

# Fun in the Sun


**F**aced with the potential of hundreds of Sun puns, I think you've probably got off a bit lightly this issue. What hasn't got off lightly is the new LX50 server from Sun, or indeed their Sun ONE strategy. We have probed and prodded the box in most of the ways imaginable to modern science, so we can bring you a detailed review of this new server and whether it is a great leap forward for Sun and Linux.

The swirling mist of vague rhetoric that has surrounded the Sun ONE platform, and what will be delivered on Linux, has also lifted a bit, so we can now bring you more detail on exactly what is promised, and how Sun themselves see it in strategic terms. In a further Sun-fest, our regular "What on Earth" Feature this month focuses on another Sun-backed initiative – Liberty Alliance – which is well worth checking out.

Another welcome release this month is *Kylix* 3. The original *Kylix* was well received, but rather disregarded by serious Linux programmers, mainly

because it used a special form of Object Pascal. This latest release includes a new IDE for developing C++ applications, which should open up new markets for this software.

We also start a new series this month, aimed at helping you get the most out of *OpenOffice.org*. We were really impressed by this in our recent roundup, and we reckon everybody should give it a try. For those of you who want to do more than just type a few letters and keep a simple spreadsheet, this series should give you some great ideas about what the software is capable of and how to achieve some impressive results. I hope you enjoy it. If you have any suggestions for other things that could be included in this series, please email, write, or phone and let us know. One of our continuing aims is to help you get more out of your Linux experience, so if you have any other ideas for tutorials or features, you owe it to yourself to let us know!



**Nick Veitch** EDITOR

**Is the LX50 any good? What does Sun ONE actually mean? All these questions and more answered p50**

**Kylix – now with added C++ support, this is really two RAD tools in one! p18**

**Get more from your office software with our new tutorial series, starting this issue p76**



# 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...



**Richard Smedley**  
In spite of regularly attending Perl Monk meetings, Rich has so far failed to become enlightened.



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



**Richard Drummond**  
As well as writing our Java series, Rich co-ordinates most of the reviews in the mag.



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



**Jon Kent**  
He scours the Net for new open source software each month, to bring you *HotPicks*.

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

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

**Simon Goodwin**  
A hardware druid in more ways than one, Simon is currently researching every emulator known.

**Andrew Channelle**  
Now studying 'culture' or some such nonsense, Andy still finds plenty time to write the news!

**Brian Long**  
Long time *Delphi* genius, Brian is also a dab hand with Borland's *Kylix*.

## Contact us

Letters for publication:  
[lxf.letters@futurenet.co.uk](mailto:lxf.letters@futurenet.co.uk)

Subscriptions/back issues:  
[subs@futurenet.co.uk](mailto:subs@futurenet.co.uk)

Technical help/Ask the Experts:  
[lxf.answers@futurenet.co.uk](mailto:lxf.answers@futurenet.co.uk)

Disc problems:  
[lxf.support@futurenet.co.uk](mailto:lxf.support@futurenet.co.uk)

General enquiries:  
[linuxformat@futurenet.co.uk](mailto:linuxformat@futurenet.co.uk)

Website: [www.linuxformat.co.uk](http://www.linuxformat.co.uk)

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

**More contact info on p114**

# Newsdesk

Audio codec brouhaha; Latin America rejects proprietary software; hyperlinks are not BT's; Intel DRM; GNOME accessibility efforts recognised; Linuxbierwandering; Xbox Linux; and Unreal Tournament 2003.

## FREWARE LICENCE RUMBLE

# MP3 not free (except they are!)

**A** row has broken out over the imposition of fees on Free/open source MP3 players. The rumble started when a Slashdot poster apparently misread the license terms imposed by Thomson Multimedia and the Fraunhofer Institute, who jointly own the technology, and informed the world that the partnership would now be charging not just for MP3 encoders, but also for the decoding technology on every device. This, of course would be disastrous for projects such as XMMS, which has no traditional revenue stream to pay licensing costs. The fee payable was quoted at US\$0.75 per decoder and minimum royalties were set at US\$15,000 annually, though vendors could opt for a one-off fee of US\$60,000.

An understandable mistake as Thomson had recently updated the page (and the royalty structure), and removed a section which read: "No license fee is expected for desktop software MP3 decoders/players that are distributed free-of-charge via the Internet for personal use of end-users." This doesn't, claim Thomson, constitute a change in licence terms. A spokesman said use of MP3 patents in freely distributed players had been, and would continue to be, granted without incurring royalties.

An unexpected problem has arisen from the confusion – the question of whether Thomson's 'freely



**Free players don't have to pay royalty fees, but their imposition has an effect on GPL software.**

distributable' model is compatible with the GPL, as it restricts a developer's right to sell the software. A Thomson Multimedia spokesman told Newsforge that if their patent license wasn't compatible with the GPL, then it never had been. Red Hat were said to be removing MP3 ripping and playback apps from their latest release, in order to comply with Thomson's stance.

### Ogg gives thanks

Xiph.org, developers of the rival, and open source, software Ogg Vorbis were quick to jump on the confusion to promote their competing codec. In an

open letter to the company, Xiph CEO Emmett Plant thanked Thomson for giving them a fantastic PR opportunity.

"Thank you for presenting a reminder to people that when they choose a patented alternative over a free one, they will eventually have to pay in one way or another. It's been difficult to send this message all by ourselves; we're glad you've decided to step up to the plate and knock it out of the park." He also asked if they could organise a similar slip up next year when Xiph plan to release their MPEG4 competitor, *Theora*. "Please be sure to threaten those who challenge your



**The royalty payments demanded by Thomson Multimedia are only applicable to commercial players.**



**Epson, one of the more Linux-friendly printer/scanner companies, found itself on the wrong end of a GPL violation.**

license fees with lawsuits and draconian collections efforts. We officially support any action you take to drive home the 'MP3 costs money' message."

The GPL was also implicated in the removal of Epson's popular *ImageScan! For Linux* and *Photo Image Print System* from the their website. A statement said that after being contacted by the Free Software Foundation, download services had been temporarily closed. "We guarantee that no problem would happen with current versions in your system, however please update to its latest versions after they are released."



President Chaves takes exception to his fourth BSOD of the day.

## LATIN LINUX

# Welcome to Free Software land

The Venezuelan government has become the latest in a long line to mandate the use of open source software to get value for tax payers' money. The government announced a policy in September that requires all software developed by or for govt uses to be licensed under the GPL, unless no open source solution exists.

The policy is allied to a concerted campaign to reduce software piracy in the country and to bring Internet access to the masses through the use of Linux-based machines controlled by a community franchise.

The Peruvian government, who undertook to enact a similar law earlier this year, found itself pressured by the American ambassador – “acting independently in the interests of both America and Peru” – to abandon the policy and stick to ‘the standard’. Bill Gates also visited the country and made a donation of \$550,000 to the nation's school system.

Recently, governments of Germany, Finland, Pakistan, India and even the UK have begun to seriously consider the security and cost implications of proprietary software solutions.



## PREPOSTEROUS PATENT

# Hyperlink loss for BT

### British Telecom's suit against

Prodigy and its parent company SBC Communications, in which BT attempted to assert its ownership of hyperlinks, has collapsed. New York District Judge Colleen McMahon ruled that neither Prodigy, nor the Internet, infringed on the company's patent, cutting off any further attempt at litigation.

BT claimed a patent applied for in 1976, and granted in 1989, covered

linking between documents on computer systems and launched a test case against Prodigy in 2000, intending to collect royalty payments on the technology. A concerted campaign to show ‘prior art’ in the case uncovered a number of previous mentions of the idea, including Ted Nelson's 1963 book *Literary Machine*, where the word ‘hypertext’ was coined, and Douglas Englbart's 1968 text-linking demonstration video.

## NEWSBYTES

■ If you're a **Mozilla/Netscape** user desperate to get a Microsoft Passport, you'll be please to know the Redmond giant has removed its block on non-MS browsers at the Passport signup page. You'll no longer be told you have to ‘update to *Internet Explorer 6*’ in order to join.



■ Open Source Evangelist **Bruce Perens**, picked up his P45 from Hewlett Packard after two years gainful employment, allegedly for ‘Microsoft baiting’. He later said the split was ‘amicable’ but inevitable after the HP takeover of Compaq.

■ **Red Hat** have seen the desktop light and are readying a package for business users. The company say they are just responding to the needs of customers frustrated by the license regime change at Microsoft. By the same token, the UnitedLinux project – despite all the partners' name and/or owner changes – should have a beta release available as you read this. The final cost per server is expected to sit below \$1,000 including support.

■ **Intel** have announced that next year's line up of CPUs will feature security features that will protect against hackers and viruses and prevent piracy. The technology, called LaGrande, will provide hardware support for Microsoft's Palladium. Intel say they won't be including any digital rights management (DRM) capabilities, but acknowledge that the LaGrande/Palladium partnership will make it easier for copyright holders to stop, for instance, CD-RW from copying protected content.

■ **IBM** are hoping to buck the downturn in the tech sector with the opening a brand new, \$2.5 billion chip fabrication plant in New York's Hudson Valley. What makes this endeavour a little different though is that Big Blue are using Linux to control the nearly 2,000 1GHz CPUs and 600 Terabytes of storage which make up the IT infrastructure of the plant. IBM and its partners evaluated both Linux and Windows control solutions, and after three months of ‘flawless performance’ opted for the open source choice. A spokesman claimed the Windows-based system had “failed after six or seven days”.

■ Graphics giant **Adobe** has requested a judgement from the US District Court in California to resolve a dispute with International Typeface Corporation (who launched an action under the DMCA) about the embedding of fonts into PDF files.

## Jono Bacon

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



## COMMENT

# United we stand

“ Once upon a time, there was a little OS called Linux. He was growing up and getting big and strong, and everyone wanted to play with him, but few wanted to play the same games. After some time there were people who played similar games with our pal Linux, but they changed the rules just slightly. After some years a new kid on the block, called UnitedLinux, stepped in to organise some standard games for our veracious young Linux to play.

Does all this sound like a fairytale? Well, not if Connectiva, TurboLinux, Caldera (now SCO) and SuSE have anything to do with it. They have clubbed together to standardise a Linux distro core to challenge the mighty Red Hat.

Fine intentions I think, but likely? Possibly. We need to refocus on the Linux *community*. At heart, Linux is not vendor dependent – anyone and their dog Barney can create a Linux distro. With this independence also comes a lack of control of the community – both a good and bad thing. UnitedLinux want to formulate a standardised core for a distro and get companies on board to develop for it. No problem, but the real challenge is to get people to play ball with it.

Linux needs standardising in package formats, user interfaces and file system layouts. The UnitedLinux technical whitepaper states it will support and use the LSB and other standards to achieve their goal. This seems honourable enough, but the real challenge is in developing the brand, market and vendor reliability. Is this possible with four different vendors with their own products and agenda's? We will have to wait and see... ”

## ENABLING USERS

# Prestigious award for GNOME accessibility

**T**he GNOME Accessibility Project has been singled out for praise in the annual Helen Keller Awards, given by the American Foundation for the Blind.

The Achievement Award in Technology was officially presented to Sun Microsystems for their 'leadership in universal design' but the judges highlighted the GNOME Accessibility architecture which, they said, "raises the bar for the computing industry and dramatically expands the options available for technology consumers who are blind or visually impaired"

Sun's Patricia Sueltz will accept the award on behalf of all those who took part in the distributed development effort including engineers from Ximian, Red Hat, Mozilla.org, Codefactory and KDE-Accessibility. A number of the engineers that Sun Microsystems have donated to the GNOME project have been working on the GNOME Accessibility project

The GNOME Accessibility project was set up to provide the tools which would bring information technology to users regardless of disability, and notable projects include *Gnopernicus*, an extension to the GNOME2 desktop for blind or visually impaired users,



Sun's Patricia Sueltz will collect the award on behalf of all those who contributed to the GNOME Accessibility Project.

and the GNOME *Onscreen Keyboard* (GOK) which provides an input method for users with limited voluntary movement. These project allow user full access to a range of standard GTK+2 and Java applications.

On a related note, the GNOME project has just finished a study into human-computer interaction which, in turn, informs the project's *Human Interface Guide* which "tells you how to create applications that look right, behave properly, and fit into the GNOME user interface as a whole." You can find the guide at

<http://developer.gnome.org/projects/gup/hig/1.0/intro.html>


## Linux Web Watch/



Tuxpaint – in rapid development.



Tuxtype – friendly typing tutor.



Debian Jr – Free kids' distro.



cbeebies – Web-based fun.

## Doin' it for the kids

Keeping the kids occupied on a rainy Sunday afternoon...

There's plenty of child-friendly Open Source stuff out there if you're prepared to look. *Linux Format* is prepared to look.

Painting is one of the best outlets for children's creativity, and if you've ever seen a three-year-old struggling with the menu structure of *GIMP*, you'll know that a child-friendly paint app is essential. One such package is *Tuxpaint* ([www.newbreedsoftware.com/tuxpaint/](http://www.newbreedsoftware.com/tuxpaint/)) which is up to version 0.4.2 and is sponsored by the Tux4Kids project. The app uses big friendly icons, humorous noises, a great

stamping tool and a save/open system which circumvents the Linux file system nicely. It's very enjoyable to use. Tux4Kids is also home to *Tux Typing*, ([www.tux4kids.org/dm/tuxtype](http://www.tux4kids.org/dm/tuxtype)) which is an educational title designed to introduce children to the keyboard and improve their letter skills, and an experimental school-friendly distribution called Omega.

For dedicated Free Software fans, there is also an edition of Debian aimed at children. The Debian Jr. project ([www.debian.org/devel/debian-jr/](http://www.debian.org/devel/debian-jr/)) has set itself the goal of

producing software for children up to the age of seven, before moving on to the far more demanding eight to 12-year-old age group. As well as making the whole UI more child-friendly, the project hopes to include packages which promote computer use.

Linux For Kids ([www.linuxforkids.org](http://www.linuxforkids.org)) is another site dedicated to bringing the best open source software to the small people in your life, and is also a member of the Schoolforge ([www.schoolforge.net](http://www.schoolforge.net)) project. The website offers a selection of downloadable applications in

categories such as maths, spelling and science, plus section for parents and teachers as well as links to articles and features on using Linux as an education tool. There are also some games there.

Finally, the Granddaddy (or should that be Auntie?) of kids websites is the BBC's Cbeebies homepage ([www.bbc.co.uk/cbeebies/](http://www.bbc.co.uk/cbeebies/)) which, while relying heavily on Flash, works fine with *Mozilla* and *Konqueror*. The visual/aural navigation system makes it ideal for children just starting out on the path to reading and the range of distractions are sufficient for any rainy afternoon.

## NEWSBYTES

■ **IBM** and **Borland** have struck a deal which will see trial versions of *Kylix*, *Delphi Studio Architect* and *C++ Builder* with the latest *Personal and Developer Editions of DB2*. The reciprocal agreement will see *DB2* bundled with the development packages from Borland.

■ **Mozilla.org** plan to continue their hectic release schedule. *Moz1.2alpha* was released as we go to press with a full release expected in early November. New features include Type Ahead Find which should improve keyboard navigation and OS-X users get a 2000%(!) improvement in IMAP mail header downloads.

■ **Netscape** hit version 7 with one notable omission from the standard *Mozilla* build: no pop up killer.

■ The **GNU project** has announced its adoption of Freenode for its official IRC network. The Free Software Foundation said the decision was just a case of formalising the arrangement. The FSF said it made sense to align with a network that was committed to the same goals of openness, innovation and the principles of free software.

■ **User Mode Linux (UML)**, which is a 'safe way of running Linux version and processes', has been integrated into the 2.5.34 development kernel. UML provides a virtual machine so you could, say, test buggy software or experiment with different kernels without affecting the underlying system. See <http://user-mode-linux.sourceforge.net/>

■ **Apple** are rumoured to be maintaining a secret x86 port of OS-X called *Marklar*. The current version is supposed to be comparable to the recent Jaguar release and, according to sources, is a backup in case Apple's usually strained relationship with Motorola reaches breaking point.



■ Computer journalist **Robert X. Cringely** has launched an open source 'computervision' show which will attempt to meet the terms of the GNU GPL. The program, *Nerd TV*, will appear in five versions, visual editions tailored to nerds and suits, audio editions in both MP3 and Ogg and, crucially, a package containing all the raw footage, which can then be edited. This is 'the source code'. The project is being coordinated with IBM. Cringely goes into 'obsessive detail' about the technology behind it at [www.pbs.org/cringely/pulpit/pulpit20020912.html](http://www.pbs.org/cringely/pulpit/pulpit20020912.html).



Driven to drink by trying to configure *sendmail*...?

### POPULAR PASTIMES

## Beer, Linux, walking

The annual **Linux Beer Hike** came

to Ireland this year and, from the pictures, it looks like a wet time was had by all. The pics come courtesy of Kirk Bollinger who travelled from San Jose to enjoy his first LBH, but says he'll try to return next year. "I really enjoyed the event, Linux brings us together." Bollinger said that as a beer loving, Linux using hiker it was the perfect event for him. "I think the beer was the most popular aspect followed by hiking and Linux. But everyone had the opportunity to have their own ratio of each."

Next year's trip will take in the wilds



Yes, some geeks can get away from the keyboard for a few hours.

of Slovakia, meaning an even longer trip for Kirk: "Perhaps I'll take two or three weeks and tour the rest of Europe while I'm there!"

### SUBSIDISED LINUX HARDWARE

## SuSE for Xbox?

After just a few months, **Xbox** hackers have succeed in bringing Linux to a (modified) Xbox, and have posted a HOWTO on the Xbox Linux Project website. The system was tested with SuSE Linux 8 and, said Michael



*TuxRacer* only manages about one frame per second, but it works.

Steil, only involved having "to change two lines in the Linux kernel, disable one init script, install two drivers for audio and networking and use the XLP X Window configuration file."

The only snag so far is that SuSE must be 'cross-installed' i.e the Xbox hard disk is initially connected to a PC for the installation and then tailored to the Xbox afterwards. Still, it works and, as a demonstration, Steil said his guide had been written on the Xbox with *OpenOffice.org* – it's not exactly great for games though as *TuxRacer* runs at a lowly 1fps.

Incidentally, when Microsoft rolls out the Xbox Live project, it's claimed that modded Xbox hardware will be excluded from connecting.

### 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

## Bold new Linux vision

“ You could see this one coming when the Red Hat 8.0 beta process was strung along forever. The delay wasn't just to produce a bug-free \*.0 release. To hear the Linux rabble tell their take on the story, 'Red Hat' replaces 'Quisling' in their vocabulary for the direction in which RH are taking one of the leading Linux distros, and in the way two favourite desktops are being modified.

A number of classic apps have disappeared because they don't work well with the UTF-8 fonts or have been superseded by newer, better apps. Red Hat doggedly refuse to offer a choice of *ReiserFS* or *JFS* filesystems for the root partition – only *ext3* is available. As well, Red Hat removed any MP3 functionality and supporting libs because of possible patent issues.

Finally, they bowdlerise the KDE and GNOME desktops; it's difficult to tell them apart. Their menus use generic descriptions ('Web Browser', not *Mozilla*).

Red Hat is now focused on producing a distro that they can afford to support ("Using non-approved apps or filesystems? Oops!"); afford to sell (no patent and trademark violations); and market to the corporate culture.

Will that jacket-and-tie approach appeal to the typical Linux hacker? Of course not, but those doOdZ aren't paying RH's bills when installing downloads, or purchasing the odd boxed set. RH is now positioned for Corporate use.

Red Hat is doing right by their shareholders and should stick to their guns. The *hoi paloi* may examine any of the dozens of other distros, do it themselves, or stick with Red Hat – that's a choice they must make. ”

## GAMING SUPPORT

# Linux client version of Unreal Tournament 2K3

**G**ood news for game fans, the sequel to *Unreal Tournament*, *UT2003*, should be available for Linux, either as a binary download or, more significantly, included on the same CD as the Windows version.

Epic's Mark Rein said the Linux versions (both client and server) of the long-awaited game would be ready to go with the demo release. In a response to questions on the Infogames forum, Rein wrote. "If space permits then we hope to put the Linux binaries and an end-user installer on the CD. We have already approved all the packaging so it won't be mentioned regardless of whether it make it in or not. There will be a Linux README file that describes how to install the game on Linux and covers potential issues users are likely to see."

"It is absolutely crucial to have Linux support for the server and seeing as we found a very competent programmer (Ryan C. Gordon) to take on the client side job we decided



**Unreal Tournament 2003 should hit the streets for Windows and Linux at the same time.**

to go full steam ahead with the client as well." He said that users who asked why the company was 'wasting time' on Linux version didn't realise that about half the servers running *Unreal Tournament* were Linux based.

"I think that if we're taking advantage of an OS for the server then we should have a client implementation as well," he said.

Look out for a review in *Linux Format*, soon after launch date.

[www.unrealtournament.com](http://www.unrealtournament.com)

## David Cartwright

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



## COMMENT

## Thinking laterally

“I can't get a Linux driver for my ADSL connection”, cried a client, toying with the idea of sticking with Windows. I have to admit that the only real drawback with using Linux day-to-day is that when new hardware or services arrive on the market, nine times out of ten the installer CD has Windows, NetWare and DOS drivers but no Linux (with a few exceptions – well done, Intel, with your RAID card setup).

As someone who, until this time last year, spent four hours a day typing on my laptop on a train, I'm painfully aware of the delay from buying a new laptop to being able to run Linux (Xfree86 has to catch up with the video card innovations, and I'm never motivated to hack the configs by hand). Generally I've set up Windows, repartitioned the drives, and waited a few weeks for the drivers to appear. Annoying if you really need Linux urgently.

Being resourceful chaps, though, us Linux people need to think laterally sometimes. OK, you may not be able to get a driver for your WhizzBang OfficePro X887.3 printer, but hey, it's got PCL5 emulation so can't we just use a generic PCL5 driver on the PC to talk to it? And while the kernel of our favourite Linux distro doesn't have support for the SpeedTouch USB ADSL modem turned on by default, why not consider buying a router instead. Not only can we plug our Linux PC into it, but also the spouse's Windows machine in the back room.

So remember: just because Linux doesn't formally support the exact widget you want to plug into it, doesn't mean you can't use it. Think laterally and you may well find the answer. ”

## Embedded Linux News

- Intel have announced an XScale-based Digital Media Adaptor designed to distribute digital media between devices in the home. A number of vendors including Dell, Mitac, Legend and Gateway demonstrated versions of the device at a recent Intel Developer Forum. The device, which uses an optimised embedded Linux system, receives data – such as MP3, streaming audio or DVD – from a PC via a 802.11 wireless link and passes it onto audio or video hardware through standard cables.

<http://developer.intel.com/technology/digitalhome/>

Intel also joined forces with Philips to create the Pronto++ reference platform. Again, it's based on XScale and Linux, and

will be used to access interactive digital content.

- Another criticism has been addressed by Trolltech with the beta release of a Linux version of Sharp's Zaurus synchronisation software. Despite being a Linux-based device the Zaurus shipped with Windows-exclusive sync software.

- Sony are entering TiVo territory with its new Cocoon product line. As well as offering the standard PVR functions, the first units will ship with 160GB hard disks (capable of recording 100 hours of TV) and will also feature Internet access over a broadband connection. The devices will also record programs it thinks you'll

appreciate and – finally – you'll be able to program it via your mobile phone (for a fee of course). Building on the Vaio strategy, Sony will release the products in Japan first to refine features and create demand before going worldwide.

- Feabas, a professional training organisation specialising in real-time systems will be holding an embedded Linux seminar on October 31<sup>st</sup>. The seminar is targeting developers with experience of traditional development tools and assumes no prior Linux knowledge. The event will take place at the Madejski Stadium, Reading, Berkshire and tickets can be booked (£65) at [www.feabhas.com](http://www.feabhas.com)



# 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: [lxf.letters@futurenet.co.uk](mailto:lxf.letters@futurenet.co.uk)

## Euro Con

I have a keyboard with a Euro symbol, I've searched the 'Net, rummaged through Usenet and tried a variety of suggestions but so far failed to work out how to get that key to give me a € symbol! Let alone how to get it to display! How about an article on those funny keyboard keys and how to use them? I've sussed the windows key, and use that for all kinds of things that it wasn't intended for, but a HOWTO and 'what-it-means' guide would be greatly appreciated!

I use *Emacs*, *OpenOffice.org* and *blackbox* and I'm sure instructions for that would differ from someone using KDE or GNOME?

Robert, via email

Why would you want to devalue your £ key! Anyway, I suspect the answers you seek are in the recent Keyboard tutorials we did, the most recent of which was in issue 28. If you don't have that issue, check out our back issues page. Alternatively, check with the Linux Documentation Project's 'Keyboard HOWTO' at [www.tldp.org](http://www.tldp.org)

## Amiga lives

Just a quick thanx for the Amiga emulator feature. Glad you gave *Worker* a mention – an absolutely essential program for any Amiga/Linux user, and being continually developed and improved. I currently run Mandrake 8.2 on a rather underpowered subnotebook (I've never had a desktop pc and nobody ever told me you didn't put Linux on a laptop three years ago so I just did it, and it worked:-) ) – and *AmiWM* takes about a tenth of the time to load that KDE or GNOME do.

A comment on KDE3 (and GNOME2) – it can't be denied that it's becoming very 'Windows-like' – whether this is a good or bad thing, I'm not sure, but I guess it's what users have come to expect...

Jen, via email



## Not so Slack

Just a quick "Thank you" for including Slackware 8.1 – first time I've seen Slackware as a coverdisc for a "long" time.

I have used it for years, although one of my home boxes also runs Red Hat, as we have to run that one at work, so I'm using it as a platform to learn the way Red Hat does network configuration, amongst other things.

Paul Duncan, via email

That's OK. Some people seem to think we have something against slackware. Not so, we hadn't included it recently because we were waiting for a good new release.

## The big debate

Can you help stifle a debate here at work please? how are you officially supposed to pronounce "linux"?

It is "lin - ix" or "ly - nix"?

Stefan Farrelly, via email

The commonly accepted pronunciation is "lin-ux" ("lin" as in "lint", "ux" as in, er, um, "crux") as Linus Torvalds, being Finnish, doesn't pronounce his first name like the Peanuts character. However, in recent years he has said his first name is mispronounced so often he isn't that fussed about it anymore...

## Double vision

I am quite a Linux fan and enjoy reading various Linux publications including *Linux Format*. I was paging through the magazine and something caught my eye. On page 51 of your article on Internet Security I see you have a diagram "Fig. 4: Dial-up host acting as firewall and desktop/workstation", which features an illustration of a firewall. I am working on a firewall project at work, and we have a *Visio* diagram which has an identical object to your illustration. So my question is, is *Linux Format* using

Microsoft products to produce the magazine?

Michael Turner Senior IT Technician  
Incepta Group plc

Well, the original drawing was supplied to us by a freelancer, and I can't say for certain what software was used to create it. Perhaps it is a commonly available clip-art object? Most of the illustration work in the magazine is done on Apple Macs, using the usual tools, simply because that's the way all the magazines are produced here.

## Doomsayers

I am a particular fan of the *WhatOnEarth* articles as they bring some of the more obscure topics to our attention, especially recent articles on the Digital Millennium Copyright Act and, in October 2002 issue, Palladium and the Trusted Computing Platform Alliance. However, both these articles were particularly gloomy in their predictions of 'the end of Open Source Software' and I would like to dispute this scenario.

I can perfectly understand the large 'Western' media and software corporations (such as Time Warner, Disney, Microsoft or anybody) wanting to protect their licensing revenues, after all that is all they really exist to do – collect licensing revenues. In order to do this, they have to produce products such as films, music and computer programs, which they have to persuade people that they want to use. They are quite good at the persuasion, part through glossy marketing, and back this up with dirty (often legal) tricks to hold on to their hypnotised marketplace users: the result is the existence of very large and powerful corporations who are able to control so much of what we in Europe and America think, use and value.



Not just nostalgia – *AmiWM* is a lightweight, capable window manager.



So if people want to use Microsoft Office, watch Hollywood movies or listen to recorded copyright music, my guess is that they should pay the appropriate licensing fees – cracking software, installing on more computers than licensed, copying CDs and DVDs and downloading illegal MP3s is quite simply theft. The founders of ‘Free Software’ always made it quite clear that if you want commercial products you should pay for them, but they offered an alternative path through collaborative working, leading to public ownership of software.

The fact that you may have to have a GNU/Linux OS capable of running TCPA-enabled software if you want to access copyright controlled media does not really change the fundamental concept of GNU Licensing, as I see it. The end user will have to pay for a certificate to run the TCPA-enabled part with the Fritz chip on their motherboard and some process/organisation will have to be established to handle this licensing. I don't see that this spells the “end of Open Source Software”, any more than having to register and pay for an Internet domain name does or running any paid for piece of commercial software does.

There will still be the alternative of not running TCPA-enabled applications. We don't *have* to use Microsoft Office, *OpenOffice.org* is every bit as good, and the same will continue to be true for many other software apps – we will be able to choose whether we use Open Source or buy commercial products.

While I don't have the same in-depth knowledge that the authors of the *WhatOnEarth* articles about DMCA and TCPA have, and I appreciate their vigilance in tracking these topics, it seems to me to be counter-productive for *Linux Format* to be pronouncing “How Microsoft's Palladium would force open source off your computer” on the cover page, and in the article to state that “idealistic young programmers will be much less motivated to write free software”. Neither of these statements is in any way a fact.

Let's take a step back for a moment and look at the Open Source Software community. I >>

## ★ Letter of the month

This month's winner receives a copy of the **MySQL Reference Manual**

### Commercial I.T.

I am the sole I.T. resource for a small company with 2 factories. I have 29 years I.T. experience mainly in project management and business analysis – i.e. I have never programmed for a living. Our main corporate system, which runs ledgers, inventory, and sales order processing is a package solution using Progress on a SCO Unix platform. It is distributed over an Ethernet LAN to dumb terminals, either green screens or Windows PCs running terminal emulation through a serial port. The number of PCs has gradually increased to around 50 and they have been networked so as to share files and printers. Apart from a few specialised applications, the PCs use MS Office (all versions), Outlook Express, typically Netscape for browsing, RamBooster to contain memory leakage, and Sophos anti-virus. There is also a DEC processor running OpenVMS. That then is the environment.

The first point to make is how dramatic is the contrast between the Unix and MS systems. I spend an unacceptable disproportionate amount of my time supporting the MS systems: OS and apps. The blue screen/freeze problems have been well documented in many letters and articles. The recent changes in MS licensing only add to the frustration. Our Unix system just runs and runs, only needing an occasional database unload/reload to restore access efficiencies.

About 18 months ago I read an article on a Sun installation that was using StarOffice 5.2. I downloaded the Windows version and began trialing. The main problem was that all modules were loaded into memory at launch time. However, its ability to read all MS Office versions was of particular interest, due to the various email file attachments we receive, which not every user can open depending on their Office version. At home my son and I became increasingly interested in

the rise of Linux, hence the subscription to your magazine.

Despite my non-technical background, I have been able to install the SuSE email package. I have also configured SuSE 7.2, Samba, Wine, and all sorts of other useful packages on some test bed PCs. I have been very impressed with the quality of documentation, HOWTOs, error messages, and comments in e.g. parameter files. More recently I have installed the Windows and Linux versions of OpenOffice.org. This has led me to the point where I am seriously considering Linux and OpenOffice.org for our standard desktops.

I would like to add to your recent review of OpenOffice.org. The main improvement is that only the chosen module is loaded, which greatly reduces launch time. The really smart touch is that there is a parameter in set up, which determines how long the module remains in memory after it is closed. The default is 10 minutes; the effect is that if the module is re-launched within this time, loading is virtually instantaneous. The other feature of note is the ease with which it can be set up under a client/server model. MS Office versions including XP can be opened

One of my colleagues has developed a number of MS Access databases, which would seem to block the introduction of Linux in some areas. However, I am in the process of researching MySQL. I would recommend looking at the HOWTO on [www.mysql.com](http://www.mysql.com), which discusses using OpenOffice.org to connect to MySQL. There is also information on linking MySQL to Access, which I have yet to test. I am intrigued by the possibilities of running MySQL on a SCO Unix platform alongside our corporate system, but that's another story.

In conclusion, I would like to put forward the following argument. If your organisation's core systems are *not* built around MS products, then any perceived



risks of moving to Linux are greatly reduced and service to users is improved. A transition path is readily available by initially using the Windows versions of OpenOffice.org and where necessary MySQL. There are superior alternatives to Internet Explorer (e.g. Opera), and Outlook Express, with all their irritating default security holes. Adobe Acrobat Reader is available on both platforms. Later on Linux can be introduced when users have gained confidence in their new application software. Although I acknowledge that setting up a PC under Linux can be a long slog (and I should know!), I believe that it is this keyboard based input, which gives Linux its stability. Once it's right, it stays right.

Thus there is already a business case for Linux, which can be made to your Board, although I suspect that the philosophy will have to be explained first. I've already been asked what happens if ‘Linux goes bust’. The benefits for me in my company are much improved stability for users, reduced support and therefore more time to implement new applications, reduced licensing costs and the benefits of everyone running the same software versions. The appearance of IBM and HP in the Linux arena makes its continued rise inevitable.

**Jim Godwin** I.T. Manager

Many thanks for this interesting and thought provoking letter. Thanks for the OpenOffice.org tips too. We are always really interested to hear what other people think of software we have reviewed in the mag, so please keep emailing us.

I'm glad to hear your Linux experience has been a positive one, and as a special prize, this month we'll give you the recently released MySQL Reference Manual, published by O'Reilly (ISBN 0-596-00265-3) MRP £28.50 – I hope you can put it to good use.

« see China, Russia, several Eastern European countries, India and lots of other countries all around the world, as well as Western Europe and America, all contributing to Open Source Software. The populations of China, Russia, India and the rest, not forgetting Africa, are vast and I am sure that they have absolutely no interest in paying Microsoft for licences for their software. Nor do I suppose that they have that much interest in the media products which DMCA and TCPA are designed to protect. So I think it safe to conclude the Open Source Software has a very assured future and perhaps we should be focusing on this instead of falling into the "End is Nigh" trap.

Tony Austin, via email

On the DMCA the question is not about preventing "theft" of "intellectual property," but the mechanisms being used to prevent people using such info fairly. E.g., preventing the development of a legal DVD movie player for Linux that can read encoded disks. The DMCA can also be used to enforce regionalisation and other restrictive licensing terms. If you buy a DVD in

## "This issue is not really about theft it's about businesses using the law in any way they can to maximise profit"

Europe, why should you have to buy another one if you move to the USA? If I buy an ebook, why shouldn't I be able to pipe that text to a screen reader and have it read aloud? The DMCA in the US also imposes ridiculous penalties on infringers. The issue isn't really about theft, it's about business using the law to maximise their profits.

With regards to the TCPA/Palladium, from what you say I think you have

missed the part about an OS using TCPA needing to be licensed itself. That means effectively, that you will no longer be able to compile a kernel and possible other key parts of the OS, and still run TCPA without relicensing. As you say, there are places where people could care less if they could view a film on their PC with Linux, but obviously I think we are right to be concerned about the erosion of users' freedom.

The underlying point about both of these pieces is that the aims of the DMCA and the Palladium are not about enabling the user, protecting them from viruses or anything other than companies trying to get you to pay them money as often as possible. There isn't anything really wrong with that, as long as there is choice, but by locking out open source, monopolists will have everything their own way.

We aren't trying to depress people, merely make them aware of what is going on in the world. As is often said, the price for freedom is eternal vigilance.

### Top Security

The article on Security was very good, but you have missed out

one major area of security – physical security.

I can have the most secure software running on my machine (and I think I have), but if I go away for the weekend anybody can put a Red Hat rescue CD into the drive, reboot, and read all my files. And that is something that I don't know how to protect against.

Yes, I could remove CD-booting



The price of freedom is eternal vigilance...

and password protect the BIOS – but remove the battery and you have access in 24 hours, or they could put the drive in another machine. I could get a removable drive, but my machine is also a gateway so it needs to be up 24/7, and it could be stolen anyway.

I guess I'm just going to have to install CCTV linked to an alarm which will ring my mobile – but if I can afford that, I might as well get Windows XP with its encrypted NTFS. At least that way I get a coffee break every time it crashes!  
Richard Kirkcaldy Bradford, West-Yorkshire

Thanks for your comments. Of course, physical security is something that's very easy to overlook. You are right that given physical access anyone can boot your system and do unspeakable things to it. The best thing to do, if this is a real concern, is to use an encrypted filesystem. Linux can mount encrypted filesystems – we covered this back in issue 21, and the article is online in the archive section of our website if you missed it.

Using encrypted filesystems an intruder may be able to boot up your

box, but won't be able to mount any of the partitions without the relevant passwords. This has been possible for some time, and the joy of Linux is that you can use any number of encryption protocols (you might for example want to use really strong encryption, or you might choose a slightly less severe version which give you the option of a shorter passphrase).

### Desktop debate


I write further to a recent letter written by Stephen Hill (issue 31, Sept 2002) and the corresponding editors reply.

I'd like to pick up on the issue of the desktop being conquered, which the editor's reply disputes. In my opinion the editor couldn't be more wrong. One only has to look back at history at OS/2. OS/2 started initially as a joint IBM / Microsoft venture. In his book *The Road Ahead*, Bill Gates gives reasons for the parting of company which includes Microsoft's desire to pursue the 'family strategy'. At the time I think it was fairly well accepted that OS/2 represented a superior operating system over

# Helpdex

BY SHANE COLLINGE

shane\_collinge@yahoo.com





### Windows 3.1.

Nevertheless workers were buying home computers with Windows 3.1, and later Windows 95, and then requesting that operating system on their work PC.

The editor also mentioned that most boxed distributions include virtually every application the average consumer could want. Perhaps a better wording would have been every "type" of application. Speaking for myself I purchased SuSE8, which I installed with ease. However the first thing I wanted to get working was my BT ADSL, which proved a major headache. Going to Linux on the desktop currently presents too many changes in one go. Not only does the operating system change, but none of my games work nor my favourite applications (like Microsoft Money). It seems to be an "all or nothing" jump, and I guess most people don't take the risk, or keep Linux as a second operating system to tinker with (I am in this boat).

Please don't take me for a Microsoft bigot. I'm not. But credit where credit is due. Microsoft understood what was required to get on every desk.

Jamie Cameron *via email*

Thanks for your opinions on the history of the desktop. Perhaps you are right, but I'm not sure the revisionist history of computing you mentioned is a reliable source! What do other readers think about this?

Perhaps "every type of application" would have made it clearer. Very few 'identical' pieces of software are available for Linux, certainly for MS applications. However, much software does run under Wine.

### Where to start

I just received your magazine a week ago and have now read

almost all of it. This is the fifth issue since I took a subscription. At first it was my intention to react to the letter of Mr Dan Franks. Like him I would be very interested in seeing a series on how to set up a Linux box from scratch. From reading the letter it was my impression that Mr Franks wants you to set up the software starting from the Kernel and building from there. Not so much the hardware. It is my experience that most new hardware works with Linux these days. I too would be interested in learning how to start with a blank system and then put on a Kernel and one software package after another the way I like them.

I would like to learn what all these hundreds of config file are for, why they are need and how to edit them. Where to find the look for the right libraries and how to install them. How to install software from source, RPMs (I still have not managed that. It may be straightforward for most of you, but so far I have been unable to add any new software to my system (Mandrake 8.2)). How to recompile the kernel and what to look for. Which modules to choose (not like you article in the last issue, just a list of modules with some explanation that makes no sense to a layman like myself. What are we looking for when we choose modules?; which do we want? and what to do when things go wrong).

Distros may be easy to install these days, but like Windows you have very little control over what ends up on your computer. five shells, five or more GUIs, ... who needs it?

This may seem a bit confusing but that is the way I feel about Linux. I know it cannot be too difficult to get - I just have not

found the key yet, and I am hoping that you might help me find it. As I said before, your magazine is a great read and does contain a lot of interesting material if only I could get some of it installed on my machine.

Have nice day,

Geert Anthonis, *Gaoxiang, Taiwan*

Thanks for your interesting letter. It can be difficult to know where to start. To gain a full appreciation of the Linux OS, projects like Linux from Scratch are very good, but obviously they assume a fair knowledge of Linux to begin with. It's rather like learning how to fly a plane by learning first how to build one.

One thing I would mention is using the 'Expert' install modes with Linux distros. With something like Mandrake, this doesn't actually require much more Linux knowledge, but does give you a lot more choice about what software gets installed.

I hope we will soon be able to announce a new series of tutorials aimed at relative newcomers and covering topics that seem to cause the most problems - installing software, kernel compiling, configuration and more. [LXF](mailto: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 question - direct those to our Q&A pages!
- Random abuse
- Nonsense rants
- 200 pages of meandering diatribe

#### WRITE TO US AT:

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



[www.linuxformat.co.uk](http://www.linuxformat.co.uk)



**Reliable, Fast,  
Effective  
Web Site Hosting**

# £59

**+Vat Per Year**

Includes the following  
(no hidden costs)

100MB WEB SPACE  
BEYOND BROADBAND SPEED  
99.99% UPTIME  
INDEPENDENTLY TESTED

Your Own Domain Name  
[www.yourname.co.uk](http://www.yourname.co.uk)  
(other domains also available)

10 x POP3 email accounts  
with forwarding  
and autoresponders

Windows and  
Apple Mac Friendly

Macromedia FLASH Compatible

 PHP4 Scripting  
Enabled

Example Form Processing Scripts

24Hr FTP Access  
Online Control Panel  
Graphical Site Visitor Statistics  
Free Technical Support

No Banner Ads -  
unless you put them there!  
Your own virtual server  
under your control

Quality UNIX Servers housed in  
a State of the Art UK Datacentre

Other Options Available:  
MySQL Database, SSH, Perl,  
FrontPage Extensions and more

Find Out More and  
Order Online at

**[www.hotfibre.net](http://www.hotfibre.net)**

 **hotfibre**<sup>™</sup>  
web site hosting

# Reviews >>>

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

## LXF verdict explained

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

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

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

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

**Value for money:** Does it have a competitive price?

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.



## THIS MONTH...

### Kylix 3

Borland bring C++ capabilities to their advanced RAD, and allow mixing of C++ and Object Pascal code, with the full CLX library **p18**

### Homebase

A Mozilla-based desktop distro with a friendly interface and an online support service that leaves .Mac behind in the value-for-money stakes **p22**

### Elx Linux >>

India's fledgling desktop distro makes a bid for the home Windows user market. But what are those server daemons doing there? **p26**

### Kohan >>

Thanks to transgaming and WineX we have another Windows game running on Linux — but is this RTS all it's cracked up to be? **p28**

### Books

As the nights draw in, we open *J2EE Professional Projects*; *Network Security with OpenSSL* and *Linux in Small Business: A Practical User's Guide* **p32**



## COMING UP SOON...

### Smoothwall

Corporate version of the popular easy-to-configure Linux firewall.

### Snapgear PRO+

We like playing with new hardware, so expect a review of this professional router soon.

### Netraverse Server

Can the server edition of the popular Win4Lin replace a network of Windows boxes?

### The big three

Red Hat, SuSE and Mandrake are all on the verge of new releases. Read about them all here.

### Lycoris Desktop/LX

Ease of use and fluffy clouds. We check out the latest update to the distro previous known as Redmond Linux.

### Xandros Linux

The project to resurrect Corel Linux nears completion.

## RAD DEVELOPMENT TOOL

# Kylix 3

Borland Kylix now reaches out to C++ developers as well as Pascal coders. **Brian Long** looks in the box to see what this latest release has to offer.

**Kylix's lengthy heritage makes it so wide-ranging and full-featured that there is little to accurately contrast it with on the Linux platform.**

- **DEVELOPERS** Borland
- **PRICE** See the box, *Kylix 3 Editions*
- **WEB** [www.borland.co.uk](http://www.borland.co.uk)

**K**ylix has had a rapid growth cycle, after a little hesitation getting started. *Kylix 1* was released in February 2001 as Borland's first attempt at a Linux-based RAD tool. It offered facilities for building GUI applications for X windows, console applications and also Web server applications (CGI applications or *Apache* modules). All this was supported through the Object

Pascal language and a rich supporting library called *CLX*.

Nine months later in November *Kylix 2* was released with a good remit of Enterprise-level support including the ability to build Web Services and CORBA applications.

In July of this year *Kylix 3* was announced. So what's new?

## Overview

The main thrust of *Kylix 3* is to finally fulfil Borland's aim of getting state of the art RAD tools on the Linux platform for both Pascal and C++ developers.

Yes that's right, *Kylix* is no longer restricted to those who want to dabble with Object Pascal (or the *Delphi* language as it has now been renamed). Naturally it still completely supports the *Delphi* language as it



**Kylix 3 brings RAD tools to C++ and Pascal developers.**

has done since its inception, but *Kylix* now contains two identical looking integrated development environments (IDEs).

One IDE is for doing *Delphi* language development (see **figure 1**) and the other is for C++ based development (**figure 2**). *Kylix 3* is effectively a Linux version of both *Delphi 6* (with some *Delphi 7* features) and *C++Builder 6*, Borland's RAD tools for the Microsoft Windows platform.

Both IDEs offer much the same features (see the *What Is Kylix?* box). They allow you to develop applications using the full *CLX* library but you have a choice of which language to develop applications in, *Delphi* or C++. It's even possible to build applications with combinations of *Delphi* and C++ source modules.

## Twin IDEs

You can see the striking resemblance between the two IDEs in the screenshots. Each IDE has the **ImageView** sample project loaded (it's available as both a *Delphi* and a C++ project) and you should be able to see the only real difference is the syntax in the code, as visible in the code editor in both cases.

Installation is straightforward, thanks to the continued use of Loki Software's installation kit, but a note of caution should be injected here. Make sure you have plenty of disk space; *Kylix 2* was a large product but *Kylix 3* is much bigger. Depending which version you install it will take between




Figure 1: The *Kylix 3 Delphi* IDE – Object Pascal at the click of a mouse button.




Figure 2: Is this the *Delphi* IDE again? No, it's the C++ IDE.

250 and 500MB for a full installation.

After installation you should have two key menu items installed on the desktop: *Kylix 3 (Delphi IDE)* and *Kylix 3 (C++ IDE)* as shown in **figure 3**. You can use these menus or launch either IDE from a terminal window with the `startdelphi` or `startcb` command.

## C++ support

*Kylix 3* represents the debut of Borland's C++ compiler for Linux. As with the *Delphi* language support you can write C++ code in the C++ IDE, compile it, link it and then run it and debug it. However there is nothing stopping die-hards still using the command line. Both *Kylix* languages offer command line tools for compiling and linking.

The C++ compiler is a fully ANSI/ISO compliant, optimising compiler and supports pre-compiled headers to increase compilation speeds. It comes with the *STLPort 4.5 Standard Template Library* as well as the *CLX* library.




Figure 3: Desktop menu items for the two IDEs in *Kylix 3*.

If you already have various tools to help your build process the C++ IDE lets you associate these tools with appropriate file types, for example you might want to use *Bison* or *Yacc* to compile .y files. You can also export C++ projects to makefiles compatible with the GNU *Make* utility. Clearly attention has been paid to ensuring that established Linux C++ developers can continue to use and leverage their existing skills and tools.

## What else?

To be fair, the main point of *Kylix 3's* release is to bring their C++ compiler to the Linux market. C++ is the programming language of choice for most Linux developers. Changes to the *CLX* component library and *Delphi* language RTL are minor, and other changes are varied but not really major in comparison to what was offered in *version 2*.

The **Object TreeView** is a new IDE window that helps whilst developing forms, frames and data modules by showing the logical relationships between visual and non-visual components in a tree view (see **figure 4**). Selecting a component tree node automatically selects them on the Form Designer and exposes their properties on the Object Inspector, whilst drag and drop can be used to change component relationships. You can also access component editors for the components by right-clicking their tree nodes.

The **dbExpress** database framework now supports *Oracle 9i* and *Informix SE* and the Indy

## What Is Kylix? Pascal metamorphosed

In a nutshell, *Kylix* is a RAD (Rapid Application Development) tool used to build X windows applications, Linux console applications, Web server applications and shared objects. All these have access to databases (*DB2*, *Informix*, *InterBase*, *MySQL*, *Oracle* or *PostgreSQL*) through **dbExpress** components.

Applications can be developed either in C++ or in the *Delphi* language (a language that's evolved over the last two decades from Pascal, which many programmers learned in college), each having their own IDE (Integrated Development Environment). Both C++ and *Delphi* are object-oriented, third-generation languages and both generate high-performance, optimised compiled code.

Applications' user interfaces can be developed quickly using the tools in the IDE, which include:

- The **Component Palette**, offering between 75 and 190 *CLX* components (dependent on version) to build your applications with. Components are the basic building blocks used to develop *CLX* applications; they are objects that represent, buttons, listboxes, database tables and so on.
- The **Form Designer**, which gives you a visual representation of a window or dialog in your application. You place components on the **Form Designer** and customise their properties.
- The **Object Inspector** to view all the properties for the selected component(s) on the **Form Designer** and to customise them as you need. It also lists all the events and will manufacture event handlers for them as required.

The IDE also offers other tools that you would expect to be in a professional

development toolkit:

- A very capable source code editor, which supports syntax highlighting for various languages, keystroke macro recording, multiple keyboard mappings, but not source code folding.
- A host of wizards to simplify the creation of GUI and console applications, shared objects, various types of **Web Server** applications, **Web Services** and **Web Services** consumers, *CLX* components, *CORBA* client and server applications and so on.
- A powerful debugger that even supports CPU level debugging for those who need it, with a disassembly window showing memory dump, CPU registers and flags, and the program stack.

The IDEs are rich with functionality and options, but command line die-hards

are not forgotten with *Kylix*. Every step of the compile and link process can be performed from a Linux prompt with the command line compilers.

Both languages (*Delphi* and C++) have their own RTL (Run-Time Library routines) but share the *CLX* library, which consists of four main areas:

- **BaseCLX** is a collection of low-level classes and routines available for all *CLX* applications.
- **DataCLX** is the set of **dbExpress** components that provide data access to applications.
- **NetCLX** is support for building Web server applications, either as CGI applications or *Apache* shared modules.
- **VisualCLX** is a set of GUI controls and graphics classes. **VisualCLX** applications have an inherent dependency on the *Qt* widget library from TrollTech, which is redistributable for **VisualCLX** applications.

components (which were also available in *version 1* and *2*) have been upgraded to *version 9*, and now include additional intercept components and I/O handler components.

## Web server apps

Web server applications (either *Apache* shared modules or CGI applications) can still be written in various ways, but debugging them is made much easier with the **Web App Debugger** (see **figure 5**). This

functions as a Web server to facilitate direct source debugging without the need for a real Web server to be installed.

*Kylix 2* introduced the ability to create new Web services and also write applications that make use of Web services (existing ones or ones that you write yourself), but this was only available in the *Enterprise* version. With *Kylix 3* the *Professional Edition* now includes support for making use of Web services, with native SOAP bindings for compile-time syntax and

## Kylix 3 Editions

*Kylix 3* has been released in three versions, just as with *version 2*, and they are priced the same as well:

- \* *Enterprise Edition* (£1399 new, £669 as an upgrade) aimed at corporate enterprise developers
- \* *Professional Edition* (£169 new, £89 upgrade) aimed at developers of commercial applications, either Web-based, X windows-based or simply command-line based, with access to various database formats if needed.
- \* *Open Edition* (free download) aimed at developers of GPL applications

If you compare these RRP's with the combined RRP's of *Delphi 7* and *C++Builder 6* on the Windows platform (which give you much the same functionality as in *Kylix 3*) you'll find yourself getting a considerably better deal on Linux. The *Personal Editions* (similar to *Kylix Open Edition*) would be £128, the *Professional Editions* £1368 (£548 upgrade) and the *Enterprise Editions* £3998 (£2398 upgrade). Borland are clearly making some concessions towards the lower price expectations of Linux consumers.

type checking (most easily done with a wizard in the IDE as shown in **figure 6**). To write Web services you still need the *Enterprise Edition*.

There's no doubt about it, this is a whopping product. C++ developers should look very carefully at this




Figure 4: **The new Object TreeView.**




Figure 6: **Kylix for C++ importing a Web service written in Kylix for Delphi.**




Figure 5: **The Web App Debugger listing out requests and responses.**

product with an open mind and, at the very least, either download the *Open Edition* or a trial version of the *Enterprise Edition* and check it out.

## Upgrade

You should like what you find. Granted there are a variety of free IDEs around for C++ development but this product has many years of history behind it; it's a mature, impressively rich and wide-ranging development system.

For users of *Kylix 2* who aren't interested in C++ there are various improvements to the IDE and to *CLX* to consider, but not that many of them. If you use *Kylix 2 Professional* then the upgrade is well worth it. If you use *Kylix 2 Enterprise* there are less compelling reasons to upgrade but you should certainly look at a trial version and see whether you like what's been added.

You should also be aware that if you are not interested in C++ you can get the *Delphi* version of *Kylix 3* free in the *Delphi 7* box (*Delphi 7 Studio Professional* comes with *Kylix 3 Professional* and *Delphi 7 Studio Enterprise* is accompanied by *Kylix 3 Enterprise*). That gives you Linux and

Windows development capabilities (as well as for .NET, since a preview version of *Delphi* for .NET is also supplied) in one product purchase.

Most newcomers to *Kylix* will be looking either at the *Open Edition* or the *Professional Edition* (the *Enterprise Edition* is more for corporate developers). There's simply no reason not to download the *Open Edition* if you have access to a high speed Internet connection and wish to create GPL applications [or try it out from our coverdisc, see page 100 for full details]. For commercial software the *Professional Edition* is hardly an expensive purchase and you get a lot for the money, although the support for C++ and using Web services is really all that sets it apart from *Kylix 2 Professional*.

## Mature product

There aren't many real gripes with *Kylix*, mainly due to its mature and long heritage on the Windows platform. The IDEs are still built mostly from the *Delphi/C++Builder* Windows code base using the *Wine* library. Yes they are real Linux ELF applications, but they still don't quite look or feel

## About Brian Long

Brian Long is a UK-based freelance trainer and problem solver for Borland's *Kylix*, *Delphi* and *C++Builder* packages. His website is at [www.blong.com](http://www.blong.com) and he can be emailed at [brian@blong.com](mailto:brian@blong.com).

like it. Before *Kylix 1* was released Borland assured the public that a proper native *CLX* IDE would be forthcoming with future releases, but so far they seem to be sticking with the *Wine* version.

There is still no help with application deployment (other than a text file describing what files you'll need to ship). Identifying all the files to deploy can be tricky, for example *VisualCLX* database applications require the application, the *Qt* library and **dbExpress** database driver to be deployed in addition to the database itself. If you build your application with runtime packages, things are rather more involved.

These downsides apart, *Kylix 3* gets the thumbs up from us. **LXF**

## LINUX Format VERDICT

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

Borland have done a good job in providing wide-ranging application dev capabilities to C++ and Pascal coders.

LINUX Format **RATING**  
**8/10**

SUPPORTED DESKTOP DISTRO

# Homebase

**Andy Channelle** finds a whole new way of working in one of three new distros aimed at those for whom computing is not a pastime.

**Browser-based desktop distro with .Mac-style online services, aimed at the Windows market.**

- **DEVELOPER:** OEone Corporation
- **PRICE:** Free download/\$19.95 per year with services
- **WEBSITE:** www.oeone.com


For most people a computer is a tool, something for emailing friends, writing job applications and resignations, surfing the Web and playing a few games. The ambitious may even manage their household budgets, or their CD library via the magic of MP3. For users like this, a thousand quid spent on a monster PC is going to be a waste, and the extensive facilities that your average Linux distro offers will be similarly left by the wayside; who wants to partition a hard drive when you could be playing *Minesweeper* or searching for pictures of Buffy?

This is where OEone's Homebase comes in. Like Lindows, and Desktop/LX from Lycoris, this is a Linux distro aimed at the average home user: unlike them, it hides everything but the essential applications and services behind a pretty and intuitive *Mozilla*-based interface – in either *Intrigue* (the blue one in most screenshots) or *Flat* (which is smart, business-like, and is now my default).

The core features of the distro cover web browsing, email, PIM facilities, word processing and instant messaging, but OEone is also offering a range of paid for services which – if you need access to your documents from more than one computer – will soon become essential. You can choose to download a free version which sits atop a standard Red Hat installation, or a standalone version (tested here) which installs from a single CD and comes complete with one year's access to the online tools, *Homebase Anywhere*. At \$29.95 this is just about a quarter of what you'd pay for similar services from Apple's new .MAC product. For downloaders, access to Homebase Anywhere is available for \$19.95 per year, and there's a free 60 day trial.

## Think really different

The first thing to note (no need to mention the installation, it's a no brainer) is that Homebase offers a completely different, but very consistent way of working within the core applications. Forget drop down menus, desktops and icons, windows, open/close/print dialogs and other common computer methodologies, this really is different. All of the navigation tools run along the bottom of the screen (or bottom and right if you're using 'Flat') and the rest of the screen is available for rendering content. As well as



Using *Homebase Anywhere*, your setting and documents are available from any 'Net connected computer.

simplifying use, this also makes the screen seem a lot bigger.


On first starting up you're asked to log in (this is a user account, not root) and go through the process of getting online, registering for OEone's services if you plan on using them, and configuring your email accounts. So far so normal. Next you're presented with the Homebase UI and your 'My Page' which can be configured to show your local weather, news from OEone (something like Yahoo! or BBC news would be a nice addition), any email waiting in your inbox and calendar events for the next five days (once you've entered them). There's also a

useful notepad for jotting down numbers, ideas, whatever. This is home.

The first thing you'll want to do is hit the 'Browser' button and explore the 'Net. The browser has OEone's front page set as home which is useful for the first few uses, but it can easily be changed in the settings page. The position of the address bar and back/forward buttons takes a little getting used to, but pretty soon the scheme makes complete sense. Bar the UI changes, browsing is the same as *Mozilla*, though most of the configuration options are absent. Tabbed browsing is here though (hit **Ctrl-T** to open a new tab) which makes up for the lack of 'windows' and can be set up in the settings page. OEone have made good use of *Mozilla*'s sidebar for things like bookmark organisation (with the excellent graphical preview), searching and 'history' browsing, and the sidebar idea even finds its way into the non-*Mozilla* portions of the distro. The first time you encounter a Flash movie or *Realplayer* content, Homebase simply downloads the appropriate player, takes a few seconds to restart the *Mozilla* engine and pushes you back to the original page. It is even easier than



The two schemes, *Intrigue* and *Flat* (above), give you lots of space for web content in the main window.



The calendar has a range of viewing option, but the sidebar always shows a list of upcoming events.

Mozilla's standard plug-in upgrade system.

So, after browsing the Web, you'll want to check your mail. Again it's all based on the standard *Mozilla* mail client which is functional, easy to configure and a breeze to use. No problems. It just works.

The third main use of home PCs after web and email access is word processing, which is covered in Homebase by a highly integrated edition of *AbiWord* (though you wouldn't know it from looks alone). *AbiWord* is ideal if your needs don't go too far beyond basic word processing as it has the usual formatting options, spell checker, image import and good integration with the email client, but doesn't get cluttered with options you may not need. The new way of working takes some getting used to – for instance, there's no need to save your document, it's saved as you go along; there's no 'open' dialog, you just select your document from the side bar; there's no 'print' option, you just hit <Printer> on the bottom toolbar. If you find yourself hitting *Abi*'s limits, there's also a full edition of *OpenOffice.org* included in the installation, just go to the applications page and select Linux Applications (this also gives you access to *The GIMP* and *Gnumeric*).

## Get organised

Next up is a brace on personal information manager (PIM) tools. The Calendar is based, as you'd expect, on OEone's very own *Mozilla* extension which will soon be part of the basic *Moz* set up. It's clean, offers a range of views (day, week, month) and has a fairly comprehensive set of reminder options, though it can't email alerts to you which is a shame. In addition to the graphical view which takes up most of the screen, the side bar has a chronological list of all events in the calendar (letting you see ahead even when concentrating on the minutiae of your daily life), and a useful search tool. The address book is simply the basic *Mozilla/Netscape* application dolled up in Homebase colours. Both applications work well and have the features the average user needs, they can also sync up with your Palm-based PDA (though it wouldn't talk to my EPOC-based one). However, these applications really come into their own once you sign up for the *Homebase Anywhere* service.

The screenshot shows the LinuxFormat.co.uk website interface. At the top, there's a navigation bar with 'Home', 'Topics', 'Downloads', 'Your Account', 'Submit News', and 'Top 10'. A prominent banner on the right says 'subscribe from just £27.45'. The main content area features several articles: 'Caldera name change' with a quote from the CEO, 'WeWantLinux.org, a Linux survey site' with a quote from the site's creators, and 'Linux takes Hollywood' with a special report on the OS phenomenon in the film industry. There's also a 'Latest Issue' section for Issue 31 - Out 14/8/02! and a 'Login' section for users. The website is displayed in a browser window with a taskbar at the bottom.

Homebase is built around *Mozilla*, so the browser works exactly as you'd expect.

Using Homebase without signing up to the *Anywhere* service is like buying a Lotus and only using the first three gears, and if you routinely use more than one PC it rapidly becomes a godsend. Once you've signed up, simply hitting the 'Synchronise' button will (if you're connected to the 'Net') send all your settings, events, contacts and, if selected, the contents of your Documents folder to OEone's secure servers. This is great for back ups, but also allows you to access your work, schedule, addresses and photos wherever you are, simply visit [www.oeone.com](http://www.oeone.com) on any PC, login and it's like having your PC in front of you. You can view or update your calendar and address book, edit documents and upload images while on the road, come home, hit the sync button and have your local copies automatically updated. It feels almost like magic and, again, it just works. Homebase promote this service as an alternative to taking a laptop away and, in many cases, it makes a lot of sense.

## Conclusion

I've not had space to mention the instant messaging (it works), the media player (ditto, though it's rather slow at ripping CDs), The games (*Othello*, *Minesweeper*, et al.), docs (I didn't

really need it), or the *File Manager* (very basic, but prevents you breaking your system), but all of these things are usually just a few clicks away from the front page. So who is this aimed at? Well, if you were considering this or Slackware, it's not for you. If, on the other hand, you've just spent £300 on a no-frills PC (or picked up a second hand bargain) and don't want to spend another hundred on an OS, it is ideal.

You get access to an uncomplicated, stable and standards compliant range of apps with very little in the way of bloat. The online services (undercutting Apple's offering by three-quarters) are a bargain and have so far worked faultlessly and would be a worthwhile purchase even if you weren't keen on the Homebase way of working. *Homebase Anywhere* is also ideal for students or home worker, who will frequently need to access documents from various locations (it's probably cheaper than buying 100MB of floppy disk space). On the two occasions I've managed to crash the system, it was only *Mozilla* that fell over, requiring a few seconds wait while the UI was loaded again. I didn't lose any work.

Homebase's biggest strength is also its weakness: it is very basic. If a feature you need isn't available it's quite a job to add it. For instance, I

wanted to download images from my camera but this involved editing *fstab* and then manually mounting the device through a terminal – not exactly beginner level. OEone told us this much needed facility is in the pipeline. Adding applications is also a non-trivial task, and if you're part of the target audience for this it could be daunting. But if your needs are undemanding and you have no desire to tinker, this is the easiest way to start using Linux. The real test was sitting my three-year-old daughter down on the PC with her favourite websites listed in the sidebar (complete with thumbnails). She was unfazed and was soon playing Tug of War with Clifford The Big Red Dog – all without intervention from me. [LXF](#)

LINUX Format VERDICT	
Installation	7/10
Documentation	9/10
Features	7/10
Value for money	8/10
An online/offline combination that deserves to be a huge success. Not for tinkerers though.	
LINUX Format RATING	

## DESKTOP DISTRO

# Elx PowerDesktop 1.0

Another desktop-orientated distribution that promises the earth, but does it deliver? **Chris Howells** investigates.

**Competing with Lycoris, Xandros and others, Elx intends to make Linux "easy to use."**

- **DEVELOPER** Everyone's Linux Ltd.
- **WEB** [www.elxlinux.com](http://www.elxlinux.com)
- **PRICE** TBA (see website)

**E**ven since the beginning, the idea of a 'distribution' has been to make it easier to install and use GNU/Linux. However, with the growing popularity of Linux as a desktop OS, some distributions have decided that they need to go further than the more mainstream distributions such as Debian, Mandrake, and SuSE, which are relatively general purpose.

The first distro aimed directly at the desktop market was Corel Linux 1.0, which featured a forked version of KDE 1.x, and a graphical installer. Since then, Corel Linux has been bought by Xandros, and other distributions such as Redmond Linux (now known as Lycoris), Ark Linux, and now Elx Linux have sprung up.

## Introduction

Originating from India, Elx Linux stands for "Everyone's Linux" and is described as a "fully featured Desktop OS". It cites ease of use as its number one goal. The distro itself is a meta-distribution based upon Red Hat which has been re-packaged, with some modifications.

Elx Linux is distributed on three CDs (surely overkill for something aiming to be an easy-to-use desktop OS?), and is supplied with two printed handbooks, a *Quick Installation Manual* and a *Reference Manual*.

## Installation

Installation starts by booting from the first CD. The default is a graphical GTK+ based installer. The first step is to choose the installation type. I chose a 'Default' installation since this is the option most likely to be chosen by a beginner. The next stage is partitioning; unfortunately custom partitioning is a fairly unpleasant experience.

Later there is the opportunity to choose the areas of packages that you wish to install – the defaults looked reasonable so I left the options as they were. As the package installation was about to start, I was told that the 1.7GB partition that I had allocated for the test install would not be sufficiently large, so was taken to a screen to choose individual packages. Unfortunately the user interface here is more gimmicky than usable. A picture of a penguin with a present is used to show whether the package is to be installed, and to change the option, you must first select the relevant penguin icon and then move the mouse across the screen to the check box. Other problems include the quality of translations, with questions such as "Are you having any other Linux OS and would like to save it?".

## Starting work

The first boot revealed that a large number of server daemons including *sendmail* (potentially making the machine an open mail relay for spammers), *Samba* (potentially allowing unauthorised access to your files), and many others were started up by default. This is very unimpressive for something which is targeted at a desktop machine, where in 99% of situations these daemons will do nothing apart from take up memory and cause security headaches.

My ATI Rage Pro 128 was detected and DRI was set up automatically, which is quite impressive. However my Creative SoundBlaster Awe64 sound card was not set up by the installer, but rather on the first boot, by the old-fashioned Red Hat utility *Kudzu*.


KDE's login manager, *kdm*, runs by default allowing a relatively easy graphical login straight into KDE, the desktop environment which is being pushed the most by Elx. The version of KDE is 2.2.2 which is extremely old, considering that KDE 3.0 has been out for a long time, and KDE 3.1 is now entering beta.



## System requirements

Pentium or compatible processor  
 CDROM drive  
 64MB RAM (128MB RAM recommended)  
 Minimum 1.2GB disk space (more recommended)  
 (If network installation is chosen, a CDROM drive is not required but an ethernet card is)

The reason for the old version of KDE is likely due to the large number of changes that Elx have made to it. For example, the desktop uses double



The **Control Panel** looks familiar, but is unfortunately lacking coherency.

## Essential Information

The distribution contains three CDROMs worth of applications, including:

**Kernel Version:** 2.4.17  
**X Window System:** XFree86 4.2.0  
**Desktops:** KDE 2.2.2, GNOME 1.4  
**Software selection:** KOffice, OpenOffice.org, AbiWord, Apache, GIMP, Mozilla, Konqueror  
**GCC:** 3.0.4  
**Glibc:** 2.2.3

click by default. The logout dialogue has been customised to allow the user to shutdown the machine straight from the logout screen (although this is implemented in KDE 3.0 anyway).

The KDE desktop has been changed a great deal to resemble Windows, with a 'Start' menu, 'My Computer' icons on the desktop, and a Windows-like window manager decoration. The KDE installation is set up to use anti-aliased fonts out of the box, which is useful for anybody looking for slightly better looking fonts.

Elx also features *launchpads*. These are small applets which are started by clicking on the appropriate icon on the KDE panel. There is a launchpad for different genres of software such as "Office", and so on. Loading the launchpad reveals a set of icons where different pieces of software can be started. Unfortunately, while these are intended to make it easier to use, I suspect that they will be found rather confusing at best. The descriptions are sadly not particularly helpful, with text



The *launchpad* applet allows access to software in similar genres.

such as "A spreadsheet application for the KDE GUI desktop" – acronyms such as "GUI" may be fine for the power user, but not for the beginner.


The quality of the RPM packaging of the distro appears to be lacking at times. In my case the *launchpads* were not installed initially due to the lack of disk space, but still the icons on KDE's panel were there. That the icons were there even though the software wasn't will cause confusion to many users. Unfortunately the installer had failed to setup the system to use a UK keyboard; when I started the KDE *Control Center* to fix this problem I discovered that none of the modules would load – clicking on an item in the tree view produced no change. This can only be attributed to packaging problems.

## General usability

The distribution is targeted heavily at someone migrating over from Windows. The default KDE installation is very Windows-like, which will make most Windows users feel at home, although unfortunately the distro is lacking in many other areas. The distro is RPM based, but unfortunately there is no easy way to install software, so the user must manually manipulate the *rpm* program on the command line, which is definitely not intuitive.

A 'Network Neighbourhood' icon on the desktop starts a network browser, which in combination with the SMB KDE *ioslave*, should make it easy to share files with Windows networks.

In terms of system administration, there is a 'Control Panel'. This features various options such as "Add/Remove



Partitioning is one of the more unpleasant stages of installation.

Software" which starts *KPackage*; 'Users', which starts *KUser*; and so on. Unfortunately the thing appears just to have been thrown together quickly, and lacks the coherency that is found in tools such as SuSE's *YaST*.

In the distribution, general office work is adequately handled with the inclusion of *KOffice1.1.1* and *OpenOffice.org 641 d*, although the distribution starts to show its age (even though it was released fairly recently), due to the fact that it includes a relatively buggy beta version of *OpenOffice.org*. *OpenOffice.org 1.0* has been available for some time.

Thankfully some attempts have been made to make Linux more suitable and accessible for the desktop user. Floppy disks and CDROMs are mounted automatically using the kernel-based *supermount*, and this arrangement seems to work fairly well.

## General

Amongst the desktop software, there is also a great deal of other software such as *Apache*, *Samba* and *BIND*. Personally I felt that this was inappropriate. For example *BIND* is a fairly complicated piece of software, and has had more than its fair share of security holes. Therefore it is not really suitable for use by beginners. It would also be prudent to only run it on a distribution which has a proven security team who can release updated packages speedily (unless you want to monitor the security mailing lists and fix any problems yourself).

The documentation is one of the better areas of the distribution, although at times it can get a little too complicated. The Reference manual is

quite comprehensive, and covers the basics of logging in, file management, customising KDE, connecting to the Internet, and setting up a printer.

## Conclusion

Elx Linux promises a great deal, but unfortunately fails to deliver in many of these areas. The installer is by no means as good as those found in other distributions, the quality of packing and general number of bugs is poor, and the distribution is not particularly well set up for a desktop machine due to the number of server daemons that are started by default.

Also lacking are some of the nice system administration utilities found in other distributions such as Mandrake and SuSE. While they are probably not needed by experienced users, they are fairly essential for non-techies to find their way around the new system.

Elx Linux shows some encouraging signs, but in its present state it is hard to recommend. The developers are currently working on an upgrade which will probably feature kernel 2.4.18, KDE 3.0.2 and others – let's hope that it is an improvement. [LXF](#)

## LINUX Format VERDICT

Performance	5/10
Ease of use	5/10
Features	5/10
Value for money	6/10

Not always well put together, and does not really meet its goal of being easy to use for the desktop user.

## LINUX Format RATING

5/10

## FANTASY REAL TIME STRATEGY

# Kohan

Reeling in horror from the sheer tackiness of the prose in the handbook, **Paul Cavanagh** steels himself to test whether there is anything new in this fantasy real time strategy.



During a campaign, other leaders will sometimes communicate with you. Mostly to taunt you or to surrender, but this time to offer a mini mission.

Classic RTS fare – you might want to consider *Heroes of Might and Magic*, or maybe *Alpha Centauri* or *Civilisation* as alternatives.

- **DEVELOPERS** TimeGate
- **PRICE** \$29.99
- **WEB** [www.transgaming.com](http://www.transgaming.com)

**W**hen it comes to fantasy real time strategy things like eternal overlords, malevolent and powerful forces and magical weapons are pretty much part of the package. They may be cheesy and outmoded, but games reviewers learn to view such things as occupational hazards. But someone's always got to come along and take things one step further. "As the masters of Khaldun, the Kohan tended their world like an immense garden. Their culture blossomed and reached a golden age of power and beauty. But Khaldun's summit of achievement also signalled its imminent decline. For beneath the gilded magnificence of Khaldun society, sinister forces were plotting to undo this celestial wonder.." This sub-*Silmarillion* tat is worthy only of contempt – thankfully for *Kohan Immortal Sovereigns*, nobody

particularly cares about narrative in computer games, but if there were prizes handed out for tacky, tired old clichés this game would collect a whole wheelbarrow full of gongs.

## What's new?

So there's these immortal types, right, and they've got to protect their world from some other immortal types who are hell bent on spoiling the crops, turning the milk sour and bringing about a terrible cataclysm. Yawn. So anyway, that's the premise, let's hope the gameplay offers something innovative to keep us interested. Let's see. *Kohan* has a number of tutorial missions, which teach you everything you need to know. So we can upgrade cities by building quarries, sawmills and blacksmiths to supply us with our material needs. Nothing new yet, but we can at least balance resource management with strategic battle, so that's good. New fighting units become available when certain upgrades have been made to towns and cities – still nothing new, but fair enough.

As the tutorial takes you by the hand and tells you all about areas of control and supply, and how to group your units together to make regiments, you'll probably begin to feel




Here's a city on the edge of a desert oasis. Sod sawmills and quarries, they ought to start an ice cream franchise and make a fortune.

comfortable with what is a classic system for a real time strategy with resource management. And yes, you right click on enemy units to attack them. But nothing new. Oh well, I guess if it ain't broke there really isn't any need to fix it.

## Recruitment drive

As you get stuck into the sixteen campaign missions, you might be a

little bit impressed with how many units the game can cope with. Huge armies can be built up and unleashed onto the enemy hordes with relative ease. And there's a good mix of individual units available to players too. You commission units from your cities, and they are composed of a frontline, support units and a captain. Your frontline troops can be foot soldiers or cavalry units, depending on what



At the end of a mission one of your fellow immortals will tell you how you're doing. This mission must have been a success with lady immortal pledging her love. Aaah.



**A sizeable army of cavalry being built up outside a full sized citadel – plenty of compost for the rose garden here.**

upgrades your city has. For example, you'll need to have a blacksmith and a barracks in order to commission mounted cavalry units. You could choose to fill your support slots with the same unit type as your frontline, but more specialised troops fit the bill better. Most support units are magical in nature (again availability depends on how well equipped your city is) and you

can use wizards, channellers, clerics, archmages and that sort of thing. Usefully, these units can heal wounded troops as well as unleash devastating magical attacks. Your captain will affect your unit's morale, and may have special abilities. You can assign one of the immortal Kohan to lead a troop if they are available, or you'll get a default captain. The Kohan can be

asked to lead a troop from a safe distance, or actively engage in battle with them (this can be effective, but normally they'll get splatted, resulting in your unit being susceptible to being routed frequently). There are different factions (royalist, council, nationalist) which control cities, and this will affect how your troops look and whether you'll get horses or beast riders, for

example. There are lots and lots of different unit types, but you'll only be able to control a few of them at any one time during a campaign game.

Campaign missions usually set you goals to defeat enemy cities, but sometimes you'll have slightly more interesting missions like repairing and protecting outposts, or ensuring that a particular character isn't assassinated. My major bugbear with the campaign game is that when you finish a mission, you start the next on a new map, so that you lose all the cities that you've so patiently nurtured, and all the troops that go with them. It's frustrating to go from a situation where you were making lots of gold, to suddenly be thrust into a deficit situation, where you have to frantically build quarries and blacksmiths all over again. The resource management aspect of the game is reasonably simplistic (if you have a deficit of materials it'll cost you in gold, if you have a surplus you can sell it off), but sometimes you just don't have enough resources and you'll have to race to finish a mission before you run out of cash.

I found that most of the missions are relatively easy to complete, and can take anything between one to six or seven hours to finish. So there's a reasonable amount of longevity there, what with sixteen missions. Occasionally the game will throw in what seems to be an impossibly difficult mission, where then enemy will just overrun you. Generally, a little



## Check Out The Competition

Gorgeous graphics



So you don't think that the pictures from *Kohan* look ancient and rubbish? Ahem. Perhaps you might want to think again.

### GENERALS

This is a screenshot from the soon to be released *Command & Conquer: Generals*. Mmmm, lovely.



### WARCRAFT

This is a closeup shot from *Warcraft III*. And if you've got a copy of *Winex 2.1* installed, you can play it straight out of the box! Who needs *Kohan*?



### NAPALM

And this, believe it or not, is a screenshot from an Amiga game in 1999. *Kohan*, old fashioned? You're not kidding.

# LinuxFormatReviewsKohan

« exploration will help out in these situations, say if you explore south, rather than east at the outset, you'll find a large citadel that will only too gladly hand control over to you, switching the balance of power in your favour. The missions are based on smallish maps, and you can choose to play on when you have accomplished the goals, so that you can use up every resource and become all-powerful (hahahahaha!!).

## Options galore

There are more gameplay options other than the campaign missions, which is just as well with only sixteen missions. The scenario options allow you to play with computer controlled teams where you can set victory conditions, choose which map to play on, decide which faction to play and so on. The AI isn't half bad and you can set the difficulty mode – these games can be fun, especially with loads of enemy teams being controlled by the CPU. When you've finished a scenario you'll be presented with graphs detailing how each player did, so that you can refine your technique next time around. Online multiplayer games are also available, and again you can tinker with victory conditions and the like. Also included is a level editor so that you can create your own maps complete with event triggers (for example you could cause the game to give the player 1000 gold pieces for discovering a ruined temple, or defeating a dragon).

## We deserve better!

This is all very well, and will keep you occupied for a good few hours. Playing the missions is reasonably satisfying, at least to begin with, and the game engine is very solid – it can be impressive watching hundreds of units marching across the map on a campaign. But if truth be told, this is a deeply disappointing release for Linux. I welcome every game release for the platform – there just aren't enough commercial games about, but we really do deserve better quality than this. Windows users (and now, owners of *Winex 2.1!*) can pop out to the shops and pick themselves up a copy of *Warcraft III* with its glorious 3D graphics, rpg type character progression and professional voice acting. In a few months they'll be able to get the latest *Command & Conquer* title, again with jaw-dropping 3D

## Download Only! Fat pipe required

**Kohan is available from [www.transgaming.com](http://www.transgaming.com), for download only. It'll cost you \$29.99, and if you're using a 56K modem, you really shouldn't be, because you'll be waiting at least 7 hours for the download. Yipes.**

visuals. And what have we got on Linux? An unoriginal, uninspiring fantasy RTS with outdated graphics – it may have won awards when it was released for Windows, but that was a far while back. You can't zoom in on the action here, or even rotate the picture. You get to view the action from the top down at a set level and



Completing a scenario brings up graphs on how well all the teams did in terms of population, military victories, economic prowess and game scores.



You can fiddle about to your heart's content using the level editor. Here I've supplied a citadel with its own swimming pool surrounded by crop circles. They'll host hippie festivals every year.

that's it. To be honest, I really have seen better graphics on an Amiga (honestly – ClickBoom's *C&C*-style masterpiece, *Napalm*). *Kohan* has no flashy 3D, no stylish cut-scenes with professional voice acting, no character development, and nothing even remotely new. You do get a playable, but ultimately rather bland RTS, and I suppose, all things considered, that's better than nothing. It's not just Linux that gets these sort of disappointing releases, mind you. Windows users have the rare privilege of having *Kohan*

*Battles of Ahmiman* available to them. It looks pretty much identical to this game, its predecessor. *PC Format* was unimpressed, gave the title a vociferous drubbing, along with a surprisingly generous 48%. We'd all like to see more games being released for Windows and Linux simultaneously, but if we are going to be forced to wait for games to be ported, then it would be great if people could make sensible choices on what titles we get. It may have been good in it's day, but modern gamers deserve more **LXF**

## LINUX Format VERDICT

Graphics	5/10
Playability	6/10
Features	4/10
Value for money	6/10

A functional, mildly diverting and challenging RTS. Unfortunately, it's also old-fashioned, graphically lacking, and deeply unoriginal.

LINUX Format **RATING**  
**6/10**

# J2EE Professional Projects

**Richard Drummond** tackles a heavyweight J2EE title and discovers its lighter side.

- **PUBLISHER** Premier Press
- **AUTHOR** Pallavi Jain and Sgadav Siddiqui
- **ISBN** 1-931841-22-5
- **PRICE** £36.99

**P**remier Press's *Professional Projects* range aims to teach developer's more advanced topics by example. *J2EE Professional Projects* adopts this tactic and presents the reader with three non-trivial Java-based server projects, with the design and implementation of each discussed in detail. The first of these is a Java-based chat server and a client with a *Swing* GUI. Although this is not really a J2EE project since it only uses APIs available in the standard class library, it does give a good insight

into socket programming and GUI development in Java. The next project gets into J2EE proper, and is an online banking system implemented with JSP, servlets and a JDBC database. The final projects is a online shopping system for a music store, and uses Enterprise JavaBeans, JDBC and JSP.

Despite the 'Professional' tag in the title, *J2EE Professional Projects* is not pitched at as an advanced a level as you might expect. It is accessible to those with no J2EE experience or indeed little knowledge of Java. It begins from first principles and the first 240 pages covers Java basics, including threading, sockets, RMI and CORBA. Even in later sections – when developing the projects – the basics of servlets, JSP, JDBC and so on are



covered. Depending on your preference, this is either the making or breaking point of this book. On one hand, it is useful to have all this information collected into one volume, while, on the other, these topics are treated more thoroughly elsewhere. Personally, I think the authors should have assumed that the reader has a more advanced level of knowledge and this would have allowed them more space to devote to what should have been the focus of the book, the

projects themselves. As it stands, little or no attention is given to practical issues such as choosing a database and application server.

## Linux Format VERDICT

A practical guide of J2EE beginners, but covers too much ground too thinly.

LinuxFormat **RATING**  
 **7/10**

# Network Security with OpenSSL

**David Coulson** reviews a book which provides a comprehensive introduction to the OpenSSL libs.


- **AUTHORS** Viega, Messier & Chandra
- **PUBLISHER** O'Reilly
- **ISBN** 0-596-00270
- **PRICE** 28.50

**T**ransferring data in plain text between client and server has become a major security concern, particularly when dealing with e-commerce, and handling transactions containing personal information. This relies on a standard protocol for encryption known as SSL, the Secure Socket Layer, which handles the connection. Of course, there is an Open Source distribution of the SSL libs for use with Linux and other platforms, *OpenSSL*.

Many people will be familiar with *OpenSSL*, as well known projects including *OpenSSH*, *mod\_ssl*, *Fetchmail* and PHP make use of *OpenSSL* for

encryption capabilities. As well as the libraries, *OpenSSL* has a fairly complex command line tool for managing certificates, which is covered in great detail within this book. The various ciphers, including DSA and RSA are looked at with some depth, and the Public Key Infrastructure (PKI) model deals with Certificate Authorities (CAs) so you can sign your own certificates.

Most of the book covers using the *OpenSSL* libraries with C, so a fair understanding of the language is needed. Fortunately, they start off with some nice and simple examples for creating servers and clients, as well as for connecting to existing servers running SSL. A major portion is dedicated to Public Key algorithms, typically the PGP alternative, S/MIME. Many of the S/MIME examples require a significant understanding of *OpenSSL*, so you'll need to read the rest of the book before skipping to that section.



If C doesn't do it for you, there is a single chapter on using SSL with Perl, Python and PHP. While it does cover the majority of the functions within each language, the actual usage of the functions are only described at an advanced level within the section which deals with C, so even a Perl hacker may want to brush up on their C if they're learning *OpenSSL* from scratch.

The non-programmer, who is using *OpenSSL* for secure TCP connections with web and mail servers will find the first third of the book to be of great use. However, unless you are especially bothered about having a nice book with pictures of sealions on the front, then you will probably prefer to find the info on the *OpenSSL* website. However,

if you are looking at using *OpenSSL* in any programming capacity, you'll be hard pushed to find anything which is as informative and easy to read as this. Even for a beginner, who is only starting to go beyond basic C, the examples are simple enough so the specific capabilities of *OpenSSL* are easily understood.

## Linux Format VERDICT

An excellent guide to using *OpenSSL*, but you'll need a solid C background.

LinuxFormat **RATING**  
 **9/10**

# Linux in Small Business: A Practical User's Guide

**David Coulson** finds a useful guide for small businesses examining Open Source software.



- **AUTHOR** John P. Lathrop
- **PUBLISHER** Apress
- **ISBN** 1-893114-46-1
- **PRICE** US\$36.95/ £27.00

**A**s Linux is significantly more cost effective than its closed-source counterparts, many business are saving money utilising a freely distributable OS, along with taking advantage of the huge number of apps which can be obtained legally from the Internet for no cost. Nothing is completely free, and Linux requires a significant amount of understanding to configure, and many services are beyond the capabilities of someone who has never even used Linux, much less installed it before.

While it does start off discussing Linux distros, it is completely based around Red Hat 7.2, so if you're using a different distribution, it will take a little imagination to be able to follow the examples in the book. As almost everything is done using GUIs, it's somewhat difficult to adapt many of the steps for system config, although if you are a RH user, it is very clear and easy to understand. Fortunately, this book manages to do a very good job of hiding the unpleasant text-based config files behind the plethora of GUI admin tools which come with Red Hat.

Everything is covered, from NFS and Samba, to Apache and sendmail.

However, if you want to seriously run any of these services, you're going to need something a little more in depth than this book. From a business perspective, it looks at whether a task can be done in-house, or if a consultant should be brought in to set everything up for you. To see this within a beginners book is refreshing, although the author may have his own agenda, as he runs his own consultancy company.

Rather than assuming that a business owner would want to replace everything with Linux, a decent amount is dedicated to integration with systems running Linux, so those still using Windows don't have to work out for themselves how to use Linux's network capabilities. Even though this book is geared towards businesses thinking of switching to Linux, it is a great book for anyone using Red Hat 7.x, and will certainly appeal to those in the home office trying to save money by moving towards Open Source software. **LXF**


## Linux Format VERDICT

Ideal for a small business trying to save some money, but only if Red Hat is the distribution of choice.

LinuxFormat **RATING**

8/10

# YOU WIN 8.1!



## NEW SuSE Linux 8.1

Putting you in the winners' circle with open standards

### For beginners: SuSE Linux Personal 8.1

- Free MS-Office compatible office suites
- Secure Internet and eMail communication
- Easy-to-install desktop solutions
- Extensive multi-media support
- Graphics manipulation tools for digital cameras, scanners etc.



### For professionals: SuSE Linux Professional 8.1

- Complete small office solution
- All you need to run your office network
- Configurable security with SuSE Firewall 2
- Additional secure file systems
- Numerous development environments and programming tools



SuSE Linux is celebrating its 10 year anniversary! We owe our success to you, thus we would like to thank you for your loyalty. Participate in the online questionnaire for your chance to win: [www.suse.co.uk/uk/81](http://www.suse.co.uk/uk/81)

Please invoice us at the following address:

Company \_\_\_\_\_  
 Surname/Name \_\_\_\_\_  
 Street \_\_\_\_\_  
 City \_\_\_\_\_ Post-Code \_\_\_\_\_  
 E-mail \_\_\_\_\_ Phone \_\_\_\_\_

Personal 8.1  
£ 39.00\*

Professional 8.1  
£ 59.00\*

\* shipping charges £5.00

Order Online Today! [www.suse.co.uk](http://www.suse.co.uk)

@ [info@suse.co.uk](mailto:info@suse.co.uk) ☎ +44(0)2083874088 📠 +44(0)2083874010

SuSE Linux Ltd.

The Kinetic Centre · Theobald St.  
Borehamwood · Herts. WD6 4PJ



# Roundup

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



## Our selection at a glance

- qmail
- Sendmail
- Postfix
- Exim
- Courier IMAP
- SuSE Linux eMail Server
- CommuniGate Pro
- GLMail
- Insight Server SE
- Volusion Messaging Server
- IEMS

## Mail servers

In an attempt to thwart Royal Mail's monopoly, David Coulson takes a look at a selection electronic mail servers.

For many people, email is as important as regular 'snail' mail or the telephone. It's quick, cheap, simple to use, and anyone with an Internet connection can make use of it. Many people see email as only being the 'client' side, such as *Outlook Express*, or one of the numerous Open Source mail clients available for Linux, and neglect to realise the poor server sitting behind the scenes. When you decide to forward a pathetic joke email (Really, has anyone ever been sent a good joke via email?) to everyone in your address book, each email will generally travel between two mail servers using a protocol known as the Simple Mail Transfer Protocol, or SMTP. SMTP is

used, as you can imagine from its name, to send mail from one mail server to another, and it's up to each individual mail server to decide how exactly it should handle delivery of messages which come in.

### Relays

Mail servers generally have two rôles. The first is as a relay, which will forward mail to another SMTP server. Usually, this is only found where you have clients connecting to an outgoing SMTP server, rather than having each user run their own SMTP delivery system. Most ISPs have an SMTP server, for customers to deliver mail through their systems.

The second use of an SMTP server is as a collection server, where incoming

mail to specific domains is distributed to mailboxes, so that users can collect it via POP3 or IMAP4, or even through a shell-based mail client such as *Mutt*. It's not impossible for a mail server to both handle mailboxes and act as a relay, so most installations will do both. Handling your own incoming and outgoing email using an SMTP server can be done very easily, although care must be taken when configuring the server, as it's not impossible to set it up such that you lose email or have messages delivered to the wrong place. Particularly if you're operating a relay, it's imperative that it does not allow everyone and their brother to relay messages through, as this type of misconfiguration, known as an 'open

relay' allows spammers to send messages through your system, meaning that when someone decides to do something about it, it's your mail server which gets blocked or filtered, rather than the spammer.

### Spam

Plonking incoming mail in a mailbox is useful, but mail servers need to do much more. With the increased use of email comes an increase in Unsolicited Bulk, or Commercial, Email, UBE and UCE, known to the man-in-the-street as spam. While no mail system can effectively remove 100% of spam, it's useful to be able to limit the amount of junk email entering the system, in order to stop your machines wasting

## Quick Reference: IPv6 IP address bonanza

The current Internet uses version four of the IP protocol, which is limited to 2<sup>^</sup>32-1 IP addresses. Over recent years, due to increased use of the Internet, the available IP ranges have been under pressure, and it is expected that IPs will run out within the next few years if prior increases are continued.

To combat this problem, version six of the IP protocol, otherwise known as IPv6, offers up to 2<sup>^</sup>128-1 IP addresses, which is significantly larger than the IPv4 range and should last for a while. IPv6 uses a completely separate IP protocol, so support is required by routers and systems which are on IPv6 networks. Fortunately, Linux has had IPv6 support for a number of years, and works happily with IPv6 networks.

As few ISPs offer IPv6 dial-ups or direct connectivity to an IPv6 network, we need to use a system known as tunnelling to route IPv6 traffic over the existing IPv4 Internet. There are a great number of different companies and organisations offering IPv6 connectivity, including [www.tunnelbroker.net](http://www.tunnelbroker.net), [www.freenet6.org](http://www.freenet6.org) and

[www.ipng.org.uk](http://www.ipng.org.uk). Each of these offer IPv6 connectivity, and route a subnet of IPv6 addresses to your network, so each system can have its own IPv6 address and be accessible from the outside world.

An IP tunnel can be setup using 'iproute2':  
# ip tunnel add ipng mode sit remote  
<remote IP>  
# ip addr add <IPv6 address>/128  
# ip link set ipng up  
# ip ro add 2000::/3 dev ipng

Public IPv6 addresses are generally within the 2001:: or 3ffe:: networks, so we don't route ::0 over the tunnel, only 2000::/3. There are many other IPv6 addresses which are not publicly routeable, and are either local to the host, the network, or are used for multicast networks. Each interface on an IPv6 enabled host has a 'fe80::' address which is constructed from the Ethernet MAC address of the device. Of course, since IPv6 has such a huge address space, we can assign IPs based on MAC address, and rather than using DHCP, there is a server known as 'radvd' which assigns IPv6 addresses based upon the MAC of the interface broadcasting for an address.

**'It sticks closely to SMTP standards and as far as security is concerned there is a \$10,000 bounty for anyone who can exploit qmail, which remains unclaimed'**

CPU time trying to deliver messages which will be instantly deleted.

A common way to avoid spam is to use a Real time Blackhole List (RBL) such as MAPS, which allows a mail server to perform DNS lookups for IPs and find out if they are a source of spam or not. MAPS is a subscription service, although there are free DNSRBLs available on the Internet which have some effectiveness against spam.

There is also the threat of email based viruses, which can attack exploitable systems, generally those running MS Windows, disrupting email for hundreds of users, and frequently causing irreparable damage to users' systems if the malicious code has such an intent. Rather than pretending that

the world is a perfect place, and everyone updates their system with the latest security patches, it's much easier to filter out nasty messages on the server. As with anything of this nature, a virus checker is only useful if it's kept up to date, and there are many commercial virus checking programs available for Linux which allow you to check queued emails.

We're going to be looking at a selection of mail servers, both those available over the 'Net free of charge, and commercial packages which come with support. As well as looking at the general mail capabilities of the servers, security and filtering play a major rôle when looking at the features of the systems.

# qmail

100% RFC-compliant, but somewhat quirky.

■ Web: [www.qmail.org](http://www.qmail.org) ■ Price: Free

```

david@niamh:~ (pts/39)
CONTROL FILES
WARNING: qmail-send reads its control files only when it
starts. If you change the control files, you must stop
and restart qmail-send. Exception: If qmail-send receives
a HUP signal, it will reread locals and virtualdomains.

bouncefrom
Bounce username. Default: MAILER-DAEMON.

bouncehost
Bounce host. Default: no. If that is supplied, oth/
erwise the literal name bouncehost, which is probably
not what you want. If a message is permanently undel-
iverable, qmail-send sends a single-bounce notice
back to the message's envelope sender. The notice is
From: bouncefrom@bouncehost, although its envelope
sender is empty.

concurrentlocal
Maximum number of simultaneous local delivery
attempts. Default: 10. If 0, local deliveries will
be put on hold; concurrentlocal is limited at confi
le time to 120.

```

Control files simplify qmail configuration.

qmail is the Marmite of the SMTP server world – You either love it, or you hate it, with nothing in between. qmail was designed to be 100% RFC compliant, and sticks closely to the standards for SMTP delivery agents. As far as security is concerned, there is a \$10,000 bounty for anyone who can exploit qmail, which remains unclaimed. Don't expect your installation to be impervious to external attackers if you're running something exploitable along with qmail.

Installing qmail, although not the standard ./configure && make which most people are used to, is simple. Config is even easier, and qmail has a simple make option to put some entries into the config for basic mail delivery to work happily from the word go. qmail simply has a selection of config files, mostly one-liners, in /var/qmail/control, and you enter the correct domains and hostnames in those files in order to have mail handled properly. Each file has a very simple name, so it's usually very easy to figure out which one you need to edit for a specific purpose. qmail can handle both mbox and Maildir delivery, and can pipe your mail through whatever app you want, such as SpamAssassin. While not having any DNSRBL or virus checking support, it's very easy to configure qmail to send all mail through a virus checker or spam filter.

One of the biggest problems with qmail is that it forks a process

for each individual mail it tries to deliver, either remotely or locally. As default, qmail will deliver 20 remote and ten local messages at any time, which should be more than adequate for most systems. As far as features go, qmail doesn't have all that many beyond basic mail delivery. However, there are more than enough patches available, offering encrypted TLS support, MySQL config and IPv6 connectivity, as well as performance improvements and improved queue storage for systems with large amounts of mail.

If you're willing to take the time to patch qmail, even with the qmail patch cocktail which most people are used to, is simple. Config is even easier, and qmail has a simple make option to put some entries into the config for basic mail delivery to work happily from the word go. qmail simply has a selection of config files, mostly one-liners, in /var/qmail/control, and you enter the correct domains and hostnames in those files in order to have mail handled properly. Each file has a very simple name, so it's usually very easy to figure out which one you need to edit for a specific purpose. qmail can handle both mbox and Maildir delivery, and can pipe your mail through whatever app you want, such as SpamAssassin. While not having any DNSRBL or virus checking support, it's very easy to configure qmail to send all mail through a virus checker or spam filter.

## LINUX Format VERDICT

Installation	6/10
Documentation	7/10
Features	5/10
Ease of Use	8/10

qmail is a great mail server, but be prepared to patch it if you want features.

## LINUX Format RATING

6/10

# sendmail

The grandfather of all \*nix mail servers.

■ **Web:** [www.sendmail.com](http://www.sendmail.com) ■ **Price:** Free

```
davide@goofy:~
# my official domain name
# ... define this only if sendmail cannot automatically determine your domain
#DOMAIN: Foo.COM

CP,

# "Smart" relay host (may be null)
DS

# operators that cannot be in local usernames (i.e., network indicators)
CO 9 % !

# a class with just dot (for identifying canonical names)
C.,

# a class with just a left bracket (for identifying domain literals)
C[[

# access_db acceptance class
C(Accept)OK RELAY

--More-- (6%)
```

Yes, it takes time to learn, but your new-found skills will be useful.

Certainly the first thing people think of when you mention email on Linux is *sendmail* and, until a year or so ago, almost every Linux distribution came with *sendmail* installed as the default SMTP server. Most people therefore learn to live with *sendmail*,

rather than actually choosing to use it. Not that *sendmail* is a mail server to be sniffed at, as it scales particularly well and has a wide array of features, including DNSRBL support and smart host capabilities. Sendmail, Inc, which has been formed by the developers of

*sendmail*, offers high-performance *sendmail* solutions for businesses.

## m4

The main problem with *sendmail* is that its configuration system is based around 'm4', which is a macro language. m4 is not the most difficult thing in the world to learn, but it does mean that you need to take a decent amount of time to figure out how to configure *sendmail* securely as well as actually make it do what you want. Of course, unless you intend doing something special with it, having to learn a new macro language in order to configure your mail server is not worth it for most people, and as one could imagine, the room for error and misconfiguration is rather large. Fortunately, the default *sendmail* configuration is actually rather useful, so even without jumping in and hacking at the configuration files, it will work as a basic SMTP server.

## Powerful

If you can master m4, through *sendmail* documentation or one of the numerous books which you can get covering the subject, *sendmail* is a very powerful and capable mail server.

As always with a popular package, there is a large and knowledgeable community of *sendmail* users on the Internet, so even a novice *sendmail* user can usually have problems fixed fairly quickly. Many distributions are moving away from *sendmail*, to simpler yet equally capable mail systems, including *Postfix* and *Exim*, as many people are moving away from *sendmail* to easier to manage servers. Should you choose *sendmail*, make sure you keep up with security updates, as *sendmail* doesn't have the greatest track record with exploits – not that one should be discouraged from using *sendmail*, as holes are plugged particularly quickly.

### LINUX Format VERDICT

Installation	7/10
Documentation	9/10
Features	7/10
Ease of use	5/10

If you can stomach m4, *sendmail* is a more than capable mail system.


### LINUX Format RATING

7/10

# Postfix

Easy-to-config drop-in replacement for sendmail.

■ **Web:** [www.postfix.org](http://www.postfix.org) ■ **Price:** Free



Discover how simple it is to configure *Postfix*.

*Postfix* originally started out as a simple alternative to *sendmail*, but has since been developed further into a fully fledged mail server in its own right. Since its conception as *VMailer*, it has become a very popular mail system, due to its simplicity and ease of configuration. *Postfix* was originally

designed to simply 'drop in' as a replacement for *sendmail*, so for systems currently running *sendmail*, or indeed any of the standard mail servers available, switching over to *Postfix* is fairly straightforward. Of course, reconfiguration of the server from scratch is required, so the

benefits of upgrading to *Postfix* must be weighed against the time it takes to set everything up.

*Postfix* can be used for a wide range of applications, from a simple mail service, to a large-scale mail server. It is rare to see *Postfix* used for a high-traffic service, although it is perfectly capable at operating as such, and the general use of *Postfix* is as a simple remailer on a workstation. *Postfix* can easily be configured so that it reads specific configuration options, including virtual domains and aliases, from a file external to the main configuration, so it is exceptionally easy to add extra domains. Unfortunately, beyond *dbm*, *Postfix* does not support the use of a database, such as *MySQL* or *PostgreSQL*, to store configuration details, so it is not ideal for real-time reconfiguration.

Mandrake Linux has used *Postfix* for some time as its default SMTP service, and its ease of configuration has done wonders for its use, as most people won't just throw it away and switch to something else because it is too difficult to manage. For simple

mailer configurations, *Postfix* is a perfectly capable system, although if you're handling a large number of domains, and want complex virtual domain configurations, other servers are somewhat more appropriate. However, for a small-scale server, handling basic mail delivery for users and for outgoing email, *Postfix* is a reasonable choice, but many prefer to choose something else, as it lacks many of the more powerful capabilities which other mail servers have. Of course, if you're looking for a mail server, there is little to lose by trying out *Postfix*.

### LINUX Format VERDICT

Installation	5/10
Documentation	7/10
Features	5/10
Ease of use	5/10

*Postfix* is a very nice and simple SMTP server, but not ideal for complex installations.


### LINUX Format RATING

5/10

# Exim

Powerful, scalable and well-documented.

■ **Web:** [www.exim.org](http://www.exim.org) ■ **Price:** Free



This powerful and secure mailer is the Debian favourite.

**Developed at Cambridge University,** *Exim* is a hugely powerful mail system, and has been designed to handle large quantities of mail, and to limit UCE/UBE on a network. Many high-traffic services, including Sourceforge, use *Exim*, as it scales particularly effectively and can be tuned specifically to the needs of the users. Whether you've five users behind a firewall, or tens of

thousands of accounts, *Exim* can handle a great range of apps and can certainly deal with pretty much everything which can be thrown at it.

*Exim* has a very comprehensive config structure, offering everything from ACLs to database support for everything from MySQL and PostgreSQL, to Oracle and IBM's DB2. However, there is no need to learn

every aspect of the configuration just to get *Exim* up and running, as there are comprehensive docs available on the *Exim* site, making setting up a basic configuration of *Exim* a breeze. Expanding server capabilities is easily done, and *Exim's* flexible configuration structure gives a huge number of options for each segment of the setup.

A number of distributions, including Debian, are using *Exim*, and in its default configuration it is particularly secure and useful. Indeed Debian, as do most other distributions, offers a comprehensive config tool allowing *Exim* to be configured to handle mail without prior understanding of the raw config files. As one would expect, once the basics of the configuration format are understood, it is simply a matter of looking up the appropriate directive for the files within the documentation and applying it to the system. As with most things 'Linux', O'Reilly have produced a book looking at *Exim*, which is a must-buy for any *Exim* administrator or, indeed, anyone who is looking at learning *Exim* seriously. The documentation on the *Exim* site is second to none, but it's often nicer to have something in print which one can pick up off a shelf when things

decide to get a little complicated.

*Exim* is certainly one of the most powerful Open Source mail systems available and, compared to *sendmail*, has a particularly straight forward configuration style, which makes it significantly more attractive to beginners. There is not really a major features which *Exim* lacks, so unlike *qmail*, one does not need to rely on third-party efforts, ensuring consistency and stability, which is always a good thing when dealing with a piece of software which is not supposed to fail. While being fairly complex, the capabilities of *Exim* are really quite astounding, and it's difficult to find a more capable system.

## LINUX Format VERDICT

Installation	8/10
Documentation	10/10
Features	9/10
Ease of use	7/10

If you're looking for an Open Source mail server, stop reading now.

## LINUX Format RATING

9/10

# Courier IMAP

IMAP and POP3 server for Maildir users.

■ **Web:** [www.inter7.com/courierimap](http://www.inter7.com/courierimap) ■ **Price:** Free



Like it says on the tin, *Courier* is for Maildir – not mbox.

**While there are many Open Source** SMTP servers, none of them have the ability for users to fetch their mail via POP or IMAP. One has to provide a separate mail server to handle mail downloading, and there are numerous POP3 and IMAP4 daemons available, including *UW-IMAP* (part of the *Pine* package), and the generic *pop3d* and

*imspd*, found in most distros. Generally, these simply authenticate against the system users, such as `/etc/passwd` or NIS, and deliver the user's mail out of `/var/spool/mail/$USER`. This is great for a small system, but users with large IMAP mailboxes are going to have problems and sysadmins are not going to be too happy when users have giant mailboxes

which gobble up system memory.

Rather than using 'mbox' storage – the simple flat-file method of storing a user's mail on Unix – mailservers such as *qmail* and *Exim*, support the use of Maildirs, which have a single file for each email. Argument for and against the use of Maildirs will continue well into the next ice age, but for users of IMAP, it means the system can quickly look through a number of simple files, rather than trawling a giant mbox. As a POP3 and IMAP4 server capable of handling Maildirs, *Courier IMAP* is a popular choice. It can authenticate out of databases, including MySQL, PostgreSQL and LDAP, making virtual mail configurations a breeze to setup. As Maildirs include all mail within their directory, users of IMAP don't end up with the sprawls of directories and messages all over their home directory as they create new folders and move messages around on the server.

*Courier IMAP* is actually part of the *Courier Mail Server*, although most people prefer to use *Exim* or *qmail* with the *Courier* mail tools. They have also made many of the components of their mail system available individually,

so one can mix and match without getting in a muddle because things don't want to work together. Using *qmail*, with the MySQL patches, along with *Courier* can quickly yield a scalable and flexible virtual hosting platform for mail, capable of handling a significant number of domains with ease.

Of course, not everyone can, or chooses to, use Maildirs, and *Cyrus* is a particularly popular POP and IMAP server for use with systems handling the mbox storage format. As with *Courier*, *Cyrus* can handle database authentication, so mail accounts need not be 'real' users on the system.

## LINUX Format VERDICT

Installation	5/10
Documentation	7/10
Features	7/10
Ease of use	5/10

If you're using Maildirs, with *qmail* or *Exim*, there is no better IMAP or POP server.


## LINUX Format RATING

6/10

# SuSE Linux eMail Server 3.1

A web interface makes this server straightforward.

■ **Publisher:** SuSE ■ **Web:** www.suse.com ■ **Price:** US\$999



SuSE slap a friendly front-end on, and try to make things easier for you.

SuSE have always been well known for their very popular Linux distro, so it's not surprising they are branching out with other products for their user-base. Like many of their other products, they are aiming their mail server

system at medium sized businesses, who can maintain their own systems without feeling the cash drain.

SuSE's system is based around an SMTP and IMAP server, along with a very simple web-based admin front-

end. Everything the admin needs to do can be done via this interface, so there is no manual config hacking, or delving into the complexities of SMTP. Of course, some understanding is required and, certainly, one could not get by without a little knowledge of how email works – but it does not require either a network or Linux boffin to make it all happen.

Along with mail, the web interface handles a calendar, address book, and a comprehensive todo list tool. Each of these are standard features available from most apps, but the fact that it is all available in one place, accessible over the network, is great for anyone who is on the move, or for a company which wants to keep everything centralised for efficiency. SuSE suggest that their system can handle around two million users, with between fifty and two-hundred people accessing the system simultaneously with a server of reasonable capacity – more than enough in most situations.

This system is actually based around *Postfix*, along with an LDAP

server for user management, so it is not impossible for one to construct a similar system without having to spend the cash on SuSE's product. As one would expect, it is the ease of configuration which will attract people to SuSE's system, although at nearly a thousand dollars, is it worth looking at Open Source alternatives, which may suit your personal needs a little better. However, since it can be installed out of a box, and be up and running pretty much instantly, it is certainly a very clean and simple choice.

## LINUX Format VERDICT

Installation	9/10
Documentation	8/10
Features	7/10
Ease of use	8/10

A simple, out-of-the-box mail solution, although it's not cheap.


## LINUX Format RATING

8/10

# CommuniGate Pro

Professional, powerful but wilfully obscure.

■ **Publisher:** Stalker Software, Inc. ■ **Web:** www.stalker.com/  
CommuniGatePro ■ **Price:** From US\$499



Enter at your peril, jargon and unnecessary features await.

CommuniGate Pro is a featureful system, offering almost all one would desire from a mailserver. It is clearly aimed at customers with an extensive technical background, as much of their feature list and manuals is well beyond

the understanding of many admins. I'm not sure if they're trying to confuse people with jargon, or include every possible protocol they can imagine, but does one really need personal websites within a mail server?

*CommuniGate Pro* gives great flexibility to the users, allowing them to collect mail via POP, IMAP or WebMail. Certainly a very useful capability, but nothing particularly special in the grand scheme of things. Everything can be handled securely over SSL, and a variety of authentication methods are supported, so compliant mail clients should have no problems using their IMAP or POP servers. For some reason, 'poppwd' and Netscape roaming is supported, neither of which are used by many clients, and indeed, it's quite unlikely that anyone is using a client these days which supports the roaming capabilities of Netscape 4. 'poppwd' was originally used by *Eudora*, although a small number of clients now use it, although there is actually no RFC or other standard stating how it should work. While there is nothing wrong with implementing non-standard features, if it is ever standardised, there is no guarantee that it will match the current usage.

At this price one would really expect a significantly more basic installation and config. Generally, anyone who can understand their instructions well enough to get it going is perfectly capable of constructing a mail server based upon Open Source projects, which for a large business who needs more than the standard fifty mailboxes, is certainly significantly more economical. While Stalker Software offer support, most companies are going to need to employ a third-party to get the show on the road with *CommuniGate Pro*, adding to the expense.

## LINUX Format VERDICT

Installation	5/10
Documentation	8/10
Features	9/10
Ease of use	4/10

A powerful system, with too many features, let down by a poor front-end.


## LINUX Format RATING

6/10

# GLMail

Powerful, easy to use and easy to install.

- **Publisher:** Gordano, Ltd. ■ **Web:** [www.ntmail.co.uk/products/GLMail](http://www.ntmail.co.uk/products/GLMail)
- **Price:** From £525.



Everything is handled through a friendly web-interface.

Gordano have produced a popular mail server for Windows NT for a number of years, but have recently seen the light and decided to port it over to Linux, and other Unix platforms. The feature list for *GLMail* is

impressive, as is the ease of installation and configuration, making it great for anyone who needs a mail server up and running, but doesn't want to waste time understanding what it's doing.

Installation is a simple process, and

*GLMail* does not rely on any specific distribution being available, which is nice if you're already familiar with a particular style of Linux and don't want to adjust in order to get your mail server installed. Everything is handled through a web-interface, allowing complete control over the use of the server, yet hiding the unpleasant configuration away from the administrator. Almost every mail configuration is supported, and expanding upon *GLMail's* capabilities by adding extra hardware at a later date is all taken care of.

## Killer app

*GLMail* is particularly popular for its anti-virus and anti-spam filtering capabilities, and certainly within our experience, one can't really argue with the level of filtering which it possesses. As good as it is at blocking mail, it also yields a low rate of false positives, giving confidence in the system and not having users worry that the mail server gobbled up their important messages thinking that they contained something unpleasant or were from a spam source.

The support available for *GLMail* is certainly very good, and a wide range

of contracts can be bought depending upon the needs of your own business. If you're interested in looking at *GLMail*, a 28-day trial is available, although we would think most people would decide well within those twenty-eight days to roll out *GLMail* on their production systems. With the ease of migration from existing mail systems, using tools available from their site, Gordano have a very capable and scalable mail platform which we're quite sure is the envy of many other mail server manufacturers. With the spam and virus filtering, which is generally not available at the same level with Open Source products, *GLMail* is close to a 'killer app'.

## LINUX Format VERDICT

Installation	8/10
Documentation	9/10
Features	9/10
Ease of use	9/10

An almost perfect mail system, even for those without experience of email servers.

## LINUX Format RATING

9/10

# Insight Server SE

Powerful, Exim-based solution.

- **Publisher:** Bynari, Inc. ■ **Web:** [www.bynari.net](http://www.bynari.net)
- **Price:** From US\$600

Bynari, Inc are known for *MDaemon*, their popular SMTP server for Windows, yet are relatively unknown within the Linux community, even though they offer large-scale mailing systems based on Linux for IBM S/390 and z-Series servers. As one would expect from a product of this calibre, it has a very impressive list of features, yet is quite happy to run on the most meagre of systems, if you don't have a spare S/390 mainframe sitting around and gathering dust.

The installation of *Insight Server SE* is done purely on the command line, which will probably put a few people off – but it is not completely cryptic, so you really don't need to understand much beyond the basic config you desire to get it up and running. Rather strangely, within their manual, Bynari explain how to add new users, groups



Yet another easy interface – for the basics, at least.

and organisations to the LDAP server using LDIF, which while being useful, seems to indicate that they're aiming their product at large organisations with very experienced administrative staff. Fortunately, if you don't care to know the intricacies of LDAP, it can all be done via the web interface.

Control mail is a breeze with *Insight*, and it handles a variety of filtering

methods, including DNSRBL and *Exim's* filtering system. Sadly, if you want the latter, you're going to have basically learn how *Exim* filters things, as they refer to it as an 'advanced option', and their web interface doesn't really help any. One should think that many businesses who are thrashed daily with spam and email attachments containing viruses and trojans would prefer this to be a basic option, with a much simpler front-end. To be completely honest, much of *Insight's* system is simply a front-end to *Exim*, which really doesn't make life any easier for the administrator, as they still need to understand how many of the options work and what capabilities *Exim* has, in order to configure their server appropriately. There is nothing wrong with the range of features, but since one has to learn how to use regular expressions in order to filter mail, this really isn't one for inexperienced administrators.

As with many other servers, *Insight* uses OpenLDAP and *Cyrus* for user accounting and POP/IMAP collection, but actually used *Exim* as an SMTP daemon, so it still has the Open

Source underbelly. Considering that they tell you how to do everything manually, one has to wonder if they've not put enough time or effort into their web-front end, or if they think that the admin staff who look after their software are happier playing with the command line, than clicking buttons and forms in a browser. They also seem to have the annoying tendency to point to the *Exim* docs within their manual, rather than describing it themselves, which does not give a particularly good impression of either their product, or their desire to inform their users.

## LINUX Format VERDICT

Installation	5/10
Documentation	5/10
Features	7/10
Ease of use	6/10

Little more than a collection of Open Source mail components, along with a web interface.

## LINUX Format RATING

6/10

# Volution Messaging Server

SMTP server with powerful front-end.

- **Publisher:** Caldera, Inc. ■ **Web:** <http://www.caldera.com/>
- **Price:** From US\$879



This is where all Caldera's work has gone.

While Caldera are only well known for their distribution, a major portion of their revenue comes from other products, including the Volution range. As with the SuSE mail server, *Volusion Messaging Server* is built around the

Open Source *Postfix* SMTP server, and Caldera have taken the use of Open Source products a step further, by making use of OpenLDAP for the storage of user data and *Cyrus* for mail access via IMAP or POP.

However, *Volusion Messaging Server* should not be seen as a simple collection of Open Source products, as it offers a very flexible and powerful web-interface, enabling the administrator to control SMTP domains, user accounts, aliases and so forth. As it is all stored under LDAP, none of the users are 'real' users on the system, so in theory, it can scale significantly further due to not being constrained by the limits of the system. Individual users can modify their aliases and auto response preferences via a cut-down version of the administration web tool.

And, well, that's about it. *Volusion Messaging Server* (which only runs on Caldera OpenLinux) does nothing outstanding, or even rudimentary, such as mailing lists, filtering and so forth. One can create the most basic of mailing lists by having an alias with multiple destinations, but compared to a powerful mailing list system, it's really quite limiting. The web interface is nice, and it certainly makes it very simple for users to modify their profile and

preferences, but beyond that, there isn't really very much else it does. No web-based schedule or address book and no web mail. All things which most other mail systems include.

For the price of *Volusion Messaging Server*, one would really expect something which does significantly more, and in its present state, it's not something which most people would choose to put on their front-end network without having another mail server sitting in front of it filtering out unwanted content.

## LINUX Format VERDICT

Installation	8/10
Documentation	7/10
Features	3/10
Ease of use	8/10

It's great at what it does, but unfortunately, it doesn't do much.


## LINUX Format RATING

5/10

# IEMS

Very powerful, but power comes at a price.

- **Publisher:** International Associates ■ **Web:** [www.ima.com/iems](http://www.ima.com/iems)
- **Price:** From US\$175.



A tremendous range of features is made easily accessible.

The Internet Exchange Messaging Server, or *IEMS*, is an all-in-one alternative to MS *Exchange* for Linux and Windows, offering an array of

capabilities and features. Rather than using Open Source products, *IEMS* has been written from scratch, and is available in a range of products, each

with differing options, so the more you pay, the more you get. It seems to make sense, until you realise that you only get *Lotus Notes* integration, which is one of *IEMS*' major selling points, distributed with *Professional* packages for more than 250 users. People with 250 users won't mind, but someone with 50 is probably going to be a little annoyed that they can't use *Notes* with *IEMS* as they publicise.

## Clusters

The feature list of *IEMS* is certainly very impressive, and it is capable of handling a significant amount of mail. While it can handle many basic tasks, including POP and IMAP, it also supports BSMTMP delivery, allowing *IEMS* clusters to be deployed without a great deal of effort. As well as handling mail for individual users, *IEMS* can handle mailing lists and, indeed, can be setup to handle massive lists with a significant number of subscribers. As with many list managers, people can subscribe to lists via email, so there is no need for the administrator to waste time adding people, unless the list is specifically configured to only permit

new subscribers who the list administrator has authorised.

Unfortunately, *IEMS*' web mail client lets the side down. It's awfully basic, not particularly pretty to look at, and looks more like an after thought than anything else. One would think that most sites using *IEMS* would install IMP or another IMAP client, which has significantly more power than *IEMS*' rather pitiful effort. If you're prepared to spend the extra cash in order to get some of *IEMS*' unique features, then it's certainly a worthwhile products and, when it comes to *Notes* and *cc:Mail* interfaces, it has little competition.

## LINUX Format VERDICT

Installation	8/10
Documentation	7/10
Features	9/10
Ease of use	8/10

A fantastic mail server, but be prepared to pay more for some features.

## LINUX Format RATING

8/10

# MAIL SERVERS

## THE VERDICT

**M**any commercial systems are, interestingly, based upon Open Source projects – hopefully an indicator of how stable and flexible systems developed openly over the Internet can be. Of course, one should not be expected to pay for something which is freely available, so one would hope that the commercial packages offer something well beyond the standard of the Open Source products. Unfortunately, in the case of *Insight Server SE* and *CommuniGate Pro*, this simply has not happened, as they seem intent on offering the bare minimum of features, other than what is provided by the underlying mail system. Maybe they think people will blindly pay for commercial variants of Open Source products without questioning what they are getting, but considering that neither Bynari nor Stalker Software are developers of the basic SMTP system they are using, there is no realistic benefit other than support, which is obtained by spending cash on their commercial product.

Of all the systems we looked at, *Volution* has one of the nicest web interfaces, and is the most straight forward to administrate. Alas, this comes at a price, and *Volution* is next

to useless for most people, due to its rather dismal lack of filter capabilities. On the flip-side, *Insight Server SE* has all the features, but a web-interface which is about as much use as a text editor, so really doesn't help anyone get to grips with the software.

Certainly, someone out there must be trying to work out how to use *Exim*, and using the web front-end to do so, but the learning curve is simply too steep for most. Maybe someone should point out that web interfaces are there to make things easier, not more difficult.

While with this type of software there is rarely a 'winner', since it is hinged on personal and business requirements, *GLMail* certainly stands out from the crowd. The only area where it was really lacking was in handling mailing lists, yet this is available as a separate package from Gordano (which we are yet to test). Considering that *GLMail* does not use Open Source components, they have taken the initiative and built a powerful system from the ground up and made it easy to understand, install and administrate. Close up behind is *IEMS*, which has more features than one can shake a stick at. With a wonderful admin interface, and connectivity with



*ccMail* and *Lotus Notes*, it is a very capable MS Exchange alternative.

However, Open Source systems should not be forgotten. *Exim* is certainly the most capable Open Source mail system out there, which is evident by its use by commercial products and inclusion as the standard mail server by a number of distributions. While not being the most simple system in the world, *Exim* is significantly simpler than *sendmail*, and is quickly becoming the SMTP

server of choice for many organisations and individuals. Each of the Open Source servers we looked at have a place, and are appropriate for different situations and for people with different levels of knowledge and experience. Most are probably going to stick with whatever is installed as default by their distribution when Linux installed and, fortunately, it's not beyond the capabilities of many to allow a standard installation to handle a significant amount of traffic. [LXF](#)

### Table of features

While many commercial mail servers are built around Open Source components, it is quite clear where the developers have gone the extra mile to create a product greater than the sum of its parts. Unfortunately, some mail servers, particularly *Volution*, really do little beyond what one can do with a basic install of a freely

distributable mail service, giving one to question the product's value for money.

Fortunately, each of the commercial systems has a web interface, so one needn't understand the internal config of the mail server. Some do this better than others, although the quality of the interface is a poor indicator of the server's capabilities.

	Qmail	Sendmail	Postfix	Exim	Courier	SuSE	CommuniGate	GLMail	Insight	Volution	IEMS
POP	N	N	N	N	Y	Y	Y	Y	Y	Y	Y
IMAP	N	N	N	N	Y	Y	Y	Y	Y	Y	Y
Notes	N	N	N	N	N	N	N	N	N	N	Y
Price	Free	Free	Free	Free	Free	\$999	\$499	£525	\$600	\$879	\$175
SSL	N	Y	N	Y	Y	Y	Y	Y	Y	N	Y
LDAP	N	N	N	Y	Y	Y	Y	Y	Y	Y	Y
Lists	N	N	N	N	N	N	Y	N	N	N	Y
Web	N	N	N	N	N	Y	Y	Y	Y	Y	Y
Filter	N	Y	Y	Y	N	Y	Y	Y	Y	N	Y
RBL	N	Y	Y	Y	N	N	N	Y	Y	N	Y
Virus	N	N	N	N	N	N	N	Y	N	N	Y
Overall	6	7	5	9	6	8	6	9	6	5	8

# HotPicks

The best new open source software on the planet!



**Jon Kent**  
A Unix veteran  
and Real time  
information  
consultant.

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](mailto:linuxformat@futurenet.co.uk)

## HotPicks at a glance

Ethereal	44
Missile Command	45
bzip 2	45
tcpdump	46
Netsaint	47
dillo	48

## HotPicks award


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



## NETWORK ANALYSER

# Ethereal

■ VERSION 0.9.5 ■ WEB SITE [www.ethereal.com](http://www.ethereal.com)



**Ethereal provides an easy to use frontend to network analysis.**

Ethereal is a GUI based network analyser which uses the *pcap* libs for packet capture, the same as *tcpdump* and many other network tools. *Ethereal* can be used to analyse live network traffic or can read previously captured network traffic saved into a file. This ability to read traffic data from a previously generated file is very useful and means that you can use a utility like *tcpdump*, or *snoop*, to capture the data and then use *Ethereal* to analyse this data off line. It is far easier to analyse a *tcpdump* file use *Ethereal* than it is to try and analyse this file with *tcpdump* itself.

When *Ethereal* is run, you are presented with three panels displaying a summary of the packets received, a detailed description of the packet and a hex display of the packet contents. Highlighting a packet within the summary panel will then display the detailed information and hex display of this packet in the associated panel. Within the detailed description you can drill down into the packet by selecting the + next to the option of interest. If

you highlight one of the lines within the detailed summary this area of the packet is then displayed within the hex display panel. *Ethereal* can also assemble packets within a conversation together in a window by selecting Tools then Follow TCP Stream.

One of the most useful features of *Ethereal* are filters you can apply to live or captured data. When you define a filter you can also set a colour highlight to make it easier to see which packets fit your filter. The filters themselves can range from very simple to the quite complex. *Ethereal* filters allow the use of standard operators, such as == or >=, along with logical expressions, such as && or !. This allows you to create the logic for the filters quickly.

*Ethereal* supports a wide range of protocols and fields contained within them. There are usually individual filters for each of these protocols which are then broken down further to the section or field you want to filter against within the packet. This is both useful and a mild hindrance as you will probably have to look at the help section each

time you want to filter a different protocol until you get used to this method. However, it does provide cleaner analysis. A very simple example of a *Ethereal* filter is :

```
ip.addr eq 172.1.243.59 &&
tcp.port == 23
```

This filter would only display telnet packets if the IP address is matched (which could be configured using a capture filter, which is covered later). With a command line utility like *tcpdump* this rule would be written as :


```
tcpdump host 172.1.243.59 and
port 23
```

This is perhaps a little cleaner for a simple analysis, however *Ethereal's* filters reveal their power with complex filters. By default *Ethereal* will capture all packets on the network being listened to. The data being captured can be modified at start up (using -f [capture filter]) or when the capture is started by adding the filter within the Filter input window. The data capture filter that can be defined at startup follows the same syntax as *tcpdump*.

Aside from standard network analysis, *Ethereal* lends itself to several other uses. Because of the configurability of filters, *Ethereal* could be used as an Intruder Detection System (IDS), by configuring the filters to look for traffic that is considered to be suspicious.

*Ethereal* has a reasonably complete set of preferences for the basic functions of the GUI. You can modify the columns of the various TCP streams, modify locations of scrollbars, and change some basic printer settings. There is also a help section available, although it does limit itself to just listing the available filters and supplies no further information. However, it is useful that it is there to refer to.

Earlier versions of *Ethereal* suffered from crashes, especially when looking at captured data, but *Ethereal* is now stable. Overall an impressive product that is extremely usable, especially when analysing capture data in detail. If you perform network analysis with *tcpdump*, this is a very useful tool and helps to speed up your analysis. Although the filtering is powerful, it is logical and you pick it up very easily.



*Missile command is a good clone of the original 80s Atari game.*

## ARCADE GAME

# Missile Command

■ **VERSION** 0.99.7 ■ **WEB SITE** <http://missile.sourceforge.net/>

**A**tari. That name brings back memories. Get home from school, plug in the Atari console, put the latest Iron Maiden LP on the record player and basically shoot anything that dares

to move. These early console games usually had a very simple concept, but were extremely addictive. A case in point is *Missile Command*, lovingly recreated for all to enjoy and waste a few hours with.

The basic—OK, very basic—idea of *Missile Command* is to protect your six cities from the missiles that are shot down from the sky by an unknown enemy. To do this you use your three missile bunkers, one each side of the screen and one in the middle. These bunkers shoot missiles to intercept and destroy oncoming missiles, hopefully protecting your cities and your missile bunkers. The easiest way to control the bunkers is by using your mouse. The left button shoots a missile from the left bunker, the middle button from the middle bunker and, yes, the right button shoots a missile from the right.

As is the usual patterns of events in this type of game, everything starts off very easy, with just a few missiles to intercept. However, as you progress through the levels, you start getting your screen filled up with missiles, which then split into two or more missiles half way down the screen, which in turn can also split into multiple missiles and things then start getting very out of hand. Your mouse will not know what's hit it.

By this point you're either remembering your wasted youth, or thinking that is sounds just too easy. But do not be deceived, this is one of those games where you find yourself saying "I'll just have one more go", over and over again. It's just plain fun. The other advantage with *Missile Command* and other games like it, is you can have a quick play for five minutes or so, unlike some games where you have to give up an entire month to play. Of course you have to be careful, you may find you have spent more time than you intended.

The game itself uses the *SDL* libs, and uses them to good effect. The game play is very smooth even when you get to the higher and more manic levels. The only drawback to this game is a total lack of any sound. Of course, if you were playing the original you wouldn't be listening to the sound effects as they were usually awful or very bland. It is a shame that this is currently missing, but it's planned to add sound before version 1.0.0 is released. So, until then chuck on your System of a Down CD and just have fun blasting anything that moves.

## COMPRESSION

# bzip2

■ **VERSION** 1.0.2 ■ **WEB SITE** <http://sources.redhat.com/bzip2/>

**Y**ou probably know and use *gzip* as it must be the most commonly known file compressor in the Linux world. It's nearly always installed by default as part of any Linux distribution, is easy to use and has a good compression ratio. As an added bonus, GNU *tar* also understands *gzip*, so you can

uncompress and untar using just one tar command, which is very useful. But this is not about *gzip*, but about its less-known cousin, *bzip2*. *bzip2* is another freely available, high-quality data compressor that you can easily install on your Linux system. Like *gzip* it is also usually installed as a default utility by most Linux distributions.

*bzip2*'s compression is usually far better than *gzip*, which is useful when downloading a large compressed tar file from the web. You may have noticed, if you get kernel sources from kernel.org, that as well as kernel source tarballs ending in *gz* for a *gzip* compressed tar file there are also tarballs ending in *bz2*. Not interesting in itself, until you look at the size differences between these two files. For example, *linux-2.4.18.tar.gz* is 28.7MB in size while *linux-2.4.19.tar.bz2* is 24.8MB in size.

To compress a file using *bzip2*:  
**bzip2 [filename]**  
This will create a compressed archive

ending in *.bz2*, that all there is to it. To uncompress you simply run:

**bzip2 -d [filename]**

GNU *tar* understands *bzip2* files, so run the following to uncompress and untar a *bzip2* compressed tar file all in one command line:

**tar -xvf - -bzip linux-2.4.19.tar.bz2**

The **--bzip** flag is the important one here. This flag informs that the tar file has been compressed using *bzip2* compression. As you can see, *bzip2* is generally better at file compression than *gzip* and is as easy to use. It is worth using *bzip2* to compress files because of the better compression ratio.

## NETWORK ANALYSER

## tcpdump

■ VERSION 3.7 ■ WEB SITE [www.tcpdump.org](http://www.tcpdump.org)

There are times when you need to go back to basics to identify and diagnose problems on your system or network. GUI based tools can help you in this, but there is still a lot to be said for command line tools.

Because *tcpdump* is one of these tools and, indeed, is really one that you should get used to using, the format *tcpdump* uses to save the network analysis to a file is considered to be a *de facto* standard. Any other network tool that wants to be taken seriously can read *tcpdump* files. *tcpdump* was originally created by Van Johnson, Craig Leres and Steven McCanne at Lawrence Berkley National Lab.

## Promiscuous

*tcpdump* allows you to dump TCP traffic to either the screen or a file for later analysis. It works by placing the network interface into promiscuous mode and, because of this, needs to be run by root.

When a network card is put into promiscuous mode, every packet that goes across the network (depending upon the network topology) is captured by the network interface, whether or not it was bound for that interface. *tcpdump* uses an underlying traffic capture library, *libcap*, which

is also used in other tools such as *Ethereal* and *Etherape*, so this needs to be installed as well if you want to use *tcpdump*.

## Defaults

By default, *tcpdump* reads all the traffic from the default network interface (usually *eth0* under Linux) and displays the output to the screen. Of course this is not necessarily going to be of much use, so *tcpdump* includes many command options to change the behaviour and produce output that is more suited to your current problem or query.

The following is an example of a typical *tcpdump* command:

```
tcpdump -s 0 -w /tmp/tcpdump.out
-vvv ip multicast and port 5428 and
dst 224.7.4.5
```

This examples breaks down in the following way. The **-s 0** informs *tcpdump* to capture the complete contents of all the packets it receives. By default *tcpdump* will only capture the first 68 bytes of data per packet, which is usually enough to aid your analysis. However, there are times when you may wish to have the entire packet.

## Raw packets

The **-w** parameter instructs *tcpdump* to write the capture out to a file as

opposed to the screen, which is the default. You may wonder why not just redirect the output from *tcpdump* to a file? While there is nothing wrong with that, it does limit your analysis. This is because, when you use the **-w** flag the raw packet data is written to the file.

You can send that file back through *tcpdump*, or another tool, adding more finely-grained parameters. You can not do this with just an ASCII text file.

The **-vvv** sets *tcpdump* in its highest verbosity mode. How useful this is depends upon the traffic you are analysing. You may be able to use a lower setting.

The next argument's **IP multicast** and **port 5439** inform *tcpdump* to only capture multicast packets originating from IP port 5439. The last argument **dst 224.7.4.5** means that *tcpdump* will only capture packets originating from multicast group 224.7.4.5.

## Operators

You may have noticed the used of the word **and**. *tcpdump* supports primitives and operators such as **and**, **or**, **not**, which allow you to create extremely complex filters for *tcpdump* to use when capturing network data.

This allows you to fine tune any data capturing you are trying to perform. Other useful operators are the **less** and **greater** parameters which allow you to define the size of packet that you wish to capture. Useful if you suspect packet loss

due to a network components dropping large packets for some reason.

Going beyond these simple primitives you can also use fairly complex expressions, such as **'ip[6:2] & 0x1fff = 0'**, which informs *tcpdump* to only capture unfragmented data or data with the fragmentation bit set to **0**. These can get even more complex and are not very readable, but in most cases using the simpler primitives should suffice.

## Documentation

How good the documentation for *tcpdump* is depends upon your attitude to man pages. *tcpdump's* man page is probably one of the longest I've seen, but as you can see there are a lot of options, parameters and operators that can be applied. Nearly everything you need to know is contained within the man page – it's just a little bit overwhelming. If you plan on using *tcpdump*, it is probably a good idea to print out the man page for further reference.

Of course *tcpdump* can be used against you as well, as it can be easily used to capture username and passwords that fly across networks that still use *r-services* (such as *rlogin*, *rsh*) and *telnet*, so you need to be wary of how you install it. It may be tempting to set the *suid* bit on the *tcpdump* binary, but this should not be done for obvious reasons as you do not want anyone being able to run *tcpdump*. Lastly, as has been mentioned in previous *HotPicks* articles, if you want to monitor a switch network you will need to ensure that the switch port you are connected to is in mirroring mode to ensure that you can see all of the network traffic going across the switch.

## Conclusion

Overall, despite its age, *tcpdump* is still one of the best, if not *the* best console network analysers available. Open Source or otherwise. Because of its *de facto* file format it is worth having a look at and familiarising yourself with, as you will probably need to use it at one point or another.


I tend to use *tcpdump* at least once at week trying to pin-point problems and I always recommend that people ensure that it is installed, ready to be used. So, install this little gem as soon as you can, you never know when you will need it.

```
14:49:18.019679 10.5.98.122.35335 > 196.7.183.91.www. . ack 42489 win 63712 <nop,nop,timestamp 1672949980 123752858> (IF)
14:49:18.020971 10.5.98.122.35335 > 196.7.183.91.www. . ack 43933 win 63712 <nop,nop,timestamp 1672949981 123752858> (IF)
14:49:18.118921 10.5.98.122.35335 > 196.7.183.91.www. . ack 45381 win 63712 <nop,nop,timestamp 1672950081 123752868> (IF)
14:49:18.119342 10.5.98.122.35335 > 196.7.183.91.www. . ack 46829 win 62264 <nop,nop,timestamp 1672950082 123752868> (IF)
14:49:18.121530 10.5.98.122.35335 > 196.7.183.91.www. . ack 48277 win 63712 <nop,nop,timestamp 1672950084 123752868> (IF)
14:49:18.125773 10.5.98.122.35335 > 196.7.183.91.www. . ack 48313 win 63712 <nop,nop,timestamp 1672950088 123752869> (IF)
14:49:18.137585 10.5.98.122.35335 > 196.7.183.91.www. . ack 49761 win 63712 <nop,nop,timestamp 1672950100 123752870> (IF)
14:49:18.138456 10.5.98.122.35335 > 196.7.183.91.www. . ack 49773 win 63712 <nop,nop,timestamp 1672950101 123752870> (IF)
14:49:18.222468 10.5.98.122.35335 > 196.7.183.91.www. . ack 51221 win 63712 <nop,nop,timestamp 1672950187 123752878> (IF)
14:49:18.222558 10.5.98.122.35335 > 196.7.183.91.www. . ack 52669 win 62264 <nop,nop,timestamp 1672950187 123752878> (IF)
14:49:18.222629 10.5.98.122.35335 > 196.7.183.91.www. . ack 54117 win 60816 <nop,nop,timestamp 1672950187 123752878> (IF)
14:49:18.223131 10.5.98.122.35335 > 196.7.183.91.www. . ack 54153 win 60816 <nop,nop,timestamp 1672950188 123752878> (IF)
14:49:18.236543 10.5.98.122.35335 > 196.7.183.91.www. . ack 55601 win 63712 <nop,nop,timestamp 1672950202 123752880> (IF)
14:49:18.236744 10.5.98.122.35335 > 196.7.183.91.www. . ack 57049 win 62264 <nop,nop,timestamp 1672950202 123752880> (IF)
14:49:18.236855 10.5.98.122.35335 > 196.7.183.91.www. . ack 58497 win 60816 <nop,nop,timestamp 1672950202 123752880> (IF)
14:49:18.237359 10.5.98.122.35335 > 196.7.183.91.www. . ack 58533 win 60816 <nop,nop,timestamp 1672950203 123752880> (IF)
14:49:18.314979 10.5.98.122.35335 > 196.7.183.91.www. . ack 61429 win 60816 <nop,nop,timestamp 1672950282 123752888> (IF)
14:49:18.356491 10.5.98.122.35335 > 196.7.183.91.www. . ack 67245 win 55024 <nop,nop,timestamp 1672950325 123752888> (IF)
14:49:18.367507 10.5.98.122.35335 > 196.7.183.91.www. . ack 67293 win 63712 <nop,nop,timestamp 1672950336 123752892> (IF)
14:49:18.430026 10.5.98.122.35335 > 196.7.183.91.www. . ack 70189 win 63712 <nop,nop,timestamp 1672950400 123752899> (IF)
14:49:18.430743 10.5.98.122.35335 > 196.7.183.91.www. . ack 70213 win 63712 <nop,nop,timestamp 1672950401 123752899> (IF)
14:49:18.437134 10.5.98.122.35335 > 196.7.183.91.www. . ack 73109 win 63712 <nop,nop,timestamp 1672950407 123752900> (IF)
14:49:18.438455 10.5.98.122.35335 > 196.7.183.91.www. . ack 73133 win 63712 <nop,nop,timestamp 1672950408 123752900> (IF)
14:49:18.467439 10.5.98.122.35335 > 196.7.183.91.www. . ack 76029 win 63712 <nop,nop,timestamp 1672950438 123752900> (IF)
14:49:18.468366 10.5.98.122.35335 > 196.7.183.91.www. . ack 76053 win 63712 <nop,nop,timestamp 1672950439 123752903> (IF)
14:49:18.527453 10.5.98.122.35335 > 196.7.183.91.www. . ack 78949 win 63712 <nop,nop,timestamp 1672950500 123752909> (IF)
14:49:18.528086 10.5.98.122.35335 > 196.7.183.91.www. . ack 78973 win 63712 <nop,nop,timestamp 1672950500 123752909> (IF)
14:49:18.554413 10.5.98.122.35335 > 196.7.183.91.www. . ack 81869 win 63712 <nop,nop,timestamp 1672950527 123752910> (IF)
14:49:18.555414 10.5.98.122.35335 > 196.7.183.91.www. . ack 84074 win 62264 <nop,nop,timestamp 1672950528 123752912> (IF)
14:49:28.746138 10.5.98.122.35335 > 196.7.183.91.www. . ack 84075 win 63712 <nop,nop,timestamp 1672950964 123753929> (IF)
```

*tcpdump* is the ultimate console based network analysis tool.

## NETWORK MONITOR Netsaint

VERSION 0.0.7 WEB SITE [www.netsaint.org](http://www.netsaint.org)



Netsaint's web front end makes it easier to monitor your systems.

When you have a large environment to support you get to a point where you need to be more pro-active about problems that can occur. The best way to achieve this is to install some form of monitoring package that will inform you when pre-defined problems or events occur. There are many software solutions available that can perform this function, and *Netsaint* is a good example of an Open Source solution to this issue.

*Netsaint* provides system and network monitoring that can report matched events to either a mail address or to a web-based front end. *Netsaint* runs as a daemon which periodically runs event tests to ensure that there are no matches. The main *Netsaint* program is actually only a control program that executes external plug-ins that perform the actual system tests. The advantages of the plug-in nature of *Netsaint* is that you can create your own plug-ins to suit your requirements, if there are no plug-ins available that perform the function you require.

The web interface is the easiest method to gauge the health of your environment. To use the web interface you need to have *Apache* installed and to have followed the installation instructions that come with *Netsaint* to configure *Apache* to serve the *Netsaint* pages. When you first access *Netsaint* via the web interface you are

presented with a title screen in the main area of the page and a menu on the left of the screen, which is split into three sections. From the General section you can access *Netsaint*'s docs, which is usually installed as part of the installation of *Netsaint*, and return to the home page of your *Netsaint* installation. The section that you would use the most is the Monitoring section. Within here you have access to all of the main monitoring reports. Lastly there is the Configuration section from which you can view the current configuration of *Netsaint*.

As the Monitoring section is the most informative, it is worth going into a little detail. The first option is the Tactical Overview which provides you with a high level view of the current status of all the system you are monitoring by summarising any errors that have occurred. This is useful as you can quickly see if there are any problems that need to be investigated. The next option, Status Details, lists all the hosts that are being monitored and the current status of the tests that have been performed (OK, Warning, Unknown or Critical). If you need to investigate further, you just click on the status you wish to get more info on, and you are then taken to another page which displays that test in more detail. From this page you can then decide if you want to acknowledge the

problem, disable the test, delay the next notification and so on, and the summary is changed as required (this usually takes a few minutes before the status is changed).


The Status Overview option simply displays a less detailed version of the Status Detail report which, again, allows you to gain more detail of the errors. This simply displays how many tests have passed and how many tests have failed and the status of that failure. The Status Summary option provides another view of the status of your environment, but this time displayed as Hosts Groups, as opposed to individual hosts that the previous report displayed. With *Netsaint*, you can configure hosts to belong to a group of host, so, if you have several servers that provide NFS, you can group all of these servers together into an NFS group, which is useful if you have a large number of servers to monitor. The remaining reports display overviews of problem trends, alert history, contents of *Netsaint*'s log file and so on.

*Netsaint* has many nice features. *Netsaint* can save data between runs (this is the default configuration) so you analyse trends over time. You can also specify whether or not to display the saved status information when the *Netsaint* page is first opened. You can also save the data produced by the status commands for future use outside of *Netsaint*. With the web-based reports the pages refresh periodically so that you do not have to worry about refreshing the page to get the latest view of your system. The *Netsaint* daemon can also be configured to accept data from outside

sources. In this mode, the remote device performs a check and writes the result in a predetermined format to a file on the *Netsaint* host machine.

Installation of *Netsaint* is fairly straightforward although you need to ensure that, as well as downloading *Netsaint*'s source, you also download the plug-ins – otherwise *Netsaint* will have nothing to do. Of course, if you use your distribution's package the plug-ins should be considered to be a dependency. The documentation is very good, and well laid out, which should ensure that you can have *Netsaint* running fairly quickly. The main configuration file is called `hosts.cfg` which is where you configure the hosts that are to be monitored, the host groups, and the tests you wish to be performed on your systems. This configuration file is not the easiest to follow, but once you get the hang of the syntax it's quite easy to add new conditions and tests as you require.

*Netsaint* is a very capable system monitoring program that is fairly easy to configure and, via the support of plug-ins, easy to extend to meet your requirements. The web front end is well laid out and easy to use, although some reports do seem to overlap to a degree. It is worth noting that the next version of *Netsaint* has had a name change and will be called *Nagios*, but, currently, *Netsaint* is the latest stable version. The name change was due to legal reasons which meant that the author of *Netsaint* felt that he should change the name, which is a shame as I like *Netsaint* as a name. All in all a very well implemented system monitoring solution.



More in-depth detail is just a click away.

## WEB BROWSER

## Dillo

■ **VERSION** 0.6.6 ■ **WEB SITE** <http://dillo.cipsga.org.br/>

As good as web browsers such as *Mozilla* or *Konqueror* are, you do pay a price in memory and CPU utilisation with these behemoths. Fairly modern systems have adequate power to cope with these browsers. However, running *Mozilla* can be too much for an under-powered system, where *Lynx* or *Lynx* is generally used as these console-based browsers require a very small footprint. Unfortunately, they are not the easiest to control – nor very pleasant on the eye as most web sites are designed to work with a graphical web browser. If you are in this situation you are now in luck, as *Dillo's* graphical web browser that also has been designed to have a very small footprint. Indeed, the memory requirements are lower than for *Lynx* at startup, and it will run quite happily on a 25MHz 486 with 8MB of RAM. *Dillo* also strictly (sometimes maybe too strictly) adheres to web standards. There are other browsers with smaller footprints than *Mozilla*, *Galeon* being a good example of such. However, unlike *Galeon*, you do not need to have *Mozilla* installed to use *Dillo*. The only requirement for *Dillo* is *GTK+* (1.2.0 or higher)

The usual browser functionality, such as forward and backward histories, home, reload, abort, is all present and correct. As with all other browsers you simply enter the website you wish to view in the URL textbox. To open files you can either click on the underlined F in the toolbar or type its full path followed after entering file: in the URL textbox. Other useful functions include a find dialog which is accessed by pressing **Ctrl-F**. Unfortunately, when a match is found, the word is not highlighted, which can make it difficult to find in the page. Other useful keyboard shortcuts include **Alt-→** and **Alt-←**, to navigate forward and backward through your history, **Ctrl-PgUp|PgDwn|Home|End** will scroll the screen, as will the space bar.

The interface has some other useful additions. These include control panels which may be removed from the main *Dillo* interface and placed elsewhere on screen. If screen space is an issue, you can double-click on the browser to toggle the toolbar between

visible and invisible, thereby freeing up some more viewing space. You are given a fair degree of scope to modify how *Dillo* looks and this works by modifying the configuration file `~/dillo/dillorc`. So, if you do not like a feature of *Dillo*, it is worth seeing if you can disable it.

*Dillo* is still being actively developed and currently supports HTML 4.01, Cookies, XHTML 1.0, PNG, JPEG and GIF image formats, almost complete HTTP 1.0 support and multiple windows. The main features that are missing are SSL support (although there is a patch available), Javascript, and there is no support for frames yet. Both frames and SSL support are being addressed and should be available soon, while Javascript and Java support are currently not planned. Lastly, file

downloads are only directly supported via the http protocol; ftp transfers are currently not supported. Unfortunately the only way to use *Dillo* to download via ftp is to paste the link to an external program such as *wget*. This is a pain, but hopefully will be resolved in later versions of *Dillo*.

Although *Dillo* has support for cookies, this is not enabled by default for 'privacy reasons', however it is easy enough to enable cookies if you want, or need, to use them. In your `$HOME/.dillo/end-bold` directory there is a file called 'cookiesrc'. To enable cookie support you simply change the **DEFAULT** setting from **DENY** to **ACCEPT**. Alternatively you can restrict URLs from setting cookies in *Dillo*. In the file `~/dillo/cookiesrc` you can specify rules for different domains by using the following syntax:


```
DEFAULT      DENY
linuxformat.com ACCEPT
.host.com    ACCEPT_SESSION
```

The first line says that we should deny all cookies from all domains by

default. The second one tells *Dillo* to save all cookies from linuxformat.com across sessions. Finally, the third line allows all subdomains of host.com to set cookies, but these will only be saved in memory until you exit from *Dillo*.

In use, *Dillo* is incredibly fast, much faster than *Mozilla* and *Konqueror*. This is very noticeable when viewing web pages as it displays the page very quickly and you can navigate around a web page with no paging glitches that some browsers sometimes do. Also, no matter how hard I tried, I could not get *Dillo* to crash and increasing the system load did not seem to have an adverse affect on it. For a relatively new project this is very impressive.

*Dillo* is a very stable browser and easy to install and use. Once SSL has been added it will be even more useable as your main browser. So if you are low on system resources, but still want to have a GUI interface to the web, *Dillo* is certainly able to meet your requirement with its low footprint and speed. **LXF**



**Dillo provides a fast and lightweight web browser – a boon to users of older machines.**

SunOne

# SunONE

cover feature



Sun's Open Network Environment sees one of the world's biggest technology companies embrace the brave new world of web services, 'end-to-end' solutions and Linux. But what lies behind this change? Where does Linux fit in? And doesn't ONE sound like something from The Matrix?

**Andy Chappelle** hits the Sun trail in search of answers.

LX50 is aimed at the 'lower-end' of the server market and ships with Linux and a range of applications.

**M**y, how times have changed. In a little over two years Sun Microsystems's attitude to Linux has shifted from barely disguised hostility (the famous "bathtub of code" remark at 2000's LinuxWorld Expo), then cautious acceptance, to chief cheerleader and self-styled biggest corporate contributor of open source code. The company's new Linux strategy is bound up in its all-encompassing Sun Open Network Environment (ONE), which brings together a range of formerly disparate products under the SunONE brand. And rather than just 'being a cog in the works', the company claims, Linux goes to the very heart of the project. Is this a case of embrace and extend – the beloved policy of another US software giant – or has Scott McNealy simply read the market and realised that there's more to life than Solaris?

The inspiration for the move (and the idea that many outfits are 'betting the company' on) can be summed up in two words: Web services. Ironically, this is something Sun have been pushing since the inception of Java, but it's only now, with the 'plumbing' almost in place, the Internet achieving a state of ubiquity, and the correct buzz-phrase identified, that the rest of the industry is starting to jump on the bandwagon. Java, of course, is a significant part of Sun's plans, but SunONE is a 'soup-to-nuts' solution that takes in everything from portal building and app servers to online authentication and wireless messaging.

As with any new technology, SunONE comes complete with a whole herd of protocols, standards and languages (more acronyms to remember!), but fortunately most of these – Java, SOAP, XML, etc – you will have already encountered (see *Acronym Soup*) and are well on the way to becoming official or *de facto* standards. The new fondness for 'open standards' and cross-platform support exhibited among the industry's many players should also ensure that this next generation of Internet technologies doesn't rest in the hands of a single entity.

'Web services' and 'services on demand' are two phrases that pop up constantly as you comb through any

of Sun's literature, but they seem such enormous, amorphous terms that pinning down a concise one line meaning is difficult. Isn't simply the act of looking at a website a service on demand? Isn't adding your email address to a database a Web service? The answer to both is obviously yes, but these two ideas encompass so much more, which is why time, money and energy is being expended by the world's biggest hardware and software companies to make sure their solution doesn't become the Betamax of web services.

### Web future – again

To simplify, Web services is the term used to cover the whole area of device integration into a single network. PCs, servers, palmtops, mobile phones, wireless access devices, applications and infrastructure services all working together transparently. Information, which is what this is all about, becomes just information once again, accessible from the most appropriate device.

Unusually, as this is a Linux magazine, we'll defer to Bill Gates who, in a 'memo' to the computer industry, came up with a pretty good breakdown of what's needed.

**1** A physical medium within which communication takes place: *i.e* the Internet.

**2** The constituent pieces of a Web service carry information about themselves (they are self-describing). There must be a universally accepted way for services to find each other on the Internet. This function is provided by UDDI (Universal Description, Discovery and Integration).

**3** A common language communicating programs can understand. XML has taken on this rôle.

**4** With the Internet being the medium of exchange, an Internet protocol is needed for this exchange to take place. This is the Simple Object Access Protocol (SOAP), created by an industry wide effort.

So, according to Gates, much of the infrastructure needed to make web services a part of every day life is either there or, at least, pretty close. This view was reiterated by IBM's Director of e-Business Standards Bob Sutor at the recent XML Web Services

One Conference in Boston, who said work on the protocols would be concluded within "six to nine months", but that standardisation through bodies such as the Organisation for Advancement of Structured Information Standards (OASIS) and Web Services Interoperability (WS-I), would be ongoing over the next couple of years. The other significant thing in the Gates proposal is that the technologies he mentions are all non-proprietary and do, in fact, form the basis of all the various Web services programmes, making almost everything interoperable – there may not even be a Betamax this time round!

The computer industry, which thrives on obsolescence, has never been reticent about shouting about 'the next big thing', but many industry insiders really do believe that Web services mark the next evolutionary step for both the Internet and personal computing. Hence the struggle to dominate.

The platform which has garnered the most mainstream publicity is .NET which developer Microsoft hopes will give them the same advantage on the web as they currently enjoy on the desktop. However, the playing field is significantly different when it comes to web services due to the nature of the Internet, and the question is: will Microsoft be able to leverage their

**“Ironically Sun has pushed Web services since Java’s inception – now the rest of the industry is climbing on the bandwagon”**

desktop dominance into this new industry or is the market wide open? Simon Tindall, Sun's UK Volume Products Business Manager, suggests that Microsoft have failed to 'articulate a coherent web services strategy' and, ironically, their platform bias is their biggest weakness.

"The marketplace – by which I include end user, customer, reseller, manufacturer, ISV and Government organisations – has learnt the lessons of the desktop, and is actively seeking alternatives to Microsoft." And the real alternative, he says, lies with open



**Sun CEO Scott McNealy has overseen a comprehensive volte-face on his company's attitude to Linux.**





standards and projects such as Linux and GNOME, not following blindly in the footsteps of Microsoft and their emulators. "Hopefully there will never be a need for Mono!" Tindall says in reference to Ximian's attempt to bring .NET to Linux, especially as Sun's own solution shows little regard for architecture loyalty. Their intention is to allow service and content providers to target people not machines.

## Many into ONE

A large part of the SunONE idea is the consolidation of a range of software lines into a single coherent brand, focused on the management and delivery of Web services now and in the future: so while there are a number of notable additions, much of the initiative involves renaming existing applications (some would argue this is the whole premise of web services). *Forté* becomes SunONE *Studio*, *Chillisoft* becomes SunONE *ASP Pages* and *iPlanet* becomes SunONE *Internet Infrastructure*, etc.

There are three main product lines that make up the family.

**SUNONE STUDIO** SunONE *Studio 4* used to be called *Forté for Java* and is the company's native Integrated

Development Environment (IDE) for Java. It is based on *Netbeans* – one of Sun's most significant donations to the open source community – and has been built specifically for developing cross-platform web services. To keep pace with developments in the sector, *Studio 4* retains *Forté's* plug-in architecture, allowing coders to integrate new features as and when they're available or needed. This range is further divided into a trio of editions, aimed at enterprise, the open source community and wireless application developers (the last two being available free of charge) and is complemented by the SunONE *Compiler Collection*, which used to be the *Forté Compiler Collection*.

The *Enterprise* and *Community Editions* of *Forté* have both been covered in these pages (and we gave you the latter on the CD) in the past, so we'll briefly mention the *Mobile Edition* which could potentially have the greatest impact on the way we access our data and services in the future. Like its cousins, the *Mobile Edition* is available for (x86) Linux, Solaris and Windows, and is used in conjunction with the Java 2 Mobile Edition (J2ME). Its primary goal is the rapid creation of low resource Java applets (MIDlets and MIDlet suites) with an integrated, two-step

compilation, pre-verification and execution process.

For testing purposes the suite includes a built-in emulator based on Sun's own Mobile Information Device Profile (MIDP) and Connected Limited Device Configuration (CLDC) reference implementation, complete with a range of skins (see pictures) which give an impression of how your work will appear on some of the most popular devices. Moreover, the Emulator SDK Registry enables developers to easily integrate third-party emulators into the system. The IDE also features automated J2ME JAD and .JAR creation.

## SUNONE INFRASTRUCTURE SOFTWARE

SunONE *Infrastructure Software* used to go under the far easier to recall name *iPlanet* and is the company's collection of middleware solutions. The suite has a comprehensive selection of servers taking in jobs such as calendar/task management, messaging, web proxies, portal design and management, directory services and web servers. The range also includes a rapid application development tool (SunONE Unified Development Server) tailored for network applications and the SunONE *Identity Server* which is the company's solution for managing user access to secure



Sun's entire software line has been repackaged and rebranded under the SunONE name. This used to be *Forté* for Java.

## Acronym Soup

Decoding the alphabet spaghetti.

**SOAP** – The Simple Object Access Protocol is an XML-based lightweight protocol designed for exchanging information in a decentralised, distributed environment. SOAP can be used to deliver services in conjunction with other protocols such as HTTP.

**XML** – Extensible Markup Language has been ratified by the World Wide Web Consortium (W3C) and is a universal format for structured documents. It can be used for designing file formats (see *StarOffice/OpenOffice.org*), UI elements (*Mozilla*) and pretty much anything else. XML is the standard for web services development and, as with SOAP, is common to the various solutions being proposed.

**LSB** – The Linux Standard Base is a specification created to simplify the deployment and maintenance of the operating system. The Free Standards Group have been consulting and devising the specification for the past couple of

years and has just certified products from three vendors: Mandrake, SuSE and Red Hat. Sun have said on a number of occasions that they are committed to upholding the LSB spec in their Linux distribution.

**J2ME** – Sun's Java 2 Platform, Micro Edition. The bonsai tree of Java runtime environments, designed for mobile phones, PDAs and other 'resource constrained' devices.

**MIDP** – This is the Mobile Information Device Profile, a set of APIs which, when used with the CLDC (see below) provides a feature-complete implementation of the J2ME runtime environment. The specification is available from <http://java.sun.com/j2me/docs>.

**MIDlet** – This is a Mobile Information Device Applet and is the basic unit of execution in MIDP. It has been optimised for the Micro edition of Java which can be seen on phones such as the Nokia

3410/6310i, Siemens SL45i/SL42i and Motorola's A008/6288. The MIDlet class extends the Java class **javax.microedition.MIDlet**.

**CLDC** – The Connected Limited Device Configuration is one of two defined aspects of the J2ME environment. Designed for small footprint devices such as mobile phones or low-spec PDAs, CLDC has a pair of VMs, *Hotspot VM* for use in devices with able to devote as much as 512k to Java and another, *KVM*, capable of piling a Java stack and applications into just 160KB.

**Java** – Not actually an acronym but vital to Sun's web services plans, Java is a cross-platform programming language that was conceived to allow developers to get the fruits of their labour working on any machine capable of running the Java Virtual Machine (JVM). Before Flash, it was the best way to add annoying content to your website, but it's capable of much more. Java was

bastardised by Microsoft (with Java for Windows), allegedly, in an attempt to prevent it achieving 'critical mass' and threatening its OS dominance.

**UDDI** – Stands for Universal Description, Discovery and Integration. The UDDI specification enables businesses to quickly, easily, and dynamically find and transact with one another. The UDDI Specification Technical Committee say they are dedicated to maintaining "a global, platform-independent, open architecture for describing businesses and services, discovering those businesses and services, and integrating businesses using the Internet"

**ASP** – Microsoft's Active Server Pages is in the vanguard of the .NET architecture and was a response to the adoption of Java. Native Linux ASP development is possible using SunONE *ASP Pages* (aka *Chillisoft ASP*).

**JSP** – Anything you can do we can do better. Sun's Java-based answer to ASP.

web-based resources (See the *Taking Liberties* box). Much of the application suite has yet to appear on Linux, but the porting process is an ongoing project with early releases including the SunONE *Web Server* (an obvious choice) and the SunONE *Application Server*.

**SOLARIS AND SUNLINUX** The third plank of Sun's strategy is their own proprietary Unix operating system Solaris, and while this remains Sun's premier OS concern, Linux is becoming increasingly important. Simon Tindall says it won't be regarded as a second tier solution. "We plan to port all of the SunONE applications to Linux eventually," he said. The high-profile LX50, Sun's low cost x86 Linux solution, ships with applications including SunONE *Application Server* (formerly J2SE), SunONE *ASP Pages* (aka *Chillisoft*) and SunONE *Web Server*. Tindall said the decision to embrace Linux was based on the expected growth of 32-bit applications on the edge of the network.

"We see the introduction of SunLinux and the LX50 as a significant extension of Sun's product offering, moving increasingly toward a full end-to-end range of solutions, from Data Centre to customer premises. LX50 is primarily positioned to address 32-bit Edge opportunities (web services, IP infrastructure etc) and acting as a compute farm node (in conjunction with *GridEngine*)." Solaris will also play in these areas, he said, but "customers buying 32-bit applications are not running them on Solaris today." The theory is that Linux and LX50 will bring the lower end

market, who traditionally would have gone to Dell, Compaq or IBM for their computing needs, to Sun's door. "Linux is a net gain for Sun," Tindall says.

At the launch of the LX50, Sun CEO Scott McNealy said the product validated the company's commitment to 'choice and an open systems.' "Today we're combining both in a move that delivers as yet unseen value to customers," he said. "Entry level systems are the fastest growing segment of our systems product line and with the Sun LX50 we're turning up the heat while expanding our available market. This unprecedented integration of software, applications and technology creates a new class of Linux and Solaris developers and lowers total cost of ownership for enterprises."

And what of the decision, announced in February and realised at this year's LinuxWorld Expo, to add another name to the ever growing mountain of Linux distributions? Though it's based – almost entirely – on Red Hat 7.2, Tindall says the decision to go it alone rather than partner with an existing product line was primarily to ensure uniformity of support.

"SunLinux allows us to control our own distribution. This means users can expect Sun quality support from a single provider. We have complete ownership of the problem – one throat to choke, if you like – and unlike our competition (IBM, HP etc.) we will not back off support to the distributor," Tindall said, adding that this approach is "already winning corporate mindshare" for Sun.

Taking control of the operating system also means having the final say on the release cycle and the speed of migration to future OS versions. Tindall points to the long migration windows common for Solaris and says SunLinux will follow a similar release schedule, concentrating on stability rather than living on the bleeding edge.

The final, and most controversial reason, is that releasing their own distribution allows the company to optimise and extend Linux for their proprietary hardware, which makes perfect sense when you're first and foremost a hardware company. It may tread, however, on




**There have been exceptions to the rebranding program. How long before *StarOffice* becomes *SunONE Office*?**

## "The unprecedented integration of apps, software and technology creates a new class of Linux and Solaris developers"

the toes of the company whose work you've based your release on.

"SunLinux is a commercial grade OS," Tindall said. "It's optimised for Sun hardware and Sun corporate customers and allows us to develop the highest degree of compatibility between Solaris and Linux." And this is where, potentially, it might get a little murky. SunLinux 5 is essentially RedHat 7.2 (which is released under the GNU General Public License (GPL)), with the addition of device drivers specific to the LX50, a tweaked installer, and a raft of alternative security measures. The company has made much of its LSB compliance but a Red Hat spokesman suggested that SunLinux users will be effectively 'locked' into their distribution as it is the only one certified for the hardware. Tindall claims this method of development is good for users and the wider Linux community. Sun's developers, he says, will take the best open source developments and combine them with the best Solaris has to offer. This could, of course, be of tremendous benefit to Linux if, for instance, something like Solaris's fault



**GridEngine** is a free download for both Linux and Solaris.



management subsystem made it into the public realm. "With 'donations' such as Netbeans and NFS, Sun is the number one contributor to the open source community," Tindall said. "Our intention is to remain in that position, and SunLinux will abide by all applicable open source licenses."

The distribution will be free to download but will only be supported on Sun Hardware.

## The Grid

Sun used their elevated position at this year's LinuxWorld Expo to unveil a pair of products. First was the LX50 server, which won one of the show's innovation awards, and second the latest version of their *GridEngine* software. Grid computing seems to be another technology whose time has come, years after "The Network Is The Computer" was trademarked, and Sun hope to capitalise on this by offering a tightly integrated combination of hardware and software under the SunONE umbrella.

You only have to wonder around the average office or university IT suite after hours to see massive computer power wasted on running screensavers. Tap into all those wasted processor cycles and, grid advocates say, you have the makings of a cheap, efficient supercomputer.

"With a grid, networked resources – desktops, servers, storage, databases, even scientific instruments – can be combined to deploy massive computing power wherever and whenever it is needed most." And this power, says Sun, is just sitting there every day waiting to be used. By Sun software, of course, running on Linux or Solaris.

**CLUSTER GRIDS** are the most common form of grid computer. Typically they will consist of one or more system united to provide a single point of access for users on a single project or department.

**CAMPUS GRIDS** are the next level up and will provide shared computing resources within a number of

departments. Campus grids can consist of physically separate workstations or servers as well as centralised resources located in multiple administrative domains.

**GLOBAL GRIDS** offer the prospect of tapping in the power of a worldwide network and will, as the name implies, consist of resources stationed at various geographical locations. Under Sun's vision, Global grids will allow individuals and organisations to offload their excess requirements onto a grid provider to make the best use of distributed computing. The world could be your renderer, prime number generator and/or impact modeller.

The combination of efficient distributed computing and web services based on open standards will put information back in to the heart of the information age, and users will be free to access their data when, and crucially how, they want to – without being tied to one platform, one architecture or one company. Sounds almost utopian. [LXF](#)

Cluster grids



Campus grids



Global grids



## Taking Liberties? Squaring up to Microsoft?

The Liberty Alliance was set up to devise an alternative to MS's Passport online authentication system. Of course the launch didn't state that goal so bluntly, but that's the aim. It's achieved support from a broad spectrum of high-profile corporations, from Novell, Mastercard and Netscape to IBM, Sony and Sun. In fact 60 outfits have signed to the project.

At the launch of the initiative, Sun's Chief Strategy Officer, Jonathan Schwartz said Liberty Alliance was vital to enable "enterprises large and small to provide their customers with a seamless and uniform interaction with all their business partners – without fear of prying eyes, or costly and proprietary single sign on solutions." It was inevitable then that the upgrade to the SunONE *Identity Server* would include support for the 'the first

set of open, non-proprietary federated network identity specifications.'

Authentication is a key area for Web services, as users, concerned about spam and stories of credit card fraud or identity theft, are wary about entrusting their data with just anybody. Meanwhile, enterprises are seeking out solutions secure, cost efficient and flexible enough to deal with internal and external communications. The Alliance claims a non-profit-making confederation of software, hardware and service vendors committed to a 'network identity infrastructure that supports all current and emerging network access devices' is preferable to single company deciding unilaterally which users and technologies can use the service.

Federated identity through open standards, they say, provides choice.



"Consumers' online IDs, their personal profiles, are self-administered, and securely shared with the organisations of

their choosing." [www.projectliberty.org](http://www.projectliberty.org)  
For more details on Liberty Alliance, turn to page 58.



## Entry level server

# Sun LX50

**Richard Drummond tests out the first product born from Sun's new Linux strategy.**

- **MANUFACTURER** Sun
- **WEB** [www.sun.co.uk](http://www.sun.co.uk)
- **PRICE** see Pricing box

**S**un's new LX50 server is a bit of a departure for the network services giant. Ignoring the server appliances from their Cobalt subsidiary, the LX50 is the first Sun server built on the IA32 architecture rather than their own UltraSPARC processors, and the first to offer Linux in addition to their own commercial

Unix flavour, Solaris. Sun's position is that while for mission critical applications, where reliability and scalability are key, the Solaris/SPARC partnership is unmatched, for so-called 'edge-of-network' applications – web serving, media streaming, firewalling – the Linux and Intel duo offers a significant performance-to-cost advantage. Supported by their Sun ONE strategy, Sun insist that the OS and hardware architecture that your server runs on will become increasingly irrelevant anyway. It is the

application and services that you run on the server that are important.

The LX50 is offered in several configurations, either with one or two Pentium III processors, and is housed in a standard, 1U, rack-mountable chassis. The unit on test was a dual-processor machine with 1GB of RAM and a 72GB disk. Since we're a Linux publication, we opted for one pre-loaded with Sun Linux 5.0 – essentially a Sun rebranding of Red Hat 7.2. The LX50 ships with a range of Sun ONE apps including Sun ONE ASP (*née* *Chillsoft ASP*), Sun Streaming Server, and Sun Grid Engine.

### Nuts and bolts

Considering the speed with which Sun got the LX50 to market, it's not surprising that it's built from the off-the-shelf components. It's based on an Intel SCB2 server board, powered by the ServerWorks ServerSet III HE chip-set. This board is a fantastic piece of kit, and with its three independent PCI buses, optimised for bandwidth.

### Pricing

#### ENTRY CONFIGURATION

1x1.4 GHz Pentium III processor,  
512MB RAM, 1x36GB UltraSCSI drive,  
Sun Linux 5.0/Solaris 8 £2,150

#### TYPICAL CONFIGURATION

2x1.4 GHz Pentium III processors, 1GB  
RAM, 1x72GB UltraSCSI drive, Sun  
Linux 5.0/Solaris 8 £3,300

An ATI Rage XL graphics controller, two Intel 10/100 Ethernet controllers and the ServerWorks chipset live on the first, 32-bit/33MHz bus, while a dual-channel Adaptec Ultra-160 SCSI controller lives on the second, 64-bit/66MHz bus. This second bus also offers a full-size PCI slot for expansion. The third, 64-bit/66MHz bus is for external expansion only via a low-profile PCI-slot and is suitable for attaching a Gigabit Ethernet controller card. The board has six slots for ECC DIMMs and supports up to 6GB of memory. >>

# SunLX50



« The LX50 chassis sports three drive bays, each of which accepts a hotplug-style drive mounting, although the hot-swapping of drives with the power on isn't actually supported. The first two bays are for LVD SCSI drives, while the third can either be used for a SCSI drive or – as in our unit – a combo low-profile ATAPI CD drive and floppy drive in a standard half-height mounting. The Adaptec Ultra-160 SCSI controller impresses, especially if you are used to ATA drives. The test unit had a 10,000 RPM 72GB drive which effortlessly achieved read speeds in excess of 50MB/s. Since this is dual-channel controller, it will scale well for RAID applications.

## “Sun Linux is based on a tried and tested foundation but the package as a whole feels rushed and lacks integration”

Four USB ports – two on the front and two on the rear of the unit – are provided for peripheral connectivity, but a PS/2 port is included for legacy mice and keyboards via a splitter cable. The LX50 has a host of features for remote management and monitoring. The mainboard includes a BMC (Baseboard Management Controller) which supports monitoring, alerting,

and logging from various hardware sensors and, in conjunction with the Intel Server Control software, provides an IPMI-compliant solution for remote management. The unit also offers dual RJ-45 serial ports for out-of-band management. Only one of these ports may be used at a time (they are essentially the same physical port). The one on the front of the LX50 is wired for connecting terminal or console management equipment, while the one at the rear is for modems, etc. The LEDs on the front of the unit include the normal power, drive activity and network activity indicators – as well as a status LED and a system ID led. These are great for locating a faulty

unit when you have racks of the beasts stacked up together.

### Sun Linux

Sun Linux 5.0 is based on the solid foundation of Red Hat 7.2. They have changed very little, apart from upgrading the kernel to 2.4.9 and optimising it for the LX50 hardware. The configuration errs on the side of

conservatism and, for example, no ACPI support is built in by default. The versions of software provided look dated, too, and – more worrying – no facility for automatically downloading and installing patches and security fixes is included.

Although Sun Linux comes pre-installed, it is supplied on three CDs should you ever need to re-install. We gave this a try, to ensure that our machine was configured to factory defaults for testing, and to evaluate the process. The installer only performs a complete re-install; if your disks have any valuable data, you'll need to back these up first. The installer is incredibly simple, being a streamlined version of the standard Red Hat installer, and only requires that you confirm the installation process. No package selection is required.

Sun Linux includes the full-range of software that makes up Red Hat 7.2. This includes valuable services such as *Apache*, *MySQL*, *BIND*, etc. but also includes software that you probably don't want on a server, such as multimedia applications and games. When you have 72 gigs of disk space to spare, wasted space isn't much of an issue, but Sun could have boiled down the essential software to one disc's worth, and reduced the install time by a factor of three.

Once installed, Sun Linux provides the usual range of Red Hat tools for administration. There's *kudzu* to take care of hardware and initial network configuration, *apacheconf* for *Apache* administration, and so on. As with Red Hat, there are no dedicated tools, such as *Webmin* – supplied for easy remote management. The Sun Cobalt *Control Station* software from Cobalt's server appliances – which provides an easy-to-use, web-based administration system – is promised for the LX50, but was not available at the time of review.

### Sun ONE platform

The suite of Sun ONE applications bundled with the LX50 gets copied as archives into the */opt* directory by the


**Despite its compact size, the LX50 offers a surprising amount of potential for expansion.**

installer, and it is up to you to install them manually. It doesn't even install a Java run-time for you. Although installation is not difficult if you follow the instructions, this is not exactly the integrated solution promised.

The LX50 provides a complete web serving solution backed by *Apache*, the *Tomcat* servlet/JSP engine and Sun ONE ASP. Sun ONE ASP is a rebranding of *Chillware ASP* and provides server-side VBScript and JScript scripting for your web applications, and support for *FrontPage* extensions. Once installed, Sun ONE ASP may be managed with the web-based management console included with the package, and this should be instantly familiar to existing users of the Cobalt range. The documentation and examples provided with Sun ONE ASP are excellent, and it shouldn't take much effort to port applications from a Microsoft server platform. The LX50 includes a full, unrestricted licence for Sun ONE in the price.

The Sun *Streaming Server*, on the other hand, is an evaluation version only. Like Sun's ASP server, this must be installed manually, but once installed it can be configured via its web interface. The Sun *Streaming Server* supports *Quicktime* and *MPEG4* media, but makes use of open standards such as RTSP (the Realtime Streaming Protocol) for delivery.

The most intriguing application of the set is the Sun *Grid Engine*. This was formerly developed by a company called Gridware and was acquired by Sun in 2000. Sun later open-sourced the *Grid Engine* project in 2001. *Grid Engine* is a distributed resource management system. Basically, it is a system for managing and using clusters of computing resources distributed over a network. Think of it as something like the distributed RC5 or SETI projects that have been



**The Sun ONE ASP server will prove a boon for those converting from Microsoft platforms.**

popular in recent years, but you can actually put it to useful tasks. Each client in the grid makes resources such as memory and computing time available to grid, and the grid management software lets you describe and schedule jobs to be executed on the grid. The Sun *Grid Engine* provides a kind of distributed shell environment for running batch jobs (although interactive jobs can also be run) and a *Motif*-based administration tool, *QMON*, is provided to submit and manage jobs. Jobs are generally submitted as shell scripts, but this doesn't limit applicability since the shell script could, say, build this source and then execute it.

## Management tools


One of the most disappointing aspects of the software provided with the LX50 is the lack of management tools. If you read Sun's marketing blurb then you see claims stating the LX50 supports open standards such as DMI (Desktop Management Interface), IPMI (Intelligent Platform Management Interface) and SNMP for remote monitoring and management. While this is true, none of this works out of the box, so to speak. An evaluation version of the Intel Server Control (ISC) software is included, but you'll have to set this up yourself (which involves building the necessary kernel module). This software is packaged very poorly and the docs vague, so all this is not as

straightforward a task as it might be. In principle, ISC should let you remotely monitor CPU temperature, voltages, fan rotation speeds, etc., and let you perform asset management – but the client part of the equation is missing. You'll need to buy a commercial server management system such as *Sun Management Center*, Intel *LANDesk Server Manager* or HP *OpenView Network Node Manager*. The *Cobalt Control Station Software* will also let you manage IPMI-capable servers such as the LX50.

While it is good to see that standards such as IPMI are supported, it would be better if Sun or Intel were to support existing open-source projects that are attempting to bring such monitoring tools to Linux, such as the *Im\_sensors* project ([www2.im-sensors.nu/~lm78](http://www2.im-sensors.nu/~lm78)) and the *MontaVista*-sponsored IPMI driver (<http://home.attbi.com/~minyad/>), rather than expecting users to rely on half-hearted and poorly maintained Linux ports.

## Unfinished

The LX50 is a hard product to judge fairly. The hardware is excellent and packs a terrific amount of computing punch into that 1U chassis. It's just a shame that it isn't backed up a similar quality in the software. Don't get me wrong. Sun Linux is based on a tried and tested foundation and should prove perfectly robust – and the Sun ONE application suite bundled with it




**QMON is the Sun Grid Engine's GUI for submitting and controlling distributed computing jobs.**

features some great technology, especially the ASP engine. But the package as a whole feels rushed and lacks integration. It's worlds away from the appliance nature of Cobalt's servers. For instance, why couldn't the Sun ONE applications be ready to roll right from installation? Why couldn't they all be distributed as RPMs? And why do they have to be installed separately? Also, apart from the ASP server, these applications don't add a significant amount of value to the package. The *Streaming Server* is only an evaluation edition and the *Grid Engine*, *Tomcat* and *JDK* are all available as free downloads anyway. The other great disappointment is the lack of management tools. While the unit features a Board

Management Controller, and can potentially support IPMI and SNMP for remote monitoring, again this is a bit of a hotch-potch and requires manual setting up. Why is the ISC software not installed and configured for you by default?

As it stands, the LX50 offers good value for money, but it doesn't really have much to distinguish it from any other similarly spec'd server. Of course, you'll be backed up by Sun's word-class technical support, but if the same hardware is available more cheaply elsewhere, you'll have to weigh up whether that support is worth the extra cost. Do remember, however, that Sun Linux is new and a certain amount of teething problems are to be expected. All of the problems I have mentioned should be easy to rectify if a little more attention is given to quality control. When the Control Station software becomes available for the LX50, and when some other niggles are ironed out, the LX50 will be a lot more attractive proposition. **LXF**



**Lurking underneath the cool, blue bezel are the drive bays, status LEDs, and USB and serial ports.**

## LINUX Format VERDICT

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

A lot of performance for the price, but Sun Linux and the software bundle feel hastily thrown together.

## LINUX Format RATING

7/10

# What on Earth is...

# THE LIBERTY ALLIANCE?

**Dr Tim Parker** forgets his passport, and signs up to the Open Source authentication service.

## »» What is Liberty?

The Liberty Alliance Project is a group composed of many industry players who are developing an open-source alternative to Microsoft's Passport system. The general concept is that a user will be able to log into one service (such as a website) with a login and password. When that user moves to another service (such as another website), the login information is transferred automatically, eliminating the need for the user to log in again. As long as the user moves from service to service within Liberty-enabled 'circles of trust', the logins are automatic. When the user logs out of a site, all the other logged-into sites can be automatically logged out as well. As a side benefit, Liberty provides authentication of the user to services, and *vice versa*, preventing identity misuse. For more information, visit the Alliance's website at [www.projectliberty.org](http://www.projectliberty.org).

## »» How does Liberty compare to Microsoft's Passport?

Microsoft introduced .NET Passport in 1999. The goals of Passport were to implement a single sign-in ability (one login and password for all Passport sites) and to provide express purchasing abilities using the Passport 'wallet' which contains billing and shipping information as well as single-click authorisations for

credit-card charges. The Liberty project includes both of these goals, implementing them with an open source as an alternative to the Microsoft-controlled Passport service. The Liberty project will collect more personal data and demographics than the Passport service, which may bother some users.

## »» So is Liberty just a non-Microsoft version of Passport?

In its early stages, yes, that would be a valid assessment. As the Liberty project proceeds, though, there will be more features added to Liberty than are currently offered by Passport (although Microsoft is bound to respond to the challenge if Liberty starts to gain popularity). Liberty is open source, open platform, while Passport is not, and that's the big difference right now.

## »» Who makes up the Liberty Alliance?

Right now, there are about 70 members in the Alliance, mostly large conglomerates who foresee an alternative to Microsoft's Passport as a good idea. Current board members include American Express, AOL Time Warner, General Motors, Hewlett-Packard, Sony Corporation, Sun Microsystems, and many other large corporations. Membership for sponsors of the Alliance is US\$120,000 per year.

## »» What is a 'circle of trust'?

The Liberty Alliance defines two circles of trust: enterprise and consumer. A circle of trust is simply a set of sites or services that use the Liberty technology and which have trust relationships with each other. The enterprise circle of trust deals with vendors and service providers, tying them together with trust relationships. The consumer circle of trust can be defined by the user, and indicates which users or services are to be trusted. Circles of trust can be established manually or automatically by the service. *E.g.*, you might create a login with one site, such as an online bookstore, which then asks if you want to create a circle of trust with all affiliated sites that sell other products or services. Now you don't need to manually establish a login with each affiliated service.

## »» What if I already have different account details on different sites? How will this affect my Liberty login?

The Liberty Alliance knows that users don't use the same account name or number and password on each site. These isolated accounts, called 'local identities', can all be amalgamated under a single Liberty login. This lets one site recognise different account details on another site. When you move to a site that has a local identity but also implements Liberty, you will have the option to tie together your local login info with your Liberty info. Where the Liberty project really becomes useful for users is when you access a service for which you have no established local identity, but your Liberty identification can be used immediately.



Personal  
**NETWORK  
IDENTITY**

Your Guarantee  
of Security

ADDRESS  
[Blurred]  
**LOCKED**

**MEMBER LOGIN**

USERNAME: [\*\*\*\*\*]

PASSWORD: [\*\*\*\*\*]

**GO**

629711300579186259983764

**SAFE**  
Online Shopping

**SAFE**  
Online Banking

UNIQUE ENCRYPTED CODE

COMBINATION OF INTERNET PROVIDERS


INDUSTRY-WIDE COLLABORATION

SECURE AND TRUSTED ARCHITECTURE

Jackie - would you like to federate your hotel account with your car rental account?

**YES** **NO**

# WhatOnEarthLibertyAlliance



## »» What kinds of information will be used in Liberty logins?

The Liberty project sees personal info as more than just a login and password. The Alliance refers to your account as a 'network identity' and it will contain many different attributes – some mandatory, some optional, and some built up with usage of Liberty-enabled sites. Inside the network identity is basic info such as your name, address, telephone numbers, and identifying numbers such as a social security number in the US. In addition, the network identity will manage specific information about your credit records, payment history, and, potentially, info about health issues, marital status, and so on. A lot of this info is not what users will want to have available to

vendors, so security and confidentiality is going to be an important aspect of Liberty. There is the potential for Liberty to become a major storehouse of personal info that could, potentially, tie many single sources of information together. Most people would not want this kind of info to be available in one location, especially one managed by public for-profit companies, so acceptance of this type of info warehousing will be an uphill battle for the Alliance.

## »» What platforms will Liberty run on?

The Liberty Alliance has set no limitations on the platforms Liberty will run on. Because the software is open source, it can be easily ported. The Liberty specs are deliberately designed to allow maximum portability regardless of OS, network infrastructure, and programming language used for service apps.

## »» Is there a set of specifications for Liberty ready now?

The Liberty Alliance released version 1.0 of their specs in July 2002. These specs deal with

interoperability between systems, with opt-in account linking and single sign-on features. The specs for version 1.0 do not include any exchange of personal info between sites, only validation of account logins. While these specifications are only a fraction of what the Liberty Alliance sees for its end product, they are a start, and the Alliance predicts the first Liberty-enabled websites will appear late in 2002. The next set of specs, which will add many of missing features to 1.0, is supposed to be released in early 2003.

## »» Why would a service provider choose Liberty over Microsoft's Passport?

Passport is essentially Windows-based, and is controlled by Microsoft. Liberty is platform independent, and is controlled by a consortium. Having said that, Liberty is still a nascent product at least a year or two away from full implementation, so most vendors are going to be slow to adopt Liberty right away. Passport, on the other hand, is in widespread use already. The open-source aspect of Liberty will appeal to some users and vendors, but public acceptance will be the real test of whether Liberty will succeed or not.

## »» How will Liberty work?

The actual process depends on how Liberty is implemented. The easiest scenario has three bodies involved: a user agent, the service provider (the site the user is trying to access) and a Liberty identity provider that establishes the user's identity.

When using HTTP redirection as a mechanism, the process involves all three bodies communicating via HTTP **GET** requests. Reduced to simple terms, when a user logs into a Liberty-enabled website an HTTP **GET** request is sent from the user agent to the service provider, which returns a redirect (status code 302 message) with an alternative URI belonging to the Liberty identity server in the Location field of the response. The user agent then sends an HTTP **GET** request to the identity server with the service provider info contained within that message. The identity service returns the validation info to the user agent, with the service provider's URI embedded. This is forwarded back to the service provider, which allows the login to proceed.

If Form-based **POST** redirection is involved, the process involves the service provider responding to the user agent through a **POST** ed HTML form, containing a parameter pointing to the identity provider and a script causing the next step of the process to be performed automatically. The form is sent to the identity provider using a **POST**, which returns the validated user information that can then be sent back to the service provider.

## »» How does single sign-in (SSI) work with Liberty?

When you visit a Liberty-enabled site, you will be provided the option to use the Liberty system. If you

choose to do so, you will provide personal info and create a Liberty login and password (which may be tied to the site's local login information as well). After completing the information, you will be offered the choice to have this information passed to all Liberty-enabled sites in the current site's circle of trust. When you visit another Liberty-enabled site, you can use the Liberty SSI information or a local sign-in to access the site. You may be presented with the option to tie your local sign-in to Liberty at that site. The way Liberty is designed, you cannot choose which affiliated sites your SSI information will be passed to: all sites in the circle of trust will receive your sign-in information.

### » Can I remove my login from a particular Liberty site if I want?

The Liberty specifications do provide for 'defederation' of a login. To do this, you need to initiate a defederation process which will remove your sign-in information from that site, without affecting your sign-ins with other Liberty sites.

### » How does a single sign-out work?

The Liberty specifications define a Single Logout Protocol that allows you to log out of all sessions that have been started since your first sign-in to a Liberty-enabled site. The logout process can be started either by the identity provider (which verifies your login) or by the service provider (such as a website). When a logout is initiated, the identity provider sends logout requests to every service provider that has a current session with the user's ID.

### » Are cookies used by Liberty?

Cookies are used to maintain HTTP state information. According to the standard that defines the use of cookies (RFC 2965), cookies are created by web servers to store local information on a client machine, and should be read only by the web server that wrote them to the client. Ensuring only the writing site can read the cookie is handled by DNS, which is easily subject to hacking. The Liberty Alliance is discouraging the use of cookies as this will require users to lower their security settings, something they deem unacceptable. Also, since many users and organisations disable the use of cookies, Liberty should work without cookies. While the Liberty specifications do not prevent the use of cookies for state information or storage of persistent details on a client, and the specifications do provide for cookie maintenance, the cookies should not be used for Liberty project-related information that could be used to gain illegal access.

### » Can I use pseudonyms with Liberty?

Part of the design specification supports the use of pseudonyms for users. You can use any pseudonym registered with Liberty services to access an account.

### » What about security?

Liberty is implementing both user confidentiality and message data integrity as part of its basic security using public-private key encryption (one of the Alliance members is RSA Security, which begin the whole PPK process). In addition, peer authentication is required as part of the Liberty implementation. Encryption and authentication protocols are used to encode data and verify sender and recipients for all data transfers.

Microsoft's .NET Passport service uses Secure Sockets Layer (SSL) and Triple Data Encryption Standard (3DES) for protecting data. To enforce information privacy, Microsoft insists that all sites that use Passport sign a contract indicating adherence to a set of privacy guidelines.

### » How does a vendor become a member of a Liberty-enabled circle of trust?

Obviously, you don't want to allow your Liberty information to be transferred to sites willy-nilly. To become a circle of trust member, a vendor has to establish bilateral agreements with other vendors, as well as arrange sharing of credentials, certificates, and trusted public keys. The process is not difficult but will be rigorous enough to prevent fly-by-night vendors from obtaining account information for misuse.

### » As a service provider, how would I implement Liberty or Passport?

To provide Passport services on your Web server, you need to sign the .NET Service Agreement (which includes a base service-level agreement or SLA). The Service Agreement specifies business guidelines intended to protect all Passport-supplied info from misuse. In addition, the SLA also requires privacy statements to be posted on all participating sites, as well as strongly recommends membership in a privacy-assurance body (such as TRUSTe). To date, the Liberty Alliance has not released an SLA of their own, specifying what a service provider must implement and how they must protect info. While SLAs are being developed by Liberty, they have not been released with the initial version of the specs.

### » How soon will we see Liberty-enabled websites?

Not right away, that's for sure. Although there are optimistic predictions of Liberty-enabled sites by the end of 2002, the first few sites may not be major service providers but smaller vendors who do not want to subscribe to Microsoft's Passport system. As of writing, the Liberty Alliance has not set up the infrastructure yet to allow a site to sign on and implement Liberty. Although the first specifications are now available, they do not include conformance info – just implementation details. In short, Liberty is not mature or tested enough for roll-out yet. Most analysts think Liberty won't be a real-world product until their next spec release is available (early 2003)

and the server implementation software is readily available. Microsoft's Passport, on the other hand, is readily available and can be implemented immediately and without too much effort.

### » What will the 2.0 specifications add to Liberty?

According to the Alliance, the next release of the specs will expand on the single sign-in features, as well as implement sharing of personal info between circles of trust. The Alliance says it will allow the user to set which information is shared, based on preferences and permissions the user creates.

### » Is Liberty designed specifically for Linux?

Although Linux servers are likely to be the major implementers of the Liberty project, the Alliance is very careful to avoid making Liberty look like a Linux project. By touting Liberty's open source code and the ability to run on any OS, the Liberty Alliance is obviously trying to not only target Linux, but also horn in on Microsoft's Passport operations.

### » Is Microsoft improving Passport in answer to the Liberty project?

Microsoft has had an upgrade plan for Passport since its introduction, well before Liberty made an appearance. However, it would be fair to say that the appearance of Liberty may make improvements and changes to Passport more likely. In the near future, Passport will implement better security features, and a new authentication system named 'TrustBridge'.

### » Can Liberty and Passport coexist on a Web server?

In theory, yes, both single sign-in services could be provided on the same server, assuming compatibility with the operating systems (Windows only for Passport). There is no conflict between the two services implemented side-by-side.

### » So, is Liberty likely to become a standard for Web single sign-ins?

It's way too early to predict the success of Liberty. It's likely to gain strong support among the Linux server community, but public acceptance is harder to gauge. The general public, who doesn't really understand Linux, and who runs Microsoft OSs, is likely to want to stay with the Microsoft Passport system, which has been used for three years. The technically savvy user may see the benefits of Liberty, and may well choose to use it simply because it's not a Microsoft product. Liberty's success will depend on a good set of specifications, assurances to the public that personal information will not be spread willy-nilly amongst participating websites, a good set of server implementation tools, and some successful site participants. Liberty's success is by no means guaranteed despite the powerful participants in the project. Only time will tell. [LXF](#)

# Emulators



# MESSy business



**Simon Goodwin rounds up his emulation series with a big MESS.**

**T**his culmination of our tour of emulators that run on Linux focuses on *XMESS*, which aims to emulate more than two hundred types of home computer, and mentions emulators as yet unreviewed. While you'll have to delve into the coverdisc – and perhaps Google – for full details of these systems, this article explores the scope and capabilities of *MESS*, a spin-off of the *MAME* arcade emulation project, and several alternatives.

## Making a MESS

Like *MAME*, *MESS* is a big package. You'll need 80MB of free disk space to compile and link it. Most Linux systems should accept

the following incantation:

```
tar -xvf xmess-0.561.tar.bz2
cp makefile.unix Makefile
make
```

A 500MHz AMD K6-2 took almost 15 minutes to make a default X11 version on my system. Makefile tweaks boost performance and enable more architectures and distros. *XMESS* targets X86, 68K, and RISC systems like Alpha, PPC, MIPS, SPARC and HP-PA, and tentatively supports Intel's C++ compiler as well as *GCC*.

*ESD*, *ALSA* and *SDL* sound are optional besides the usual *OSS*, and various joystick options. The display can use a flexible but slow X11 window; or *Glide*, *SVGALIB* on PC hardware, XFree86 DGA, *SDL* or *OpenGL* drivers. If you get *zlib* error reports, uncomment **ZLIB=1** to use a version packaged with the emulator.

The superuser command **make install** copies the *xmess.x11* binary into `/usr/local/bin/`. That's close to 6MB long, though less than half the size of its sibling *xmame.x11*. The man page installed is for *MAME* and doesn't mention *MESS*, though the command line interface is very similar. *MAME* (*LXF26*) is the focus of development and *MESS* is a poor relation.

Linux loners may prefer to override the default ROM location `/usr/local/lib/xmess/` via the config file. This is not made automatically, though **make install** creates a directory for it. To extract settings in the required format, type:

```
xmess.x11 -showconfig > ~/.xmess/xmessrc
```

Then edit 'xmessrc' to specify where the ROMs on your system live – I used **rompath /home/simon/mess**.

Past columns enthused about emulators for mainstream 6502 and Z80 micros. While *MESS* supports those, most such systems are better supported by dedicated emulators, so the



**MESS got this far into TI emulation before falling over.**



**MESS menu overlay above these Colour Genie UBoat graphics.**

accompanying boxes focus on more eclectic *MESS* features, like the Motorola 6809 and Texas TI9900 processors, sparsely supported by Linux.

## Key points

*MAME* hijacks computer keys for arcade setup options but those are needed for computer control in *MESS*, so the **ScrollLock** key switches between partial emulation, when the *MAME* setup keys work, and the rest are sent to the emulation and full emulation when the user interface is blocked. For instance the **P** key pauses – as in *MAME* – in partial emulation when menus are available, and enters a letter if the menus invoked with **Tab** are disabled.

Keys are configurable but you're liable to end up with inconsistent setups; if you chance to press **Enter** twice after selecting the long scrolling 'general input' menu you clobber the assignment that got you there, and hence access to all *MESS* menus. At last I resolved this by deleting the `~/xmess/cfg/default.cfg` file.

Wrapping between top to bottom makes the scrolling menus easier to use, but I'd prefer shortcuts or the chance to type the path or filename I wanted rather than scroll and select from clumsy menus.

Input is annoyingly modal – if you press **Esc** while in menu mode you return to the emulation but if **ScrollLock** has toggled since then, **Esc** closes down the emulator and dumps you back in the launch shell, losing the setup and anything entered.

## Vectrex

If you're looking for one reason to run *XMESS*, it should be the Vectrex driver. This unparalleled system offers vector graphics almost as smooth and not as flickery as the real thing.

The Vectrex home vector graphics system featured a built in mono screen, with coloured filter overlays like early arcade games, and its games are now freely available. It is a throwback to the very first computer games like *SpaceWar* on the PDP1 – emulated by *MESS* if you can guess how to start it – updated to run on a 1.5MHz 6809 with AY8910 and DAC sound.

Most computer displays are built up like TV pictures, scanning the entire screen systematically line by line, displaying dots at fixed positions in a rectangular grid or 'raster'. Vector graphics computers control the display beam directly, sending it around the screen in any sequence like a laser beam, drawing continuous outlines rather than discrete dots. Vectors need little display or pattern memory, and objects can be moved or scaled without jumps in position and resolution. Rotation, reflection, stretching and squashing are trivial. But the more vectors there are the longer the system takes to re-paint them all, so screens flicker and dim as images becomes more complicated.

You can simulate vector graphics on a high resolution raster, though the results are never quite as smooth or subtle as a real vector scan image. *MAME* emulates the classic arcade vector games like *Asteroids*, *Battlezone* and *Tempest*, but *MESS* offers more through its support for Vectrex and Raaspec.

## AdventureVision

The bizarre Entex Adventurevision console used 40 LEDs and a spinning mirror to generate a low resolution flickery red display. This is worth seeing, even in emulation, if not something you'll return to often.

*XMESS* emulates this in the centre of an X window, and the eight-bit 8048 processor, similar to the QL I/O chip, but without

## Enter the Dragon

### 6809 systems

Motorola's 6809 is the ultimate eight bit microprocessor, but development took so long that the 16-bit 68000 overtook it and it missed the 1970s boom. Even so it was used in many arcade systems like *Defender*, the matchless Vectrex console, and home micros from Tandy, Dragon and Thomson. This box explains how to emulate the latter on Linux.

*XMAME* has a good 6809 core and Motorola video chip emulation is a key *MESS* component. Motorola did not manufacture micros themselves – at least till the StarMax Mac clones – but published a reference design based on the 6809, 6847 Video Generator, 6883 memory controller and two 6821 ports. This was adopted by Tandy, seeking a colour followup for their TRS-80, and Dragon – a Welsh company founded by toy-makers Mettoy, bought by GEC, then sold to Eurohard in Spain.

The most significant difference between Dragons and Tandy 'CoCo' Color Computers is that Dragons have a parallel port. Programs are easily converted though keymap and BASIC tokens vary, along with the amount of RAM and ROM. Tandy's original CoCo had a tiny BASIC and just 4KB RAM, but later models boosted both, and graphics too.

The first Dragon had 32KB RAM and 16KB ROM; Dragon 64s doubled the RAM to run OS/9, the Unixoid system that revealed the true strengths of the 6809. *MESS* also emulates the original CoCo, the expanded models 2 and 3, and Tandy's cut-down MC10, all with authentic six-bit DAC sound. Press ~ with emulator menus enabled, then adjust the volume in 1 dB steps with arrow keys.

### Loading

*MESS* emulates DragonDOS and slow loading from '.cas' files, using menu selection and an on-screen tape counter. I was not able to get any '.pak' ROM cartridges to work, but the **CLOAD** command loaded BASIC games like *Transylvanian Tower* or *Mined Out*, and **CLOADM** read Dragon or CoCo code files like *Scarfmán*, *Shock Trooper*, *Phantom Slayer* and *Return of the Jeti*; **EXEC** starts code after loading.

Select Tape Control and the Rewind item repeatedly to wind back if you get a filetype error when you try to load with **CLOADM**, and try **CLOAD**, or vice versa. Files like *Arcadia* use a short BASIC loader, then **CLOADM** the code and start it with **EXEC**. Tape Control doesn't work till



The Dragon 32 was a Welsh relative of Tandy's Color Computer.

you have selected a '.cas' image with the File manager.

If code tapes crash after loading, try again without the *ddos10.rom*, only needed

for disk images. Reset (**F3**) is a bit erratic and sometimes you must reload the emulator before you can load another tape.

*MESS's imgtool* manipulates cassette and disk images. It has good Dragon and CoCo support, and can also detokenise BASIC.

There are no dedicated Dragon or Color Computer emulators for Linux so this *MESS* setup fills a significant gap, but it is crude compared with many emulators dedicated to particular systems; *MESS* emulation is rarely the best choice if there is a bespoke alternative.

### Les Thomsons

There are several dedicated 6809 emulators for Linux, but these mimic more obscure systems. *Sim6809* is a processor-level simulation, while *FlexEmu* emulates a DOS that is to the 6809 roughly what CP/M was to the Z80. Both are on the coverdisc, plus an open-source 6809 assembler and *SBC09*, another cross-development system.

The French appreciate good design and got into micros a bit later than English-speakers – the Dragon and CoCo lack lower case letters, let alone French accents – but the big French electronics firm Thomson made a neat range of 6809-based micros through the 1980s, and Linux emulators for those are on the coverdisc.

The Thomson 9000 or T07 occupied much the same position in French schools as BBC Micros did in their British counterparts. *EmuT07* emulates this in an X window, with the mouse standing in for the 'crayon optique' or lightpen, a standard feature used in Thomson software. The 24KB T07 was followed by the T0770, with more RAM and colours, and has a derived emulator, *EmuT070*.

A later open source project brought about *THOM* and *TE0*, capable T0770 and T08 emulations which can use an 8, 16, 24 or 32-bit display with sound, joystick, cartridge, tape and disk support. Finally *XemuL5* emulates the Thomson M05 home system. Source for this is derived from *EmuT070*, and it comes with some English, as well as French, HTML documentation.

sound. The flickery red display was barely readable and appeared on top of the menus rather than in place of them. The **F4** key clears up the small part of the middle of the screen where the action occurs, and I got a helicopter variant of horizontal scroller *Scramble* to work reasonably well. Adventurevision expects the cartridge name at the end of the launch command:



# Emulators

## Texan Instruments

From TTLs to 16-bit LSI

### TI99/4

US chip fabricators Texas Instruments invented the TTL (Transistor Transistor Logic) integrated circuits which were the mainstay of small computers before the microprocessor, and the glue that held early micros together before custom-chips. Even Motorola's 6809 processor and Amiga custom chips were prototyped in TTL, spread across several tables.

But as Sinclair's ZX81 replaced a *smorgasbord* of ZX80 TTL with one Ferranti gate array that did far more in less space and more cheaply TTL became obsolescent. It's still used for prototyping and small projects, but PROMs, PALS, FPGAs and LSI chips dominate modern computers.

When microprocessors and other Large Scale Integrated chips arrived, Texas produced rivals for Motorola parts. They developed 16-bit minicomputers, and licensed sound and graphics chips but, unlike their rivals, TI set out to make and sell home computers directly to customers at the end of the seventies.

The TI99/4 was the result – the first 16-bit home computer, with exceptional colour graphics and sound for its day. There are three TI99/4 emulations for Linux.

### TI99sim

*TI99sim* comes as an X86 binary or source; **make** produced code to suit my Debian 2.2 system, after editing the installation paths at the top of 'Makefile.linux'. *TI99sim* eschews the raw '.bin' ROM files that suit *MESS* and other emulators, but a tool converts TI memory contents into its hybrid '.ctg' format in the roms directory. So **convert-ctg 994a.bin** reads the 24KB *994agrom.bin* and associated 8KB *994arom.bin* from a *MESS* ROM set, making a combined file about 32KB long.

*TI99sim* can use *SDL*, full screen or in a window, or run as a console application showing processor registers alongside text from the original computer's screen. You can edit register values when this executable is paused, and explore the TI9900 processor. This elegant minicomputer architecture was far in advance of other micros from the seventies, but RAM bottlenecks and losses in the 1983 micro market meltdown stifled development.

The *SDL* binary has graphics and runs in a small 256x192 pixel window that can be scaled up to 512x384 pixels or full-screen. Games can overlay up to four sprites on top of any line of the 16 colour attribute graphics playfield, with collision detection and a fast text mode for editing and listings. TI also made Video Display Processors for MSX, Memotech and Tatung Einstein micros and influenced later console and Spectrum graphics.

Emulated sound uses three beepers plus random noise, and a speech synthesiser. 32KB expansion RAM, floppy disk images and joysticks are also supported. Speed is limited to match the original, which is authentic but a pity as layers of interpretation made the TI99/4A frustratingly slow.

*TI99sim* has no GUI, and just four control keys: **Esc** quits, **F2** and **F3** save and load the computer's state and **F10** reboots. You can also save and load to disk images as if using an original TI floppy controller and ROM.

*TI-99sim 0.0.7* comes with utilities and a helpful *readme.html* file. **dumpcpu** disassembles code in TI memory images, with optional symbol and label assignment to make the result more readable. **dumpgrom** does a similar thing with GPL (no relation) pseudocode and strings in GROM cartridges, while **list**



The Texas TI99-4A was the original 16-bit home computer.

detokenises TI-BASIC and Extended BASIC programs. **decode** tries to extract TI tape file contents from WAV files, and **disk** converts and extracts files from AnaDisk and 'V9T9' images or raw tracks from another emulator, *PC99*.

Speech synthesis, using linear predictive coding and an add-on cartridge, was an early strength of the TI, as you might expect from the makers of *Speak-and-Spell*. You can use original code and data from *spchrom.bin* or make a 'stripped down' ROM if you don't have the original, and potentially add new samples, though the emulator has no LPC-10 encoder to do this for you.

### TI MESS

*XMESS* promises to emulate the 3MHz 16-bit minicomputer-style 9900 CPU, TI graphics and sound, with single sided disk images, V9T9 carts and 'partial' speech. *MESS* drivers match international hardware releases; TI99/4AE uses the European ROMs and display format, launched with: **xmess.x11 ti99\_4ae**

Alas, the colourful startup screen was as far as it got. Pressing a key caused a segmentation fault. Again, *MESS* currently promises more than it can deliver. The third contender, *TI4Linux*, is text-only and needs a lot more work to be useful. Despite quirks, *TI99sim* wins out.

### ◀ xmess.x11 advision -cart cobras.bin

This also loads 'avbios.rom' from the advision directory on the *XMESS* path, configured in '~/.xmess/xmessrc', or 'usr/local/lib/xmess/messrc' for all users. As for all *XMESS* drivers, special configurations for one or many users of the X11 build go in the corresponding 'xmess-x11rc', and ones for a specific system in 'rc' sub-directories like '~/.xmess/rc/advisionrc'.

272 computers, described as 'games' in documentation hacked from *MAME*, are listed as 'supported' by *MESS*, but slight


variations get their own entry and 60 of the systems don't work at all. Others have missing features, keys or sound. Dozens of tests confirmed that *MESS* needs more work. It's let down by poor documentation as well as patchy code. Every driver includes a 'Machine usage and history' menu item, but they all came up 'System Info not available'. Online documents contains large gaps and basic errors.

## Consoles

Eight-bit Sega and Nintendo console drivers are respectable though slow, but the SNES one only runs one early *Mario* game. The Zilog-based Amstrad PCW16 emulation has no EPP printer port, which makes emulation of this dedicated word processor rather pointless. Atari Lynx emulation has blitter, timing and audio bugs, so it's no competition for *Handy* (LXF30). The eponymous Computers Lynx is not emulated, though if that was buggy it might be hard to tell as the original was so dodgy! The Intellivision emulator lacks collision detection, scuppering many games; prefer *Bliss* and *Jzintv* (LXF30).

*MESS*'s Amiga driver lacks sound and only runs a subset of programs for the oldest Amigas, so it's no competition for *UAE* (LXF32). The MacXL or Lisa2 emulation is ropy and *vMac*, *Basilisk* and *Executor* (LXF18) are far more useful. *Bochs* (LXF23) offers better emulation for old x86 computers, especially as *MESS* skimps on platform-specific hardware.

The *Kemulator* front-end for *MAME* also supports some *MESS* drivers.



Even with appropriate ROMs and an emulation listed as working it might not do anything in practice, as I found with the Memotech MTX driver. Exidy Sorcerer and Mattel Aquarius emulations can't generate equals key-presses, so they're barely programmable. Sorcerer and Z88 drivers warn of bugs and are glacially slow. As the name hints, *MESS* is a curate's egg.

## Bare Boards

*MESS* development for early bare-board micros like the SYM-1, Aim-65, and Cosmac VIP stalled after basic work on the display and keyboard because "it is too cryptic to operate and therefore [sic] not interesting." But the driver for Chuck Peddle's milestone KIM-1 runs tiny but authentic programs, including a calculator-style *Lunar Lander*, and the 2KB ROM and its original source are available for those who fancy a taste of 1975-vintage microcomputing.

The Compukit UK101 was a PAL clone of the Ohio SuperBoard II 6502 system, reworked by Dr. AA Berk for a *Practical Electronics* magazine series, much promoted in the early eighties. *XMESS* assigns a resolution of 256 by 400 pixels, implausible for TV text, and I got no further than the initial DCWM prompt when I launched this driver. *Emu8* is another UK101 emulation, written with old C++ and Tk installations on Red Hat 4, and won't run without changes on current Linux systems. Can a reader make the best of both?

## Orics

Oric Computers, derived from Microtan65 boards, once rivalled the Spectrum but lack decent Linux emulation. The giftware *Euphorix* just gives a segmentation fault, and I've found mentions but no source for Jean-Francois Fabre's even older *Oric48K* emulator for X, so it's a pity the *XMESS* offerings are broken.

*XMESS* has drivers for the Oric 1, Atmos, Telestrat, and Pravetz 8D, a Bulgarian Atmos clone. But bugs mean some tapes and disks don't load. In fact the Atmos emulation doesn't seem to load any tapes at all, and a fault in Telestrat disk emulation means that won't even boot.

## System 80

*MESS* TRS-80 disk emulation only suits NewDOS80 version 2, one of a dozen TRSDOS clones. *XTRS* is a better choice unless you really want the EACA System80 (also known as Video Genie or PMC-80) variant which *XMESS* lists beside Tandy originals, and may yet be closer to the system you were expecting. In *XMESS* only the left **Shift** button works, which is authentic for some later Genies with missing TRS-80 keys in the right **Shift** position, but annoying otherwise.

## Colour Genie

*XMESS* offers the only Colour Genie EG2000 emulator for Linux. This follow-up to EACA's Video Genie uses Hitachi graphics based on Motorola's 6845 VDP, giving 16 colour 40x24 character text, 128 redefinable 8x8 character patterns, and four-colour 160x96 block graphics rather than the mono 128x48 blocks or 64x16 characters of the EG3013 and TRS-80 TTL logic.

It also has an AY8910 beeper like the Spectrum 128, which *MESS* emulates well, though without serial and parallel ports. But input does not match my US PC keypad, or the original I reviewed 20 years ago for *Computing Today* and *Micro Choice*.

Only the left **Shift** key works. Right **Alt** emulates the 'Mode Sel' key to the left of the spacebar on the real thing, **F1** to **F3**

work like the original function keys, but **F4** only works if the *MESS* menus are disabled. You can reassign the keys in the GUI but the defaults are typically badly chosen. **F5** toggles a clearer display which uses a wider font but sometimes leaves narrow vertical lines on the screen.

Neither of the original two reset keys appear to be implemented so there seems to be no way to break into a tape load that goes wrong. BASIC is loaded with **CLOAD** and code with **SYSTEM** followed by a name, as on TRS-80s. DOS uses the TRSDOS **CMD**, but **CMD "IO"** gives a directory of drive 0, **CMD "L NAME/BIN** loads and **CMD "S NAME/CMD"** runs **NAME/CMD**, while **CMD "S FIREBIRD/CMD:3** loads from drive 3 – closing quotes are optional.

The *XMESS* driver offers 16KB or 32KB RAM, optional 8KB disk and 4KB extension ROMs or RAM in the same space, and tape support. Between 4 and 10KB of RAM is lost to the system, depending on the setup.

The Z80 was clocked at 2.2MHz, 9/7 the speed of earlier Genies but still only 2/3 the speed of a Spectrum. With the limiter removed by **F10**, *MESS* runs Colour Genie programs around 4 times faster than the original on my K6. It's stable and not a bad emulator apart from the keyboard bodes.


## Swansong

Home Computers used many common parts, so *XMESS* can share input, sound and graphics support as well as processor cores and a generic interface – but what works for *MAME* does not work as well for *MESS*.

Many of the drivers are buggy or limited, though the CPU cores seem solid. But Colour Genie, Kim1, Dragon, CoCo and especially Vectrex emulation mean that *MESS* fills significant gaps. Any emulation completist should try it, which is why it terminates this series, after testing more than 100 emulators that run on Linux.

If you plan to write an emulator, plenty of scope remains. I've yet to see an emulation of the MK14, Sinclair's first micro, or the ComX35; their SC/MP and 1802 processors make them esoteric, though the 1802 controls NASA Space Shuttles and Voyager probes. I'd also enjoy an emulator for Belgium's DAI, a fast UK-programmed 8080 micro born of European TI distribution problems.

Your coverdisc includes *VMIPS*, a virtual R3000 processor (as used in PlayStations and SGI workstations), emulations of Intel's seminal 4004, several DEC minicomputers, and MIX, the virtual machine documented in Donald Knuth's matchless *Art of Computer Programming* books. There are also TI and HP calculator emulators, Xtel and Oric resources, *Emu8* and *GalEmu* archives. Extended and updated articles from this series will continue to appear online at <http://simon.mooli.org.uk/LXF/> and I welcome correspondence via [linux@studio.co.uk](mailto:linux@studio.co.uk). [LXF](#)



The main menu of the well-presented French emulator *Thom*.



The built-in Vectrex game *Minestorm* is a smooth variation on *Asteroids*.

## Links

LXF Emulators online: [www.simon.mooli.org.uk/LXF](http://www.simon.mooli.org.uk/LXF)  
 MESS CVS: <http://cvs.mess.org:6502/cgi-bin/cvsweb.cgi/>  
 MESS files: <http://pkg.lugbs.linux.it/PACCHETTI/Mess-roms/>  
 MESS home: <http://mess.emuverse.com/>  
 MESS ROMs: [www.geocities.com/messroms](http://www.geocities.com/messroms)  
 Thom and Theo: <http://thomson.rcroms.com/archives/>  
 Thomson Emul5: <http://emul5.thomsonistes.org/>  
 TI99sim: [www.mrousseau.org/programs/ti99sim](http://www.mrousseau.org/programs/ti99sim)  
 XMAME/XMESS: <http://x.mame.net/>

# 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](mailto:linuxformat@futurenet.co.uk)) or log on to our website and post your suggestions in our special forums? ([www.linuxformat.co.uk](http://www.linuxformat.co.uk)). Hope to hear from you soon!

**Nick Veitch** EDITOR

## THIS MONTH...

### Low resource Linux

Tailoring Linux for that old laptop that you'd thought was redundant **p68**

### System programming

Delving under the bonnet of the Linux kernel to look at the world of system calls **p72**

### OpenOffice.org >>


A new series on getting the most out of 'the greatest Open Source project ever' **p76**

### Perl

This month a quick project to glue news headlines into your website with RSS **p80**

### Java

Continuing our index searching applet project, this month we construct a GUI with the AWT widget set **p82**



### Kylix

It's not just for GUI fans! We conclude our series with *Kylix* programs that can be run from the command line **p84**

### PHP

Having introduced SQL, we now look at getting PHP to query *MySQL*, and round off with a look at normalisation **p90**

## How code is represented

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap
between lines:
begin
mniWordWrap.Checked := not
end;
Usually, you'll find the code on
our CD/DVD too.
```

## TIP OF THE MONTH!

The **find** command is one of the most useful of the shell tools supplied with every distro. It's not just there for locating files, although it is handy for that. With its ability to execute an arbitrary command for each file or directory that it locates, the *find* command becomes a recursive pattern-matching tool.

As an example, suppose you wished to change the permissions of every file in the directory tree rooted at the current directory to readable and writable by the owner and group, but just readable by others. You can do this with

## Not just for finding

```
find . -type f -exec chmod
0664 {} \;
```

Here the switch **-type f** tells find to match only files while the escape characters make sure the special characters in the **-exec** clause are not interpreted by the shell.

Perhaps you wanted to change the permissions of all the directories in the tree rooted at *somefolder/* so that they are accessible only by the owner. How about this:

```
find somefolder/ -type d
-exec chmod og-x {} \;
```


The **-type d** switch here matches only directories.

Another great use for *find* is when you want to search a tree of files. Suppose you wanted to look in the kernel source tree and find all the source files which contain the word 'acpi'. You could do this with:

```
find /usr/src/linux -name "*.c[h]"
-exec fgrep -l acpi {} \;
```

Here the **-l** switch to *fgrep* disables the normal output and instead makes *fgrep* print only the filename of a file in which the pattern matches.

For more information on *find* consult the man page or, even better, the texinfo pages, since these contain lots of example uses.



## EFFICIENT DESKTOPS

# Low resource Linux

**PART 2** This time, Marco Fioretti presents some specific programs and methods to work faster on limited computers.

Last month, in our bid to get the most out of low resource machines, we examined the general system configuration. This time we will focus on the end user space, *i.e.* on the actual apps which deliver the greatest real functionality for home and small office with the smallest possible hardware needs. This includes text processing, email, news, browsing, file management, accounting, and several other things. Games and multimedia activities (with the exception of MP3 playback) will not be covered for the simple reason that they inherently require newer and more powerful hardware, regardless of the operating system, falling therefore outside the scope of this tutorial.

It should be explicitly noted that the solutions and programs presented here, even if optimal or mandatory on older computers are in no way restricted to them. The fact that a processor has a 2GHz clock is not a valid reason to not make it even faster, not

to mention the fact that several times the most efficient programs also turn out to be the most customisable.

## Prerequisites

The only real prerequisites to set up an efficient desktop are the right attitude and respect of netiquette. There is no need to be a programming guru, but, exclusively in the name of efficiency, GUI applications should be put on the back shelf in favour of the command line or, to be more specific, of some shell scripting.

Remember that, even though several of the programs described below might not fit today's standard 'user friendly' dogmas, they normally have much more features than their multi-button counterparts – or at least only those really useful – and run faster. This is *real* user friendliness, isn't it?

Another thing to take into account is that many of these programs are still (unfairly) used mostly by Linux experts. Newbies should not be scared by this, just be conscious of the environment: the users of these tools are very competent and active in the respective mailing lists, but also have a lower tolerance of improper behaviour. They will give plenty of help to

newbies, but: do read the FAQ and any relevant documentation first and do give meaningful subjects and relevant details in your messages – “problem with *sometool*” on the *sometool* list will be ignored in the best case, mightily flamed otherwise.

## System interaction/administration

### Driving from a GUI

Shell scripts and commands give the maximum power, efficiency and flexibility, but remembering them all is not easy, and opening a terminal just to type six or seven characters doesn't make much sense. When we closely look at it, this is the main, if not the only reason, why many people become addict to full fledged mouse-based environments, and limit themselves to the ready made ones. The way to keep the best of both worlds is (as we mentioned in the first part) any window manager which accepts, via user configuration file(s), the binding of a shell script to its root menu. All scripts and commands presented in the rest of this tutorial can (and should) be used in this way, to give a friendly interface with low (human and silicon) memory requirements.

### The window manager

It is obvious, both in general and from the previous paragraph, that the right choice of window manager is critical to guarantee the right mix of ease of use and performance. There are a lot of really minimal and light solutions, from *twm* to *qwm*, but most of them cannot easily support the menu flexibility that we require. *XFCE* and *WindowMaker* can, look much better, and they are indeed quite fast, but remember what was said in the first part about older PCs usually having less pixels and colours to spare on icons and docks. This is one more reason to try *Blackbox* (<http://blackboxwm.sf.net>), which seems to give the best compromise between speed, minimal dependencies and clean look on one side, versus completeness and flexibility on the other.

### Internet Connectivity

Some simple shell script will happily replace all the ‘connection managers’ available. Unfortunately, too many modem, *pppd* or email transfer options vary from country to country, if not from ISP to ISP, to make the same scripts always work. The only generally valid tip that can be given in this space is that, before asking anything, one should check if PPP support was included at install time, and if the modem is supported by Linux. After that, rather than wasting time in Internet wide searches, or queries on high traffic lists, the right thing to do is to search in your ISP's online help for Linux-specific settings, or in the local LUG home page. If nothing comes out of that, just subscribe to the local LUG list, and ask for PPP/ADSL scripts, specifying your ISP, where you live, Linux distribution and modem model.

### Mounting devices

Floppy, CDROM, and external iomega drives can all be mounted and unmounted through a GUI as normal users, without heavy file managers, or custom icons. First of all, they should be configured to be mountable by normal users, with lines similar to this (note the ‘user’ option):

```
/dev/cdrom /mnt/cdrom iso9660 noauto,user,ro 1 1
```

in the */etc/fstab* file. The man pages of *mount* and *fstab* give all the details. The commands to mount and unmount devices are all similar to these:

```
mount -t iso9660 /dev/cdrom /mnt/cdrom # mounts cdrom
```

```
umount /mnt/cdrom # unmounts cdrom
```

and can obviously be bound to root menu entries.

### Monitoring

There is a class of icons that cannot be substituted by menus: we refer to all the dashboard-like faces showing, in real time, parameters like connection speed, CPU and RAM utilisation, battery or mailboxes status, or, very simply, time and date. When this visual feedback is needed, *gkrellm* (see [www.gkrellm.org](http://www.gkrellm.org)) is relatively efficient, and works using only the *Gtk+* libraries. Several window managers, namely *Blackbox* and *WindowMaker*, come with their own auxiliary applications.

### Backup

Backup is another activity where it is possible to go from typing commands by hand every time to network wide, self scheduling commercial tools. Even at the single desktop level, however, there are quite a lot of script based utilities that automate most of the dirty work, sometimes even dumping the backup straight to CD, and require almost no resources. One example for all is the *axlbac* tool, available online at <http://tuxpowerf2g.net/axlbac.php>

### Text processing

Writing tools, memos, diaries and programs is an excellent use for a limited workstation. *TeX* typesetting is not covered here, for the simple reason that it would fill by itself older drives. *Abiword* is probably the lightest WYSIWIG word processor available, and can even import *Word97* files. For all other needs, there is no reason to start yet another holy war between *Vi(m)* and *Emacs*: both should be known, installed and used. The first is often the only editor usable in rescue situations, and can do a lot anyway. The latter, already omnipotent by itself, can be extended with autonumbering (see <http://tinytools.sourceforge.net>), PIM and calendar modules, and much more: just remember, as explained last month, to remove all the unneeded files.

### File management

One thousand generic computer users will more or less agree about what a computer file is, but give (at least) one thousand different definitions of file management: this is obvious, since even the same person needs different “file management” solutions depending, every time, on the type of files with which they work (source code vs plain text vs images...), on their quantity, and on what needs to be done (find differences, move or rename things, etc.).

### Managing differences and archives

Browsing directories, or looking inside tar or compressed archives can still be done with *mc*, the *Midnight Commander*. When the problem is to find files with some specific properties (newer than N days, belonging to user X, with a .jpg extension...) nothing beats (predictably) the *find* command. An hour spent reading its man page can speed thing ups more than a faster processor. Moving and renaming large quantities of files can be achieved either with *find* or with the *mmv* (*multiple move*) command. Differences in directories (up to five) are *dirdiff*'s domain (see pic, overleaf).

### Images galleries

There is at least one case where even command line fanatics must start up a graphical file manager, and that's handling large collections of images. Let's assume that there is a CDROM worth >>



Keeping your system under control with *gkrellm*.

◀◀ of family pictures scanned and saved on disk without any criterion: they are all in one folder, named in ten different ways (2001\_Holidays.jpg, Johnny\_playpen.gif, Nevada.png.) To sort out something like this, thumbnail galleries are essential: with them, one can select immediately all the pictures taken last Christmas, and move them to a new folder. The *ROX* filer is light, cute, and perfect for these tasks and many others. Its homepage, including links to icon collections and the rest of the *ROX* desktop environment, is <http://rox.sourceforge.net>.

To browse galleries of already sorted images instead, two solutions that combine satisfactorily speed and features (including slideshow support) are *gqview* (<http://gqview.sourceforge.net>) and *display*. The latter is part of the *ImageMagick* suite (available at [www.imagemagick.org](http://www.imagemagick.org)) which also includes several tool to resize, convert and process in many other ways images on the command line.

## Browsing

Web surfing is probably the activity that is considered most penalised by a low power computer. There is no doubt that a lot of fashionable gadgets and activities, from movies embedded in webpages to online gaming, require the latest hardware and a broadband connection whatever software is used. For this reason, as already mentioned, this use of the Internet is outside the scope of this tutorial. Everything else, however, can be accomplished even with very modest means.

The first step to make web surfing lighter on your CPU has nothing to do with software. Prepare a polite and objective letter which says (without even mentioning Linux or any specific browser) "Your website is made slow and unusable by too many pointless animations, gratuitous JavaScript, and non standard HTML. It bored me and I will never return". Once prepared, *always* send it to the webmaster of *every* page that requires twenty different plug-ins just to read a phone number.

On a more technical note, there are indeed several ways to make an old computer surf faster (and, since they are typically attached to a phone line billed per minute, cheaper too!)

It goes without saying that a lot depends upon the choice of the browser. The text based ones (choose any among *w3m*, *links*, and *lynx*) besides being intrinsically lighter, make surfing many

sites faster by ignoring everything but real information.

Even for images and JavaScript support one can usually do without heavyweights like *Konqueror* or *Mozilla*. The 4.7x version of *Netscape* (yes, *Netscape*...) looks lean and mean in comparison, even if something like *Galeon* ([galeon.sourceforge.net](http://galeon.sourceforge.net)) is still relatively light and makes much more sense in terms of support of modern standards. Products like *Browsex* ([www.browsex.com](http://www.browsex.com)), *Dillo* (<http://dillo.sourceforge.net/>) or *Skipstone* ([www.muhi.net/skipstone](http://www.muhi.net/skipstone)) are even lighter, although somewhat more limited and/or still less stable. In any case, before starting the browser it is mandatory to make sure that it will have to work and wait as little as possible. To make this happen, one should first of all set up a DNS proxy to diminish queries to your ISP's DNS server (try *pdnsd*, <http://home.t-online.de/home/Moestl/>). Right after that, *privoxy* (which offers banner blocking and several privacy protection functions) should be installed.

Depending upon connection speed and upon the nature of the sites visited most frequently, a real time saver may be offline browsing, which can mean one of two things. The first is caching static pages on one's PC and sending those copies to the browser, instead of downloading them each time. The second is to mirror the whole set of pages of interest (for example the headline page of a news portal plus all the links it points to) and browse that local copy when finished. Programs like *wwoffle* and/or *wget* are the answer in this case.

## Email


Sometimes email is even more needed than surfing, isn't it? To manage it effectively, however, especially when subscribed to one or more high volume mailing lists, and under different rôles (work, private, and so on) there are some basic features to shop for. An efficient email environment is one that at least lets the user:

- define many criteria to sort incoming email automatically
- catch and destroy as much spam as possible, possibly without ever downloading it
- add and edit addresses and aliases easily
- recognize and open correctly all common attachments
- always have the From and Reply-to headers set according to the context, *i.e.* depending from the address to which the message will be sent, or that to which the original message was delivered.
- sign and encrypt messages with different keys

There are several big caliber email programs, like *Kmail* and *Evolution*, that do to some extent everything listed above, and quite a bit more. In a low resource environment, they are not an answer: the user of such systems needs to install several smaller tools, and make them talk to each other with proper configuration files. This takes some time, of course, but it's only initial setup, and the end result is not simply faster: it also has many more features than anything else, which is the reason why even many sophisticated GUI users stick to similar solutions.

*Sylpheed* is one of the lightest GUI based clients, but one even fastest and endlessly customisable, down to and including the colours to use for each part of each message, is *Mutt* ([www.mutt.org](http://www.mutt.org)). A lot of its power comes from its 'hook' mechanism: *Mutt* can recognize several events, from moving to a specific mailbox to sending a message to a known address, and change accordingly its internal settings, from the sender address to the mailbox to which messages should be saved.

**Spotting differences between directories. This handy tool shows them visually (green signals the newest files, black the missing ones) and can also merge and patch directories. More info is available at [http://freshmeat.net/projects/dirdiff?topic\\_id=45](http://freshmeat.net/projects/dirdiff?topic_id=45).**



All types of attachments (and HTML email) can be opened straight from Mutt, but this only happens if the mailcap file is configured properly: see examples in the *Mutt* pages listed below. There is also a very small and simple address book utility, *abook* (<http://abook.sourceforge.net/>), written explicitly for *Mutt*. Compared to its much more visually appealing competitors *Mutt* misses only one functionality: It cannot send email by itself, just pass it to an SMTP server on the same machine. This is not going to change, so it is useless to ask for it, because "it's not a bug, it's a feature": SMTP support is not an email client task, but something that would make the program slower and trickier to maintain.

The practical consequence is that even a single user has to install some SMTP server to send email. The good news are that the single user configuration of *Postfix* or *Exim* are quite easy to set up, and that there are even lighter alternatives for limited computers. These range from simpler servers like *nullmailer* (<http://untroubled.org/nullmailer/>), to nothing more than Perl scripts. The last option is really quite limited and dirty, but sufficient in simpler cases: a good read for the pros and cons of these solutions is the thread "sending with Perl instead of MTA" in the August 2002 archive of the *Mutt* users list.

So, *Mutt* needs external programs to work with but what are they? Apart from any SMTP server, the default solutions proposed around are valid even in this case: *fetchmail* to download all messages, even from multiple accounts, and *procmail* to sort them, send automatic replies, and stop spam (possibly with help from *SpamAssassin*, <http://spamassassin.org/>). A welcome addition on dial up connections, which save call costs if not CPU cycles, are programs like *animail*, which delete large unwanted messages on the server instead of downloading them.

The best online resources to know how to integrate *Mutt* and its friends are Gary Johnson's *Mutt* Page at [www.socom.com/users/gjohnson/mutt](http://www.socom.com/users/gjohnson/mutt), the list of third party applications for *Mutt* at [www.symonds.net/~prahladv/mutt.php](http://www.symonds.net/~prahladv/mutt.php), and Dave Pearson's *Mutt* Page, [www.davep.org/mutt](http://www.davep.org/mutt).

As far as newsgroups are concerned, *Mutt* itself may read them too, but only if patched and compiled from source. A ready made alternative, with a similar interface too, is *SLRN* (<http://slrn.sourceforge.net/>).

## Audio

Recording and playing back music is possible even on very slow machines (if there isn't anything else going on, and top quality doesn't matter). Apart from actually listening to songs, the creation of song catalogues and many other "music administration" tasks require very little resources. For example, *mp3conv* (<http://tuxerf2g.net/mp3conv.php>) converts MP3 files to Ogg format, while *Mp3CdSpeler* (<http://tuxpowerf2g.net/mp3cdspeler.php>) creates CDs properly formatted for the MP3 Discman: no more than forty songs per directory, and eight characters file names.

The screenshot shows the Midnight Commander interface in a terminal window. The left pane displays the contents of a file named 'CHANGELOG', which lists various updates and fixes, including 'Attempt to fix segfaulting seen since 2.3.2', 'Moved the re-enable service after cps vi', and 'Fixed the Filelog problem of printing ga'. The right pane shows a directory listing with columns for Name, Size, and Permission. Files listed include 'COPYRIGHT', 'INSTALL', 'README', 'sample.conf', and 'CHANGELOG'. The bottom status bar shows the current directory as 'CHANGELOG' and the file size as '27105'. A hint at the bottom suggests using 'M-l' to look at command output.

*mc* – the *Midnight Commander* – with its ever-useful, twin-pane view.

Those are only two of a myriad of utilities demonstrating that one can manage multimedia on a hardware shoestring: Freshmeat will list many more, and, as an exercise for the reader, we leave to you to find any of the several webpages explaining how to transform a leftover Pentium box into a juke-box.

## CDROM burning

Making CDROMs is not a luxury, but a necessity on any computer (network) that can afford it, both financially and physically. A very nice comparison of both GUI and command line frontends to CD burning tools is at <http://sites.inka.de/~W1752/cdrecord/frontend.en.html>, which also lists the external dependencies of many programs. People who really burn several CDs per week, each one in a different way, are probably better off with a GUI.

The rest of us, however, those who just need to do one or two things always in the same way, can certainly live with command line interfaces like *cdrx* ([cdrx.sourceforge.net/](http://cdrx.sourceforge.net/)), or just bind to a menu entry some scripts with the proper calls to the *cdrecord* and/or *mkisofs* commands (especially as almost none of the fancy frontends support "all" their options).

## Faxes and photocopies

This is yet another field where scripting is more than enough. On the author system, the following three lines script and the *SANE* packages are all that is needed to make copies of documents:

```
scanimage --mode LineArt --resolution 300 LineArt >
/tmp/tmp_copy
pbmtolj --resolution 300 /tmp/tmp_copy/lpr
rm /tmp/tmp_copy
```

calling the *fax* command in the second line with the proper options would directly fax the document just scanned. The man pages of *scanimage* and *fax* provide all the details for colour printouts, multi page faxes, and so on.

## Conclusion

We hope to have shown that, with the right tools and a bit of script-based cheating, Free Software can still do more with less, and that this can be a lot of fun too. Portals like Freshmeat continuously announce small, unknown utilities to perform every possible task faster and more lightly: keep an eye on them and, above all, don't be afraid of experimenting. **LXF**

## SYSTEM CALLS

# Under The Hood

**PART 1** In a new series, **Chris Brown** gives us the low-down on that smooth-running engine, the Linux kernel.

Cars appeal to people in different ways. My son views his as an acoustic chamber in which to confine 400 watts of audio; the fact that it happens to be on wheels being of secondary importance (I suspect the power output of the sound system exceeds that of the engine). My wife, bless her, fills ours with little old ladies whom she ferries around tirelessly, handing it over to me only when the MOT is about to expire or, as happened recently, it starts overheating. We called out a mechanic who forced me to re-acknowledge the existence of parts of the car I try to ignore. The parts that make it go.

Linux also appeals to people in different ways. Some see it as a set of network services. An end user (whatever *they* are) might think mainly about the desktop tools. A third approach, perhaps rather narrow and pedantic, and more likely to be held by a developer, views Linux as a set of system calls that programs can use to get their job done. In this series of tutorials we'll be looking 'under the hood' of Linux to examine the engine compartment – the services provided by the Linux kernel. We should 'fess up' right now that this stuff is not new. In fact most of it was inherited from Unix and has been around for 20 years or more. We're also going to stick with the traditional approach and show our examples in C. Although the system services in Unix are potentially available from any language, the system calls (in section 2 of the man pages) have always been documented in C and it is in some sense the 'natural' language to discuss them in.

## What does the kernel do, exactly?

The kernel is the first piece of software that gets loaded into memory when a Linux system boots. In a sense, the kernel and the set of system services it provides, *is* Linux. All the other good stuff, like the shell and the GNU utilities, the X window system,


the window managers and desktop tools, the network services and so on, are mere applications which happen to run on top of Linux. **Figure 1** shows how the kernel fits in to the picture.

The general picture painted in **fig 1** is one of increasing levels of abstraction as one moves up through the layers. At the bottom layer lies the hardware. Nasty, complicated stuff controlled by fiddling around with bit fields in status registers. The kinds of things you can tell raw hardware to do are pretty basic. Take a disk drive, for example. You can select a cylinder, a track and a sector and command the drive to read the sector to a specified address in memory. Messy. Moreover, the way you do this varies from one type of disk to another. Not many programmers want to muck about at this level. So instead, the Linux kernel provides an abstraction, (a figment of the kernel's imagination, if you will), called the file system. It lets programs deal with named files within an hierarchical structure of files and directories, and to perform operations on those files such as open, read, write and close. Behind this lie multiple implementations – SCSI disks, IDE drives, CDs, DVDs, even remote filesystems accessed via network protocols such as NFS. The Linux kernel provides drivers for each of these implementations, providing a single, consistent system call interface to application programs.

Let's consider another example – memory. Pretty simple stuff, you'd imagine. An array of addressable locations that a program can store data in without any help from the kernel, right? Well, yes and no. Memory starts to get more complicated in a multi-process environment like Linux where there are many programs competing for it. For a start, all those programs would like to imagine that they have the memory to themselves – that they can simply access a 'logical address space' (a range of memory addresses) starting at address zero and extending up as far as they need. In reality they are all sharing a physical address space (the addresses which determine the actual locations in RAM where the data will be stored). Also, the amount of memory physically present in the machine may be less than the total

memory requirement of the currently active programs. To support all this, the kernel works in concert with the memory management hardware of the computer to establish a mapping between the logical addresses that the programs use and the corresponding physical addresses. Some pieces of the logical address space may even be 'paged out' (parked in the 'swap space' on disk) and retrieved on demand – see **fig 2**. The mechanics of this are called demand paged virtual memory, and it's another important kernel service. Another vital task of the memory management system is to keep the address spaces of each process separate. That is, one process cannot

Figure 1: **Linux in Layers.**



(either maliciously, or through errant code) scribble into the memory of another.

The virtual memory services are a lot less 'visible' to application programs than the file system services. That is, you rarely need to explicitly call on the memory management services of Linux from within a program, as you do for file I/O. It's a bit like breathing whilst driving a car. You focus consciously on the traffic, not on inflating your lungs. That doesn't mean that breathing is unimportant, just that it's mostly automatic.

The kernel, then, provides a set of system services to the applications that run on it. The specification of this set of services (known technically as system calls) defines a 'Linux Virtual Machine' which is independent of the underlying hardware. The standardisation of this interface ensures the portability of applications from one implementation of Linux to the next. It's this virtual machine that we're concerned with here.

## No cheating allowed

As well as providing system services, the kernel also prevents application programs from 'cheating' by bypassing the kernel and talking directly to the hardware. Think about it – if a program could manipulate the control registers of a disk drive directly, it could read or overwrite anything, making a mockery of any file-level access control that might be in place. This protection, whilst configured and monitored by the kernel, is ultimately provided by the processor, which traps specific privileged instructions (including those that access the device registers) and allows them to proceed only if the processor is running in a privileged mode. When a program makes a system call into the kernel, the processor switches into a privileged mode. We talk about moving from 'user space' to 'kernel space'. This is a tightly controlled entry point into the kernel and is achieved using a special instruction called a trap.

## Library functions and system calls


Now, it turns out that ordinary programs do not issue the trap instructions directly, but do so by calling what we might describe as 'wrapper' functions in the standard runtime library. It is these wrapper functions that actually execute the traps.

Look back at **figure 1** and at the C program listing below. This is a silly program whose only purpose is to illustrate some points about library functions and system calls.

```
#include <stdio.h>
#include <math.h>

main()
{
    double x;
    x = sqrt(8.0);
    printf("Square root of 8.0 is %f\n", x);
    write(2, "hello world", 11);
}
```

First, the program calls a library function called **sqrt()**, which calculates square roots. The calculation is done entirely within user space; that is, no system call is made into the kernel, and the result is simply returned to the caller. Next, the program calls **printf()**, a formatted output function which writes to the 'standard output' stream. The **printf()** function does all the formatting in user space (that is, in the library) but will ultimately need to write out the formatted string. For this, it will make a **write()** system call to the kernel. Finally, the program calls



**write()**, to send a simple text string to output stream 2 (the standard error stream.) This **write()** function, which exists within the standard library, is just a wrapper which makes the appropriate system call into the kernel and passes the parameters through.

To summarise: programs obtain services by calling library routines. Sometimes the library routine does the entire job unaided. Sometimes it will make its own called into the kernel. And sometimes, the routine is little more than a wrapper which executes the trap into the kernel and passes the appropriate parameters through.

If you want to look at the system calls executed by a process, you can use the program *strace*. Suppose you have created the little program shown above in a file called *calltest.c*. You can compile and link it with the command

```
$ gcc calltest.c -lm -o calltest
```

Then run it and trace the system calls with:

```
$ strace calltest
```

You will have to look carefully in the output because there is a great deal of rather fearsome-looking system call activity as the program starts up. However if you look closely at the last few lines of output you should see the **write()** system call resulting from the **printf()** call, and a second **write()** call corresponding to the last line of the program. There is no output from *strace* from the **sqrt()** calculation. Remember, *strace* is only monitoring system calls into the kernel, not calls to the runtime library.

## Processes and Programs

Having established, hopefully, some context as to what system calls actually are, we're going to spend the rest of this article looking at the calls which allow us to create processes and run

Figure 2: Logical to physical address mapping.




Figure 3: **The memory layout of a process.**

programs. To cut to the chase, the calls are **fork()** and **exec()**. The way these calls work, more than any others perhaps, are unique to the Unix process model. No other operating system does it quite the same.

First, we need to understand the difference between a program and a process. The concept of a program is fairly easy to grasp. It's a list of instructions saying what's to be done. If you asked me to take the cover off my computer and show you the programs, I could, in principle at least, point to the sequences of machine code instructions in memory and say "those are the programs". A process is a more abstract concept. If you asked me to point to the processes inside my computer, I'd have a harder time.

One way of looking at a process is as the entity responsible for getting a program executed. You might draw an analogy between the script of a play (the program) and an actor (the process) who is reading the script (executing the program). Just as an actor has an identity (a name), each process has a numeric identifier, called the process ID, or PID. You can see all the processes running in your machine with the command

```
$ ps ax
```


The first column (headed PID) shows the process ids.

A process also keeps track of the resources, or context, needed to run a program. These include environment variables, file descriptors, and memory. The memory space of a process is divided into four regions called the code segment, the data segment, the stack and the heap, as shown in **figure 3**. This figure shows the logical address space of the process. We've already hinted at how Linux maps this onto physical addresses and the swap space.

### Creating processes

New processes are created with the **fork()** system call. Now, **fork()** is a very simple call. It doesn't take any arguments, and it returns a simple integer value. Nonetheless, **fork()** is confusing the first time you meet it because although only one process calls **fork()**, two processes (the original 'parent' process and the newly created 'child') return from it. We now have two concurrently executing process.

Figure 4: (below) **Life cycle of a process.**



Let's get a bit of help from our analogy. A lone actor stands on the stage, reading a script. He reads a stage direction "summon new actor" so he puts in a call to a theatrical agent to hire a second person. The new actor (the child) joins him on stage. The original actor (the parent) and the new actor both continue reading whatever immediately follows the "summon new actor" directive in the script. The analogy isn't perfect, the child process wasn't "waiting in the wings" like some out-of-work actor, it was born as the **fork()** call was executed.

When a programmer uses **fork()** to create a child process, it's inevitable he'll want it to do something different from the parent. So that you can tell which process is which, the **fork()** returns different values in the parent and the child. In the parent, it returns the PID of the newly created child, and in the child, it returns zero. So you always see the **fork()** call embedded inside an **if()** test, like this:

```
if (fork()) {
    /* Do the parent stuff */
}
else {
    /* Do the child stuff */
}
```

The following example is worth typing in and running because it nicely illustrates the concurrent execution of two processes.

```
int main()
{
    int i;
    if (fork())
        for (i=0; i<10000; i++)
            write(1, "*", 1); /* Parent */
    else
        for (i=0; i<10000; i++)
            write(1, ".", 1); /* Child */
    return 0;
}
```

To run the example, put the code into a file called *twoproc.c* and compile and run it like this:

```
$ cc twoproc.c -o twoproc
$ twoproc
```

The example works best on a console terminal (type **Ctrl-Alt-F1** to get one) rather than in a terminal window on a graphical desktop. You'll see a bunch of stars and dots mixed up in some unpredictable fashion. What we're witnessing here is the scheduling performed by the Linux kernel in an effort to give both our parent and child processes a turn at actually running. If you run *twoproc* repeatedly, you'll get a different sequence of stars and dots each time.

Another thing you'd notice about the output (if you took the time to count) is that there are always exactly 10000 stars and 10000 dots. That tells us that the parent and the child have quite separate copies of the variable *i*, each counting up independently in its own loop. What happens is that as the new process is created, a copy of the parent process's data segment and stack (the regions of memory where variables are stored) is created for the child.

### Executing programs

Creating child processes which continue to execute (different parts of) the same program as the parent is fairly common. For example, the *Apache* web server creates multiple child processes so that it can simultaneously field requests from many browsers.

Each child runs the exact same service loop in the *Apache* server. This “child per client” model is very common. However, another reason for creating a child is because you want it to go off and execute a completely different program. This is what the shell does, for example, when you ask it to run a command.

The system call **exec()** is used by a process to run a new program. To return to our actor analogy, a call to **exec()** is like a stage direction in *King Lear* that says “Go and read *Hamlet*”. The actor sets aside the old script and begins to read at the beginning of the new one. He’s the same actor, but the old script has been discarded, and a new one taken up. Similarly, when a process executes **exec()**, it remains the same process, but the memory image and data of the old program are discarded, a new code and data segments are established for the new program, and execution begins at the beginning of the **main()** function.

It turns out that there are six versions of **exec()**. Since we’re trying to have fun learning the concepts here, rather than getting bogged down in details, we’re going to look at just one – **execlp()**. (There are many books on Unix and Linux systems programming which would be delighted to fill you in on the details, if you’re interested.)

Here’s a typical call to **execlp()**:

```
execlp("sort", "sort", "-n", "foo", 0);
```

The first argument specifies the name of the program we want to execute. (It’s equivalent to the *Hamlet* in our “go and read *Hamlet*” analogy.) Linux will search for this program in all the directories listed in the **PATH** environment variable of the process, just as it does if you typed in the name of a program at a shell prompt. The remaining arguments specify the “command line” which will be passed to the new program, with the **0** argument marking the end of the list. So this particular example is equivalent to entering the shell command:

```
$ sort -n foo
```

Here’s a complete program that does an **exec**:

```
int main()
{
    write(1, "The date is ", 12);
    execlp("date", "date", 0);
    write(1, "Have a nice day", 15);
    return 0;
}
```

and here’s the output from it:

```
The date is Fri Aug 16 13:08:12 BST 2002
```

What happened to the “Have a nice day” message? Well, unless the **execlp()** call fails, control never returns back into this program. The process goes off and executes the *date* program, and that’s that. Once our actor has tossed aside *King Lear* and picked up *Hamlet*, there’s no going back.

So what’s the right way of doing what this example was trying to do ... print a message, run a program, wait for it to finish, then print another message? The solution is to create a new child process to run the *date* command, and have the parent wait until the child has finished before printing the second message.

The code looks like this:

```
int main()
{
    write(1, "The date is ", 12);
    if (fork())
        wait(0); /* Parent */
    else /* Child */
        execlp("date", "date", 0);
}
```

```
write(1, "Have a nice day", 15);
return 0;
}
```

The only new piece here is the call to **wait(0)**, which waits for the child process to finish. Without this call, there would be no guarantee that the “have a nice day” message would appear after the output from *date*.

## The shell in 13 lines

We are now in a position to write our masterwork for this month ... a tiny shell. Here, in pseudo-code, is what a shell does:

1. Print out a prompt
2. Read in a command line
3. Do a whole lot of substitutions and other mucking about with the command line, which we’re going to completely ignore here
4. Create a child process to execute the command
5. Wait for the child to finish
6. Loop back to step 1.

Here’s the code for our minimalist implementation:

```
int main()
{
    char command[80];

    while (putchar('>'), gets(command)) {
        if (fork())
            wait(0); /* Parent */
        else { /* Child */
            execlp(command, command, 0);
            printf("command not found\n");
            exit(1);
        }
    }
}
```

The test expression in the **while()** loop is one of those disconcerting pieces of C code that programmers like to show off with by demonstrating how much they can accomplish in a single line. It prints a prompt (the **>**), calls **gets()** to read a line from the keyboard into the array ‘command’, and tests to see if **gets()** encountered an ‘end of file’. The process then forks, and the child tries to execute the command. If this succeeds, control never returns to this program in the child. If the **exec** fails (maybe the command could not be found?) the child prints a message to say so, and exits. (**exit()** is another system call. It causes a process to terminate.) Meanwhile, the parent simply waits for the child to finish then loops round to prompt for another command. **Figure 4** summarises the normal process life cycle.

If you compile and run this example it will happily accept simple commands like **date** or **hostname** or **ls**. If you try anything fancy, like **ls -l** it will fail, because it will try to execute a program called, literally, *ls -l*. Our simple-minded implementation isn’t even able to parse the command line out into a command name and arguments. Well, what do you expect from 13 lines? You can exit from this program, just as you can from a regular shell, by entering your ‘end of file’ character (probably **^D**). This causes the **gets()** call to **return 0** and the **while** loop to terminate.

If you understand how this program works then you have pretty much mastered the Linux process model. [LXF](#)

## NEXT MONTH

Next month we’ll look at the system calls which allow us to read and write data in files, and how to access a file’s attributes such as its access permissions and timestamps. We’ll see how to write an important class of programs in Linux called filters, and develop a simple program which takes a “snapshot” of a file so that we can tell if it has been changed at a later date. (This is the basis of security tools such as *tripwire* which monitor the filesystem for unauthorised tampering.)

## WORKING WITH OPENOFFICE

# OpenOffice.org for power users

Over the next few months, **Neil Lucock** will be looking at the different tools in **OOo** and the features that make it so good, and showing you how to get the most from **OOo**.


The procedures described were done on a 700MHz Pentium 3 running KDE 2 under Mandrake 8.2. This tutorial assumes that you are using an Intel/AMD machine and know how to do the basics, such as navigate around the filesystem, become Root and make a new directory.

## What you get

You get tongue-tied, because they chose an awkward name. The phrase "open office" perfectly describes the suite. Unfortunately they can't call it that because someone else has already grabbed that name for their suite. Rather than think of something new, (*Office Libre, Buro Frei*, anyone?) they called it after their organisation, so we all have to call it the *OpenOffice.org* suite to avoid legal problems. We call it *OOo* in print to avoid repetitive strain injury from excessive typing.

Although it does not include the *Adabas* database that is included in *SO*, *OOo* has nearly every productivity tool you would need in an office environment. For something that's a free download, it's really the most amazingly good piece of software I've used. It isn't perfect, there are still remnants of *StarOffice 5.2's* interface design that have not yet been replaced, but there is now an effective free office suite that will allow a business to read its *Word* documents and *Excel* spreadsheets. You can still use that CD of clipart too. *OOo* uses windows metafiles (wmf) for its own clipart, it imports *Photoshop* psd files and has import filters for many common file types. However, it's one thing to list a file type under the "import" or "open" dialogues, I'll be looking to see how well the files created by the original programs open in the appropriate tutorials. I was impressed when I wrote the July

**Home user setup.**  
**Choosing the right file to install always helps.**  
**Extract everything to a temporary directory.**



*LXF* review of the presentation software element in the suite. If the *PowerPoint* import filter is representative, then most of the arguments about not being able to read your old correspondence and use your old work under Linux are no longer valid. A word processor that imports *Word*; a spreadsheet that knows *Excel*; a presentation tool that reads *PowerPoint* files, all contained in the same interface, makes a very good reason to look hard at *OOo*.

## Single user setup

Before we describe this, if more than one person uses your Linux box, or if you want to install to a user account, you will be better off using the "Network" install described later.


Very few Linux programs seem to have the one click setup procedures that you find under Windows, but it's really not difficult. The file you need is on the *LXF* July 2002 CD/DVD. You can also download it from [www.openoffice.org](http://www.openoffice.org) (if you have a fast connection or a lot of time) or buy it on a cheap CD from various distributors. Find the directory and the file called 'OOo.1.0.0\_LinuxIntel\_install.tar.gz'.

Find a place on your hard drive (I have a 'tempfiles' directory for this sort of thing) and copy the file there. It is useful to split the file manager window into two view-panes for this. Once it is on your hard drive, right-click on it and use the 'extract here' option. (If this does not work, you should navigate to where your system keeps the tool you use for making or extracting compressed archives, launch it and dump the contents of the tar.gz file into your temporary directory). You will then have a lot of files in the directory. Ignore them all except for the file called 'Setup'. Click this and it launches the installation program. From then on, just follow the prompts and click on the 'next' buttons. I did the 'standard installation' which put 183.6MB of files onto my drive. At the end of the process, it tells us that *OOo* is now on our hard drive, but how do we launch it?


There's two easy ways to put an icon on the desktop. If we right click, one of the options is to create a new "link to an



Find this file, click it and follow the wizard.



The wizard makes installing simple. Choose custom if you are short of space and don't want some parts of the suite.



Making a link for your desktop. Make sure you type the correct path into the execute box. The path to the file is displayed at the top of the 'Open' dialog box.

application". Once we have told it what name we want to use, we tell it where to find the program. OOo should have made a directory with the same name in your Home directory. Inside this are two files that you can link to. The one you want is called 'soffice' but the one called 'spadmin' will also launch the program. Make the 'soffice' file the target of the link. The other way to make an icon on the desktop is to drag and drop a copy of the 'soffice' file onto the desktop. Left-click and hold the mouse button to grab the file, release it over the desktop and choose 'copy here' from the small menu that appears.

It's also nice to have a *Kicker* menu entry. Mandrake provides a utility called *MenuDrake* which allows you to add entries to GNOME/KDE menus, – other distributions provide different tools. OOo provides an icon (found it in OpenOffice.org1.0/ share/icons. It's called 'ooo\_gulls.png') but I'm going to make my own.

*MenuDrake* shows a list of icons, but does not tell you where on the filesystem they are found and does not give you an option to navigate to anywhere else, so I'll need to find out where it keeps the icons it can access. First find an icon (from those

## History Lesson


### A star is born

*StarOffice* was originally produced by Stardivision, a German company, during the 1980s. Sun Microsystems bought the software during 1999 and released it as *StarOffice 5.2* during June 2000. It was a huge suite, it included a web browser interface, it could also make toast and also do the ironing. I think. It was also really slow. It came with a calendar (so you could see how slow it was). Nobody really liked the interface which took over the entire desktop and served to launch the various parts of the suite. Sun had another look at the code, adopted an open source development model and incorporated many improvements. The saved files are XML based, the file import filters are improved, it was a very different product that Sun released as the *StarOffice 6* and *OpenOffice.org* suites during 2002.

The major difference is that the integrated desktop interface has gone. It no longer acts as

a web browser/email client or calendar. The word processor, spreadsheet, drawing package and presentation software are there, there are a lot of other nice touches too. The business world uses *MS Office*. Sun's amazing contribution to the Open Source world might very well change that. Both suites open files from *Word*, *PowerPoint* and *Excel*, reliably and sometimes better than Windows opens them, in my experience.

*StarOffice* is Sun's commercial suite. It's based on the same code as OOo, but has lots of proprietary code in it that Sun can not release as Open Source because it does not own the code itself. With SO you get lots of time saving templates, OOo has most of the major functions (certainly all the vital ones) but you'll have to make the templates yourself. Having said that, it's free (as a download) so you can't complain. It's supposed to be the largest Open Source project yet.



offered) whose name is easy to remember, then search for the file to find out where it lives. Mandrake keeps the icons for KDE programs in `usr/share/icons`. That's where I'll put my new icon (or copy the supplied one).


I started OOo and took a screenshot of the splash screen as it started up. I trimmed it in *The GIMP* and resized it to 32x32 pixels. This was saved as a .png and copied into the `usr/share/icons` directory (you might have to be Root to do this). I then used *MenuDrake* to pick up the icon I made. I named the file (*OpenOffice.org*; it's the name that is displayed under the icon) and told *MenuDrake* the path to where the program lives. Once you've done that look on your *Kicker* menu and launch the program to prove that it works. If it does not, check that you have typed the path to the program correctly. Of course, this procedure can apply to any program that does not create menu entries and desktop icons.

## Network setup


The web site tells you to unpack the archive (again, use a temporary directory) and install the software in the `/opt` directory by becoming Root and typing

```
/install -prefix=/opt
```


Drag and drop a copy of the file onto your desktop.



Using **MenuDrake** to add an icon to the **Kicker** menu in KDE.



The **Kicker** menu with the new icon for **Ooo**.



The icon to make a new file. You can also use **File>New** to get here.

Return to your own user status and navigate to the OOo directory in `/opt`. Type

```
./setup
```

and ask it to do a workstation installation. It might ask you where Java lives on your machine, you can use the *whereis* tool in the shell to find the path to type in. The graphical *find files* will also give you results, but not as quickly. It listed Java in a few places, including in `mnt/windows`, not the path you want!

The last task is to delete the files that were copied to the temporary directory, just to free up some hard drive space. Installation is now complete.

## Working with OOo

Launch the suite and you are presented with the word processor. If that's not what you want, you can use the menu: `file>new` or click the icon (see screenshot) that gives the same options menu. OOo will then make a new spreadsheet (or whatever you chose) for you. Before you start using OOo, it's worth while having a look around the interface and setting a few options to make it easier to work with. Under `Help` turn on the extended tool tips. Hold the cursor over an icon and it tells you briefly what it does. (Some programs call this 'bubble help'). The extended tool tips give a short paragraph telling you how to use the tool. It's a great idea when you are learning new software. Once you know how to use it, you can revert to the standard.

Some OOo icons have a small green arrow to tell you that there's a submenu. Hold the mouse button down over it (long click). If you grab the title bar at the top of the menu that appears it detaches. If you are going to be repeatedly using a couple of tools, leaving the submenus on the workspace as

floating toolbars really speeds up the work. Some other icons have a long click function but no green arrow. The 'font color' and 'font background' icons change the highlighted text to the displayed colour with a single click. You need to do the long click to change to the colour you want.


OOo keeps the same basic interface but displays icons giving functions appropriate to the kind of document that you are creating, so if you create a spreadsheet, it displays cells in the work area and offers you spreadsheet functions on the toolbars.

## What it also does

In the articles that follow this one we will be looking at the major parts of OOo: the word processor, spreadsheet and presentation software. OOo is a cut down version of *SO*, so you get the same high quality tools that Sun sells as a rival to *MS Office*. What you don't get is all the extras that are in *SO*. What we are going to do now is show some of the unexpected qualities of OOo, make a business card to hand out and learn a little about OOo's way of working.

## Calling card

OOo can make business cards for you. Nearly every computer shop of any size sells attractive perforated cards in A4 sheets. Buy some that you like the look of. OOo includes a nice range of templates to create a suitable card. First start a new document and choose 'Business Card'. It gives you a tabbed dialogue box to set all the options. I chose Avery A4 as the Brand, even though the cards I bought were Decadry. It doesn't matter, everyone uses the same sizes of business cards. You'll have to hunt for the right format under 'Type'. Mine were 85mm x 54mm and C2364 seemed to be about right. You have to select one to see the details you need. Once you think you've got the right size, have a look under 'Format' to make sure that the margins at each side correspond to the margins on the cards you've bought. Go to "Options" to set your printer and tell it to print over the entire page. You can fill in the details that you want on your cards under the 'Business' or 'Private' tabs. Your last stop should be 'Business Cards'. This gives you a preview of what the card will look like. Once you have found one you like, click the 'New Document' button at the bottom. You are returned to the OOo interface with all your business cards across the screen. We now need to make a few changes. I wanted a picture on the cards, so I used the menu `Insert>Graphics>From File` to choose a picture. The dialogue box gives you a preview. I chose a Mandrake logo (what else is there to shout about except Linux?) and dropped it in the top left hand card. I had not put in a fax or telephone number in the fields under 'Business' but the 'Tel No.' and 'Fax' headings were still displayed. The picture had also moved the URL out of



How to put a picture on your business card.

## Made In Linux

### Blending in to the heterogeneous environment


OOo offers you the chance to get Linux onto the office desktop where you work. Tell people how good Linux is, ask the IT people how they like having to cope with paying for licences for software and also keeping records. I burn Mandrake GPL CDs and offer them to anyone who is interested. Mention that Linux is free (no paperwork is a good argument in every company, high quality and free is an even better one.), does everything that a Windows machine does, has no problems with viruses, unauthorised software installations or people

“borrowing” the licensed CDs for their machines at home. It’s designed to be secure; compare the number of Windows vulnerabilities described on The Register or Slashdot websites to the number of Linux ones.

OOo runs on Windows and Solaris. (The betas for Mac OS X have been released for testing at the time of writing). You can run it on your existing Windows machines at work, (ask permission first) just to show them what you are talking about. Get a couple of people interested in it, ask them for their opinions.

Show them what it can do. I do Railtrack *Powerpoint* presentations at home using OOo and email them to myself at work. Mention the fact that it was “Made in Linux.” Can they tell the difference?

Once a couple of people know how to use OOo, have seen what it can do, why not ask if they want to try running it on a Linux machine? Tell everyone you are trying to save the company money. Lots of it. Of course, the IT people might hate you. As long as they have Windows, there’ll always be work for them to do.



Ensure that the margins are correct. Holding the mouse over any item gives a “tooltips” help bubble.

sight. I used the View menu to zoom in on the card and edited it to make it look better. When you’ve finished return to the 100% page size. Click the ‘Synchronise Labels’ button and they all look like the one you’ve just edited. Put the perforated card in the printer and print out a masterpiece. You want to save the file. When I did this OOo put up a light bulb to tell me that the help system had something for me. I clicked it and found it was displaying a page of useful information on the different file formats it saves. If you use the word processor, it will tell you that it has corrected a spelling mistake. They have struck the right balance here.

The Help system tells you it’s got something for you, but does not interrupt your work. In many Open Source projects, the programmers make a fine program that works wonderfully, but never get round to writing a decent help system. If you can write, learn the software and volunteer to write help files. It’s a way to give something back to the Linux community. OOo has a very good help system, comparable to anything you’ll find commercially. I liked the bookmarks idea, it’s the first time I’ve seen that under Linux. If you need to return to a useful page, right click on the page and add it to your bookmarks. It also allows you to print out the page as a reference. If you want to delete a bookmark, select it in the Navigator window on the left (if it is not displayed, right clicking gives you the option).

## OOo Basic

This was a surprise. OOo contains a Basic programming language! There’s plenty of information in the Help system. I did not try this out, (I’m not a programmer) but you can write macros in the IDE (Integrated Development Environment) that is provided. With the



‘Synchronise Labels’ makes all the others like the one you were working on.

whole thing being modular, you can alter OOo to suit your tastes or needs. If you regularly need a particular function for your business, you can have someone write it and integrate it into your desktop. There’s also a developer kit available at the OOo website.


This is yet another argument for Open Source. Not only do you get more, it builds on the same principle that ID Software used with *Doom*. Here’s the game, here’s the tools, make some levels and release them. It adds value to the software which becomes more capable with time. There have been suggestions that the XML file standard used in SO and OOo should be adopted as a universal standard.

## OOo math

If you want custom formulas, there’s a nice module to make them.

## Conclusions

The more I learn about OOo, the more I’m impressed with it. There are a couple of odd behaviours, sometimes the menus don’t close if you don’t choose anything. You have to click somewhere on the desktop to deselect them. If you put one of the 3D objects into a business card as a picture, it does not copy it to the other cards when you try to synchronise. I’ll discuss any problems as well as all the great features over the next few months. [LXF](#)



Made in Linux and Indistinguishable from *PowerPoint*. This will be used (without the slogan) in Railtrack NW Safety Briefings.

## PERL MODULES

# Headlines

Make it easy to add headlines to your site or share yours with others. **Dave Hodgkinson** demonstrates the Perl modules that make this easy.



**S**yndication of your website content is a wonderful thing. It means that links to you get spread all over the Internet and it makes you look really knowledgeable and important. Maybe you have a website and you want to spruce it up a bit with headlines from somewhere else. Or maybe you have a site with sparkling, original and witty writings you want to share with the world. Or you just want to know how Slashdot.org gets all those headlines in the 'slashboxes' down the right column.

The answer is 'RSS', and in this article, I'll show you how it's done and which CPAN modules to use to make the process easy. RSS was introduced by Netscape as a way of getting headlines onto my.netscape.com. The principle of broadcasting news headlines on the 'Net goes back much further to the days when "push" was going to rule the world through Pointcast and CDF (Channel Definition Format) was being offered by MS. These days it's used much more to pull the news you want, when you want it.

So what is RSS? To quote: "Rich Site Summary (RSS) is a lightweight XML vocabulary for describing metadata about websites, ideal for news syndication"<sup>1</sup>. More accurately, it is actually the "RDF<sup>2</sup> site summary", where RDF expands to "Resource Description Framework". RDF is important since it is the foundation of the "Semantec Web"<sup>3</sup> which will make the web understandable to machines as well as humans.

A major user of RSS is the blogging community. Blogging started out as a system for linking to interesting resources and commenting on them. It has become a personal journal phenomenon. With something like 750,000 people keeping blogs, including some names like Moby and Wil Wheaton<sup>4</sup>, you know, Ensign Crusher from *Star Trek*.

For our purposes, that's pretty much all we need to know. I'm assuming you have a web site and you'd like to pull some headlines from another site and put them on yours. Later, we look at generating an RSS file so that your news can find its way elsewhere.

## Grabbing an RSS feed

Taking headlines from [www.slashrock.com](http://www.slashrock.com) as an example, we can very easily grab an RSS file and examine it. The important stuff is near the end with the ten latest headlines:

```
<item rdf:about=
"http://www.slashrock.com/article.pl?sid=02/08/21/2155206">
<title>lan Paice DVD now out</title>
<link>http://www.slashrock.com/article.pl?sid=
02/08/21/2155206</link>
<dc:creator>danielb</dc:creator>
<dc:subject>lanPaice</dc:subject>
<dc:date>2002-08-21T16:53:06+00:00</dc:date>
<slash:department>thunderous</slash:department>
<slash:section>articles</slash:section>
```

```
<slash:comments>0</slash:comments>
<slash:hitparade></slash:hitparade>
</item>
```

What is so great about these paragraphs of information is how utterly simple they are, and indeed it would be fairly easy just to use a regular expression to extract things we need from the RSS feed. However, we have some Perl modules we can use. For this example, since all I want to do is generate a perl data structure that maps roughly to the layout of the RSS feed, I will use **XML::Simple**. For dealing with more complex XML you might want to look at **XML::SAX** and **XML::DOM**.

The final part of our program will be to render the data structure out in some kind of simple HTML format. For this I'll use Andy Wardley's extremely powerful Template Toolkit<sup>5</sup>.

Note that I'm using **XML::Simple** to parse the data not **XML::RSS**. They both do the same thing. You would use **XML::RSS** if you wanted to do more manipulation on an RSS data set.

```
#!/usr/bin/perl -w

use strict ;

use LWP::Simple;
use XML::Simple ;
use Template ;

# get the rss data
my $content = get('http://www.slashrock.com/slashrock.rss') ;

if (defined $content) {
    my $data = XMLin($content) ;
    my $config = { INCLUDE_PATH => './include' };
    my $template = Template->new($config);

    my $vars =
    {
        item => $data
    };
```

## Links

- [1] RSS: [www.webreference.com/authoring/languages/xml/rss/1](http://www.webreference.com/authoring/languages/xml/rss/1)
- [2] RDF: [www.w3c.org/RDF](http://www.w3c.org/RDF)
- [3] Semantec Web: [www.w3.org/2001/sw](http://www.w3.org/2001/sw)
- [4] Wil Wheaton: [www.wilwheaton.net](http://www.wilwheaton.net)
- [5] Template Toolkit: [www.template-toolkit.org](http://www.template-toolkit.org)
- [6] MwForum: [www.mawic.de/mwforum](http://www.mawic.de/mwforum)
- [7] The Highway Star: [www.thehighwaystar.com](http://www.thehighwaystar.com)
- [8] Syndic8: [www.syndic8.com](http://www.syndic8.com)
- [9] Amphetadesk: [www.disobey.com/amphetadesk](http://www.disobey.com/amphetadesk)

```
$template->process('headlines.tt',
                  $vars,
                  "../include/slashrock.html")
|| die $template->error();
} else {
die "Couldn't get slashrock";
}
```

Rather than embed the HTML in the code, it's good practice to separate it out into a file and use one of the many templating systems. In this case, we will use the Template Toolkit:

```
[% USE date %]
<TABLE BORDER="0">
<TR><TD><B>[% item.channel.title %]</B></TD></TR>
[% n = 0 ; FOREACH i = item.item %]
<tr>
<td VALIGN="TOP">
<img SRC="/hs-pics/news.gif" WIDTH="18" HEIGHT="18"
BORDER="0" ALT="">
</td>
<td VALIGN="TOP">
<font size="-1">
<a href="[% i.link %]">[% i.title %]</a>
</font>
</td>
</tr>
[% n = n+1 ; %]
[% LAST IF n >= 5 %]
[% END %]
</TABLE>
```

Without a crash course in Template Toolkit's syntax, it should be fairly obvious that TT commands are delimited by **[% %]**, that the 'item' structure we passed in is iterated through in the var **i**, and that hash elements are dereferenced like **i.title**. So, as a result of running the above code, we should have some HTML suitable for inclusion on a webpage, say via server side includes, starting:

```
<TABLE BORDER="0">
<TR><TD><B>Slashrock</B></TD></TR>
<TR>
<td VALIGN="TOP">
<img SRC="/hs-pics/news.gif" WIDTH="18" HEIGHT="18"
BORDER="0" ALT="">
</td>
<td VALIGN="TOP">
<font size="-1">
<a href="http://www.slashrock.com/article.pl?sid=
02/08/25/2125242">Candice Night goes solo</a>
</font>
</td>
</tr>
```

## Creating an RSS feed

Now let's say we want to create an RSS file. Of course, there's a module to do that for us: **XML::RSS**. We could try doing this ourselves, possibly even using a template as above, but this module will be kept in step with any changes made to the RSS standard and therefore so will our code.

For this example, I am going to extract the ten most recently changed topics from the forums<sup>5</sup> on the Highway Star<sup>7</sup> website and make an RSS feed from that.

```
#!/usr/bin/perl -w
```

```
use strict ;
use DBI ;
use DBD::mysql ;
use XML::RSS ;

my $dbh = DBI->connect("DBI:mysql:database=purple1",
"username", "password");

# get the 10 most recently modified topics
my $statement = "SELECT
id, subject, lastPostTime
FROM
topics
LIMIT 10" ;
my $sarts = $dbh->selectall_hashref($statement,lastPostTime) ;

# set up the channel object
my $rss = new XML::RSS ;
$rss->channel
(
title => "The Highway Star Forums",
link => "http://www.thehighwaystar.com",
description => "Most Recent Topics from the Highway Star
Forums",
dc => {
date => scalar localtime,
subject => "",
creator => 'daveh@hodgkinson.org',
publisher => 'daveh@hodgkinson.org',
rights => 'Copyright 2002, The Highway Star',
language => 'en-gb',
},
syn => {
updatePeriod => "daily",
updateFrequency => "1",
updateBase => "1901-01-01T00:00+00:00",
}
);

# iterate down the list of articles in reverse article_id order
foreach my $art ( sort { $b <=> $a } keys %$sarts ) {
$rss->add_item
(
title => $sarts->{$art}->{subject},
link => "http://www.thehighwaystar.com/cgi-bin/mwf/
topic_show.pl?tid=" . $sarts->{$art}->{id},
description => $sarts->{$art}->{subject}
)
}

# and output the RSS file
$rss->save("../include/forums.rss") ;
```

## Syndicate!

You now have the basics with which to generate and read RSS. What you need to do know is let the world know you are out there. The way to do this is to let Syndic8<sup>8</sup> know about it. This is a database of some 16000 RSS feeds and is used by headline readers like *Amphetadesk*.<sup>9</sup> Now have the tools with which to grab the headlines you want and to disseminate your own. [LXF](#)

## About Dave Hodgekinson

Dave has been noodling on the Internet since before it was the Web. While not doing strange things with *Apache*, *Perl* and *MySQL*, he is editor-in-chief of *The Highway Star*, the one and only Deep Purple fan site. [daveh@hodgkinson.org](mailto:daveh@hodgkinson.org).

## NEXT MONTH

Charlie Stross jets back in to talk about Parrot, the virtual machine for executing Perl 6 (& other language's) bytecode.

## LIGHTWEIGHT COMPONENTS

# Speaking Java



When AWT doesn't provide the widgets you need, make them yourself. **Richard Drummond** explains.



You can find the source code to accompany this article on the coverdisc in the folder Magazine/Java.

We will at last get around to constructing a GUI for the index searching applet we've been building. The problem is, however, how to display the results of a query to the user. Ideally, we would generate a page of HTML and display that in the host browser to present the results in a manner similar to any of the popular Internet search engines. But, while the Applet API will let you redirect the host browser to a new page, we cannot write files from within the secure sandbox that an applet runs in – not unless we've been given additional permissions, and that is not a simple task for the user to set up. We must therefore consider how to create a GUI to display the results, and here we are up against the simplicity of the AWT widget set (we can't use *Swing* widgets, because we want the applet to work on open source JVMs such as *Kaffe*, which don't supply an equivalent to *Swing*).

The solution is to create our own widget or widgets based on the AWT framework, and include them with the applet. This solution may not be the simplest, but it is elegant and portable – and has the added bonus that we can re-use any components that we create in other projects.

The problem, then, is how to present a list of search results – that is, package names, descriptions, locations on the disc, etc – and allow the user to select a result to bring up that entry's HTML page in the browser. Possibly the most desirable solution would be to create a widget for rendering HTML, but that is just far too involved for such a small project. The approach I will adopt is to create a multi-column list widget.

## DIY widgets

First, how do we create our own AWT widgets? Well, strictly speaking, we can't. The AWT widgets are 'heavyweight' components, meaning each widget has a native counterpart or peer. When you instantiate an AWT button, for example, the underlying, native toolkit (which is usually *Motif* on Unix platforms) creates its own button – the native peer – which does all the work of on-screen rendering and communicating with the windowing system; the AWT button is simply a Java wrapper for that native widget, if you like. What we can do is create 'lightweight' or 'peerless' widgets that fit in with the AWT framework (and this is exactly what the *Swing* components are).

To do this, we simply subclass the AWT **Canvas** class and override the necessary methods to implement the functionality of our widget. So, for example, we would override the **paint()** and **update()** methods to draw our widget physically on screen, and we would override **processKeyEvent()**, **processMouseEvent()**, **processMouseEvent()**, etc. to handle any input events we are interested in. For example, in our **processMouseEvent()** we would interpret any mouse clicks within our widget, update the widget's state accordingly, possibly schedule any redrawing required, and generate any appropriate 'output' events, such as

**ItemEvent** events if our widget supports selection or **ActionEvent** events if we support any kind of semantic actions, such as button presses. (We must remember to manually enable any classes of event we wish to listen to, though, since these are not enabled automatically.)

The following is a rough skeleton of a custom, lightweight widget.

```
class OurWidget extends Canvas
{
    public OurWidget() {
        super();
        // handle our initialisation;

        // we want key to listen to events
        enableEvents( KEY_EVENT_MASK );
    }

    public paint( Graphics g ) {
        // draw our widget
    }

    public update( Graphics g ) {
        // re-draw our widget
    }

    public processKeyEvent( KeyEvent e ) {
        // handle key presses
    }

    // etc.
}
```

That's all there really is to it. The devil is in the details of course.

## The MultiList widget

For this project, we wish to create a multi-column list widget to display the search results from our index applet. You should be familiar with the kind of thing I mean; just take a look at your email client. This widget is made up of three components. The list body – which we implement as the **MultiListBody** class – displays the data in the list in a tabular form, and lets the user select a row (multi-selection is for another day); the list header – the **MultiListHeader** class – displays the column titles and lets the user re-size the column widths; and horizontal and vertical scrollbars allow the user to scroll around the list body when it contains more data than can be displayed on screen and for these we can use standard AWT scrollbars. The **MultiList** class itself will simply be a container (using a **BorderLayout**) which joins all these pieces together and oversees the passing of events between them.

The other piece of the puzzle is the interface between the widget and the actual data we wish to display. We want to make the list widget as re-usable as possible, so we don't want to restrict the type of data we can display in it. Therefore, rather than allowing the list widget direct access to some data structure containing the list data, we only allow the list access via a wrapper provided by the **MultiListModel** interface. This interface provides methods to return the number of rows and columns in the list and to access the data in a particular row and column. Thus the list data can be stored in any form – an array of arrays or a **List** of **Vectors**, say. We simply need to create a class which implements **MultiListModel** so that the list widget can use it. (At the moment, our list and **MultiListModel** assume that the underlying list won't be modified while being displayed in the widget. Strange things will happen if it is.)


For now, we will concentrate on the **MultiListBody** class, the part of the widget that does the hard work of drawing the list data on screen. Have a look at the code supplied on the coverdisc.

The actual job of rendering the list is straightforward. What makes it more complex is the need to optimise for speed. Thus we don't want to redraw the whole list every time it is scrolled or the selection point is moved. We just want to draw the parts that have changed. Thus our **update()** method makes use of a member variable **updateStatus** to check what needs redrawing and acts according. When the list is scrolled vertically, for example, **updateStatus** is set to **UPDATE\_VERTICAL\_SCROLL** and an update scheduled with a call to **repaint()**.

The next time the list's **update()** method is invoked by Java, we know that the list needs to be scrolled vertically and we know by how much, so we can perform the scroll operation. First we blit the area that has moved, and then redraw only the new content that has been made visible by the scroll operation (by installing a suitable clipping region with the **setClip()** call). The following code fragment illustrates this.

```
// do we need to do a vertical scroll?
if( (updateState & UPDATE_VERTICAL_SCROLL) != 0 )
{
    int dy=yOffset-newYOffset; // how much to scroll by

    if(dy<0)
    {
        // scroll up//
        g.copyArea( 0, -dy, getWidth(), getHeight()+dy, 0, dy );
        g.clearRect( 0, getHeight()+dy, getWidth(), -dy );
        g.setClip( 0, getHeight()+dy, getWidth(), -dy );
    }
    else /* dy>=0 */
    {
        // scroll down
        g.copyArea( 0, 0, getWidth(), getHeight()-dy, 0, dy );
        g.clearRect( 0, 0, getWidth(),dy );
        g.setClip( 0, 0, getWidth(),dy );
    }
    yOffset=newYOffset;
    // reset update flag
    updateState &= ~ UPDATE_VERTICAL_SCROLL;
    // redraw table body
    paintTableColumns( g );
}
```



This shows our list widget running on Sun's JDK 1.4 (left), and the Kaffe VM (right).

Another tricky aspect of writing a custom widget is maintaining the geometry of the widget. When the widget gets resized – either because it has just been created and displayed, or because the user has resized the window it lives in, or whatever – you need to react and adapt to that change in size. While you could override all the various **setSize()**, **setBounds()** and **reshape()** methods, an easier method is simply to listen to component events by overriding **processComponentEvent()**. Here's how we do it.

```
public void processComponentEvent( ComponentEvent e )
{
    super.processComponentEvent( e );
    if( e.getID() == ComponentEvent.COMPONENT_RESIZED ) {
        // Cache our new size - remember JDK 1.1 has
        // no Component.getHeight()/getWidth
        Dimension dim = super.getSize();
        this.height = dim.height;
        this.width = dim.width;

        // recalculate our geometry here . . .
    }
    else if( e.getID() ==
ComponentEvent.COMPONENT_SHOWN ) {
        updateState |= UPDATE_REFRESH;
        repaint();
    }
}
```

## Work in progress

You can get a better idea of how to create a custom **AWT** component by looking at the **MultiListBody** class supplied on the coverdisc. This is unfinished, though, so expect a few bugs and unimplemented features. As it stands, the event handling is not complete and the resizing of columns is not included at all. I'll present the completed source code to the widget on next month's coverdisc.

If you want to try it out now, give the **DemoList** and **DemoFrame** classes a whirl. The **Ant** build script will build the example as a jar which you can execute. The build script specifies the **-target** switch to the Java compiler with the parameter **1.1**. This is to ensure that bytecode suitable for JDK 1.1-compliant JVMs will be generated. We need this if we compile under JDK 1.4, but wish to use the compiled code on an earlier JVM. [LXF](#)



## BUILDING A PROCESS VIEWER


# Who needs a GUI anyway?

**PART 15** In the last instalment in this tutorial series **Brian Long** looks at building console applications.

**S**adly this is the last part of this series on *Kylux Open Edition*; we hope you've found the subjects useful. This month we are finishing by looking at applications that don't have a GUI, ones that can be run from any shell prompt, including those outside X.

Command line applications are often written in C, C++ and Perl but *Kylux* is perfectly capable of joining the ranks of tools that can be used to build command line utilities.

Figure 1: The **File | New...** dialog.



## Let's get started

The first step to getting a command line application in *Kylux* is to choose **File | New...** to invoke the **Object Repository** dialog (see **figure 1**). Select the **Console Application** icon and press **OK**. This gives you a small project file, with no accompanying units and no classes defined in it.

```
program Project1;
```

```
{$APPTYPE CONSOLE}
```

```
begin
```

```
end.
```

Of course as you develop your console applications it will be appropriate to add units to modularise it, but the fresh project doesn't have any.

This is the bare bones console app with just a project file and no units at all. If you run it from the IDE and *Kylux* was launched from a console window (rather than a desktop menu), then the application's output will appear in that console. If not, the output will go off into the ether.

If you didn't launch *Kylux* from a console then you can either run your programs manually from a console prompt or alternatively try an IDE option intended to resolve the issue. The **Run | Parameters...** dialog has an option that runs your console program via a launcher – in short it launches a copy of *xterm* and runs your program from there (see **figure 2**). To be fair we have had difficulties getting this option to have the intended effect in any version of *Kylux* (1, 2 or 3) but that may be our installation setup. We recommend starting *Kylux* from a console window.

Alternatively you can run it from a command prompt. GUI applications require you to run the *kyluxpath* script from the *Kylux* bin directory (as explained in the online article <http://bdn.borland.com/article/0,1410,28057,00.html>) to ensure

you get access to the *Qt* library, required for *VisualCLX* controls. However console applications have no dependency on *CLX* so this is not necessary.

The following slightly modified project includes one output statement and you can see what it produces in **figure 3**.

```
program Project1;
{$APPTYPE CONSOLE}

begin
  WriteLn('Hello world')
end.
```

Notice that *Kyl1x Open Edition* applications always try to advertise their development tool by default. GUI applications display a temporary non-modal window on startup with a little message, which can be disabled by passing the **-ns** command line switch. Console applications display a little *Kyl1x* advert as their first line of output unless you remove the **\$APPTYPE** compiler directive from the project file.

The directive actually does nothing useful anyway so it is safe to remove (just delete the whole line). It has more of a point in *Delphi* on the Windows platform where the compiler has to do special things to make a working console application. With it gone we get only what we specifically write out, as shown in **Figure 4**.

## Standard file handles

Linux (and C/C++) uses three standard file handles for dealing with the keyboard (stdin) and screen (stdout and stderr). *Kyl1x* surfaces these through three text file variables **Input**, **Output** and **ErrOutput**. You can read from **Input** to read keyboard input and you can write to **Output** or **ErrOutput** to display text on the screen as we shall see.

These variables support shell file redirection, so the user could run your application with the stdin file redirected to a file on disk. Similarly the output could be redirected to another disk file, as can the error output.

## Reading and writing

Now on to the more interesting stuff, actually doing the reading and writing to the standard file variables... The Pascal routines we need to be familiar with are **Read**, **ReadLn**, **Write** and **WriteLn**. These are just a few of the available file manipulation routines in the Object Pascal language (though you should note that as of *Kyl1x 3* and *Delphi 7* the Object Pascal language has been renamed the *Delphi* language – *Kyl1x 3* allows development in either the *Delphi* or C++ languages).

These routines all operate on a file variable. If the file variable is specified, then things are straightforward, but if no file variable is specified then **Read** and **ReadLn** assume they should read from **Input**, whereas **Write** and **WriteLn** default to writing to **Output**. The code listing above used **WriteLn** with no file




Figure 2: Ensuring console apps work from the IDE.




Figure 3: A trivial (and small) command line application.

variable. It wrote the specified text to stdout and followed it by a new line (that's what the **Ln** part implies). To show the idea, here are a number of equivalent statements written in various ways.

```
WriteLn('Hello world');
WriteLn(Output, 'Hello world');
Write('Hello world'#10);
Write('Hello world', #10);
Write(Output, 'Hello world', sLineBreak);
Write('Hello', #32, 'world', sLineBreak);
Write('Hello world'); WriteLn;
```

Each line achieves the same thing. Note that **Write** and **WriteLn** are special routines in *Kyl1x* as you can pass as many values to them as you wish (of varying types, such as strings, characters, integers, floating point numbers or Booleans). Each value is written in the order passed and a line feed is automatically output if **WriteLn** has been called. You can read more about these routines in the online help. You can also find




Figure 4: A console app without a nag line.




Figure 5: Where to learn about the file handling system in the Delphi language.

« more about file processing in general in the *Object Pascal Reference Guide* (choose it from the Help menu). A suitable help topic is highlighted in **figure 5**.

This code represents a simple program that reads from stdin and writes to stdout (and sometimes to stderr). It keeps track of how many lines have been read and simply prefixes them with a

little description before outputting them. If the string happens to be longer than some arbitrary limit (30 characters) a message is sent to stderr instead of stdout.

```
program Project1;
var
  Line: String;
  LineNum: Integer;
const
  MaxLineLength = 30;
begin
  LineNum := 0;
  while not Eof(Input) do
  begin
    ReadLn(Input, Line);
    Inc(LineNum);
    if Length(Line) <= MaxLineLength then
      WriteLn(Output, 'Line ', LineNum, ' is: ', Line)
    else
      WriteLn(ErrOutput, 'Line ', LineNum, ' is too long
        (more than ', MaxLineLength, ' characters)');
  end;
end.
```

If the program is run from a prompt you can type in lines and they will be spat back out at you, modified as per the source.

**figure 6** shows a sample session. You can see that most lines were accepted and prefixed appropriately, except the longer fourth line. The program was terminated by pressing **Ctrl-C**.

As you might expect, console redirection is supported so you can feed input to the program from a file, redirect output to a file and even redirect error output to a different file. Assuming the existence of a text file *test.in* this command would take input from that file and would create the file *test.out* to store the results:

```
./Project1 <test.in >test.out
```

Anything written to **ErrOutput** would still appear on the screen as before. This version of the command would redirect both stdout and stderr to the same output file, *test.out*:

```
./Project1 <test.in &>test.out
```




Figure 6: Reading from and writing to the console.

## Code Templates

### At your fingertips

Note that when implementing classes, or entering any other common code snippets, you can save on the typing by using *Kyl* Code Templates. We looked at **Code Completion (Ctrl+Space)** and **Code Parameters (Ctrl+Shift+Space)** some while ago, but **Code Templates** is another editor feature.

There are a number of predefined **Code Templates** that can be entered by pressing **Ctrl-J** and selecting from the list (see **fig 10**). The text of the **Code Template** will then be entered into the editor with the cursor left

in the most likely place you will need to type something.

You can get a shorter list by typing the first letter(s) of the **Code Template** you need. For class declarations you could enter **cl** and then press **Ctrl-J** to see only the pertinent ones listed (see **figure 11**).

You can set up your own **Code Templates** on the **Code Insight** page of the editor options dialog (Tools | Editor Options...), which looks like **figure 12**. Any commonly entered text, such as calls to **MessageDlg**, **ShowMessage** or **Format**, or even a simple

**begin...end** statement would be useful to make new **Code Templates** for.

Press the **Add** button in the **Code Templates** group box and enter a name and description of your new **Code Template** (see **figure 13**). Then you can enter the code snippet you want to be represented by the **Code Template**. An example is shown in **figure 14**. Notice the use of the pipe sign to specify where the cursor should be left (in this case between the quotes of the string parameter).




Figure 10: Selecting a Code Template from the list.




Figure 11: Selecting a Code Template from a smaller list.




Figure 12: Customising the editor productivity options.




Figure 14: A simple Code Template for a call to MessageDlg.

Figure 13: Adding a new Code Template.



◀◀ Finally, this command would redirect stdin to read from *test.in*, stdout to *test.out* and stderr to *test.err*.

```
./Project1 <test.in >test.out 2>test.err
```

You can see these commands being executed in **figure 7**.

## Checking for parameters

For a console application to be flexible it must support command line parameters (after all, there is no menu available to choose from). There are several ways to access the command line used to invoke your program and identify the various passed arguments.

The recommended approach is to either use the **ParamCount** and **ParamStr** functions, which give you access to each individual argument as a string, or to use **FindCmdLineSwitch**, which lets you search for a specific parameter you are expecting.

Let's take the latter option first. You'd use **FindCmdLineSwitch** if you are expecting specific individual command line options rather than lots of command switches that can be merged together in various permutations. So if you want to respond to **-help** and **-version** options in a case-sensitive fashion you could write code like this:

```
uses
  SysUtils;
...
if FindCmdLineSwitch('help', ['-'], False) then
  WriteLn('Help message, blah, blah');
if FindCmdLineSwitch('version', ['-'], False) then
  WriteLn('Version 1 and a bit');
```

However there are two additional overloaded versions of this routine. The first one assumes case sensitivity and so takes just the first two parameters. The second one assumes case sensitivity and the use of **-** as the command line switch

prefix character. So you could abbreviate this to:

```
if FindCmdLineSwitch('help') then
  WriteLn('Help message, blah, blah');
if FindCmdLineSwitch('version') then
  WriteLn('Version 1 and a bit');
```

Of course it is quite typical for Linux applications to take various single letter command line switches and in most cases these can be combined into "compound" command line options. For example, if the program supports **-i**, **-v** and **-o** command line options then **-iv** or **-oi** or **-ivo** would also be perfectly acceptable. Due to the number of possible combinations, **FindCmdLineSwitch** is not so suitable here so we should look at the other options.

The code below represents a simple project that checks all its command line parameters for validity. **ParamCount** tells you how many parameters have been passed and **ParamStr(n)** returns the *n*<sup>th</sup> parameter. The code loops across each parameter and checks that it starts with the appropriate **-** character and that there are additional characters after it before proceeding to look at each character individually (a nested loop). Recognised command line characters are reflected by messages sent to stdout, but a bad character causes a message to go to stderr and the code bails out.

```
program CmdLineSwitches;
uses
  SysUtils;
var
  ParamNo, SwitchNo: Integer;
  Param: String;
  BadSwitch: Boolean;
```

```
Konsole - blong@Pyramid-Mandrake81.Polyhedra: /home/blong/tmp/exe_bin - Konsole
[blong@Pyramid-Mandrake81 exe_bin]$ ./Project1 <test.in
Line 1 is: First string
Line 2 is: 2nd string
Line 3 is: String no. 3
Line 4 is too long (more than 30 characters)
Line 5 is: A shorter string
[blong@Pyramid-Mandrake81 exe_bin]$ ./Project1 <test.in >test.out
Line 4 is too long (more than 30 characters)
[blong@Pyramid-Mandrake81 exe_bin]$ ./Project1 <test.in >test.out 2>test.err
[blong@Pyramid-Mandrake81 exe_bin]$
```

Figure 7: Console I/O supports redirection of input, output and errors.

```
Konsole - blong@Pyramid-Mandrake81.Polyhedra: /home/blong/tmp/exe_bin - Konsole
[blong@Pyramid-Mandrake81 exe_bin]$ ./CmdLineSwitches -o -itv
-o switch detected
-i switch detected
-t command-line option not supported
[blong@Pyramid-Mandrake81 exe_bin]$
```

Figure 8: Detecting command line switches, with **ParamCount** and **ParamStr()**.

```

Konsole - blong@Pyramid-Mandrake81.Polyhedra: /home/blong/tmp/exe_bin - Konsole
[blong@Pyramid-Mandrake81 exe_bin]$ ./pslst
pslst Version 1.0
Written in Borland Kylix by Brian Long

Syntax:
  pslst [options]

options can be one of:
  k PID  - kill specified process (PID is decimal Process Identifier)
  a      - list all processes
  n      - list processes not owned by root
  r      - list only running processes
  u      - list only user processes
[blong@Pyramid-Mandrake81 exe_bin]$ ./pslst r
PID      User State  VSZ  RSS Command
8354     blong      R 28900 21640 /etc/X11/X :0 -deferglyphs 16
8657     blong      R  2028   856 ./pslst r
[blong@Pyramid-Mandrake81 exe_bin]$ █

```

```

begin
BadParam := False;
for ParamNo := 1 to ParamCount do
begin
  Param := ParamStr(ParamNo);
  if (Length(Param) > 1) and (Param[1] = '-') then
  begin
    for SwitchNo := 2 to Length(Param) do
    case Param[SwitchNo] of
      'i': WriteLn('-i switch detected');
      'o': WriteLn('-o switch detected');
      'v': WriteLn('-v switch detected');
    else
      begin
        WriteLn(ErrOutput, '-', Param[SwitchNo],
' command-line option not supported');
        BadParam := True;
        Break
      end
    end;
  end
  if BadParam then
    Break
  end
  else
    begin
      WriteLn(ErrOutput, Param, ' command-line option not
supported');
      Break
    end
  end
end
end.

```

Of course, this is simple code and you would need to extend it to set appropriate flags for the options found (instead of writing out messages) and in many cases these command line switches are followed by values. Depending on your needs you will need to

butcher this basic framework to a greater or lesser extent. You can see an invocation of this program in **figure 8**.

## A sample console app

Last month we built a GUI process viewer. This month's project uses most of the code from that application to build a command line utility that can list processes (in varying ways) and kill a specified process.

The project is called *pslst.dpr* and can be found on this month's disc along with last month's GUI version, *pslst.dpr*. Since most of the code is the same the only interesting part to show is the command line parameter analysis.

In this case the app takes one of five possible command line options and this simple application expects only one of them to be passed. If this is done the program will either kill the process identified by the PID (the **k** command line option), list details of all processes (**a**), the current user's processes (**u**), running processes (**r**) or processes owned by users other than root (**n**). If not, you get a help screen explaining the syntax.

You can see a couple of invocations of the program in **Fig 9**, one to get the help screen and one to list running processes. As you can see the utility lists out the process PID, the owning user, its state, its virtual memory size, resident memory size and the command line that invoked it.

## Summary

We have seen that *Kylix* is just as suited to building command line utilities as it is for building full-scale graphical X applications.

Please keep an eye out for *Kylix 3*, which is out now. It includes not only the next version of the *Delphi* language version but also a completely compatible C++ version. You can use your knowledge of the *CLX* library in C++ apps if you require, and can also mix Pascal and C++ source files in the same project.

Anyway, that concludes this tutorial series, we hope you enjoyed it. Good luck with your *Kylix* programming. 

Figure 9: The command line process list application.

## About Brian Long

Brian Long is a UK-based freelance trainer and problem solver for Borland's *Kylix*, *Delphi* and *C++Builder* packages. His Web site is at [www.blong.com](http://www.blong.com) and he can be emailed at [brian@blong.com](mailto:brian@blong.com).

## MYSQL AND PHP

# Practical PHP programming



Continuing last month's SQL introduction, Paul Hudson shows you how PHP easily brings databases to the web.



As seen last issue, Structured Query Language (SQL) is a very high-level language that lets you manipulate data with the minimum of fuss. Also, if you've read this far, you probably also agree with me in thinking that PHP is the best choice for a multi-purpose web development language. This month we'll be looking at how to bring these two wonders together to form an unbeatable development team, and also see how you can make your code more flexible by normalising your data.

As a point of note, if you followed last month's article you should be fairly proficient in using the *MySQL* monitor to work with your databases. Simple as it may seem, that's definitely a skill you'll want to keep, trust me – there is no better SQL debugging environment than the plain old *MySQL* monitor.

## Connecting to your server

Before you can work with your data, you first need to connect to your *MySQL* server using PHP. This is accomplished using the function `mysql_connect()`, which takes three important parameters and returns a link to your server. The three parameters in order are the name/IP of the server to connect to, the username to provide, and the password to provide. If the function fails, the result PHP sends back is false, otherwise your connection succeeded.

Once your server connection is open, it will almost certainly stay open for the life of your script, and you can proceed to send queries. Once you're finished with the server, it is good practice to call the function `mysql_close()`. Note that if you do not call `mysql_close()`, PHP will automatically terminate the connection for you when your script ends. Here is what we have so far:

```
<?php
$link = mysql_connect("localhost", "usernamefoo",
"passwordbar");
// amazing web site here...
mysql_close($link);
?>
```

For fool-proof code, it is best to check the result of `mysql_connect` – this can be done simply with `if ($link) {`, as `$link` will be false if the connection failed, the `if` statement will evaluate to false.

It is very common practice to follow up your `mysql_connect()` function with a `mysql_select_db()` function. If you think back to last month, you will remember we had to execute a `USE` command before we could operate on a database. This is precisely what `mysql_select_db()` does – you specify the name of a database to use as the first parameter and a *MySQL* server link as the second parameter, and all subsequent queries to that connection will be on your chosen database.

The second parameter, the link identifier to your *MySQL* server, is entirely optional – if you don't specify it, PHP will automatically use the most recently connection. This is normally the behaviour you want, so feel free to omit the second parameter here.

## Fixing Errors

### Three easy ways to debug SQL and PHP

Once you're away from the *MySQL* monitor, it's very common to see messages like **0 is not a valid MySQL resource index**, which basically translates to "Your query is wrong – sort it out!" Naturally, this isn't very informative, and there are three basic ways to learn more about the problem.

#### 1. `mysql_error()`

Use `echo mysql_error()`; to have PHP print out the last error message received from *MySQL*. This is likely to be something similar to **Unknown field FOO in table BAR**, and is usually enough to help you locate the problem.

#### 2. `echo "SELECT * FROM..."`

Just printing out your SQL from your script can

be very helpful. This shows you what was passed to *MySQL*, because PHP replaces any variables used with their values, showing you the complete SQL statement.

#### 3. Use the *MySQL* monitor

If you've got no idea why something isn't working, give it a try from the *MySQL* monitor (see last issue for instructions). Very often you might just have gotten the syntax slightly wrong, or forgotten a quote.

Remember, if you ever encounter a problem you can't fix, *MySQL* has almost as good channels of support as PHP has. The *MySQL* docs are complimented by its excellent mailing lists. Visit [www.mysql.com](http://www.mysql.com) for more information.

## Starting off with simple queries

Connecting is the easy stuff, and, because you have pretty much learnt all there is to know about it already, we can move onto the core *MySQL* function in PHP – `mysql_query()`.

`mysql_query()` does exactly what it says on the tin – it sends a query to your *MySQL* server, and, where appropriate, returns data for you to read. It takes one parameter of interest, which is the SQL you wish to execute, and it returns a resource identifier. If your query returns no data, which is quite normal if it was **INSERT INTO**, **DELETE FROM**, or the like, then `mysql_query()` will merely return true if the call succeeded, or false if it failed.

The return value, a *MySQL* result resource identifier, can be compared to a summary of the query that also includes the query information you requested. As well as just reading the matching data from the tables you queried, you can also perform a variety of other functions on a *MySQL* resource to find out more about the data – one of which is `mysql_num_rows()`. This

function takes just one parameter, the resource to work with, and it returns the number of rows that matched your query. Here's an example for you to see how it works:

```
<?php $link = mysql_connect("localhost", "usernamefoo",
"passwordbar");
mysql_select_db("lxfdb");
$result = mysql_query("SELECT * FROM staff;");
echo "Query returned " . mysql_num_rows($result);
mysql_close($link);
?>
```

In the above example, we execute an SQL query we used last month – **SELECT \* FROM staff**, or "List all data from all rows in the staff table". The result of this query is placed into **\$result**, which we then query using **mysql\_num\_rows()**.

We could just as easily have used another helpful function, **mysql\_num\_fields()**. Can you guess what it does? :-)

Of course, very often the most important thing to do with a result index is to actually read the content it returned. There are several ways that this can be done, but my personal preference is to use the function **mysql\_fetch\_array()**, as it provides the most functionality.

**mysql\_fetch\_array()** takes one key parameter, which is your result index, and it returns an associative array of the next row of your data, or false if there are no (more) rows. Every time it is called, it increments its row count so that it cycles through all the rows in its result as you call it. It's much easier if I show you how this works in code, so here goes!

```
<?php
$result = mysql_query("SELECT Name, Age FROM staff;");
while ($r = mysql_fetch_array($result)) {
    echo $r['Name'], ' is ', $r['Age'], ' years old.<BR>';
}
?>
```

We're using the same query as before, so it's safe to ignore the first line. However, notice how a **while** loop is used with **mysql\_fetch\_array()**, and the condition for the loop is **\$r = mysql\_fetch\_array(\$result)**. Basically, this boils down to calling **mysql\_fetch\_array()** again and again, assigning the array it returns to **\$r**, until there are no more rows, at which point the loop terminates.

**mysql\_fetch\_array()** returns your data as an associative array where each field name for a row is an array key, and each field value is an array value. So, we can read **\$r['Name']** to see the **Name** field for the current row.

As you can imagine, this makes life very easy – including connecting, selecting a database, and closing a connection, we can echo out the results of a query in just seven lines of code!

```
mysql> DESC staff;
+-----+-----+-----+-----+-----+
| Field | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+
| Name  | varchar(255)  | YES  |     | NULL    |      |
| Age   | int(11)       | YES  |     | NULL    |      |
| Job   | varchar(255)  | YES  |     | NULL    |      |
| Pay   | int(11)       | YES  |     | NULL    |      |
+-----+-----+-----+-----+-----+
4 rows in set (0.01 sec)

mysql>
```

If you missed last issue, here's the schema for our staff table.

The screenshot shows a web browser window with the URL [http://www.mysql.com/doc/en/String\\_functions.html](http://www.mysql.com/doc/en/String_functions.html). The page content includes:

```
strings to be concatenated:

mysql> SELECT CONCAT_WS(", ", "First name", "Second name")
-> 'First name,Second name'
mysql> SELECT CONCAT_WS(", ", "First name", NULL, "Last name")
-> 'First name,Last name'
```

Other functions listed include:

- LENGTH(str)**
- OCTET\_LENGTH(str)**
- CHAR\_LENGTH(str)**
- CHARACTER\_LENGTH(str)**

Returns the length of the string **str**:

```
mysql> SELECT LENGTH('text');
-> 4
mysql> SELECT OCTET_LENGTH('text');
-> 4
```

Note that for **CHAR\_LENGTH()** and **CHARACTER\_LENGTH()**, multi-byte characters count as one.

- BIT\_LENGTH(str)**

Returns the length of the string **str** in bits:

```
mysql> SELECT BIT_LENGTH('text');
-> 32
```

- LOCATE(substr, str)**
- POSITION(substr IN str)**

Returns the position of the first occurrence of substring **substr** in string **str**.

You may find it easier to use the **extract()** function to convert all the variables from the array **\$r** into standalone variables in their own right, but this is not mandatory – it just makes reading your code easier on the eye.

Let's take a look at a slightly more complicated solution, this time incorporating **mysql\_num\_rows()**:

```
<?php
$result = mysql_query("SELECT Name, Age FROM staff;");
$numrows = mysql_num_rows($result)
if ($numrows) {
    echo "Your query matched $numrows records:<BR><BR>";
    while ($r = mysql_fetch_array($result)) {
        extract($r);
        echo "$Name is $Age years old.<BR>";
    }
} else {
```

Whilst not quite as good as the PHP documentation, the **MySQL** manual is still very thorough.



```

<< echo "Your query matched no records";
    }
?>

```

With the added `mysql_num_rows()`, we are able to output a message to our visitors if no rows are matched. Remember that our `while` loop only works when `mysql_fetch_array` doesn't return false – if your query returns no rows, the while loop will not execute even once, so users would see nothing unless you used `mysql_num_rows()` to check beforehand!

Also note how I've used `extract()` to change `$r['Name']` into just `$Name`. I'm doing this merely for the sake of clarity – you go your own way.

In SQL queries, you can use the special keyword **AS** to define the name with which you wish to refer to a field. This may not seem important immediately, but it really is very helpful when using *MySQL's* built-in functions. Later on we'll take a short look at the extra functionality *MySQL* can provide for you through its built-in functions, but for now I want to show you how you can format your database results in a more attractive manner.

## Formatting query results

Very often the most desirable thing to do with database results is to tabulate them neatly on your page to make for the easiest reading. Here is a complete example of how you can connect to a database, read data, output it with formatting, and tidy up:

```

<HTML>
<BODY>

<?php
$link = mysql_connect("localhost", "usernamefoo",
"passwordbar");
mysql_select_db("lxfdb");
$result = mysql_query("SELECT Name, Pay FROM staff
ORDER BY Name ASC;");
if (mysql_num_rows($result)) {
?>

Results:
<TABLE BORDER="1"><TR>
<TD ALIGN="CENTER">Name</TD>
<TD ALIGN="CENTER">Pay</TD>
</TR>

<?php
while ($r = mysql_fetch_array($result)) {
extract($r);
echo "<TR><TD>$Name</TD> <TD>$Pay</TD></TR>";
}
echo "</TABLE><BR><BR>";
} else {

```

## Normalisation

### Get to grips with this key topic

This is a key topic in databases, and, if you want to get even remotely serious, is something you need to understand. Not everyone sees its importance first time, so don't worry if you find yourself having to re-read this a few times before it sinks in.

Newcomers to databases tend to put all related data in one big table. For example, the golf course staff table we created last month contained four fields: **Name**, **Age**, **Job**, and **Pay**. What would you do if you wanted to store information about their job description? I have seen countless developers in this situation who, on the supposed grounds of "it's easier this way", have simply added another field **JobDesc TEXT** which includes a whole chunk of text about the person's job.

Our staff table last month had three security people added into it, each with the same title – **Security Man**. If we had a two-hundred word description of what the role of a security man was in the company, and had followed the thinking that just adding a field was OK, we'd have to have the same job description three times in our table. Not only is that a chronic waste of space, but it's also very slow, and exceptionally hard to maintain.

The solution here is normalisation. Normalisation is the process of producing a set of tables with pretty much the same properties we would have had in our large table, except split neatly up into grouped elements.

Having added a **JobDesc** field, our staff table currently looks like this: **ID**, **Name**, **Age**, **Job**, **Pay**, **JobDesc**.

Looking over that list, you can see that **Age**, **Job**, and **Pay** all relate directly to each staff member. **JobDesc**, however, is the odd one out, because it relates to the

**Job** field, and it is only *through* the **Job** field that it has any relation to the staff member.

So, we could split our one large table into two smaller tables – **staff** and **jobs**. The **staff** table would have the fields **ID**, **Name**, **Age**, **Job**, and **Pay** (although **Job** would no longer be a **VARCHAR(255)**, it would be an **INTEGER**), and the **jobs** table would have the fields **ID** (this needs to be auto incrementing), **Title**, and **Description**. In our **jobs** table, we would enter a row for each job in our golf courses, and give it a description. Then, in our **staff** table, we would set the **Job** to be the ID number of that person's job from the **jobs** table. That is, if **Security Guard** had **ID 4** in the **jobs** table, we would set the **Job** field in our **staff** table to **4** wherever someone worked as a security guard.

Now if in the future we decide to change **Security Guard** to **Peace Officer**, we only have to change one field in one table – **ID 4** in the **jobs** table.

By breaking up tables this way, we remove all duplication (or at least as much as is possible without getting too nit-picking), and in the process save a lot of disk space, and increase the speed of our queries. Everyone's a winner!

How you normalise data might not be immediately apparent to you in your own projects, but it can make a big difference if done properly. Put simply, every

individual thing which uniquely has attributes of its own in your database should have its own table of data – that is, if you were shipping computers, you might have a table for monitor types, a table for CPU types, a table for graphics card types, etc, and then a master table of PCs, which contained IDs into the other tables – e.g., **PC #1** might contain monitor **#4** (which in the **Monitors** table might be a Mitsubishi Diamond Plus 200), **keyboard #6**, **graphics card #2**, etc.

Normalisation is a really big topic, and this only just covers the first few steps into it. If you're not sure how to take data from both tables rather than one big table, then you should realise there are two options: either use two queries – in-between `connect()` and `close()`, you can execute as many queries as you like – or you can make your query smarter.

The first option is very simple, but not very powerful. The second option is the absolute opposite! Using the above example, try this out:

```

SELECT s.Name, s.Pay, j.Title, j.Description FROM staff
AS s, JobDesc AS j
WHERE j.ID = s.Job;

```

As you can see from the screenshot, the above query takes data from two tables at once (using the format **table.field**), and matched the **Job ID** field against the **staff Job** field to ensure that the correct title and description are shown for each person.

Name	Pay	Title	Description
Joseph Smith	29000	Security Guard	Keeps the peace
Harold Barnes	29000	Security Guard	Keeps the peace
Carmen Hobbes	30000	Security Guard	Keeps the peace

Querying several tables at once makes your life much easier.

```

?>
Sorry, there are no staff members currently on the system.
<?php
}
mysql_close($link);
?>
</BODY>
</HTML>

```

## MySQL functions

MySQL, and indeed most database managers, give you a variety of functions for use in your queries, to help you query more intelligently. Perhaps the most popular function is **Count()**.

**Count()** returns a field that contains the number of rows which matched your query. For example:

```
SELECT Count(ID) FROM staff;
```

That query would return the number of staff members in the staff table. Yes, an alternative is to use **SELECT \* FROM staff** and then run **mysql\_num\_rows()** on the result, but why force MySQL to do a whole lot of work it doesn't need to do?

You can also try using **MAX()**, **MIN()**, and **AVG()** to get the maximum value, minimum value, and average value of a field. You can make these more complicated again by passing in the result of another operation first as opposed to just a field – for example, to get the average age, use **AVG(Age)**, but to get the average of age plus pay, you can use **AVG(Age + Pay)**.

If you try executing these queries in the MySQL monitor you will notice that MySQL names the field **AVG(Pay + Age)**. Try reading *that* from an associative array in PHP! Instead, you can use **SELECT AVG(Pay + Age) AS AvgPayPlusAge** to get something more manageable – if you're using **mysql\_fetch\_array()** as above, this can be read by using **\$r['AvgPayPlusAge']**.

MySQL also provides excellent date and time functions for example **SELECT UNIX\_TIMESTAMP(now());** returns the current Unix timestamp on your server. You can use this directly inside queries, for example:

```
SELECT * FROM News WHERE DateAdded >
(UNIX_TIMESTAMP(now()) - 86400);
```

This would extract all matching fields from a hypothetical news table, where the **DateAdded** field (presumably an **INTEGER** that has been filled with the Unix timestamp each article was added) is greater than the current time minus **86400** (one day, in seconds). Notice how the **now()** function (which returns a MySQL timestamp) is embedded inside a **UNIX\_TIMESTAMP** function. You can try out more function stacking in the golf staff table by using a query like

```
SELECT AVG(LENGTH(Name)) FROM staff;
```

## Pushing your SQL boundaries

Before I finish this two-part mini-series on SQL with PHP, there are a few final notes that are important for me to mention.

Firstly, don't be afraid to wander from MySQL – many people out there swear by PostgreSQL (check their website, [www.postgresql.org](http://www.postgresql.org), for a pronunciation guide!), and it's not hard to understand why – it has many features that MySQL lacks, although it doesn't come close with regards to overall performance. PHP has top-notch support for other databases,

Revision	\$Revision: 1.42.2.2 \$	
<b>pcre</b>		
PCRE (Perl Compatible Regular Expressions) Support	enabled	
PCRE Library Version	3.4.22-Aug-2000	
<b>mysql</b>		
<b>MySQL Support</b>		
Active Persistent Links	0	
Active Links	0	
Client API version	3.23.39	
MYSQL_MODULE_TYPE	builtin	
MYSQL_SOCKET	/tmp/mysql.sock	
MYSQL_INCLUDE		
MYSQL_LIBS		
<b>Directive</b>		
<b>Directive</b>	<b>Local Value</b>	<b>Master Value</b>
mysql.allow_persistent	On	On
mysql.default_host	no value	no value
mysql.default_password	no value	no value
mysql.default_port	no value	no value
mysql.default_socket	no value	no value
mysql.default_user	no value	no value
mysql.max_links	Unlimited	Unlimited
mysql.max_persistent	Unlimited	Unlimited

and you should remember that MySQL isn't the best choice by default – I've even found myself using Oracle and, gulp, Microsoft SQL Server in the past!

Secondly, if you find you're having problems getting help with MySQL, don't be afraid to ask other PHP developers – MySQL questions are very common in PHP help circles, so give it a try.

Thirdly, MySQL is a very fast changing system, so keep up to date with it to make sure you're taking advantage of all it can do. Future versions promise a slew of new features offering such long-awaited wonders as nested subqueries, stored procedures, foreign key integrity rules, and more. By the way, these are all already available in PostgreSQL.

Fourthly, SQL is a very, very large topic. I have tried to pass on the key elements here, especially those relating to PHP. However, if you really want to take advantage of all that SQL can offer your site – and really I would recommend you at least research the possibilities – you need to do quite a bit more SQL learning. If you have £40 around, I cannot recommend highly enough the Connolly & Begg book mentioned last month, *Database Systems*. I keep it in my top drawer, if that's any indication of how good it is.

Fifthly and finally, spare some time to read the MySQL manual section on optimising MySQL. It gives literally hundreds of golden information nuggets to allow you to push your MySQL server to the limit and beyond, going from how to best compile your server, through table and database design, all the way up to how to optimise your queries.

## Conclusion

Well, that about wraps it up for using SQL with PHP. It is, by the way, something that we'll be making use of regularly, so keep this article safe for future reference.

The SQL functions are very often the most commonly-used part of PHP, and you would do well to experiment with them to get the best understanding of how they work. Remember that PHP supports a wide variety of database systems as well as MySQL, and you can port your code across with little fuss – just be careful of little differences in the SQL dialect used.

With a little trial and error, you will find even two or three hours spent toying with MySQL very productive indeed – don't be afraid to get your hands dirty! [LXF](#)

If you experience serious problems working with MySQL, check the configuration settings in [phpinfo\(\)](#).

## About Paul Hudson

Paul Hudson is a London-based web developer specialising in PHP and Perl. He can be emailed at [hudzilla@php.net](mailto:hudzilla@php.net)

## NEXT MONTH

Next month we'll be looking at how PHP helps you easily parse XML documents, including how to work XSLT into the mix. Note that if you didn't follow my original PHP installation instructions, you may need to configure your PHP installation for XML and XSLT – see the PHP manual online for more information. If there is something you want to see covered here, drop us an email and we will try our best to incorporate it into a future instalment.

# 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!

## Experts this month

Whatever your question is, we can find an expert to answer it – from installation and modern 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.



**Richard Drummond** is an experienced programmer who can answer queries on a variety of subjects. A keen Debian user, he's also our resident Java guru.



**Nick Veitch** is the editor of the magazine, and answers your easy questions! Or indeed anything to do with *Grub*, *LILO*, *netatalk*, *vi*...



## Tar gzip

**Q** I am having a problem installing tar.gz applications from the Aug 02 issue.

**My system is a standard PC with an Athlon 850, 512MB RAM, DVD drive.**

I have recently upgraded my Linux installation from Red Hat 7.2 to Red Hat 7.3. This works fine and I have loaded several RPMs with no problems.

I have now tried to install *OpenOffice.org* from the August DVD and got the following response:

```
tar xvf /mnt/cdrom/Office/
OpenOffice/OOo_
1.0.0_LinuxIntel_install.tar.gz
```

```
tar: This does not look like a
tar archive
```

```
tar: Skipping to next header
```

```
tar: Archive contains obsolescent
base-64 headers
```

```
tar: Error exit delayed from previous
errors
```

I then tried to install *Rosegarden* with a similar result:

```
tar xvf /mnt/cdrom/Sound/
Rosegarden/rosegarden-0.1.5.tar.gz
```

```
tar: This does not look like a
tar archive
```

```
tar: Skipping to next header
```

```
tar: Archive contains obsolescent
base-64 headers
```


```
tar: Error exit delayed from previous
errors
```

Although I have played around with various brands of Linux I am really much of a newbie and don't know what to do next.

Can you help?

**Alan Blundell**

**A** Both the files you are attempting to untar are also gzipped, so you will need to ungzipped them before being able to untar them. Thankfully, *tar* can do this all in one go with the **z** flag, so you will need to use **tar xvzf** rather than



URLs in the format **http://user:pass@domain.tld/** perform HTTP authentication on the server.

**tar xvf.** *tar* really doesn't know what you're feeding it, although you can use the *file* command to figure out what format a particular file is in. Beware that you are also trying to untar the files to the CD, which obviously won't work – you may like to copy them to your hard drive first.

## Plug in camera

**Q** I read in the June 2002 issue your reply regarding the use of *GIMP* to acquire scanner and digital camera images. I have *GIMP1.2.2* installed but it does not have this facility.

**File/acquire** only has 'Screen shot'.

Do I need a later version or a plug-in.

**Mike Davies**

**A** A plug-in is required for collecting images from a scanner or digital camera.

Depending upon what you're using, a plugin should be available. For digital cameras, a plug-in is available as part of the *gtkam* distribution, although it does depend upon the particular model of camera you have.

As you didn't say what device you had, we can't really point you in the right direction. However, in either case, the relevant website has plenty of information for integration into the *GIMP*.

## Apache security

**Q** I've found that I can access a protected directory on a webserver running *Apache* by placing a valid username and password in a URL on a webpage.

The webserver is using basic authentication, and a user usually gets the standard username/password dialog box when accessing a protected directory. All I need to do is insert the username and password followed by **@** between the **http://** and the **www** in the URL.

Here's an example:

```
http://username:password@www.any
where.com/cgi-secure/anyscript.pl
```

Having something like this as the URL of a webpage link and then clicking on it takes me straight into the protected directory and runs the script which was called.

In can't find anything about this 'feature' in the *Apache Bible*. I would like to know if *Apache* can be configured to deny this method of access and trigger the standard username/password dialog instead.

**Jim Anderson**

**A** This is actually a feature of the web browser, rather than *Apache*, and *Apache* will never see the URL in that format. The web browser will expand the URL, and perform HTTP authentication

via the web server as part of the HTTP transaction.

One really has to beg the question as to why you would want to stop people using this feature for authentication in the first place? It is no less secure than authenticating via the login box, not that HTTP authentication is very secure in the first place, as it is done over plain text.

It's actually very convenient for book marking secure sites which require authentication, or for linking to pages which would otherwise ask for login details.

## Disk workout

**Q** 1. How do I change from *kdm* to *xdm* (I'm using Mandrake 8.2)?

2. On both Red Hat 7.2 and Mandrake 8.2, it seems that about one minute after I've logged in, the hard disk decides to do a 'workout' for about five minutes.

It's probably just some filesystem clean up stuff and scripts that are being executed. But how do I stop it? It's almost impossible to run other programs when this is going on.

André Sørensen

**A** Switching your system to start up using *xdm*, rather than *kdm*, is straight forward. All you need to do is edit `/etc/sysconfig/desktop` and change the **DISPLAYMANAGER** entry to:

```
DISPLAYMANAGER=xdm
```

The next time you restart the system, *xdm* will start up.

The other thing you are seeing is *updatedb*, which is executed daily by the *cron* daemon. If you don't leave the system on all the time, then *cron* will execute **updatedb** at the next possible instant, so when you come to start the system up again, *updatedb* will trawl your hard disk updating its database, which will slow things to a crawl. *updatedb* is started out of a script in `/etc/cron.daily/` directory, which you can simply delete from there.

However, once *updatedb* is not executed, commands such as *locate* will fail to work properly. Of course, if you never use *locate*, then it's probably not a major problem to disable *updatedb*. Alternatively, you could move it from `cron.daily` to `cron.weekly`, so it is only executed once a week, which should be more bearable.

```
david@niamh:/etc/cron.daily (pts/13)
niamh 10:05pm Mon Sep 02
david:/etc/cron.daily# ls
calendar logrotate modutils ntp-simple sysklogd
find      nan-db      netkit-inetd standard tetex-bin
niamh 10:05pm Mon Sep 02
david:/etc/cron.daily# cat find
#!/bin/sh
#
# cron script to update the 'find.codes' database.
#
# Written by Ian A. Murdock <imurdock@debian.org> and
# Kevin Dalley <kev in@aimet.com>

if [ -f /etc/updatedb.conf ] ; then
    . /etc/updatedb.conf
fi

cd / && updatedb --localuser=nobody 2>/dev/null
niamh 10:05pm Mon Sep 02
david:/etc/cron.daily#
```

`etc/cron.daily/` contains the *find* script which runs *updatedb* every day.

## PCMCIA

**Q** Thanks for a great read – your mag educates and amuses me and keeps me out of mischief. The problem most people have with getting Linux up and running is lack of drivers for the extensive range of hardware.

For example I use a Thinkpad and have an IBM 10/100 ethernet card which I cannot get working as there isn't a driver for the card. Now I want to buy another

ethernet PCMCIA card but want to know what makes Linux have drivers for. Is there a website with a list of all supported hardware that Linux supports? This would ensure that anyone attempting to load Linux would have a good idea if they are going to succeed and if someone wants new hardware would know what to purchase.

I think Linux is a great Operating System and am amazed at how much time people give voluntarily to progress it. All the whiners that keep complaining it is not as easy as Windows should give it more respect, as when you do get it up and running is more stable and, here is the twist, you learn a lot getting it up and running – but when it's running forget, as what you don't use you lose!

Tony

**A** As far as PCMCIA is concerned, you should really take a look at <http://pcmcia-cs.sf.net/>, as that has a comprehensive list of devices which work with the Linux PCMCIA sub-system.

Generally, 3com Ethernet cards are well supported, as are those from a number of different manufacturers. You've not said which distribution you are using, or what version of *pcmcia-cs* you are using, but upgrading *pcmcia-cs*, or even your entire distribution, should help your IBM card to work correctly, as it is supported.

There are a number of sites pertaining to the use of Linux on Thinkpad notebooks, and a quick search on Google produces a whole list of them. Not knowing which Thinkpad model you have, we can't

## A QUICK REFERENCE TO: IPv6

The current Internet uses version four of the IP protocol, which is limited to  $2^{32}-1$  IP addresses. Over recent years, due to increased use of the Internet, the available IP ranges have been under pressure, and it is expected that IPs will run out within the next few years if prior increases are continued.

To combat this problem, version six of the IP protocol, otherwise known as IPv6, offers up to  $2^{128}-1$  IP addresses, which is significantly larger than the IPv4 range and should last for a while. IPv6 uses a completely separate IP protocol, so support is required by routers and systems which are on IPv6 networks. Fortunately, Linux has had IPv6 support for a number of years, and works happily with IPv6 networks.

As few ISPs offer IPv6 dial-ups or direct connectivity to an IPv6

network, we need to use a system known as tunnelling to route IPv6 traffic over the existing IPv4 Internet. There are a great number of different companies and organisations offering IPv6 connectivity, including [www.tunnelbroker.net](http://www.tunnelbroker.net), [www.freenet6.org](http://www.freenet6.org) and [www.ipng.org.uk](http://www.ipng.org.uk). Each of these offer IPv6 connectivity, and route a subnet of IPv6 addresses to your network, so each system can have its own IPv6 address and be accessible from the outside world.

An IP tunnel can be setup using *iproute2*:

```
# ip tunnel add ipng mode sit
remote <remote IP>
# ip addr add <IPv6 address>/128
# ip link set ipng up
# ip ro add 2000::/3 dev ipng
```

Public IPv6 addresses are generally within the `2001::` or `3ffe::` networks, so we don't route `::0` over the tunnel, only `2000::/3`.

```
david@niamh:~$ ip route
Kernel IP routing table
to: * :::0 dev: eth0
to: :::1 dev: eth0
to: :::2 dev: eth0
to: :::3 dev: eth0
to: :::4 dev: eth0
to: :::5 dev: eth0
to: :::6 dev: eth0
to: :::7 dev: eth0
to: :::8 dev: eth0
to: :::9 dev: eth0
to: :::a dev: eth0
to: :::b dev: eth0
to: :::c dev: eth0
to: :::d dev: eth0
to: :::e dev: eth0
to: :::f dev: eth0
to: :::10 dev: eth0
to: :::11 dev: eth0
to: :::12 dev: eth0
to: :::13 dev: eth0
to: :::14 dev: eth0
to: :::15 dev: eth0
to: :::16 dev: eth0
to: :::17 dev: eth0
to: :::18 dev: eth0
to: :::19 dev: eth0
to: :::1a dev: eth0
to: :::1b dev: eth0
to: :::1c dev: eth0
to: :::1d dev: eth0
to: :::1e dev: eth0
to: :::1f dev: eth0
to: :::20 dev: eth0
to: :::21 dev: eth0
to: :::22 dev: eth0
to: :::23 dev: eth0
to: :::24 dev: eth0
to: :::25 dev: eth0
to: :::26 dev: eth0
to: :::27 dev: eth0
to: :::28 dev: eth0
to: :::29 dev: eth0
to: :::2a dev: eth0
to: :::2b dev: eth0
to: :::2c dev: eth0
to: :::2d dev: eth0
to: :::2e dev: eth0
to: :::2f dev: eth0
to: :::30 dev: eth0
to: :::31 dev: eth0
to: :::32 dev: eth0
to: :::33 dev: eth0
to: :::34 dev: eth0
to: :::35 dev: eth0
to: :::36 dev: eth0
to: :::37 dev: eth0
to: :::38 dev: eth0
to: :::39 dev: eth0
to: :::3a dev: eth0
to: :::3b dev: eth0
to: :::3c dev: eth0
to: :::3d dev: eth0
to: :::3e dev: eth0
to: :::3f dev: eth0
to: :::40 dev: eth0
to: :::41 dev: eth0
to: :::42 dev: eth0
to: :::43 dev: eth0
to: :::44 dev: eth0
to: :::45 dev: eth0
to: :::46 dev: eth0
to: :::47 dev: eth0
to: :::48 dev: eth0
to: :::49 dev: eth0
to: :::4a dev: eth0
to: :::4b dev: eth0
to: :::4c dev: eth0
to: :::4d dev: eth0
to: :::4e dev: eth0
to: :::4f dev: eth0
to: :::50 dev: eth0
to: :::51 dev: eth0
to: :::52 dev: eth0
to: :::53 dev: eth0
to: :::54 dev: eth0
to: :::55 dev: eth0
to: :::56 dev: eth0
to: :::57 dev: eth0
to: :::58 dev: eth0
to: :::59 dev: eth0
to: :::5a dev: eth0
to: :::5b dev: eth0
to: :::5c dev: eth0
to: :::5d dev: eth0
to: :::5e dev: eth0
to: :::5f dev: eth0
to: :::60 dev: eth0
to: :::61 dev: eth0
to: :::62 dev: eth0
to: :::63 dev: eth0
to: :::64 dev: eth0
to: :::65 dev: eth0
to: :::66 dev: eth0
to: :::67 dev: eth0
to: :::68 dev: eth0
to: :::69 dev: eth0
to: :::6a dev: eth0
to: :::6b dev: eth0
to: :::6c dev: eth0
to: :::6d dev: eth0
to: :::6e dev: eth0
to: :::6f dev: eth0
to: :::70 dev: eth0
to: :::71 dev: eth0
to: :::72 dev: eth0
to: :::73 dev: eth0
to: :::74 dev: eth0
to: :::75 dev: eth0
to: :::76 dev: eth0
to: :::77 dev: eth0
to: :::78 dev: eth0
to: :::79 dev: eth0
to: :::7a dev: eth0
to: :::7b dev: eth0
to: :::7c dev: eth0
to: :::7d dev: eth0
to: :::7e dev: eth0
to: :::7f dev: eth0
to: :::80 dev: eth0
to: :::81 dev: eth0
to: :::82 dev: eth0
to: :::83 dev: eth0
to: :::84 dev: eth0
to: :::85 dev: eth0
to: :::86 dev: eth0
to: :::87 dev: eth0
to: :::88 dev: eth0
to: :::89 dev: eth0
to: :::8a dev: eth0
to: :::8b dev: eth0
to: :::8c dev: eth0
to: :::8d dev: eth0
to: :::8e dev: eth0
to: :::8f dev: eth0
to: :::90 dev: eth0
to: :::91 dev: eth0
to: :::92 dev: eth0
to: :::93 dev: eth0
to: :::94 dev: eth0
to: :::95 dev: eth0
to: :::96 dev: eth0
to: :::97 dev: eth0
to: :::98 dev: eth0
to: :::99 dev: eth0
to: :::9a dev: eth0
to: :::9b dev: eth0
to: :::9c dev: eth0
to: :::9d dev: eth0
to: :::9e dev: eth0
to: :::9f dev: eth0
to: :::a0 dev: eth0
to: :::a1 dev: eth0
to: :::a2 dev: eth0
to: :::a3 dev: eth0
to: :::a4 dev: eth0
to: :::a5 dev: eth0
to: :::a6 dev: eth0
to: :::a7 dev: eth0
to: :::a8 dev: eth0
to: :::a9 dev: eth0
to: :::aa dev: eth0
to: :::ab dev: eth0
to: :::ac dev: eth0
to: :::ad dev: eth0
to: :::ae dev: eth0
to: :::af dev: eth0
to: :::b0 dev: eth0
to: :::b1 dev: eth0
to: :::b2 dev: eth0
to: :::b3 dev: eth0
to: :::b4 dev: eth0
to: :::b5 dev: eth0
to: :::b6 dev: eth0
to: :::b7 dev: eth0
to: :::b8 dev: eth0
to: :::b9 dev: eth0
to: :::ba dev: eth0
to: :::bb dev: eth0
to: :::bc dev: eth0
to: :::bd dev: eth0
to: :::be dev: eth0
to: :::bf dev: eth0
to: :::c0 dev: eth0
to: :::c1 dev: eth0
to: :::c2 dev: eth0
to: :::c3 dev: eth0
to: :::c4 dev: eth0
to: :::c5 dev: eth0
to: :::c6 dev: eth0
to: :::c7 dev: eth0
to: :::c8 dev: eth0
to: :::c9 dev: eth0
to: :::ca dev: eth0
to: :::cb dev: eth0
to: :::cc dev: eth0
to: :::cd dev: eth0
to: :::ce dev: eth0
to: :::cf dev: eth0
to: :::d0 dev: eth0
to: :::d1 dev: eth0
to: :::d2 dev: eth0
to: :::d3 dev: eth0
to: :::d4 dev: eth0
to: :::d5 dev: eth0
to: :::d6 dev: eth0
to: :::d7 dev: eth0
to: :::d8 dev: eth0
to: :::d9 dev: eth0
to: :::da dev: eth0
to: :::db dev: eth0
to: :::dc dev: eth0
to: :::dd dev: eth0
to: :::de dev: eth0
to: :::df dev: eth0
to: :::e0 dev: eth0
to: :::e1 dev: eth0
to: :::e2 dev: eth0
to: :::e3 dev: eth0
to: :::e4 dev: eth0
to: :::e5 dev: eth0
to: :::e6 dev: eth0
to: :::e7 dev: eth0
to: :::e8 dev: eth0
to: :::e9 dev: eth0
to: :::ea dev: eth0
to: :::eb dev: eth0
to: :::ec dev: eth0
to: :::ed dev: eth0
to: :::ee dev: eth0
to: :::ef dev: eth0
to: :::f0 dev: eth0
to: :::f1 dev: eth0
to: :::f2 dev: eth0
to: :::f3 dev: eth0
to: :::f4 dev: eth0
to: :::f5 dev: eth0
to: :::f6 dev: eth0
to: :::f7 dev: eth0
to: :::f8 dev: eth0
to: :::f9 dev: eth0
to: :::fa dev: eth0
to: :::fb dev: eth0
to: :::fc dev: eth0
to: :::fd dev: eth0
to: :::fe dev: eth0
to: :::ff dev: eth0
to: :::100 dev: eth0
to: :::101 dev: eth0
to: :::102 dev: eth0
to: :::103 dev: eth0
to: :::104 dev: eth0
to: :::105 dev: eth0
to: :::106 dev: eth0
to: :::107 dev: eth0
to: :::108 dev: eth0
to: :::109 dev: eth0
to: :::10a dev: eth0
to: :::10b dev: eth0
to: :::10c dev: eth0
to: :::10d dev: eth0
to: :::10e dev: eth0
to: :::10f dev: eth0
to: :::110 dev: eth0
to: :::111 dev: eth0
to: :::112 dev: eth0
to: :::113 dev: eth0
to: :::114 dev: eth0
to: :::115 dev: eth0
to: :::116 dev: eth0
to: :::117 dev: eth0
to: :::118 dev: eth0
to: :::119 dev: eth0
to: :::11a dev: eth0
to: :::11b dev: eth0
to: :::11c dev: eth0
to: :::11d dev: eth0
to: :::11e dev: eth0
to: :::11f dev: eth0
to: :::120 dev: eth0
to: :::121 dev: eth0
to: :::122 dev: eth0
to: :::123 dev: eth0
to: :::124 dev: eth0
to: :::125 dev: eth0
to: :::126 dev: eth0
to: :::127 dev: eth0
to: :::128 dev: eth0
to: :::129 dev: eth0
to: :::12a dev: eth0
to: :::12b dev: eth0
to: :::12c dev: eth0
to: :::12d dev: eth0
to: :::12e dev: eth0
to: :::12f dev: eth0
to: :::130 dev: eth0
to: :::131 dev: eth0
to: :::132 dev: eth0
to: :::133 dev: eth0
to: :::134 dev: eth0
to: :::135 dev: eth0
to: :::136 dev: eth0
to: :::137 dev: eth0
to: :::138 dev: eth0
to: :::139 dev: eth0
to: :::13a dev: eth0
to: :::13b dev: eth0
to: :::13c dev: eth0
to: :::13d dev: eth0
to: :::13e dev: eth0
to: :::13f dev: eth0
to: :::140 dev: eth0
to: :::141 dev: eth0
to: :::142 dev: eth0
to: :::143 dev: eth0
to: :::144 dev: eth0
to: :::145 dev: eth0
to: :::146 dev: eth0
to: :::147 dev: eth0
to: :::148 dev: eth0
to: :::149 dev: eth0
to: :::14a dev: eth0
to: :::14b dev: eth0
to: :::14c dev: eth0
to: :::14d dev: eth0
to: :::14e dev: eth0
to: :::14f dev: eth0
to: :::150 dev: eth0
to: :::151 dev: eth0
to: :::152 dev: eth0
to: :::153 dev: eth0
to: :::154 dev: eth0
to: :::155 dev: eth0
to: :::156 dev: eth0
to: :::157 dev: eth0
to: :::158 dev: eth0
to: :::159 dev: eth0
to: :::15a dev: eth0
to: :::15b dev: eth0
to: :::15c dev: eth0
to: :::15d dev: eth0
to: :::15e dev: eth0
to: :::15f dev: eth0
to: :::160 dev: eth0
to: :::161 dev: eth0
to: :::162 dev: eth0
to: :::163 dev: eth0
to: :::164 dev: eth0
to: :::165 dev: eth0
to: :::166 dev: eth0
to: :::167 dev: eth0
to: :::168 dev: eth0
to: :::169 dev: eth0
to: :::16a dev: eth0
to: :::16b dev: eth0
to: :::16c dev: eth0
to: :::16d dev: eth0
to: :::16e dev: eth0
to: :::16f dev: eth0
to: :::170 dev: eth0
to: :::171 dev: eth0
to: :::172 dev: eth0
to: :::173 dev: eth0
to: :::174 dev: eth0
to: :::175 dev: eth0
to: :::176 dev: eth0
to: :::177 dev: eth0
to: :::178 dev: eth0
to: :::179 dev: eth0
to: :::17a dev: eth0
to: :::17b dev: eth0
to: :::17c dev: eth0
to: :::17d dev: eth0
to: :::17e dev: eth0
to: :::17f dev: eth0
to: :::180 dev: eth0
to: :::181 dev: eth0
to: :::182 dev: eth0
to: :::183 dev: eth0
to: :::184 dev: eth0
to: :::185 dev: eth0
to: :::186 dev: eth0
to: :::187 dev: eth0
to: :::188 dev: eth0
to: :::189 dev: eth0
to: :::18a dev: eth0
to: :::18b dev: eth0
to: :::18c dev: eth0
to: :::18d dev: eth0
to: :::18e dev: eth0
to: :::18f dev: eth0
to: :::190 dev: eth0
to: :::191 dev: eth0
to: :::192 dev: eth0
to: :::193 dev: eth0
to: :::194 dev: eth0
to: :::195 dev: eth0
to: :::196 dev: eth0
to: :::197 dev: eth0
to: :::198 dev: eth0
to: :::199 dev: eth0
to: :::19a dev: eth0
to: :::19b dev: eth0
to: :::19c dev: eth0
to: :::19d dev: eth0
to: :::19e dev: eth0
to: :::19f dev: eth0
to: :::1a0 dev: eth0
to: :::1a1 dev: eth0
to: :::1a2 dev: eth0
to: :::1a3 dev: eth0
to: :::1a4 dev: eth0
to: :::1a5 dev: eth0
to: :::1a6 dev: eth0
to: :::1a7 dev: eth0
to: :::1a8 dev: eth0
to: :::1a9 dev: eth0
to: :::1aa dev: eth0
to: :::1ab dev: eth0
to: :::1ac dev: eth0
to: :::1ad dev: eth0
to: :::1ae dev: eth0
to: :::1af dev: eth0
to: :::1b0 dev: eth0
to: :::1b1 dev: eth0
to: :::1b2 dev: eth0
to: :::1b3 dev: eth0
to: :::1b4 dev: eth0
to: :::1b5 dev: eth0
to: :::1b6 dev: eth0
to: :::1b7 dev: eth0
to: :::1b8 dev: eth0
to: :::1b9 dev: eth0
to: :::1ba dev: eth0
to: :::1bb dev: eth0
to: :::1bc dev: eth0
to: :::1bd dev: eth0
to: :::1be dev: eth0
to: :::1bf dev: eth0
to: :::1c0 dev: eth0
to: :::1c1 dev: eth0
to: :::1c2 dev: eth0
to: :::1c3 dev: eth0
to: :::1c4 dev: eth0
to: :::1c5 dev: eth0
to: :::1c6 dev: eth0
to: :::1c7 dev: eth0
to: :::1c8 dev: eth0
to: :::1c9 dev: eth0
to: :::1ca dev: eth0
to: :::1cb dev: eth0
to: :::1cc dev: eth0
to: :::1cd dev: eth0
to: :::1ce dev: eth0
to: :::1cf dev: eth0
to: :::1d0 dev: eth0
to: :::1d1 dev: eth0
to: :::1d2 dev: eth0
to: :::1d3 dev: eth0
to: :::1d4 dev: eth0
to: :::1d5 dev: eth0
to: :::1d6 dev: eth0
to: :::1d7 dev: eth0
to: :::1d8 dev: eth0
to: :::1d9 dev: eth0
to: :::1da dev: eth0
to: :::1db dev: eth0
to: :::1dc dev: eth0
to: :::1dd dev: eth0
to: :::1de dev: eth0
to: :::1df dev: eth0
to: :::1e0 dev: eth0
to: :::1e1 dev: eth0
to: :::1e2 dev: eth0
to: :::1e3 dev: eth0
to: :::1e4 dev: eth0
to: :::1e5 dev: eth0
to: :::1e6 dev: eth0
to: :::1e7 dev: eth0
to: :::1e8 dev: eth0
to: :::1e9 dev: eth0
to: :::1ea dev: eth0
to: :::1eb dev: eth0
to: :::1ec dev: eth0
to: :::1ed dev: eth0
to: :::1ee dev: eth0
to: :::1ef dev: eth0
to: :::1f0 dev: eth0
to: :::1f1 dev: eth0
to: :::1f2 dev: eth0
to: :::1f3 dev: eth0
to: :::1f4 dev: eth0
to: :::1f5 dev: eth0
to: :::1f6 dev: eth0
to: :::1f7 dev: eth0
to: :::1f8 dev: eth0
to: :::1f9 dev: eth0
to: :::1fa dev: eth0
to: :::1fb dev: eth0
to: :::1fc dev: eth0
to: :::1fd dev: eth0
to: :::1fe dev: eth0
to: :::1ff dev: eth0
to: :::200 dev: eth0
to: :::201 dev: eth0
to: :::202 dev: eth0
to: :::203 dev: eth0
to: :::204 dev: eth0
to: :::205 dev: eth0
to: :::206 dev: eth0
to: :::207 dev: eth0
to: :::208 dev: eth0
to: :::209 dev: eth0
to: :::20a dev: eth0
to: :::20b dev: eth0
to: :::20c dev: eth0
to: :::20d dev: eth0
to: :::20e dev: eth0
to: :::20f dev: eth0
to: :::210 dev: eth0
to: :::211 dev: eth0
to: :::212 dev: eth0
to: :::213 dev: eth0
to: :::214 dev: eth0
to: :::215 dev: eth0
to: :::216 dev: eth0
to: :::217 dev: eth0
to: :::218 dev: eth0
to: :::219 dev: eth0
to: :::21a dev: eth0
to: :::21b dev: eth0
to: :::21c dev: eth0
to: :::21d dev: eth0
to: :::21e dev: eth0
to: :::21f dev: eth0
to: :::220 dev: eth0
to: :::221 dev: eth0
to: :::222 dev: eth0
to: :::223 dev: eth0
to: :::224 dev: eth0
to: :::225 dev: eth0
to: :::226 dev: eth0
to: :::227 dev: eth0
to: :::228 dev: eth0
to: :::229 dev: eth0
to: :::22a dev: eth0
to: :::22b dev: eth0
to: :::22c dev: eth0
to: :::22d dev: eth0
to: :::22e dev: eth0
to: :::22f dev: eth0
to: :::230 dev: eth0
to: :::231 dev: eth0
to: :::232 dev: eth0
to: :::233 dev: eth0
to: :::234 dev: eth0
to: :::235 dev: eth0
to: :::236 dev: eth0
to: :::237 dev: eth0
to: :::238 dev: eth0
to: :::239 dev: eth0
to: :::23a dev: eth0
to: :::23b dev: eth0
to: :::23c dev: eth0
to: :::23d dev: eth0
to: :::23e dev: eth0
to: :::23f dev: eth0
to: :::240 dev: eth0
to: :::241 dev: eth0
to: :::242 dev: eth0
to: :::243 dev: eth0
to: :::244 dev: eth0
to: :::245 dev: eth0
to: :::246 dev: eth0
to: :::247 dev: eth0
to: :::248 dev: eth0
to: :::249 dev: eth0
to: :::24a dev: eth0
to: :::24b dev: eth0
to: :::24c dev: eth0
to: :::24d dev: eth0
to: :::24e dev: eth0
to: :::24f dev: eth0
to: :::250 dev: eth0
to: :::251 dev: eth0
to: :::252 dev: eth0
to: :::253 dev: eth0
to: :::254 dev: eth0
to: :::255 dev: eth0
to: :::256 dev: eth0
to: :::257 dev: eth0
to: :::258 dev: eth0
to: :::259 dev: eth0
to: :::25a dev: eth0
to: :::25b dev: eth0
to: :::25c dev: eth0
to: :::25d dev: eth0
to: :::25e dev: eth0
to: :::25f dev: eth0
to: :::260 dev: eth0
to: :::261 dev: eth0
to: :::262 dev: eth0
to: :::263 dev: eth0
to: :::264 dev: eth0
to: :::265 dev: eth0
to: :::266 dev: eth0
to: :::267 dev: eth0
to: :::268 dev: eth0
to: :::269 dev: eth0
to: :::26a dev: eth0
to: :::26b dev: eth0
to: :::26c dev: eth0
to: :::26d dev: eth0
to: :::26e dev: eth0
to: :::26f dev: eth0
to: :::270 dev: eth0
to: :::271 dev: eth0
to: :::272 dev: eth0
to: :::273 dev: eth0
to: :::274 dev: eth0
to: :::275 dev: eth0
to: :::276 dev: eth0
to: :::277 dev: eth0
to: :::278 dev: eth0
to: :::279 dev: eth0
to: :::27a dev: eth0
to: :::27b dev: eth0
to: :::27c dev: eth0
to: :::27d dev: eth0
to: :::27e dev: eth0
to: :::27f dev: eth0
to: :::280 dev: eth0
to: :::281 dev: eth0
to: :::282 dev: eth0
to: :::283 dev: eth0
to: :::284 dev: eth0
to: :::285 dev: eth0
to: :::286 dev: eth0
to: :::287 dev: eth0
to: :::288 dev: eth0
to: :::289 dev: eth0
to: :::28a dev: eth0
to: :::28b dev: eth0
to: :::28c dev: eth0
to: :::28d dev: eth0
to: :::28e dev: eth0
to: :::28f dev: eth0
to: :::290 dev: eth0
to: :::291 dev: eth0
to: :::292 dev: eth0
to: :::293 dev: eth0
to: :::294 dev: eth0
to: :::295 dev: eth0
to: :::296 dev: eth0
to: :::297 dev: eth0
to: :::298 dev: eth0
to: :::299 dev: eth0
to: :::29a dev: eth0
to: :::29b dev: eth0
to: :::29c dev: eth0
to: :::29d dev: eth0
to: :::29e dev: eth0
to: :::29f dev: eth0
to: :::2a0 dev: eth0
to: :::2a1 dev: eth0
to: :::2a2 dev: eth0
to: :::2a3 dev: eth0
to: :::2a4 dev: eth0
to: :::2a5 dev: eth0
to: :::2a6 dev: eth0
to: :::2a7 dev: eth0
to: :::2a8 dev: eth0
to: :::2a9 dev: eth0
to: :::2aa dev: eth0
to: :::2ab dev: eth0
to: :::2ac dev: eth0
to: :::2ad dev: eth0
to: :::2ae dev: eth0
to: :::2af dev: eth0
to: :::2b0 dev: eth0
to: :::2b1 dev: eth0
to: :::2b2 dev: eth0
to: :::2b3 dev: eth0
to: :::2b4 dev: eth0
to: :::2b5 dev: eth0
to: :::2b6 dev: eth0
to: :::2b7 dev: eth0
to: :::2b8 dev: eth0
to: :::2b9 dev: eth0
to: :::2ba dev: eth0
to: :::2bb dev: eth0
to: :::2bc dev: eth0
to: :::2bd dev: eth0
to: :::2be dev: eth0
to: :::2bf dev: eth0
to: :::2c0 dev: eth0
to: :::2c1 dev: eth0
to: :::2c2 dev: eth0
to: :::2c3 dev: eth0
to: :::2c4 dev: eth0
to: :::2c5 dev: eth0
to: :::2c6 dev: eth0
to: :::2c7 dev: eth0
to: :::2c8 dev: eth0
to: :::2c9 dev: eth0
to: :::2ca dev: eth0
to: :::2cb dev: eth0
to: :::2cc dev: eth0
to: :::2cd dev: eth0
to: :::2ce dev: eth0
to: :::2cf dev: eth0
to: :::2d0 dev: eth0
to: :::2d1 dev: eth0
to: :::2d2 dev: eth0
to: :::2d3 dev: eth0
to: :::2d4 dev: eth0
to: :::2d5 dev: eth0
to: :::2d6 dev: eth0
to: :::2d7 dev: eth0
to: :::2d8 dev: eth0
to: :::2d9 dev: eth0
to: :::2da dev: eth0
to: :::2db dev: eth0
to: :::2dc dev: eth0
to: :::2dd dev: eth0
to: :::2de dev: eth0
to: :::2df dev: eth0
to: :::2e0 dev: eth0
to: :::2e1 dev: eth0
to: :::2e2 dev: eth0
to: :::2e3 dev: eth0
to: :::2e4 dev: eth0
to: :::2e5 dev: eth0
to: :::2e6 dev: eth0
to: :::2e7 dev: eth0
to: :::2e8 dev: eth0
to: :::2e9 dev: eth0
to: :::2ea dev: eth0
to: :::2eb dev: eth0
to: :::2ec dev: eth0
to: :::2ed dev: eth0
to: :::2ee dev: eth0
to: :::2ef dev: eth0
to: :::2f0 dev: eth0
to: :::2f1 dev: eth0
to: :::2f2 dev: eth0
to: :::2f3 dev: eth0
to: :::2f4 dev: eth0
to: :::2f5 dev: eth0
to: :::2f6 dev: eth0
to: :::2f7 dev: eth0
to: :::2f8 dev: eth0
to: :::2f9 dev: eth0
to: :::2fa dev: eth0
to: :::2fb dev: eth0
to: :::2fc dev: eth0
to: :::2fd dev: eth0
to: :::2fe dev: eth0
to: :::2ff dev: eth0
to: :::300 dev: eth0
to: :::301 dev: eth0
to: :::302 dev: eth0
to: :::303 dev: eth0
to: :::304 dev: eth0
to: :::305 dev: eth0
to: :::306 dev: eth0
to: :::307 dev: eth0
to: :::308 dev: eth0
to: :::309 dev: eth0
to: :::30a dev: eth0
to: :::30b dev: eth0
to: :::30c dev: eth0
to: :::30d dev: eth0
to: :::30e dev: eth0
to: :::30f dev: eth0
to: :::310 dev: eth0
to: :::311 dev: eth0
to: :::312 dev: eth0
to: :::313 dev: eth0
to: :::314 dev: eth0
to: :::315 dev: eth0
to: :::316 dev: eth0
to: :::317 dev: eth0
to: :::318 dev: eth0
to: :::319 dev: eth0
to: :::31a dev: eth0
to: :::31b dev: eth0
to: :::31c dev: eth0
to: :::31d dev: eth0
to: :::31e dev: eth0
to: :::31f dev: eth0
to: :::320 dev: eth0
to: :::
```


## FREQUENTLY ASKED QUESTIONS: USB

### FAQ How can I find out what devices work with USB under Linux?

Everything USB related for Linux is documented over at [www.linux-usb.org](http://www.linux-usb.org). While they do not have a comprehensive list of things which will, or won't, work with Linux, they do link to the sites for the various components of the USB subsystem which have far more detail and generally specifically say which products will work happily with that driver. Of course, people sometimes get kit working which is not listed, so checking <http://groups.google.com/> and seeing how other people fared with the same item is probably a good idea.

### FAQ What do I need to do to get USB up and running?

Assuming you have everything working on the hardware front, allowing Linux to use USB devices is simply done by compiling USB support into the kernel, or indeed, as a loadable module. There is a USB section of the kernel configuration tool, allowing you to



USBView displays a hierarchy of USB devices connected to the system.

select a number of different USB subsystem modules, depending upon the chipset your motherboard uses. Most systems use the UHCI standard, so the *usb-uhci* module needs to be available. Alternatively, there is the *uhci* module, which works on the same hardware, but was written separately. As always, try one and see if it works. If it doesn't, try the other.

Once you've got the correct

module, you can build it all and reboot the machine. The module can then be loaded, as root, with:

```
modprobe usb-uhci
```

If you check *dmesg*, there should be some useful debugging output telling you what it found and how many USB ports are available.

### FAQ Will that work every time I reboot?

You will need to configure your system to load the *usb-uhci* module when you boot up. This is different depending upon which distribution you are using. For example, Debian users need only list *usb-uhci* in */etc/modules* and it will be loaded when they reboot. Of course, *usb-uhci* will automatically load whenever any USB modules are loaded, so if you configure your system to load modules when their */dev* entries are accessed, then there is no need to load *usb-uhci* at boot time.

### FAQ I plugged my device in. Now what?

Hopefully, the kernel will notice that you've plugged something in, and produce some useful debugging output so you can check that it's recognised properly. Before you can actually use it, you will need to load the appropriate module so that user-space tools can interact with the device. Many devices, including printers, are built into the kernel's USB support, but others, including webcams, mice and so forth need additional modules.

### FAQ Can I make Linux load the kernel modules

« offer much further information, but [www.linux-laptop.net](http://www.linux-laptop.net) is a good place to start.

## Explore NTFS

**Q** In the June 2002 issue Matthew had a question about transferring data from Linux to Windows 2000. You suggested to create a FAT partition.

I would like to suggest another approach. I have had the same problem a couple of years ago when I was using a NT4/Linux in dual boot. Back then I discovered two programs which could read data from Linux from within Windows. The first one is *Captain Nemo* which can be found at [www.runtime.org/captain.htm](http://www.runtime.org/captain.htm). This one is not free. It costs €104.40.

It has however a downloadable try out. The other one is called

*explore2f* and it is freeware, inclusive of the source code. It can be found at <http://uranus.it.swin.edu.au/~jn/linux/explore2fs.htm>.

**Michiel Mol**

**A** *explore2fs* is a particularly popular program, and many have reported success with it. However, it's all dependent upon how much you trust such an application not to gobble up your Linux partition and spit it back out again into the floor. Linux certainly has a decent track record with NTFS, so accessing Win2k partitions from within Linux is hopefully a safe prospect.

In either case, creating a separate partition for copying files to and from another OS is preferable, so if the worst does happen, then you won't destroy either system.

## Wireless access

**Q** I am a subscriber of your magazine and I have read the wireless networking article in April 2002 issue. I have a question: what brand of access point base station will work connected to a Linux server which is operating as a Samba server, and also as a router to the Internet. In other words I need to buy a wireless access point base that serves as a wireless connection to Windows-based laptops (that communicate with wireless technology), and hooks up to the Linux server that presently works as networking OS (*Samba*, actually) and also as router to Internet connection (fibre-optic connection).

Thank you very much for your assistance.

**Filippo Del Favero**

**A** Wireless access points act as ethernet bridges, so there is no requirement for any operating system support on the network, beyond basic ethernet connectivity.

Depending how you architect your network, you simply need to have your Linux router route packets from the network handled by the wireless access point to the Internet. If you just decide to hook it up with your existing wired LAN, there is no need for any reconfiguration of the Linux system.

Of course, it is possible for a Linux box to function as an access point in its own right using a Prism2-based wireless card. It does require a firmware upgrade, but it will work out considerably cheaper than buying an access point in a separate box.

## automatically when I plug things in?

Many distributions install *hotplug* as standard and, if not, the rest should have it available as a package. *hotplug* is notified whenever something is connected or disconnected from the USB system, allowing you to run scripts and so forth when a device becomes available. All one needs to do is set up a USB map in `/etc/hotplug/usb.usermap`, and put a script in `/etc/hotplug/usb`. As always, **man hotplug** has all the information you'll ever need.

## FAQ I use gphoto2 with my digital camera, but it keeps complaining about permissions. How do I fix this?

*gphoto2* accesses raw devices via `/proc/bus/usb`, and those will need to have the appropriate permissions set so that non-root users can access the appropriate file. *gphoto2* comes with a number of scripts which make it easy for the system to set the permissions appropriately using *hotplug* when the camera is connected, so that you won't need


to be root to download the images from the device.

## FAQ I reconfigured X to use /dev/input/mice, but it refuses to start unless I have my mouse plugged in when I boot.

As `/dev/input/mice` is a device file, it is important that the *mousedev* module is loaded, as this will allow XFree86 to open `/dev/input/mice`. One can either have *mousedev* load on startup or, if you are using *devfs*, an alias from 'input/mice' can be created in `/etc/modules.conf`, enabling the *mousedev* module to be loaded whenever `/dev/input/mice` is accessed.

## FAQ How can I check what USB devices are connected to my system?

The easiest way to check USB devices is with the *USBView* utility, which displays the current USB devices in a hierarchy based upon the hub they are connected to. Most distributions include *USBView* as part of the standard packages provided.



If you want to access *ext2* filesystems from within Windows, *expore2fs* is a good choice.

## Athlon boot

**Q** My problem is this: I got a free issue of *LXF* when *Amiga Format* closed.

The coverdisc had *Definite 7* on it, which I recently tried to install. I'm a complete Linux newbie, so when the installer failed, I was lost. Assuming there was no chance that you would offer any support for such an old disc, I bought issue 31 for the full distro, thinking that that one would probably work instead, and if it didn't there was at least a chance of support.

Both installers fail at the same point, giving exactly the same final message:

VP\_IDE: not 100% native mode:  
will probe IRQ's later

plus some status reports on the two IDE channels such as

ide0: BM-DMA at 0xc000-0xc007,  
BIOS settings: hda:DMA, hdb:DMA  
then nothing.

I've questioned the *LXF* forums, and been advised to try command line parameters like `mem=nopentium` and `ide=nodme`, I've tried a couple of bootdisk images, I've even overridden cable select and specified master/slaves on my motherboard's IDE channels. None of this has helped. I can't even boot the DVD, although this seems to be a wider problem than just me.

My system is as follows:

QDI Kinetiz 7B mobo (BIOS v3.8), Athlon T'bird 1.3GHz  
256MB PC133 RAM (single DIMM)

Primary master: Seagate 40GB HD (can't remember the model number off the top of my head)  
Secondary master: Maxtor Diamond Max 15GB HD  
Primary slave: Matshita SR8583A DVD-ROM drive  
Secondary slave: Acer 6432 CD-R/RW drive

I can boot the discs past the lockup point on an older K6-2 400MHz, VIA MVP3-based PC I have, but not on the Athlon.

Please help. I want to try Linux out, but my every effort fails miserably before I even get to the first hurdle.

Thanks.  
Neil Morford.

**A** Definite Linux is a dead duck, and it is neither maintained nor distributed anymore. You can safely throw that CD in the trash, or fashion fancy designer jewellery from it and attempt to make a small profit.

Since your boot process fails so early in the kernel initialisation, it is quite likely to be a hardware issue. Assuming you are booting using a recent kernel, it may be worth trying to boot the machine with the bare minimum of hardware in order to figure out what's going on. There are known issues with Athlon CPUs under Linux, so you may wish to contact the distribution's support staff directly, or post to their newsgroup, in order to find out if there is a patch or alternative boot image available in order to get your system up and running.

You should really try the system out with a recent 2.4 kernel, and you can download boot images from the Internet and write them to a boot floppy to see if you have any more success before looking at the entire distribution.

## Hard Slack

**Q** I used to have SuSE installed for a bit of light relief but that was a few years ago – so when I saw your magazine, with Slackware 8.1, I thought "Yes! it is the year 2002 I'll give this a whirl, it cannot be too hard!" But yet again it appears that Linux folk want Linux for ... Linux folk.

Your DVD disk does not appear to be bootable as claimed in your magazine. And yes I have got BIOS set to enable booting from CDROM. I thought OK, I'll try booting from floppy – but this did not work. Made a boot disk (from .s or .i?? Who knows? It would not be Linux I guess if meaningful instructions were made available), but after this at first appeared to work it then prompted me for something it could not find on the floppy... Lets try bootable CD – ah, but the README only gives instructions for Linux users: excuse me but is this just a little bit daft?? Or is the Linux community not looking to get converts from Windows?

While there is merit in working things out for oneself, for those who have a life beyond playing with >>>



**Debian has different releases on the go simultaneously – be sure that you’re running the right one.**

◀◀ the insides of PCs and simply want to use the PC as a tool for a job, I do not think it is unreasonable to expect that an OS to be readily installable. As you are increasing sales by providing the software I believe you have a responsibility to ensure that it is adequately documented.

Rather your brief comments refer readers to poorly presented READMEs that are of little value.

**Mike Scott**

**A** We’ve been assured that the DVD was indeed bootable, although if you think you have a dead one, then you can contact our CD department and request a replacement. Booting off a floppy disk should work for systems which are unable to boot from a DVD.

Slackware is not the simplest of Linux distributions available, so you may want to look at something like Mandrake or Red Hat if you are a complete beginner. Slackware is certainly better than it has been, but is certainly not something which those who are not familiar with Linux will want to try out. If you would like to give Mandrake or Red Hat a go, head on over to [www.linuxemporium.co.uk](http://www.linuxemporium.co.uk), or download the ISOs from their FTP servers and burn them to a couple of CDs.

## Unstable X

**Q** Hi. I’m running Debian unstable/testing on an AMD Athlon, with an nVIDIA GeForce2 MX. Whenever I try to start X I get the following error message:

**Fatal server error:  
could not open default cursor font  
‘cursor’**

**Nick Wilson**

**A** Debian unstable is a very different barrel of monkeys from testing, as it contains far more packages which are less functional than those in the current testing distribution. If you’re tracking Debian unstable, a.k.a. sid, you really need to subscribe to Debian mailing lists to keep up with things.

The debian-x list had a number of posts about this problem, which seems to be rectified by reinstalling the *artwiz-cursor* package. Hopefully, this should fix your problem. If not, there are many mailing lists for Debian users, which will hopefully yield some answers.

## Samba path

**Q** After issuing the command `/usr/sbin/rcsmb start` I get the following error messages:

```
Jul 23 13:30:10 Mars smb[2437]: param/loadparm.c:service_ok(2156)
```

```
Jul 23 13:30:10 Mars smb[2437]: No path in service HP710C-PPA - using /tmp
```

```
Jul 23 13:30:13 Mars nmbd[2434]: [2002/07/23 13:30:13, 0] nmbd/nmbd_responserecordsb.c:find_response_record(237)
```

```
Jul 23 13:30:13 Mars nmbd[2434]: find_response_record: response packet id 14560 received with no matching record.
```

```
Jul 23 13:30:13 Mars nmbd[2434]: [2002/07/23 13:30:13, 0] nmbd/nmbd_responserecordsb.c:find_response_record(237)
```

```
Jul 23 13:30:13 Mars nmbd[2434]: find_response_record: response packet id 14561 received with no matching record.
```

**Could anyone tell me what’s wrong and how to repair this?**

**Derk Drukker**

**A** The first problem we see is that you’re missing a ‘path’ option within the *Samba* section for your HP701C-PPA share. While we can’t be 100% sure, we’re guessing that HP710C-PPA is a printer, rather than a filesystem share, so you should not need to create a separate share for it, as the **[Printers]** section will handle it for you.

As you didn’t send us your `smb.conf` file, we can’t check it thoroughly for misconfiguration or other mistakes, so we obviously can’t be certain that this is the problem you’re experiencing. If this does not fix your problem, sending us your

`smb.conf` file will give us more clues in order to deal with the issues you’re experiencing.

## Mozilla with style

**Q** Peter Harrison had problems getting *Mozilla* to recognise his external style sheet. It is possible that it is being sent with an incorrect content type, *Mozilla* doesn’t accept style sheets that the server claims contain anything other than `text/css`.

An easy way to check is using `wget`:

```
wget http://localhost/default.css
```

```
...
```

```
Length: 409 [text/css]
```

If it returns anything other than `text/css` (many servers are unfortunately set up to send as `text/plain` or even `application/x-pointplus`) then you should contact your web host and ask them to fix it. If you are hosting yourself then the file to edit is most likely `/etc/mime.types`

**David Dorward**

*Mozilla* is particularly fussy about the MIME type of the CSS file it reads. In addition to the suggestions you made, it is also possible to add a MIME type to Apache using the **AddType** directive, which may be included as part of a `.htaccess` file.

```
AddType text/css .css
```

Of course, it is preferable for a web hosting service to set MIME types correctly, but if all else fails, a `.htaccess` file does the job.

## Webmin bug

**Q** I’m the reader of *Linux Format* from Malaysia. I’m facing some problem configure DHCP server that comes with Mandrake Linux 8.2. I’ve installed all the DHCP packages into the machine from all 3 CDs.



**Mozilla supports CSS, but you need to make sure your CSS document has the correct MIME type.**

**After I’ve configured the DHCP server using Webmin 0.99 (I upgraded the Webmin from 0.90), I get the following error message:**

```
=====
You must add a ddns-update-style statement to /etc/dhcpd.conf to get the same behaviour as the 3.0b1pl11 and previous version, add a line that says
“ddns-update-style add-hoc; “
=====
```

I’ve add the line

```
ddns-update-style add-hoc;
```

in the `dhcpd.conf` file, but when I used *Webmin* to restart the DHCP server, I still get the same message. What thing did I configure wrongly? **Ee Beng**

**A** It seems like *Webmin* has a typo, as the appropriate line you need to use is;

```
ddns-update-style ad-hoc;
```

You’ve got the latest release of *Webmin* already, so you might want to take a look at the *Webmin* mailing lists to see if there have been reported cases of this problem. Hopefully there is a fix already, or you can just continue as you are once you have the appropriate line in your `/etc/dhcpd.conf` file. **LXF**

## Submission advice

We are happy to answer all sorts of Linux related questions. If we don’t know the answer, we’ll find out for you! But in order to give you the best service, it helps a lot if you read the following submission advice.

- Please be sure to include any relevant details of your system. ‘I can’t get X to work’ doesn’t really mean anything to us if we don’t know things like what version of X you are trying to run, what hardware you are running on.
- Be specific about your problem. Things like ‘it doesn’t work’ or ‘I get an error’ aren’t all that helpful. In what way does something not work? What were you expecting to happen? What does the error message actually say?
- Please remember that the people who write this magazine are NOT the authors or developers of Linux, any particular package or distro. Sometimes the people responsible for software have more information available on websites. Try reading the documentation!

We will try and answer all questions. If we don’t answer yours specifically, you’ll probably find we’ve answered one just like it. We can’t really give personal replies to all your questions.

**WRITE TO US AT:**  
Linux Format, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or email: [lxf.answers@futurenet.co.uk](mailto:lxf.answers@futurenet.co.uk)

# missed one?

## LINUX FORMAT BACK ISSUES

Every month *Linux Format* brings you the best tutorials, the essential reviews and the latest news. But if you've missed out on a must-read feature or a vital 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 below for more details).



**October 2002**

Product code:  
**LXFB0032**(cd)  
**LXFD0032**(dvd)

**DVD HIGHLIGHTS:**  
Knoppix, Drip, Squeak, extra FlightGear maps, Ogg Vorbis 1.0, Knoppix

**CDs HIGHLIGHTS:**  
FlightGear (runs from disc), Aglaophone, UAE, Clam Antivirus, Perl 5.8, Quanta Plus, Netclipboard, Mah-Jong, HTML-Mason, WebSuck, Epsutil

**MAGAZINE FEATURING:**  
Building better databases, 'Trusted Computing' – beware Palladium, USB 2.0, firewall roundup, Amiga emulation, Gentoo review



**September 2002**

Product code:  
**LXFB0031**(cd)  
**LXFD0031**(dvd)

**DVD HIGHLIGHTS:**  
Slackware 8.1, Cinelerra, Ogle & Zine, Gnome2, MultiCD, Phobia III, Grip, Zinf, TKVoice

**CDs HIGHLIGHTS:**  
Slackware 8.1, Fluxbox, Kallers, Torcs, SaveMyModem

**MAGAZINE FEATURING:**  
Linux goes to Hollywood, HTML Editors roundup, Internet security special, Ruby scripting language



**August 2002**

Product code:  
**LXFB0030**(cd)  
**LXFD0030**(dvd)

**DVD HIGHLIGHTS:**  
Intel's C++ & Fortean compilers, CPAN, Gnome2, Gentool Linux, Boson, Mozilla, Normalize, EarCandy

**CDs HIGHLIGHTS:**  
Highlights: Intel's C++ & Fortean compilers, Gnome2, Transcode, Album Cover Grabber, Office Suites

**MAGAZINE FEATURING:**  
Ultimate Office: Every current office solution on test, plus Quantum computing, Red Hat 7.3, PHP tips, Eden motherboard test



**July 2002**

Product code:  
**LXFB0029**(cd)  
**LXFD0029**(dvd)

**DVD HIGHLIGHTS:**  
Highlights: OpenOffice.org 1.0, Mozilla 1.0rc3, Netscape 7.0, Opera 6.0, Mac on Linux

**CDs HIGHLIGHTS:**  
Highlights: OpenOffice.org 1.0, Evolution, Omnis Studio 3.01, Clam AntiVirus, Python 2.21

**MAGAZINE FEATURING:**  
Customise your kernel. WineX latest, Inside IPv6, Astronomy applications roundup



**June 2002**

Product code:  
**LXFB0028**(cd)  
**LXFD0028**(dvd)

**DVD HIGHLIGHTS:**  
KDE3, Mozilla 1.0, Gentoo, Beehive, PixiePlus, Funk You, NetComics, Fluxbox

**CDs HIGHLIGHTS:**  
Highlights: KDE3, Mozilla 1.0, Linux from Scratch, Tomsrtbt, GKrellm, evoBB

**MAGAZINE FEATURING:**  
Mozilla special issue, Video editing, Networking with Macs, KDE3 uncovered, Crossover Office



**May 2002**

Product code:  
**LXFB0027**(cd)  
**LXFD0027**(dvd)

**DVD HIGHLIGHTS:**  
Mandrake 8.2, knoda, Wine, Gspy, Morphon CSS-Editor, Robotcop, WebCalendar

**CDs HIGHLIGHTS:**  
Mandrake 8.2 (featuring Kernel 2.4.18, XFree86 4.2, KDE 2.2.2, KOffice 1.11, USB2 support)

**MAGAZINE FEATURING:**  
HP's vision of a Linux future, Linux Format Awards, Security update, Web know-how tutorials, Making Linux accessible

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 1454 271100 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.

**TO SUBSCRIBE from as little as £27.45\*** call 0870 4448645<sup>†</sup> or log onto [www.linuxformat.co.uk](http://www.linuxformat.co.uk)

\*for 6 months paying by Direct Debit saving £22.97 per year (CD issue)

<sup>†</sup> or call +44 1458 271100 from overseas

**LINUX**  
FORMAT  
[www.linuxformat.co.uk](http://www.linuxformat.co.uk)

# Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format* CDs. This month a developer's bonanza.

## On the CD



Wherever you see this logo it means there's related stuff on the CD

### Essential info

On page 000 we have grouped together essential info on the different types of packages on your coverdiscs – along with instructions for installing source packages.

### Important notice

Before you even put the CD or DVD in your drive, please make sure you read, understand and agree to the following: The *Linux Format* CD 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 CD 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.

## READ ME FIRST

**W**e have been busy in the past month, lining up some special treats for this month's cover CDs and DVD. The first CD contains a comprehensive development kit, all you need to start coding (with the exception,

that is, of an idea for a program and some programming knowledge).

The second CD contains the usual mix of utilities, games, network software and system tools. DVD users get a rare treat, a new Debian GNU/Linux release. In

contrast to other distributions, major Debian releases happen once in a blue moon. Read the DVD pages for the full details of Debian 3.0, installable directly from the DVD but also able to be burned to separate CDs for installation on other machines.

### Development/Kylix3

Our series of *Kylix* tutorials has been running for quite some time. While you have been developing your *Kylix* skills, Borland have been developing *Kylix* itself, it has now reached version 3. Amongst the changes and improvements for version 3, *Kylix* now includes a high-performance 32-bit optimising C/C++ native-code compiler. It can now build applications in C++ as well as *Delphi*.

Before you can use *Kylix 3*, you will need to register with Borland to get a key. Pop along to [www.borland.com/products/downloads/download\\_kylix.html](http://www.borland.com/products/downloads/download_kylix.html) – if you don't want to type in the URL, click on the link in the *Kylix* description on the CD. Click on the first "Open edition" link, in the Downloads section, the Keys Only section only has a link for version 2. After submitting the form, you will be taken to the download page and mailed a key. Naturally, you don't need to download it (it's 300+MB despite what the webpage says) because it's on the CD. Download your mail and put the key in your home directory, now you can unpack the tarball from the CD and install with

```
tar xzf /mnt/cdrom/Development/
Kylix3/kylix3_open.tar.gz
cd kylix3_open
./setup.sh
```

You should read the PREINSTALL and INSTALL files before the final step, although on my Mandrake system no pre-installation was needed. If you install on an RPM based system, the installer will update your RPM database with details of the files and packages installed. However, that doesn't mean you should uninstall with *rpm*. Use the uninstall script that is installed into the *kylix3* directory if you ever need to remove it from your system.

### Development/AnjutaIDE

*Anjuta* is a versatile Integrated Development Environment (IDE) for C and C++, written for *GTK* and *GNOME*. *Anjuta* has a wide range of features to ease and speed up the process of writing and developing programs in C


or C++. These include a powerful source editor, project management facilities, an interactive debugger and various application wizards. The source editor has a number of features to improve productivity, such as; syntax highlighting, code formatting, auto-completion, code folding/hiding and function prototypes.

*Anjuta* is on the CD in various formats. Installation from the source tarball is done with the usual **./configure && make && make install** process. The binary RPM is for Red Hat, if you are using a different RPM based distribution it would be safest to create your own binary RPM using `rpm --rebuild anjuta-0.9.99-1.src.rpm` and then install this RPM as usual. Debian users can install from the `.deb` package.

### Kylix 3 licence restrictions

This software can only be installed onto one PC. It is NOT shareware. Applications developed using the software may not be deployed. The software is for personal use only by software developers and may not be used for development or teaching in a commercial or educational establishment. Programs and applications that have been constructed with the

software may not be distributed. The software is provided only with the aim of allowing the user to learn the use of this software. For distribution rights of owner generated applications, the owner will have to purchase a copy of additional software or a package designed for this purpose. No resale of the software is permitted. No free support is available with this software.



**Anjuta is a versatile Integrated Development Environment for C and C++.**

## Development/ KDevelop

*Anjuta* is aimed squarely at GNOME development, but KDE advocates needn't worry as it has its own IDE, *KDevelop*. We have included the slightly older version 2.1.2 as this one works with KDE 2 as well as KDE 3, whereas the later 2.1.3 only works with KDE 3. Both produce code for KDE 3 so you lose almost nothing but gain some flexibility.

Development of *KDevelop 2* stopped in August, all effort is now being concentrated on *KDevelop 3.0*, so this version will remain current until that is released. That is likely to be some time as 3.0 will be a complete rewrite.

Although designed for creating KDE applications, *KDevelop* will run on GNOME provided you have the KDE libraries installed. Similarly, *Anjuta* runs on KDE if GNOME is installed.

*KDevelop* has a comprehensive range of tools and features. This probably accounts, at least in part, for the amount of KDE software available. These features include; the *KWrite* editor with syntax highlighting, an integrated debugger, a dialog editor to simplify the creation of interface code, class tools to show the inheritance and methods of classes, browse classes graphically and add new methods and attributes.

*KDevelop* comes with comprehensive manuals and a "Quickhelp" online help system. There is no RPM or Debian package on the CD, installation consists of unpacking the tarball and running the standard compile commands.

## Development/Qt

*Qt* is a platform-independent C++ toolkit. It provides a standard API for all functions related to the GUI, networking, database access and file handling. This means that a program written on one platform using *Qt* can be ported to another platform far more quickly than if you had accessed the various APIs of that platform directly. The CD (and the DVD) contains a tarball of *Qt 3.0.5* as well as packages for the major distributions. Install it using whichever method suits you best. After all, *Qt* is about making your life easier.

The first CD also includes a few extra items including; *autoconf* and *automake* for creating configure scripts and makefiles, the latest version of the *GCC* compiler suite and some HOWTOs on C/C++ programming

It all looks a bit serious so far but don't worry, there's lighter stuff on the second CD. We know what they say about all work and no play, so we have some playthings on this disc.

## Games/VegaStrike

"Here you can explore star systems, engage in trade, make cash and meet your destiny!" is the promise from the *Vega Strike* website. You get to shoot people too. This is a 3D space action role playing game where you make your fortune by a combination of trade and bounty hunting. It is in the spirit of *Elite*, but is by no means a port or copy of *Elite*. This is a complex game with many choices to be made, hopefully getting as many right as


## BX Pro 6

The professional GUI builder for Motif

We reviewed ICS's *Builder Xcessory Pro 6* last issue, and most impressed we were with it too. If you weren't convinced by the review, now's your chance to test it out yourself, with the special evaluation version provided on this month's coverdisc.


*BX Pro* lets you develop *Motif* and Java AWT GUIs quickly and accurately, and it can generate source code in C, C++ or a mixture of C and IDL (IDL is *Motif*'s interface definition language, which lets you separate the interface design from your code, and so aids maintainability). *Motif* is the industry standard GUI toolkit for Unix platforms, and *BX Pro* will work with *Motif 2.1* or the open-source *OpenMotif 2.1*. *OpenMotif* can be freely distributed with open-source operating systems, so is very attractive for development on Linux. *OpenMotif 2.1* is included on the coverdisc in source code and binary forms and should be installed before *BX Pro*

Full instructions for installing *BX Pro* can be found on the coverdisc. Before



**Easy GUI creation with BX Pro – the professional Motif builder.**

you can use the software, you'll need to obtain a license key. Visit [www.scl.com/linuxformat](http://www.scl.com/linuxformat) for more details. Further instructions and tutorials on using *BX Pro* can be found on ICS's web site at [www.ics.com/products/bxpro](http://www.ics.com/products/bxpro).



**KDevelop is a powerful IDE for developing KDE2 and KDE 3 applications**

possible. It's also pure escapism, which is exactly what you need after a hard day's work. If you don't work, or don't work hard, I'm sure you can find another excuse for playing it.

## Desktop/ MenuConvert SystemToE

Having a choice of so many different window managers is great, it lets you change the way your computer works to suit your needs, instead of changing

the way you work to suit your computer's needs. However, there are drawbacks to this level of customisation. After spending time setting up your current favourite window manager to work the way you want, you then have to do it all over again if you want to switch to a different one. This isn't such a big deal if you tend to stick with the defaults, but then people that stick with defaults are less likely to switch window managers (or even use Linux).

As its name suggests, this is a menu conversion program. More



## LinuxFormatCoverdiscCD



**Vega Strike is a complex 3D space combat and trading game**

« specifically, it scans your system for any KDE and GNOME menus and adds them all to an *Enlightenment* sub-menu, giving instant access to your KDE and GNOME menus from another window manager.

### Internet/GTransfer Manager

Downloading large files from the Internet can be a real pain in the neck if you don't have a broadband

connection. Even with unmetered dialup you still have to deal with timeouts, or someone wanting to use the phone line. The *GNOME Transfer Manager* is one solution to this problem.

It splits a file into pieces, downloads each piece separately and then reassembles them. You can download each piece in a separate session, without needing to stay connected continuously. It can also speed up the process when downloading from a slow or overloaded server. By downloading two or three segments at a time you can often get a faster overall download than from a single connection to the server.

There is a GNOME applet supplied with the program. You can drag and drop URLs from a browser onto this applet to have them passed straight to *GNOME Transfer Manager*. You can also transfer URLs directly from *Galeon*.

### System/Parted

No matter how much thought you put into partitioning your hard drive, sooner or later your needs are likely to change and you will want to change the space allocated to various partitions. Alternatively, you may decide to buy a larger and/or faster drive and need to transfer partitions between the two.

GNU *Parted* is a tool to create, destroy, resize and copy partitions. It works with the most popular filesystems, *ext2/3*, *ReiserFS*, *XFS* and Windows' *FAT16* and *FAT32* partitions. A partition shouldn't be mounted when you resize it, you don't want something trying to write to it during the resize process, so *Parted* comes as floppy disk images as well as the normal source code tarball. These images can be copied to floppy disks with the *dd* command and booted from. This enables you to work on even your root partition.

## »» CD CONTENTS AT A GLANCE

### Disk A

#### Development

<b>AnjutaIDE</b>	Versatile Integrated Development Environment for C and C++
<b>Autoconf</b>	Produces shell scripts to configure source code packages
<b>Autodepend</b>	Tracks files used during a build to generate dependencies
<b>Automake</b>	Generates Makefiles compliant with the GNU Coding Standards
<b>GCC</b>	The GNU Compiler Collection
<b>Highlight</b>	Source code to HTML, XHTML, RTF, TeX, or LaTeX converter
<b>HOWTO</b>	HOWTOs on C++ Programming
<b>KDevelop</b>	C/C++ development environment for KDE2/3
<b>Kylix3</b>	Kylix 3 - Open Edition
<b>QT</b>	GUI toolkit for software developers

### Disk B


#### Magazine

<b>Emulators</b>	All the files mentioned in this month's Emulation article.
<b>HotPicks</b>	All the programs covered in this month's HotPicks section.
<b>Java</b>	The files to go with this month's Java tutorial
<b>Kylix</b>	Example files from the Kylix tutorials
<b>Perl</b>	Example scripts from this month's Perl tutorial
<b>PHP</b>	PHP source code and example scripts

#### Desktop

<b>CDRecBox</b>	Set of tools to burn CDs under a Linux shell
<b>CDrecord</b>	Creates home-burned CDs with a CD-R/CD-RW recorder
<b>EviIWM</b>	Minimalist window manager for the X Window System
<b>Gnetload</b>	GNOME applet showing the network load
<b>GnomeRIG</b>	GUI to the Ham Radio Control Libraries using GTK+/GNOME
<b>GnomeToaster</b>	Full CD creation suite for X11 and GNOME
<b>GnomeXsu</b>	Gnome Xsu is an interface for the su command
<b>GQradio</b>	Interfaces with FM radio cards through the video4linux API
<b>Ion</b>	Keyboard-friendly window manager
<b>Katoob</b>	Small text editor based on GTK+ 2.0.
<b>KBirthday</b>	Panel applet to remind you of birthdays and anniversaries
<b>KExchange</b>	Convert between 70+ currencies using up-to-date rates
<b>Klat</b>	LaTeX editor for KDE
<b>KPSK</b>	PSK31 digital radio communications application

<b>KRemoteControl</b>	Graphical irexec for KDE
<b>Krusader</b>	An "old-school" twin-panel file manager
<b>MenuConvert_SystemToE</b>	Converts KDE and Gnome menus to Enlightenment
<b>Nightfall</b>	Interactive application to simulate eclipsing binary stars
<b>Openbox</b>	Window manager with no external library dependencies
<b>PackageDataBaseView</b>	Generates an HTML view of an RPM or a dpkg database
<b>SGutils</b>	Variants of the Unix dd command
<b>Synergy2</b>	Share a single mouse and keyboard between computers
<b>TuxTyping</b>	Educational typing tutorial game starring Tux
<b>VisualREGEXP</b>	Design and debug regular expressions
<b>ZIPTool</b>	Control IOMega JAZ and ZIP drives from Linux
<b>Distros</b>	
<b>a-Linux</b>	x86 floppy-based mini-distribution
<b>Games</b>	
<b>AdaptivePoker</b>	Complete Texas Hold'em style poker
<b>AdvanceSCAN</b>	ROM manager for MAME, MESS, AdvanceMAME, AdvanceMESS,
<b>Raine</b>	
<b>BalloonChase</b>	Hot air balloon game
<b>Civil</b>	Cross-platform, turnbased, networked strategy game
<b>CSBuddy</b>	Tool for Counter-Strike server owners
<b>DominoOnAcid</b>	Solitaire variant of Dominoes with weird colourful tiles
<b>Glaxium</b>	OpenGL-based space-ship "shoot-em-up" styled game
<b>JumpNBump</b>	Funny multiplayer game about cute little fluffy bunnies
<b>SDLBlock</b>	3d Tetris game, a clone of BlockOut and XBlockout
<b>VegaStrike</b>	Action 3D space combat simulator
<b>XBubble</b>	X Window-based clone of Bust-A-Move/Puzzle Bubble
<b>Graphics</b>	
<b>Freevo</b>	Turns a PC into a multimedia jukebox/VCR/PVR/HTPC/DVR/STB
<b>Gspy</b>	Processes images from a v4l device into a daily MPEG movie
<b>ImageMagick</b>	Automated and interactive manipulation of images
<b>ImageViewer</b>	Viewer with real time interactive pan and zoom viewing
<b>Kino</b>	Record, create, save, edit and play movies from DV cameras
<b>njbWorld</b>	Edit and interact with VRML and other 3D formats
<b>MT2Thumbnailer</b>	Generate HTML photo albums from pictures
<b>PNG2ico</b>	Converts .PNG files to Windows .ICO icon resource files
<b>Internet</b>	
<b>Balsa</b>	Gnome e-mail client supporting POP3, IMAP and local mail



**Transfer your KDE and GNOME menus to Enlightenment with MenuConvert**

We should warn you that there is always an element of risk when altering disk partitions. For example, a power cut (or tripping over the power lead) during a resize would be most likely trash at least one of your partitions.

For this reason, you should backup all data before running a program like this. This means the important data like the main system partitions, plus the data you think is unimportant... until you lose it (best backup your whole /home to be safe).

## System/Advisor

Everyone agrees that security is important, as is the need to keep your system up to date with the latest packages. Running old software makes you a prime target for hackers who know all about the insecurities in old versions.

That's the theory, in practice, trawling through security advisories to see which apply to your system is not most people's idea of fun. *Advisor* will do the job for you, at least it will if you use Red Hat or Mandrake. This program will monitor the security advisories for your distribution, comparing them with your currently installed packages. When it sees that one of your packages should be upgraded for security reasons, it sends an email with links to the updated packages. The program does all the hard work, all you need to do is click your mouse, which is exactly as it should be.

## Sound/XMMS-alarm

I have a problem getting up early in the morning (actually I have a problem getting up at any time, being early just compounds it). Alarm clocks make a raucous sound designed to awaken you as rudely as possible. Because of Murphy's Law, any radio station you tune a clock radio into will be playing something even worse than the alarm sound at the time you set it for.

Now you can have an alarm that plays when you want when you want, at the volume you want. *XMMS-alarm* turns *XMMS* into an alarm clock, with programmable playlist, sleep times and volume. You can even have it start quiet and fade the volume up slowly, to ease you into the day gently. It also has the advantage that you can't throw a desktop PC across the room when you really don't want to get up. All you need now is a way to get to sleep with those CPU, PSU and case fans running. [LXF](#)

<b>CaffeineMonkey</b>	Ping, traceroute, nslookup and whois with a Web interface
<b>FastFileSearch</b>	Indexes FTP servers and SMB shares
<b>GTransferManager</b>	Retrieve multiple files from the Web
<b>I-Spy</b>	Grab and compare contents of FTP directories and Web pages
<b>Melon</b>	Mailbox flag for X
<b>Opera</b>	Alternative, lightweight, X11-based Web browser
<b>Pan</b>	Newsreader, loosely based on Agent and Gravity
<b>Pushmail</b>	Push email messages to an SMTP server without a local MTA
<b>SpamProbe</b>	Spam detection program
<b>StaticChargeGTK</b>	GTK is a front-end to update your StaticCling hostname
<b>WebSamba</b>	Samba client for the Web
<b>Yahoo2mbox</b>	Retrieves messages from Yahoo! Groups into a local mbox

## Mobile

<b>AppletManager</b>	Enable and disable Zaurus taskbar applets
<b>Discoverer</b>	Small 802.11b network detector for Linux
<b>e3</b>	Full featured text editor
<b>Embeddedkonsole</b>	Console (Shell) with tab support
<b>UnofficialZaurusFAQ</b>	The Zaurus FAQ
<b>zbtattleship</b>	A battleship implementation for the SHARP Zaurus

## Office

<b>CK-Ledger</b>	Double-entry ledger accounting system for PHPGroupWare
<b>GnuCash</b>	Track bank accounts, stocks, income and expenses
<b>jGnash</b>	Personal finance application written in Java
<b>KBudget</b>	Budgeting and money management program for KDE
<b>MrProject</b>	Project management program
<b>SWX2html</b>	swx2html converts OpenOffice Word (.sxw) files to HTML
<b>SWX2txt</b>	Converts Open Office Word files (.sxw) to ASCII text

## Server

<b>ArrowHeadASPServer</b>	Java Servlet that supports the ASP syntax and VBScript
<b>AWLinks</b>	PHP/MySQL-based link catalog administration system
<b>BoxModelAutoHacker</b>	Generates CSS rules that work around IE5 bugs
<b>FreeLinuxCDProjectSystem</b>	The code that makes FreeLinuxCD.org possible
<b>IFtop</b>	Real-time bandwidth usage information on an interface
<b>mod_savi</b>	integrates the Sophos anti-virus engine into Apache
<b>nWorks</b>	Web-based network management system (NMS)
<b>Oinkmaster</b>	Update your Snort rules

<b>Privoxy</b>	Web proxy based on Internet Junkbuster
<b>Shoutstats</b>	Shoutcast server statistic analysis program
<b>Sysmon</b>	Network monitoring tool
<b>VServer</b>	Run many independent servers simultaneously in one box
<b>ZoneMaster</b>	Zone file and name server management tool

## Sound

<b>AiceriMultimediaPlayer</b>	Simple multimedia player using GTK+ 2
<b>BEAST_BSE</b>	GTK+/GNOME-based frontend to the Bedevilled Sound Engine
<b>EQGraphicalEqualizer</b>	Realtime graphical equalizer for XMMS
<b>GnomeMP3</b>	Renames and tags MP3 files
<b>MP3blaster</b>	Interactive text-based player for several audio file types
<b>MusicControlCenter</b>	Organize and expand your sound file collection
<b>RezasRioReceiver</b>	OS replacement for the client software on the Rio Receiver
<b>Sweep</b>	Sound wave editor, recording and playback tool
<b>SynthesisToolKit</b>	Audio signal processing C++ classes and instruments
<b>XMMS-alarm</b>	Use XMMS as an alarm clock
<b>XMMS-fc</b>	XMMS plugin to play back Amiga Future Composer modules

## System

<b>Advisor</b>	Monitors a security advisory database and sends alerts
<b>AGT</b>	Powerful console frontend for iptables
<b>BackPackUSBLoader</b>	Upload firmware to Micro Solutions' BACKPACK drives
<b>BackupByLoop</b>	Automatic and complete system backups
<b>CatweaselLinuxDeviceDriver</b>	Driver for the Catweasel Advanced Floppy Controller
<b>GAG</b>	Graphical Boot Manager with a lot of features
<b>Gromit</b>	Automated system configuration tool
<b>IOzone</b>	Filesystem benchmark tool
<b>IPtablesControl</b>	Fast and easy iptables filter configurator
<b>OpenSSL</b>	Robust, commercial-grade, fully featured SSL toolkit
<b>Parted</b>	Create, destroy, resize, and copy partitions
<b>ReadyExec</b>	Speed up the startup of programs like procmail
<b>RubyUpdate</b>	Keep track of packages installed from source
<b>Stress</b>	Impose certain types of stress on a POSIX system
<b>UPMS</b>	Universal Package Management System, based on BSD ports
<b>Watchfolder</b>	Watches specified folders for incoming files

# Coverdisc



**Neil Bothwick** is your guide through the wonders of this month's jam-packed *Linux Format DVD*. This could be your last installation...

**D**ebian's testing distribution is updated continually, but it takes a long while for these updates to be considered stable and put into a new distribution. So a major version change is a rare event in the Debian world, and we have version 3.0 on the DVD for you. Debian's package management and installation system is generally reckoned to be the best around, but their installer is less friendly than many. So we'll walk through the main steps of setting up a Debian system from the DVD.

## Partitioning the disk

You will need to create at least a swap and root partition on your hard

disk. You can do this beforehand or from within the installer using *cfdisk*. *cfdisk* is fairly self-explanatory, press **H** at any time for help. Note that when you select **Write** to commit your changes to the disk, you have to answer **yes** to the "are you sure?" question, entering **y** is not enough. Once you have your partitions set up, the installer will initialise and mount them. It is generally safe to accept the default options here. When you reach the section on configuring device driver modules, you probably won't need to do anything if you have fairly standard hardware. There's no option to skip this section, so enter it and select **exit**.

If you have a network connection, you should configure it next. If the network also connects you to the Internet, you'll be able to use this to download updated packages towards the end of the installation. The base system will now be installed. Unlike other distribution installers, the Debian installs only the core packages to start with and then reboots the system using this core before installing optional packages. When asked where you want to install the base system from, press **Enter** to accept the default.

## Booting

Once the base system is installed, you need to be able to boot from it. This can either be from the *Lilo* boot loader on your hard disk, or from a floppy disk. If you choose to boot from hard disk, you should still make a boot floppy, for emergency use. Don't forget to remove the DVD before rebooting, otherwise you'll boot from the DVD again and be right back at the start of the installation.

The installation continues automatically after rebooting. After a few setup questions, it will scan the DVD for available packages. This will take a few minutes. When asked to

```

Welcome to Debian GNU/Linux 3.0!

This is a Debian CD-ROM. Keep it available once you have installed
your system, as you can boot from it to repair the system on your hard
disk if that ever becomes necessary (press <F3> for details).

For a "safe" installation with kernel 2.2.20, you can press <ENTER> to begin.
If you want additional features like modern hardware support, specify a
different boot flavor at the boot prompt (press <F3> to get an overview).
If you run into trouble or if you already have questions, press <F1>
for quick installation help.

WARNING: You should completely back up all of your hard disks before
proceeding. The installation procedure can completely and irreversibly
erase them! If you haven't made backups yet, remove the CD-ROM
from the drive and press <RESET> or <Control-Alt-Del> to get back to
your old system.

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law. For copyright information, press <F10>.

Press <F1> for help, or <ENTER> to boot.

boot: _

```

This is what you first see when booting from the DVD. Press **Enter** to begin the installation, or press **F1** to read the help screens before installing.

```

Debian GNU/Linux Installation Main Menu

Debian is now installed (although not yet fully configured).
You should build a custom boot floppy, in case Debian
GNU/Linux doesn't boot directly from the hard disk. Please
select "Next" to build the custom boot floppy.

Next      : Make a Boot Floppy
Alternate: Reboot the System

Configure the Keyboard
Preload Modules from a Floppy
Partition a Hard Disk
Initialize and Activate a Swap Partition
Activate a Previously-Initialized Swap Partition
Do Without a Swap Partition
Initialize a Linux Partition

```

Once the base system is installed you must reboot, but make sure you have set up *Lilo* and/or a boot floppy before you do this.

give a name for the disk, you can use anything that will let you identify it later, LXF33 seems a good choice. At this point you get the option to add another source of packages. Adding an ftp or http source means that it will be able to update packages from the Internet. You also get to add a specific source for security updates.

## Package selection

Now we come to the fun part, choosing the packages to install. The *Taskel* program lets you select groups of packages to install, such as X

Windows, desktop environment, games etc. You can fine tune these choices later if you wish. If you are new to Debian, it is probably best to go with the groups from *Taskel*, you can always add or remove packages later.

The packages you selected will be installed now. If you configured a security source and are online, some updates will be downloaded first. The package installation may take a while, but you may be prompted for information during the process, so you can't go to the pub while it's running. Assuming you select both X window



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.**

## Running Flightgear from last month's DVD

Some people have reported problems running *Flightgear* from last month's DVD. This is because the files are in a different location on the DVD from on the CD. To run from the DVD, type the following, assuming your DVD drive is mounted at /mnt/cdrom


```
cd /mnt/cdrom/Games/Flightgear/bin/
sh Flightgear.sh
```

*Flightgear* should now start up and you can follow the rest of the instructions from last month's magazine.

and desktop environment from the package groups, you will need to set up the X server. The installer can take care of this for you. Don't use the advanced option in Xserver setup unless you have the specifications of

```
Setting up libx11 (1.0.17-2) ...
Setting up libgimp1 (4.2.0-4) ...
Setting up libpaper1 (1.1.8) ...
Setting up xutils (4.1.0-16) ...
Setting up gsfonts-x11 (0.16) ...
/usr/sbin/update-fonts-scale: warning: absolute path /usr/lib/X11/fonts/Type1
s provided.
/usr/sbin/update-fonts-alias: warning: absolute path /usr/lib/X11/fonts/Type1
s provided.
Setting up liboaf0 (0.6.10-2) ...
Setting up oaf (0.6.10-2) ...
Setting up bsdgames (2.13-7) ...
Setting up bzip2 (1.0.2-1) ...
Setting up c2man (2.41-14) ...
```

Installing and configuring packages can take a while, but don't go away, your input may be needed.



Transfer movies from DVD to hard disk, cropping, zooming and re-encoding them on the way.

your monitor to hand, particularly the various operating frequency ranges.

## What else?

That's it, your Debian system should now be installed. Debian seems to have used up most of the available space for discussing the DVD this month, although it has only used half of the space on the DVD. See the CD pages for details of the CD contents, all of which are on the DVD too, and the *DVD At a Glance* box for a complete listing of the disc's extra content. There is just about enough space to highlight a couple of goodies.

*DVDRip* does what it says, it rips movies from DVDs. It also has options to process those movies. For example, you could resize them to something like 320x240 for viewing on a portable device. Re-encode them at a suitable bit rate, and you can fit an entire movie onto a Compact Flash memory card. You'll find *DVDRip* in the Graphics directory. As always with DVD

## Creating Debian CDs from the DVD

Share them with your friends

As part of our ongoing effort to please all of the people all of the time, we have included software to enable you to create ISO images of Debian install CDs from the files on the DVD. This means you can install on other machines without DVD-ROM drives, without setting up a network install. You will of course need a CD burner to burn the ISO images to CD.

Everything you need is in the *jigdo* directory of the DVD. *Jigdo*, an abbreviation of "jigsaw download", was written to allow people to download an ISO image in parts and reassemble it after download, but it works equally well with local files. These instructions assume your DVD drive is mounted at /mnt/cdrom. The first thing you must do is copy the *jigdo* directory to your hard disk, so you'll need 2GB of free space to create them all, or 700MB to make them one at a time. Either type

```
cp -a /mnt/cdrom/jigdo ~/
```

or drag the directory from the DVD to your home directory (or any other directory you can write to that has sufficient space). Now you need to run the *jigdo-lite* program from a shell.

```
cd ~/jigdo
./jigdo-lite woody-i386-1.jigdo
```

```
./jigdo-lite woody-i386-2.jigdo
```


```
./jigdo-lite woody-i386-3.jigdo
```

```
./jigdo-lite woody-i386-4.jigdo
```

When you run *jigdo-lite* the first time, it will ask you if you have an old version of the CD you are creating. enter the path to the DVD here, say /mnt/cdrom. Subsequent runs will offer this path as an option. *jigdo-lite* will find all the files it needs on the DVD, so will never reach the download stage. The first time you run it, *jigdo-lite* scans the contents of the DVD, this can take quite a while so leave it to get on with it. The good news is that it caches the results of this scan, so the following runs are much faster.

This process will create four ISO images, ready for burning to CD. Obviously, if you don't have the 2GB of free space needed for all four, you can create and burn them one at a time. Once you have burned the CD, boot from the first one and install the same as from DVD.

There is also a version of *Jigdo* for Windows, in case your DVD-equipped computer doesn't have Linux installed (yet). This hasn't been tested with the DVD but should work the same. The program and documentation is in a zip file in the *jigdo* directory. After unzipping it, you'll need to copy the *jigdo* and *.template* files to its directory.



OpenOffice.org again, this time with some new features and fixed bugs.

reading programs, encrypted commercial DVDs will require the use of *libdvdcss*, which we are unable to include on the coverdiscs. Check the program's home page for a link or you can find them at [www.dtek.chalmers.se/~dvd](http://www.dtek.chalmers.se/~dvd)

Office/  
**OpenOffice.org**

We included the initial 1.0 release of *OpenOffice.org* on the DVD earlier this year. Since then, a number of bugs

have been discovered and fixed, especially with respect to printing. It's a large package for anyone with a dialup connection to download, so we have the fixed and improved version 1.0.1 of *OpenOffice.org* on the DVD this month. Installation is the same whether you had the original version installed or not. Unpack the archive, **cd** to the install directory and type **./setup**

As always, read the docs before installing, to avoid any surprises. [LXF](#)

## »» DVD CONTENTS AT A GLANCE

### Desktop

<b>Alist</b>	Collects hardware and software information in a database
<b>CVSfs</b>	Mount a CVS project like any file system
<b>Feta</b>	Frontend to the Debian package management system
<b>Freud</b>	Standalone C++ port of the Emacs Psychiatrist
<b>Gcompressor</b>	GUI file compression interface for GNOME
<b>Genmenu</b>	Generates menus for some common window managers
<b>GkACPI</b>	GkrellM plugin to display battery statistics
<b>GKrellM</b>	GTK and Imlib-based stacked monitor program
<b>GLgraph</b>	Interactive OpenGL based function grapher
<b>GNUTypist</b>	Interpreter of typing tutorials
<b>Goats</b>	Post-it note applet for the GNOME panel
<b>GPGKeys</b>	GUI frontend to GPG
<b>HamFax</b>	Sending and receive amateur radio facsimiles
<b>Hopify</b>	Tools and documentation to teach you how to speed read
<b>HTMLtoPDF</b>	Generates PDF files from a Web pages
<b>Karenn</b>	Ship hull design software
<b>Katalog</b>	Create/list/search an index of your CDs
<b>Navigatrix</b>	Navigatrix is a GPS navigational system for KDE.
<b>Paralogger</b>	Tails the system logs in borderless transparent Eterm(s)
<b>PlextorTool</b>	Change various properties of Plextor CD and CD-R(W) drives
<b>Slackports</b>	BSD-like ports system for Slackware
<b>TrashCan</b>	Command line recycle bin
<b>VTWM</b>	Virtual window manager with adjustable graphical complexity
<b>Wmpasman</b>	Stores passwords and makes them available for pasting
<b>Xmldiff</b>	Show the differences between two similar XML files
<b>Xplore</b>	Powerful and highly configurable Motif file manager
<b>XWine</b>	Graphical user interface for the WINE emulator
<b>ZuretafaceMon</b>	Displays the activity of your network interfaces

### Development

<b>Barter</b>	Code generator for AspectJ
<b>Botan</b>	Library of cryptographic algorithms
<b>Buildtool</b>	Make programs more portable and easier
<b>Dialog</b>	Use dialog boxes from a shell script
<b>DiaSCE</b>	C/C++ code editor for GNOME
<b>dietLibc</b>	Create small, statically-linked binaries
<b>Eboxy</b>	Builds simple user interfaces for entertainment PCs
<b>ExOr</b>	Creates self-extracting and self-installing tarballs
<b>Execline</b>	Light, non-interactive scripting language
<b>Fenris</b>	Multipurpose tracer, debugger, and code analysis tool
<b>GnomePython</b>	Set of interfaces to gnome-libs
<b>GPGME</b>	Makes access to GnuPG easier for applications
<b>GTK2ForPascal</b>	Translation of the GTK 2 headers to Free Pascal and Kylix
<b>Hibernate</b>	Object/relational persistence and query service for Java
<b>ID3lib</b>	Manipulate meta-information in digital audio files
<b>ImLib3D</b>	Library and visualization system for 3D image processing
<b>OpenCM</b>	Secure, high-integrity replacement for CVS
<b>PerITidy</b>	Perl script indenter and beautifier
<b>Pounder</b>	Utility for testing Java GUIs
<b>PyGame</b>	Python extension modules designed for writing games
<b>PyGTK</b>	Set of Python bindings for the GTK widget set
<b>PyKyra</b>	Fast game development framework for Python
<b>PyMAD</b>	Use the MPEG Audio Decoder library with Python
<b>PyUID</b>	Unique ID (UID) generator in Python
<b>SGE</b>	Graphics library for the Simple Direct Media Layer
<b>SLOccount</b>	Count physical source lines of code (SLOC) in large systems
<b>SQLite</b>	C library that implements an SQL database engine

### Distros

<b>Debian</b>	Debian 3.0, installable from the DVD
---------------	--------------------------------------

### Games

<b>Enigma</b>	Version of Oxyd, a puzzle game popular on the Atari ST
<b>Gnono</b>	GNOME version of the Windows card game WUNO
<b>KsirK</b>	Computerised version of a well known strategy game
<b>LBreakout</b>	Breakout game with nice effects, graphics and sounds
<b>LM-Solve</b>	Solver for several types of puzzles on the Logic Mazes site
<b>MindlessAutomaton</b>	Play collectable/trading card games over a network
<b>Openglad</b>	SDL port of an old DOS game called Gladiator
<b>Six</b>	A game with very simple rules and deep tactical complexity

### Graphics

<b>AstroCamNetwork</b>	Webcam network for everybody
<b>DVD-rip</b>	Simplifies the whole process of making copies of your DVDs
<b>EXIFtags</b>	Extracts EXIF data from digital camera images
<b>Kavi2SVCD</b>	Generatesg MPEG files from an AVI file
<b>Motion</b>	Uses a video4linux device for detecting movement
<b>Namp</b>	GTK2 frontend for mplayer

### Internet

<b>DownloaderForX</b>	Downloads files via both HTTP and FTP
<b>Gnarwl</b>	Email autoreply tool
<b>GrabCartoons</b>	Comic-grabbing utility
<b>L2TP</b>	Layer 2 Tunneling Protocol VPN client/daemon
<b>Netrik</b>	Full-featured, text mode WWW browser
<b>NetworkQueryTool</b>	Get information about a domain or IP address
<b>Werkmail</b>	Browser-based mail client written in PHP

### Mobile

<b>GNUchess</b>	IPKG of GNUchess by spiralman.
<b>mfighter</b>	Street Fighter style game for the Zaurus
<b>QPhoto</b>	Photo viewer
<b>TheEliminator</b>	Eliminate all balls on the board
<b>WirelessApplet</b>	Displays the status of a wavelan card in a taskbar icon
<b>ZRally</b>	Race against the Computer, Race against the Clock

### Office

<b>AWOL</b>	In/out board
<b>MantisHelpDesk</b>	Help desk application, designed to be easy to use
<b>OpenOffice.org</b>	Bug-fixed update to OpenOffice.org

### Server

<b>AIMSniff</b>	Monitors and archives AOL Instant Messenger messages
<b>CyrusIMAPServer</b>	IMAP server to be run on sealed systems
<b>Drall</b>	Secure remote access
<b>FuseCMS</b>	Content Management System
<b>Gcount</b>	Solid Web-based counter in PHP
<b>LABE</b>	Web application to administer a centralized LDAP directory
<b>PHPSquidPass</b>	Change user authentication files for the squid Web proxy
<b>PLP</b>	Perl embedder, primarily for HTML documents
<b>SingleHoneyPot</b>	Simulates many services like SMTP, HTTP, shell and FTP
<b>Spasm</b>	Spam filter for sendmail 8.12+
<b>WebCyradm</b>	Web-based tool to admin the Cyrus IMAP Server
<b>WebPackageSurfer</b>	PHP-based package management system

### Sound

<b>aRts</b>	Framework for developing modular multimedia applications
<b>BurnMP3</b>	Tcl/Tk GUI for burning audio CDs from MP3 files
<b>Ecamegapedal</b>	Realtime effect processor
<b>GAIM-XMMS</b>	GAIM plugin
<b>Julie</b>	Configurable music jukebox that supports MP3 and Ogg
<b>MP3CD-tools</b>	Packs directories of MP3 files in a space-efficient way
<b>Normalize</b>	Adjust the volume of audio files to a standard volume level
<b>Orpheus</b>	Text-mode menu- and window-driven player for CDs and MP3s
<b>Zina</b>	Graphical interface to your MP3 collection

### System

<b>Atop</b>	ASCII full-screen performance monitor
<b>Dnotify</b>	Execute a command when the contents of a directory change
<b>Duplicity</b>	Incremental backup tool
<b>FBlogo</b>	Replace the penguin boot-logo you see when booting
<b>Fcron</b>	Periodical command scheduler
<b>iBackup</b>	Automate the backup of system configurations
<b>Keychain</b>	Manage RSA and DSA keys
<b>LILO</b>	Boot loader for Linux/x86 and other PC operating systems
<b>Mondo</b>	Simplistic health monitoring daemon
<b>MonModules</b>	A collection of modules for mon
<b>RPMupdate</b>	Automates the download and installation of errata RPMs
<b>Stitch</b>	Makes backups from many computers to a large storage array
<b>StoreBackup</b>	Backup utility that stores files on other disks
<b>TTF2PT1</b>	Converter of various scalable font formats

# Essential disc info

Read this important information before you use your Linux Format coverdisc – CD or DVD.  
We've collated some helpful info to help you get the most from these jewels of data!

## Finding the essentials

### Missing something?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases the latest libraries and

other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

### Package formats

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

### Documentation

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

## Creating install CDs with cdrecord

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord -scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v  
/path/to/image.iso
```

You can simplify the command by saving some default settings in */etc/default/cdrecord*. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the "Burn" tab and the "ISO 9660 Image" gadget near the top of the window. Put the path to the image file in the gadget and press "Combust!". Now put on the kettle while the CD is created for you.

### Other OS?

You 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 info.

## 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 xvzf --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 (e.g. */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.**

# User Groups

Your local Linux User Group needs you! LUGs worldwide are full of members keen to help with your problems, discuss ideas and generally natter about all things Linux. We have collected a load of information here so you can find the LUG closest to you. You can find lots more information online at: [www.lug.org.uk](http://www.lug.org.uk) or <http://lugwww.counter.li.org/groups.cms>

## 1 Hampshire

URL [www.hants.lug.org.uk](http://www.hants.lug.org.uk)  
Contact Hugo Mills

## 2 Bristol & Bath

URL [www.bristol.lug.org.uk](http://www.bristol.lug.org.uk)

## 3 Scottish

URL [www.scottish.lug.org.uk](http://www.scottish.lug.org.uk)

## 4 Oxford

URL [www.oxford.lug.org.uk](http://www.oxford.lug.org.uk)  
Contact Alasdair G Kergon

## 5 Kent

URL [www.kent.lug.org.uk](http://www.kent.lug.org.uk)  
Contact John Mills

## 6 Brighton

URL [www.brighton.lug.org.uk](http://www.brighton.lug.org.uk)  
Contact Johnathan Swan

## 7 Worcestershire

URL [www.worcs.lug.org.uk](http://www.worcs.lug.org.uk)  
Email [info@thirdeyedevlopment.com](mailto:info@thirdeyedevlopment.com)

## 8 Northants

URL [www.northants.lug.org.uk](http://www.northants.lug.org.uk)  
Contact Kevin Taylor

## 9 Anglian

URL [www.anglian.lug.org.uk](http://www.anglian.lug.org.uk)  
Contact Martyn Drake

## 10 Milton Keynes

URL [www.mk.lug.org.uk](http://www.mk.lug.org.uk)  
Contact Denny De La Haye

## 11 Doncaster

URL [www.doncaster.lug.org.uk](http://www.doncaster.lug.org.uk)  
Contact Andy Smith

## 12 Moray

URL [www.moray.lug.org.uk](http://www.moray.lug.org.uk)  
Contact Stewart Watson

## 13 West Wales

URL [www.westwales.lug.org.uk](http://www.westwales.lug.org.uk)  
Contact Dan Field

## 14 Wolves

URL [www.wolves.lug.org.uk](http://www.wolves.lug.org.uk)  
Contact Jono Bacon

## 15 Peterborough

URL [www.peterboro.lug.org.uk](http://www.peterboro.lug.org.uk)  
Contact Steve Gallagher

## 16 Edinburgh

URL [www.edinburgh.lug.org.uk](http://www.edinburgh.lug.org.uk)  
Contact Alistair Murray

## 17 Tyneside

URL [www.tyneside.lug.org.uk](http://www.tyneside.lug.org.uk)  
Contact Brian Ronald

## 18 Leicester

URL [www.leicester.lug.org.uk](http://www.leicester.lug.org.uk)  
Contact Clive Jones

## 19 Greater London

URL <http://gl.lug.linux.co.uk/>  
Contact John Southern

## 20 Surrey

URL [www.surrey.lug.org.uk](http://www.surrey.lug.org.uk)  
Contact Jay Bennie

## 21 Cambridge

URL [www.cam-lug.org](http://www.cam-lug.org)

## 22 Devon & Cornwall

URL [www.dclug.org.uk](http://www.dclug.org.uk)  
Contact Simon Waters

## 23 Falkirk

URL [www.falkirk.lug.org.uk](http://www.falkirk.lug.org.uk)

## 24 Manchester

URL [www.manlug.mcc.ac.uk](http://www.manlug.mcc.ac.uk)  
Contact John Heaton, Owen Le Blanc

## 25 Hertfordshire

URL [www.herts.lug.org.uk](http://www.herts.lug.org.uk)  
Contact Nicolas Pike

## 26 West Yorkshire

URL [www.wyflug.lug.org.uk](http://www.wyflug.lug.org.uk)  
Contact Jim Jackson

## 27 Sheffield

URL [www.sheflug.co.uk](http://www.sheflug.co.uk)  
Contact Richard Ibbotson

## 28 Staffordshire

URL [www.staffslug.org.uk](http://www.staffslug.org.uk)

## 29 North East

URL [www.shofar.uklinux.net/NELUG](http://www.shofar.uklinux.net/NELUG)

## 30 London

URL [www.lonix.org.uk](http://www.lonix.org.uk)

## 31 Thames Valley

URL [www.sclug.org.uk](http://www.sclug.org.uk)

## 32 Liverpool OpenSource

URL [http://linux.liv.ac.uk/\\_liv\\_linux\\_ug/](http://linux.liv.ac.uk/_liv_linux_ug/)  
Contact Simon Hood

## 33 Deal Amiga Club

Email [superhighwayman@hotmail.com](mailto:superhighwayman@hotmail.com)  
Contact John Worthington

## 34 Chesterfield

Email [spirelug@yahoo.co.uk](mailto:spirelug@yahoo.co.uk)  
Contact Robin Needham

## 35 South Derbyshire

URL [www.sderby.lug.org.uk](http://www.sderby.lug.org.uk)  
Contact Dominic Knight

## 36 Belfast (BLUG)

URL [www.belfastlinux.cx](http://www.belfastlinux.cx)  
Email [russell@belfastlinux.org](mailto:russell@belfastlinux.org)

REVISED  
DETAILS

## 37 Wiltshire

URL [www.wiltshire.lug.org.uk](http://www.wiltshire.lug.org.uk)  
Contact Jason Rudgard

## 38 South London

URL [www.sl.lug.org.uk](http://www.sl.lug.org.uk)  
Email [ben@ilovephilosophy.com](mailto:ben@ilovephilosophy.com)

## 39 Cheshire

URL [www.sc.lug.org.uk](http://www.sc.lug.org.uk)  
Contact Anthony Prime – [enquiry@sc.lug.org.uk](mailto:enquiry@sc.lug.org.uk)

## 40 North Wales

URL [www.northwales.lug.org.uk](http://www.northwales.lug.org.uk)  
Contact Jonathan Cole

## 41 Midlands

URL <http://midlandslug.port5.com/>  
Contact Pete Thompson

REVISED  
DETAILS

## 42 Cumbria

URL [www.cumbria.lug.org.uk](http://www.cumbria.lug.org.uk)  
Contact Jamie Dainton

## 43 Dorset

URL [www.dorset.lug.org.uk](http://www.dorset.lug.org.uk)  
Contact John and Mat

## 44 Shropshire

URL [www.shropshire.lug.org.uk](http://www.shropshire.lug.org.uk)  
Email [shropshire@lug.org.uk](mailto:shropshire@lug.org.uk)

## 45 South West

URL [www.southwestlug.uklinux.net](http://www.southwestlug.uklinux.net)  
Email [southwest@lug.org.uk](mailto:southwest@lug.org.uk)

## 46 South Wales

URL [www.sw.lug.org.uk](http://www.sw.lug.org.uk)  
Contact Tim Bonnell

## 47 North London

URL <http://www.kemputing.net/lug/anlug-aims.html>

## 48 Malvern

URL [www.malvern.lug.org.uk](http://www.malvern.lug.org.uk)  
Contact Greg Wright

## 49 Huddersfield

URL [www.hud.lug.org.uk](http://www.hud.lug.org.uk)  
Contact Adam Brookes

## 50 Nottingham



URL [www.nottingham.lug.org.uk](http://www.nottingham.lug.org.uk)  
Contact Godfrey Nix


## 51 St Albans & Luton

URL [www.lust.lug.org.uk](http://www.lust.lug.org.uk)  
Contact Michael Culverhouse – [mike@easily.co.uk](mailto:mike@easily.co.uk)

## 52 Wrexham

Contact Paul Kersey-Smith  
Email [paul@pkls.fsnet.co.uk](mailto:paul@pkls.fsnet.co.uk)



**53 Preston & Lancs**

URL [www.preston.lug.org.uk](http://www.preston.lug.org.uk)  
 Contact Phil Robinson

**54 Derry**

URL [www.derry.lug.org.uk](http://www.derry.lug.org.uk)

**55 Isle of Wight**

URL [www.iow.lug.org.uk](http://www.iow.lug.org.uk)  
 Contact David Groom – [info@iow.lug.org.uk](mailto:info@iow.lug.org.uk)

**56 Scarborough**

URL [www.scarborough.lug.org.uk](http://www.scarborough.lug.org.uk)  
 Contact Jamie Adams

**57 Blackburn**

Contact Matt Wright  
 Email [matt@consultmatt.co.uk](mailto:matt@consultmatt.co.uk)

**58 York**

URL [www.york.lug.org.uk](http://www.york.lug.org.uk)  
 Contact Ewan Mac Mahon

**59 Lincs**

URL [www.lincs.lug.org.uk](http://www.lincs.lug.org.uk)  
 Contact Chris Lingard

**60 Hull**

URL <http://hull.lug.cjb.net>  
 Contact [hull.lug@jaxxnetwork.co.uk](mailto:hull.lug@jaxxnetwork.co.uk)

**61 Walton-On-Thames**

Contact William Mutch  
 Email [rael@freeuk.com](mailto:rael@freeuk.com)

**NEW  
DETAILS****62 Gloucestershire & Cotswolds**

URL [www.gloucs.lug.org.uk](http://www.gloucs.lug.org.uk)  
 Contact Francis Barton

**63 West of Scotland**

Contact Ray Gillian  
 Email [raygillan@fire-bug.co.uk](mailto:raygillan@fire-bug.co.uk)

**64 South Staffordshire**

URL [www.staffs.lug.org.uk](http://www.staffs.lug.org.uk)  
 Contact Oliver Keenan

**65 Mansfield**

URL [www.mansfield.lug.org.uk](http://www.mansfield.lug.org.uk)  
 Contact Brent Vardy

**66 Borders**

URL [www.linux.bordnet.co.uk](http://www.linux.bordnet.co.uk)  
 Contact Welby McRoberts

**67 South Birmingham**

URL [www.sb.lug.org.uk](http://www.sb.lug.org.uk)  
 Contact Tim Williams

**68 Coventry**

Contact Darren Austin  
 Email [info@coventry.lug.org.uk](mailto:info@coventry.lug.org.uk)

**69 Newark & Lincoln**

URL [www.newlinc.lug.org.uk](http://www.newlinc.lug.org.uk)

**70 Bedfordshire**

URL [www.beds.lug.org.uk](http://www.beds.lug.org.uk)  
 Contact Neil Darlow

**71 Lincoln**

URL [www.lincoln.lug.org.uk](http://www.lincoln.lug.org.uk)  
 Contact Jon Shamash

**72 Loughborough**

URL [www.loughborough.lug.org.uk](http://www.loughborough.lug.org.uk)  
 Contact Martin Hamilton

**73 Exeter University**

Contact Nicholas Murison  
 Email [N.J.Murison@exeter.ac.uk](mailto:N.J.Murison@exeter.ac.uk)

**74 Sunderland**

Contact Thomas Croucher  
 Email [thomas.croucher@sunderland.ac.uk](mailto:thomas.croucher@sunderland.ac.uk)

**75 East Yorkshire**

Contact Daniel Gallacher  
 Email [sharkonline@whatemail.com](mailto:sharkonline@whatemail.com)

**76 Cleveland Open Source Group**

Contact Haniff Din  
 Email [openlug@digitalmedia.co.uk](mailto:openlug@digitalmedia.co.uk)

**77 Beverley**

Contact Vladimir Lukyanov  
 Email [vladimir\\_lukyanov@hotmail.com](mailto:vladimir_lukyanov@hotmail.com)

**78 Dundee & Tayside**

URL [www.dundee.lug.org.uk](http://www.dundee.lug.org.uk)  
 Contact Duncan Gauld

**79 Sussex**

URL [www.phpworld.co.uk/~swlug](http://www.phpworld.co.uk/~swlug)  
 Contact Gareth Ablett

**80 Wigan & St Helens**

Contact Paul F. Johnson  
 Email [paulf.johnson@ukonline.co.uk](mailto:paulf.johnson@ukonline.co.uk)

**81 Brixton**

URL [www.communitytechnology.org.uk/~linuxhome](http://www.communitytechnology.org.uk/~linuxhome)  
 Contact R.M. sanchez

**82 St.Andrews, Fife**

Contact Stuart Anderson  
 Email [stuart@nx14.com](mailto:stuart@nx14.com)

**83 Nuneaton**

URL [www.nuneaton.lug.org.uk](http://www.nuneaton.lug.org.uk)  
 Contact S Prosser

**84 Isle of Man**

Contact John Mylchreest  
 Email [john@dpm.co.im](mailto:john@dpm.co.im)

**85 Aylesbury**

URL [www.aylesbury.lug.org.uk](http://www.aylesbury.lug.org.uk)  
 Email [drbond@educational-computing.co.uk](mailto:drbond@educational-computing.co.uk)

**NEW  
DETAILS**

# LinuxUserGroups

## LUG OF THE MONTH!

### Northants

The Northants Linux Users Group held its first meeting in August 1998, at the house of the founder, Kevin Taylor. As one of the first LUGs in the area, we were never quite sure how it would take off, but we are still going strong, with regular monthly meetings and an active mailing list.

We have a range of members, from experienced, unstable Debian users (you know who you are :), through to people installing for the very first time. All are welcome, and

no question is too silly to ask.

Every three months we have a formal meeting at a local college, consisting of a talk or demo by one of the group members, coffee and chat, finishing up in a local pub. Between these, we have informal sessions at someone's house, where members can bring machines along for practical help and advice. Online facilities include mailing lists, our own IRC channel and webpages – supported by [www.lug.org.uk](http://www.lug.org.uk).

Members' interests range from Linux in education, music (playing and composition), coding, networking and security, and the use of Linux by blind and visually impaired users.

Through regular book reviews we have built up a comprehensive book and CD library. Members are also involved in writing LDP HOWTOs, kernel hacking and supporting advocacy initiatives.

[www.northants.lug.org.uk](http://www.northants.lug.org.uk)



## Worldwide Linux User Groups

Free Software users across the globe

### Australia

#### ADELAIDE

URL [www.linuxsa.org.au](http://www.linuxsa.org.au)  
Email [mtippet@anu.edu.au](mailto:mtippet@anu.edu.au)

#### ALICE SPRINGS

URL [www.aslug.org.au](http://www.aslug.org.au)

#### MELBOURNE, VICTORIA

URL [www.luv.asn.au](http://www.luv.asn.au)  
Contact [luv-committee@luv.asn.au](mailto:luv-committee@luv.asn.au)

#### PERTH

URL <http://plug.linux.org.au/>

### Europe

#### AUVERGNE

URL [www.linux-arverne.org](http://www.linux-arverne.org)  
Email [Cyril.Hansen@wanadoo.fr](mailto:Cyril.Hansen@wanadoo.fr)

#### COSTA DEL SOL (English speaking)

URL [www.fuengirola.lug.org.uk](http://www.fuengirola.lug.org.uk)

#### DENMARK

AIssund [www.alslug.dk](http://www.alslug.dk)

Esbjerg [www.eslug.dk](http://www.eslug.dk)

Fyns [www.flug.dk](http://www.flug.dk)

Midt-og Vestjylland [www.mvjlug.dk](http://www.mvjlug.dk)

Nordjylland [www.njlug.dk](http://www.njlug.dk)

Skåne Sjælland [www.sslug.dk](http://www.sslug.dk)

Trekantsområdet [www.tlug.dk](http://www.tlug.dk)

Vest-fyn [www.haarby-net.dk/vflug](http://www.haarby-net.dk/vflug)

Århus [www.aalug.dk](http://www.aalug.dk)

#### EIRE

URL [www.linux.ie](http://www.linux.ie)

Email [root@linux.ie](mailto:root@linux.ie)

URL [www.dilu.org](http://www.dilu.org)

Email [glossary@dilu.org](mailto:glossary@dilu.org)

#### GOTHENBURG

<http://nain.oso.chalmers.se/LUGG/>

UK: Don't forget the distro-specific lists:

URL [www.lug.org.uk/maillist.html](http://www.lug.org.uk/maillist.html)

### India

URL [www.linux-india.org](http://www.linux-india.org)

Email [newsmaster@linux-india.org](mailto:newsmaster@linux-india.org)

#### TRIVANDRUM

URL [www.river-valley.com/tux](http://www.river-valley.com/tux)

Email [anil@river-valley.com](mailto:anil@river-valley.com)

### Middle East

#### EGYPT

URL [www.linux-egypt.org](http://www.linux-egypt.org)

Contact Hesham Bahram

### North America

#### ALASKA

URL [www.aklug.org](http://www.aklug.org)  
Email [deem@wdm.com](mailto:deem@wdm.com)

#### ALBERTA

URL <http://calgary.linux.ca/>

#### BATON ROUGE

URL [www.br lug.net](http://www.br lug.net)  
Email [dpuryear@usa.net](mailto:dpuryear@usa.net)

#### BAY AREA

URL [www.balug.org](http://www.balug.org)  
Email [aftyde@balug.org](mailto:aftyde@balug.org)

#### CLARKSVILLE, TN

URL [www.cllug.org](http://www.cllug.org)  
Email [tux@cllug.org](mailto:tux@cllug.org)

#### DENVER

URL <http://clue.denver.co.us/>

#### FLORIDA

URL [www.flux.org](http://www.flux.org)

#### LOS ANGELES

URL [www.lalugs.org](http://www.lalugs.org)  
Email [dank@alumni.caltech.edu](mailto:dank@alumni.caltech.edu)

#### NORTH COLORADO

Email [nclug@nclug.org](mailto:nclug@nclug.org)

#### OTTAWA CANADA

URL [www.oclug.on.ca](http://www.oclug.on.ca)

#### TAMPA

URL [www.suncoastlug.org](http://www.suncoastlug.org)  
Email [president@suncoastlug.org](mailto:president@suncoastlug.org)

#### UHACC Normal, IL

URL [www.uhacc.org](http://www.uhacc.org)  
Email [lug@uhacc.org](mailto:lug@uhacc.org)

#### VIRGINIA TECH

URL [www.vtlug.org](http://www.vtlug.org)  
Email [nega@vt.edu](mailto:nega@vt.edu)

### South America

#### BUENOS AIRES

Email [dcoletti@impost.com.ar](mailto:dcoletti@impost.com.ar)

#### CHILE

URL [www.linux-chile.org](http://www.linux-chile.org)

#### MONTEVIDEO

URL [www.linux.org.uy](http://www.linux.org.uy)

#### PARAGUY/ ASUNCION

Email [rolgiati@conexion.com.py](mailto:rolgiati@conexion.com.py)

#### SAO PAULO

URL <http://gul.ime.usp.br/>  
Email [gul@ime.usp.br](mailto:gul@ime.usp.br)



In the first of a new series of columns, **Jono Bacon** says if you like Linux, don't just keep it to yourself.

**Well, welcome to our first** mini-article on advocating and promoting the use of Linux. In this first article we will look at advocacy, why we do it, and how we can best use it.

Put simply, advocacy is the fine art of promoting something that you like, enjoy or find useful. Take for example your favourite band/musician/group. If you like their music, and you meet someone who you think might like them also, you may suggest them to this person. This is an act of advocacy. Now, the person you suggested the band to may go away and listen to them and really like them, buy their CDs, see them live and otherwise support the group. Now something special has happened – your advocacy of the group has directly helped the band which may in turn help them become a better band. This is the aim of these articles – to get Linux into more places and used by more people.

Advocacy however, may not be as straightforward as it seems. It is not

just a simple case of saying "Hey, use Linux...it's cool". Granted, you may get some guy nicknamed Crazy Sam to install Mandrake but, for the greater good, it is not the best technique. The first step is to understand what you are promoting. Sure, Linux is stable, it is free, it is capable, but do you know how many different architectures it will run on, what companies support it and where people can get it on CD?

Next step is in understanding your audience. If you are going to advocate to schools, learn the educational and cost benefits of Linux. Businesses need to know the support benefits and home users need to know how pretty and easy to use it is.

Next month we will look at how we can best identify the connections between Linux and our audience, and until then I suggest you choose an audience to target (business, charities, schools etc) and look into the needs of the audience. We can then begin to draw a battle plan...

## 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](mailto:linuxformat@futurenet.co.uk)

# NEXT MONTH

## Issue 34 on sale Monday 4 November

### EDITORIAL

**Editor** Nick Veitch [nick.veitch@futurenet.co.uk](mailto:nick.veitch@futurenet.co.uk)

**Art Editor** Julian Jefferson [julian.jefferson@futurenet.co.uk](mailto:julian.jefferson@futurenet.co.uk)

**Reviews Editor** Richard Drummond [richard.drummond@futurenet.co.uk](mailto:richard.drummond@futurenet.co.uk)

**Production Editor** Richard Smedley [richard.smedley@futurenet.co.uk](mailto:richard.smedley@futurenet.co.uk)

### Editorial Contributors

Jono Bacon, Neil Bothwick, Chris Brown, David Cartwright, Paul Cavanagh, Andy Channele, David Coulson, Hoyt Duff, Marco Fioretti, Simon Goodwin, Dave Hodgkinson, Chris Howells, Paul Hudson, Jon Kent, Brian Long, Neil Lucock, Tim Parker

### COVER CD PRODUCTION

**CD Editor** Neil Bothwick

**Senior CD Editor** Kate Hadley

### ART CONTRIBUTORS

**Photography** Philip Sowels, Louise Parker, DaveCaudery, Eyewire

**Illustration** Paul Bateman, Chris Winn, Shane Collinge

### ADVERTISING SALES

Matt Stanley 01225 442244

[matt.stanley@futurenet.co.uk](mailto:matt.stanley@futurenet.co.uk)

### MARKETING AND PROMOTIONS

**Marketing Manager** Fiona Tully

**Subscriptions Executive** Nick Skardon

### PRODUCTION

**Production Co-ordinator** Diane Ross

**Group Production Manager** Clare Tovey

**Advertising Production Co-ordinator** Jo Crosby

### MANAGEMENT

**Publisher** Dave Taylor

**Group Advertising Manager** Simon Moss

**Group Publisher** John Weir

**Circulation Director** Sue Hartley

**Managing Director** Colin Morrison

### DISTRIBUTION AND CIRCULATION

**Circulation Manager** Jamie Malley

[jamie.malley@futurenet.co.uk](mailto: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

Simon Wear [simon.wear@futurenet.co.uk](mailto:simon.wear@futurenet.co.uk)

Tel +44(0)1225 442244 Fax +44 (0)1225 732361

### Contact Details

Linux Format, 30 Monmouth Street, Bath BA1 2BW

**Tel** +44 (0)1225 442244 **Email** [linuxformat@futurenet.co.uk](mailto:linuxformat@futurenet.co.uk)

### Subscriptions and Mail Order

**Phone** +44 (0)1458 271178. See page 99

**Email** [linuxformat.subs@futurenet.co.uk](mailto: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 recognises all copyrights and trademarks. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight.

**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. We cannot be held responsible for mistakes or misprints. Linux Format recognises all copyrights in this issue. Where possible we have acknowledged the copyright holder. Please contact us if we have failed to credit copyright.

**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 © Future Publishing Ltd 2002

**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

[www.thefuturenetwork.plc.uk](http://www.thefuturenetwork.plc.uk)



BATH • LONDON • MILAN •  
NEW YORK • PARIS •  
SAN FRANCISCO • WROCLAW



Yes, it's that time of year again, the UK Linux Expo show in London. If you'll be travelling along to Olympia, we'll see you there, but if not, next issue will give you the lowdown on the launches, the speeches, the stands and the editor being heckled during the Great Linux Debate (probably). This year there should be some interesting new products, and with the big three distros all on the verge of a new release, there's sure to be plenty going on.

If you aren't interested in (ugh) people, but prefer the pure joys of maintaining a server, or designing a dynamic website, you simply must read our extensive feature on **PHP Accelerators** which will uncover just how much faster your website could work.

On review we're hoping to have the very latest **Red Hat 8.0**, **Netraverse Terminal server**, more server hardware, the new version of **Arkeia** backup software and plenty more.

## DON'T FORGET

There'll also be another lavish helping of Linux Pro, complete with more tales of Linux deployments, tips and security advice.



## DON'T MISS YOUR COPY OF LINUX FORMAT!

Get it delivered to your door every month – subscribe on page 99