to their PCs now though. The great thing there is we offer a desktop, but we also offer access to all the university admin systems. In the past getting access to these systems involved filling in a lot of forms, then someone had to come to your PC and install a client or whatever. Now we can just add them to the user group and they can be running almost immediately instead of waiting for days.

LXF: So it makes policy management easier?

JW: Much quicker. It boils down, using terminals or PCs turned into terminals, thin clients are much easier for management and support. For example, even though we used the same model of PCS in the open access area, we had to maintain an image of that system and go round doing updates quite regularly. That's all gone now. Whatever we need to change we can do it on the server and push it out to the clients.

LXF: These have been a success for you. Any areas were you didn't find it went the way you expected?

JW: Well, we did have some teething problems. We had problems with printing initially, but that turned out to be a driver issue. The actual terminals themselves went in without any bother. We just took them out of the box and they just worked.

LXF: Was there any resistance to swapping PCs for thin clients?

JW: A lot of the users didn't like the idea that they were losing their PC, because they thought that they would be losing out in some way. But then when you put in this nice little silver device, with a TFT monitor, and they can reclaim loads of desk space, and then it works as advertised, that's a big win for us. A lot of them thought when we said we were going to put terminals in that they were going back to the old green-screen terminals. Until they actually see it and use it...

LXF: Many people may remember their first experience of computing as being on those old terminals. That sort of idea was prevalent in the day, then it sort of went out of fashion. But now it seems to be back.

JW: I think you have to look at what was feeding it. The original terminals weren't really graphical, so the PC on the desk became a status symbol. You had a graphical interface that was easier to use, and looked smarter.

Then companies like Citrix have taken that environment back to the server. So the management is coming back to the people that should be managing and looking after the devices, and people using them can get back to their work. When you have a PC on the desk you effectively become a sub-administrator for that device. Users had the ability to

change and modify their environment, and also install software, and maybe fix some problems, but usually when something went wrong the support issue, the demands on time became extremely high. Now we have put the block on people loading software or whatever, they are getting back to doing what they should be doing. It may seem a bit harsh that, but should people really be downloading games or fancy screensavers that eventually end up putting their computer out of action for two days.

LXF: The longevity of thin client devices is obviously increased over PCs.

JW: We've had the Capios for about 2 years. We have about 350 Capios, and we have only had one failure.

“With a PC on your desk, you effectively become sub-administrator for that device. Thin clients free you to be productive.”

LXF: And how often were you getting through PCs?

JW: It's a hard comparison to make because we had a lot more PCs than thin clients. It's difficult to quantify in comparative terms, but we do have an engineer from one of our suppliers working on site here. It's usually the moving parts that fail, the Capios don't have any so it isn't surprising they are more resilient.

We do have a few failures on floppy drives that we have connected to the clients. That's usually because someone sticks an old floppy in and it comes out without the metal shim. The record for the most number of shims found in a drive here, by the way, is ten.

LXF: Ten??!! We're surprised they managed to get the tenth disk in! Has the volume of support calls dropped on the whole though?

JW: Yes, definitely. By a considerable amount. There was one open access area called Trinity. We used to get four or five support calls a day there, based around problems with software or things that had happened to the devices. Now, if we get one a week, that would be a lot.

LXF: So now you can all sit around with your feet up?

JW: I have plenty to do! But it means that the people in first-line support can now get on with the things they were meant to be doing – instructing and helping people.

LXF: It frees up resources to tackle other issues?

JW: Yes. In the open access areas, these places are full 9 'til 9, with people queuing to get on. We want to be able to keep all these device working 99% of the time. If one does go down from a hardware fault we can just change it and plug another one in.

LXF: Is the thin client route one you will be happy to continue with?

JW: We have installed around 700 thin client devices in the last 2 years. Whenever we are looking at a new deployment, we now look at thin client as an option. ■