

# Worried? You should be...


**W**hile the *MySQL/Apache/PHP* combination seems to power more and more websites these days, and it's easier to set up than ever, that doesn't mean that you are getting the most out of the combination. Database construction is often seen as something of a black art, and generally people seem happy that it works, rather than being concerned about how appropriate the configuration is. We're trying to help you leverage those extra percentage points of performance (indeed, we'll even show you how to measure the performance) from your database with a range of tips and advice from *MySQL* expert Jeremy Cole. You'll find the advice covers all sorts of topics from basic database setup to interacting with Perl and PHP, and I think you'll find something worthwhile here, no matter how much of a database guru you think you are.

We also have a special 'What on Earth' this issue, with an extended look at Microsoft's 'Palladium' and the whole TCPA gambit. If you've

never heard of either of these, you should definitely take a read of this piece and find out how and why other people want to determine what you can and can't do with your computer.

Given that most of the implications for the consumer are detrimental, you might wonder why anyone would buy such hardware – but the scary thing is that you may end up with no choice, especially if you aren't aware of it. As outlined here, it seems doubtful that TCPA will be able to do any of the 'good' things it claims (prevent viruses, reject spam, etc) but be very good at the things that it doesn't make any great claims about (stop you playing DVDs out of region, implement pay-per-use software and, possibly, prevent you from running an OS you have built yourself). Worried? Check it out on p56.

But cheer up, there's all sorts of good stuff in the mag too, including a look at USB 2.0 and USB 2.0 devices, Debian 3 on test, the firewall roundup and plenty more!



**Nick Veitch** EDITOR

# 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.

**Don't play in the traffic, build a better database p44**

**Upgrade to the latest incarnation of USB with Linux, and play with some nice, fast, new peripherals p50**

**A trusted computer is one that won't do what you want it to... p56**



## Meet Linux Format's team of writers...



**Richard Smedley**  
A refreshing infusion of dried weeds gives Rich2 the strength to sub almost a dozen emulation features.

**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.

**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.

**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!

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



**Charlie Stross**  
Master of Perl, Charlie has been writing about Linux for more years than anyone can remember.

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

## 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**

# Contents

# LINUX


FORMAT

LXF32 October 2002

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


## USB 2.0 and Linux

We round up the fast connection technology which is now supported in the Linux kernel, and test the latest USB 2.0 devices under Linux.



## What is Palladium?

Not the old variety theatre, but a troubling new intrusive technology which could prevent you doing many things on your own computer.



## Soar with FlightGear

Discover the joys of flying across Europe (and North America, with the DVD version) with this top Linux flight simulator, on our coverdisc.



**COVER FEATURE**

# BUILDING BETTER DATABASES

MySQL expert Jeremy Cole spills the secrets of efficient databases

**44**



Save money and subscribe to *Linux Format*. See page 96 or phone 0870 444 8645

## »» REVIEWS

**20 DEBIAN**  
Debian's 'testing' distro might be updated every day, but a new stable release is cause for celebration.

**24 DHCP TURBO**  
Does the fast, scalable and easy to set up DHCP server really cut the mustard? Read on to find out.

**26 BX PRO**  
A point-and-click GUI builder for the industry standard widget toolkit, *Motif*. Maybe there's no need to switch to *Qt* yet.




**28 GENTOO**  
The ultimate in configurability – a source-based distro that can give Linux From Scratch a run for its money.



**30 BOOKS**  
*JSP Examples and Best Practices; TCP/IP Network Admin and Web Database Applications with PHP and MySQL.*

## »» TUTORIALS

**68 LOW RESOURCE LINUX**  
How to optimise standard GNU/Linux distros for low memory and small disks.



**72 PERL**  
Our look at the newly-released Perl 5.8 – learn about the benefits and gotchas.

**74 JAVA**  
Time for a spot of housekeeping as we tidy up our build processes with *Ant*.

**76 KYLIX**  
In our penultimate *Kylix* tutorial we use what we've learned to build a Linux process viewer.

**82 CD WRITING**  
This month we get GUI as we round off our look at CD writers with *Gnome Toaster* and KDE's *KreateCD*.

**86 PHP**  
Get your website talking to your database as we introduce SQL, in the first of a two part look at *MySQL* and PHP.



## »» REGULARS

NEWSDESK .....	6
MAILSERVER .....	14
HOT PICKS .....	38
ANSWERS .....	90
BACK ISSUES .....	95
USER GROUPS .....	110

## »» FEATURES

<b>Roundup</b> .....	32
Top Firewalls on test.	
<b>HotPicks</b> .....	38
ElectricSheep; Fluxbox; GRUB; nano; ntop	
<b>Emulators</b> .....	62
The long-awaited Amiga feature.	
<b>User Groups</b> .....	110
Now with added advocacy!	

# Coverdiscs

A DVD or 2 CDs packed full of the latest Linux goodies **98**

**FlightGear** runs from the CD!; **UAE** Amiga emulator; **Clam Antivirus** as used by *LXF*; **Aglaophone** real-time audio processing; **Perl 5.8** latest version; **Quanta Plus** HTML editor; **On the DVD: Knoppix** install-free, Debian-based demo distro; **Ogg Vorbis 1.0** patent-free, quality audio compression; **Drip** store your DVDs on your hard drive; **Transcode** video stream processing; **Squeak** powerful Smalltalk-based language; and more...



Please read the coverdisc instructions starting on page 98 before installing from the coverdiscs!



# Newsdesk

Sun go Linux; IBM fight back; Lindows PCs cheaper than a night at the opera; nVIDIA open Cg compiler; RealNetworks' open source community; Gobe GPL'd; Cinellera premieres; Pocket cinema (with Zaurus).



**SUN RELEASE NEW SERVERS AND NEW LINUX DISTRO**

## Sun 'legitimise' Linux in the server space

**A**fter much speculation, Sun's Linux strategy is finally becoming clearer following recent product announcements and a keynote speech delivered by CEO Scott McNealy at the recent Linuxworld conference in San Diego.

The first announcement was a new range of Sun servers aimed at the edge-sector volume markets. The Sun

LX50 server is a 1U rackmount box with a variety of processor options based around Pentium PIII processors, and will run Sun's own version of the Linux operating system, Sun Linux 5.0.

Further probing reveals that Sun Linux 5.0 is a reworking of Red Hat 7.2 (a good version to choose), with Sun optimisations for the hardware.

This means Sun are now actively supporting two operating systems –

the new Linux distribution and the venerable Solaris OS. Solaris has been one casualty of the steady growth of Linux over the last few years, as even normally conservative customers seek a cost effective server solution. This not only favours Linux, but also commodity Intel based hardware, as Solaris has only been available on SPARC processors for some time. But Sun execs insist that there is still a

space for Solaris in vertically scaling systems, and even the new LX50 server will include a Solaris option.

However, the other message coming out of the Sun camp is that they view the choice of OS as being increasingly redundant, with the introduction of the Sun One platform. The idea with Sun Open Net Environment is that applications/services/whatever designed to run on Sun ONE will work

## Lindows low low price

**Hot on the heels of Walmart's** \$299 Linux-powered machine, Lindows boss Michael Robertson has said his company has inked a deal with major hardware and retailing partners to bring a sub-\$200 Lindows based machine to market. "For less than the price of most

handheld devices you'll be able to browse the Internet, check email and run a variety software products," said Robertson.

Each system would, of course, be pre-installed with Lindows, and a trial membership of the company's online software library.

### LX50 at a glance

**Processor:** Single/Dual Pentium III 1.4GHz  
**Memory:** 256MB to 6GB  
**Ethernet:** Dual embedded 10/100 BaseT  
**Drive interface:** Dual channel Ultra 160 SCSI  
**Drive Bays:** Two SCSI bays, CD/floppy/HDD bay  
**Video:** ATI Rage XL  
**Ports:** 2x RS232 serial (one at front)  
 4x USB (two at front)

**Software:** Linux /Solaris OS  
 Linux distribution (Sun Linux 5) to include: Java2SE, Sun ONE ASP, Tomcat, MySQL, Apache, wu-ftp, sendmail, BIND. Further Sun applications to follow.



on any Sun ONE implementation. The Sun ONE vision has already been partially extended to Linux in the form of the Java 2 Platform and *Grid* engine software with Sun ONE *Studio* (a further development of *Forté*), and other components to follow.

The backing of Sun, particularly the addition of their own supported version of the OS, must certainly be seen as another strong boost to the credibility of Linux for 'real' computing tasks, should anyone still be in doubt. The fact that Sun will be providing technical support for their own version of Linux may be a key factor in winning more converts to the OS.

Sun's acquisition of Internet appliance firm Cobalt last year gave them an entry to the low-end server market, and while sales of Cobalt devices have apparently been sluggish, it gave Sun valuable experience of supporting Linux systems. While the Cobalt Qube may survive for a while, it seems likely that the RaQ will be phased out in favour of the Sun-badged LX50 server and its descendants. Sun seem to have big ambitions for the volume server market, which in the long term will mean an even larger Linux install base.

## NEWSBYTES

■ **IBM** have created an affordable high-performance solution for parallel and distributed rendering. The Deep View System contains an octet of Linux nodes and two network switches plus a Scalable Graphics Engine (SGE) connected to the main system by eight Gigabit ethernet links. Deep View's Linux cluster performs computations that produce 3D geometry, renders the geometry to produce 2D pixels, and then transfers the pixels to be displayed on a display or video wall. [www.research.ibm.com/deepview/index.html](http://www.research.ibm.com/deepview/index.html)

■ **Bruce Perens**, an employee of Hewlett Packard, was planning a live DMCA-baiting hack which would demonstrate how to make a region-specific DVD player able to play discs from any part of the world. However, after garnering a lot of publicity, Perens backed down from his demo. Later it emerged that HP were attempting to use the controversial law to crack down on SnoSoft, a 'loosely organised research collective' who publicised a security hole in the company's Tru64 OS on BugTraq. In a missive sent to SnoSoft, HP VP Kent Ferson said the collective could "be fined up to \$500,000 and imprisoned for up to five years" if it went ahead and published the specifics of a buffer overflow exploit which could give a cracker full administrator access to a Tru64 system. However, in the face of a 'significant' reaction from customers, HP later called off the lawyers.

■ Are **Dell** selling non-Windows PCs or not? The Register (and other sources) acquired an email sent to various Dell departments suggesting that Microsoft had 'told them' they could no longer offer computers without Windows. The email began ominously: "New Microsoft contract rules stipulate that we can no longer offer the 'No OS' option to our customers beyond September 1<sup>st</sup>" and went swiftly down hill, ending with: "FYI - this effects all of our competitors as well." No-one from the [definitely not monopolistic] Microsoft was available for comment. Dell will get round the problem by offering systems with FreeDOS as an option, though it won't be pre-installed.

■ David Wheeler has written a paper which takes a detailed look at a 'typical' GNU/Linux distribution and comes up with a few surprising results. The first is that Red Hat 7.1 (his test distro) contain some 30 million physical lines of source code, required 8,000 man years of development and, if created under conventional proprietary terms, would have cost somewhere in the region of \$1000 million to develop. Another interesting point Wheeler makes is that the three biggest contributions to his tested distro (Kernel, Mozilla, Xfree86) are not GNU software. [www.dwheeler.com/sloc/redhat71-v1/redhat71sloc.html](http://www.dwheeler.com/sloc/redhat71-v1/redhat71sloc.html)

### David Cartwright

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



## COMMENT Wireless security

“ More and more of us Linux types are going wireless when it comes to networking. My office LAN is partly wireless, for instance, mainly because it was far easier than mucking about drilling holes in walls while standing in the rain on a wobbly shed with power drill in hand, and I was impressed when I set up the kit (gizmos from NetGear's range, for those who are interested) because it asked me as part of the install process whether I wanted to turn encryption on, made me define an encryption key, and so on. But I wonder how many of us just click through that screen, thinking "Ah, it doesn't matter, it's only a private network"?”

In the past we could sit and smugly contemplate how well off we were in the Linux world, with secure shells, *IPChains* and all that jazz thrown into our distros by default. Now we must remember that Big Brother may well be watching us from three doors up the street – you'd be surprised how far your wireless LAN (WLAN) can reach even on a bad day.


I've always had snooping in the back of my mind as an issue with WLANs, but it was only when a friend thrust me his AirMagnet (the coolest new toy this year – check out [www.globaltech.co.uk/products/airmagnet.htm](http://www.globaltech.co.uk/products/airmagnet.htm)) at a trade show recently that I realised just how many people's networks I could scan – without them even knowing I was doing so. When you see a list of dozens of servers appearing on your list, with guest accounts and all sorts of scary stuff turned on, you really start to realise that we're in a scary new world, and that those of us using wireless for a living need to be more careful than ever. ”

## OPEN SOURCE WINS

# Boost for Open Source projects

**A**s a host of top companies – including RealNetworks, nVIDIA and Gobe – launch major new open projects, it's been a good few months for the open source community.

First up is RealNetworks and their Helix initiative which, the company says, will 'transform the digital media industry'. The project consists of three elements. The Helix Platform itself promises to deliver feature rich media and applications to any device regardless of its underlying operating system. Echoing previous open source projects such as *Mozilla* and *OpenOffice.org*, the second plank of Real's strategy for Helix is the creation of a 'community' providing a central point of contact for industry partners, institutions and individual developers to access and license the Helix source code to create their own encoder, server and client applications. The third part of the release, and perhaps the most significant, is the Helix Universal Server which is capable of streaming any media to any device capable of receiving it. "No longer it is necessary to maintain three different delivery infrastructures to reach the largest audience," a spokesperson said.



**GobeProductive is available in beta form for Windows (above) and in Pre-alpha for Linux.**

"When you have a Helix Server, you have the capability to deliver media to any individual with a media player." Real's figures also suggest that the Universal Server outperforms Microsoft's *Windows Media Server* (serving WMP native files) to the tune of 400%.

Helix is seen as another advance in the long-running battle with Microsoft who are making significant gains in the media serving sector while Real has seen sales fall by 36 per cent in the second quarter of 2002. Hot on the



heels of the Helix launch, Real announced a restructuring program which will see 11% of staff picking up their P45s. Linux evangelist Bruce Perens said the Helix project was welcome, but didn't go far enough. "While RealNetworks is making a significant contribution to Open Source, today's release does not include the 'crown jewels' – their 'codex', the encoding and decoding software for their proprietary RealAudio and RealVideo formats."

Real are also working with Xiph.org to integrate the Ogg Vorbis format and codec into Helix.

An announcement by FreeRadicalSoftware will see the available number of open source productivity suites increase by one. FreeRadical, a new company started by former Gobe employees, has acquired the rights to *GobeProductive*

– a formerly BeOS only product that has recently come to Windows and Linux – and plans to relicense it under the GPL. *GobeProductive* has received praise for its document-centric interface and svelte size so should be one to keep an eye on.

Cg, nVIDIA's new programming language for 3D development, will get an open source compiler. The source code, available from [www.cgshaders.org](http://www.cgshaders.org), contains the parser that reads the language and creates intermediate code for compilation and a generic back-end. Dan Vivoli said the positive response from developers to Cg encouraged them to open up the compiler code. "We're open sourcing this compiler code to further accelerate the transition to an era of advanced real-time effects through Cg."


John Carmack, *Quake* guru, said the decision by nVIDIA to opt for a very liberal license, should prevent some of the 'petty licensing disputes' that have marred previous attempts at industry-wide cooperation. "This is a positive step which I hope other vendors will follow," he said.

**Helix – rich media, regardless of underlying operating system.**

## NEWSBYTES

■ The UK government has published its open source initiative white paper which sets out its commitment to open standards and suggests that government will 'explore further the possibilities of using OSS as the default exploitation route for Government funded R&D software.'

[www.e-envoy.gov.uk/news/newsonsite/oss-policy.htm](http://www.e-envoy.gov.uk/news/newsonsite/oss-policy.htm)



Run with the Dog ... pre-installed!  
[Yellow Dog Linux](#) | [The Terra Soft Store](#)

■ In an attempt to boost hardware revenue, Apple reseller **Terra Soft** are to offer Macintosh hardware with YellowDog Linux preinstalled. Terra Soft will position the systems at high-end users building clusters around Apple's new Xserve product who'd rather have Linux than OS-X.

■ You may have seen the joke about Luxury Linux Airlines ("You had to do what with the seat?"), but **Chris Stevens** has taken this fable to heart and created Penguin Airlines; and Penguin's relationship to Linux is more than just having a Tux-a-like as a mascot.

■ **UnitedLinux** said it would ship the first beta of its newly minted distribution at the end of August, and would be on track for the full range launch, in accordance with its initial roadmap, by the end of the year.


■ **The Free Standards Group** which oversees management and certification of the Linux Standard Base (LSB) has certified four Linux distributions as LSB compliant. Caldera, Red Hat, Mandrake and SuSE are the first to achieve the LSB mark which was launched at New York's LinuxWorld Expo back in January.

■ Those impressed with **PHP Accelerator** on last month's coverdisc may be interested to know that it has been updated to version 1.3.2, and is very competitive with *APC* and *Zend Accelerator*. Look out for future releases on our coverdisc. More details at:  
[www.php-accelerator.co.uk](http://www.php-accelerator.co.uk)

■ **optimoz**, a project to bring gesture recognition to the Mozilla browser, has started on phase II of its work: pie menus. Pie menus "leverage muscle memory for enhanced speed and accuracy in menu selections." Debate is raging over whether XUL or DHTML will give the best results.  
<http://optimoz.mozdev.org/r>

## FILM EDITING

# From the ashes of B2000... Cinellera



**Cinellera brings low-cost real time compositing and effects to Linux.**

**C**inellera, Heroine Virtual Ltd's replacement for the much-loved *Broadcast 2000* film editing system, has finally hit its premiere official release, and on first looks, it seems very capable indeed. *Cinellera* is a significant development in this sector and brings high-quality editing facilities within the grasp of anyone with a digital camcorder and Linux box. "Users can now expand their render farms with a minimal bottleneck due to software development costs. At the same time, users are not restricted by the usage license and corporate agenda of their

operating system vendor. *Cinellera* gives users complete ownership of the means to production," the company said. *Cinellera*'s features include realtime effects generation, 16-bit compositing and SMP and render farm optimisations. It is capable of capturing via IEEE1394, *Video4Linux*, Motion JPEG, and screenshots.

Defiantly professional, *Cinellera* demands a Dual 1.6Ghz Athlon, 512MB RAM for standard definition (1 GB RAM for high definition), 200 GB storage and a Gigabit ethernet adaptor. It is available from <http://heroinewarrior.com>

## Movies on the move

### Two reel Zaurus

It's been promised for a while, but finally you can watch *The Matrix* (or whatever) on the bus. You just need a Sharp Zaurus and *tkcVideo* from TheKompany. This application is capable of decoding MPEG, DivX, Realvideo, AVI and MJPEG videos on the move, while encoding options include MP1, 2 and 3 making it a useful audio player as well. Using MPEG4 (aka DivX 4/5) theKompany says you can squeeze one hour of 320 x 240 (Zaurus screen resolution) onto a 128MB CompactFlash card.



**You can watch movies on your Zaurus full screen or in a Window.**

## 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

# Finding the lice in Linux

“ My young daughter arrived home from school one day and informed us that there had been an outbreak of head lice in her class. Horrified, upon checking her head we found the tiny lice eggs, known as 'nits'. Treatment was easy: wash our hair in a medicated shampoo to kill the lice, but the almost microscopic nits were 'glued' to the hair shafts and had to be manually removed, or 'picked', with the use of a fine-toothed metal comb. It was both traumatic and tedious, but the end result was worth it. That day, we learned what 'nitpicking' really was.

Recent discussions on the Red Hat beta mailing list and elsewhere have fallen into the category of such nitpicking, especially with regards to the names used in the menus, the icons associated with them, the shapes of dialog boxes, the nuances of colour and selection of background images. Interestingly, the default desktop image (for the beta2 version) is a dragonfly, an insect unfamiliar to many Europeans who were repulsed by the image; it will not be in the final version.

The changes of great import (deprecating 'classic' packages, changing the compiler and so on) are discussed in depth, but without the emotion of the nitpicking that surrounds the details. This release is an important step for Red Hat because this attention to desktop details signals Red Hat's response to desktop-focused distros.

It also means that Linux is moving closer to being a desktop contender in the mainstream marketplace. When you have overcome the major hurdles, only the nitpicking remains and once the lice are gone from Linux, what a delight it will be. ”

## DANCING TO AN OSS BEAT

# Dance music ReBorn. EU's multimedia Linux

The Roland TB303 bassline synth and 808 drum machines have provided the foundation – or inspiration – for about 80% of dance music produced over the past two decades. Thanks to the average PC's processing power, and some clever coding, *ReBorn* has brought both these machines (via The Propellerheads' seminal softsynth *ReBirth*) to a Linux box near you... well nearly. Mere days after announcing the app's release, developer David Singer was contacted by Propellerheads and ordered to pull the product as the interface infringes on their software rights. Singer wrote that his work on *ReBorn* wouldn't go to waste as he was reworking the app and integrating the audio core into a

brand new softsynth 'with an all-new interface and 100% incompatibility with all Propellerhead products.'

Singer also asked those who'd already downloaded the application to respect the wishes of the ReBirth developers and not distribute *ReBorn*.

Also on the music front is a new dev effort, which is ongoing over the next two years and is supported by the EU, to create a Red Hat-based distribution specifically for the production of music. ReHMuDi is supported by a number of acoustic research bodies including Institut de Recherche et de Coordination Acoustique/musique (IRCAM) in Paris, Kungl Tekniska Högskolan (KTH) in Stockholm, Universita Pompeu Fabra in Barcelona and also FSF-Europe.



## Agnula hosts the Red Hat and Debian multimedia development efforts.

Franz Meyer, Red Hat's Southern European Director, said the intention was to develop a release "specifically designed for professionals in the music industry, Red Hat wants to enable [musicians] to free themselves from

technological and cultural constraints." The ultimate aim though, is to "expand the global nature of music even further and to extend the concept of Open Source Software to Open Source Music." [www.agnula.org](http://www.agnula.org)

## Linux Web Watch/



Linux Networks – x86 cluster.



Netraverse – Win4Lin 4.0.



Sun – new commitment to Linux.



Shaolin – remote desktop tools.

## And the winners are...

Starting up *Emacs* to write their acceptance speeches...

August's LinuxWorld conference ([www.linuxworld.com](http://www.linuxworld.com)) demonstrated the ascendancy of Linux in the enterprise sector as Linux luminaries, hackers and, curiously, Microsoft touched down in San Francisco to ply their trade.

Part of the excitement of the bi-annual exposition is all the awards that get dished out at these things, and one of the more interesting gong-shows was the Enterprise Evolution

Awards. The first of five winners were Linux NetworkX ([www.linuxnetworx.com](http://www.linuxnetworx.com)) and their Eclipse Database Cluster which combines hardware and software cluster management with off-the-shelf x86 hardware.

Netraverse ([www.netraverse.com](http://www.netraverse.com)) was praised for the stability and compatibility of *Win4Lin 4.0*. LinuxWorld.com editor Mark Cappell had no hesitation recommending it

to enterprises "that want to move to Linux and have custom-built Windows applications."

*OpenOffice.org* was third in line, thanks to its price, capabilities and efforts to offer robust support for the ubiquitous Microsoft *Office* file formats.

Sun's ([www.sun.com](http://www.sun.com)) high profile launch of LX50 secured fourth place in the awards. The LX50 is a small-form server built around a pair of 1.4GHz Pentium III processors, Linux, Solaris 9

and the Sun ONE *Application Server*, among other soft goodies. Although the judges were underwhelmed by the hardware spec, LX50 was awarded as a demonstration of one of the world's biggest computer companies commitment to Linux.

Last of all *Aptus 2.0* from Shaolin Microsystems ([www.shaolinmicro.com](http://www.shaolinmicro.com)) which provides a raft of tools for installing and maintaining Linux desktops remotely across a network.



IBM add a new model to their burgeoning xSeries range.

## BIG BLUE FIGHTS BACK

# IBM gunning for Sun

### IBM have responded to Sun's

Damascene conversion to Linux by sprucing up a couple of its Intel-based servers, launching a new eServer Cluster line and publicising a list of the top ten customers – including Air New Zealand and Satellite Records, the biggest dance music retailer in the US

– dumping Sun for Big Blue hardware and software. IBM are also making much of their 'long term' experience and commitment to Linux and team of experts well versed in Linux and open source deployment.

The new eServer, x335, is a 1U box based around Intel's Xeon processors

and featuring support for Ultra320 hard drives, Gigabit ethernet and 64-bit/100Mhz PCI-X slots. The x335 is designed specifically for serving up web pages. Further up the line is the 1350, an integrated cluster complete with management software, storage and network infrastructure.

## WIN ON LIN

# Serving up Office for all

### Building on the success of

*CrossOver Office* and *CrossOver Plugin*, Codeweavers used San Francisco's LinuxWorld Expo and Conference to announce the server edition of their *Office* app which is a direct competitor to similar solutions from Citrix and Microsoft's own *Terminal Server* system. Jeremy White, Codeweavers' president and founder, said that the *Server Edition* was a logical extension of their current range and would extend the product's core marketplace into areas such as engineering. "These organisations have a strong need for Windows software running over a thin client. With *Server Edition*, these workgroups will have a powerful incentive to standardise their desktops on a single Unix environment, rather than having both a Unix and a



### *CrossOver Office Server Edition* offers cost big saving per user.

Windows system on each user's desk." This, he said, would offer significant cost savings, especially with the advent of Microsoft's licensing 6.0 program.

Codeweavers have also announced version 1.2 of *CrossOver Office* which


now includes support for *Quicken* and *Microsoft Visio*.

Last minute bugs apparently thwarted a rumoured beta appearance at the show of a *CrossOver* release bringing *Photoshop* to Linux.

## PROCESSOR-OPTIMISED LINUX

# AMD unveil Hammer-based CPUs

Red Hat and AMD (who reportedly had the most impressive stand at LinuxWorld Expo) used the event to announce their intention to collaborate on bringing Red Hat Linux Advanced Server to forthcoming Opteron and Athlon processors based on AMD's next generation Hammer technology. While retaining support for the current range of 32-bit processor, RH will provide native optimisations for the x86-64 architecture with future Linux products.



AMD's next generation Opteron (right) and Athlon (left) processors get native support from Red Hat.

AMD's Richard Heye said the combination would bring "enterprise-class servers and workstations" within the price range of more users. "A range of customers with data-intensive applications are ready for the performance of 64-bit computing," he said. "And the AMD-Red Hat combination is intended to provide a mainstream solution as new apps become available across consumer and enterprise customer segments."

## Jono Bacon

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



## COMMENT

# Migration train

“ Here at my humble Comment on these lovely pages of *LXF*, I try to remain fully buzzword compliant as often as I can, and share my thoughts on these terms and technologies. Admittedly, buzzwords are to be used with caution and are not for impressing your mother-in-law, but this month I am going to focus on a stonker of a buzzword – migration.

OK, I hear the groans at the back and mumbles that migration is not nearly as buzzword compliant as the daddy (XML), but migration has an important message – with care and feeding, we can use this message to help promote Linux.

Migration is all about moving from one thing to another, be it web servers, databases or, indeed, OSs. Many IT managers across the land are put in a position of making decisions on which technologies to use, their costs, and how much coffee is needed implement them. Many of these decision makers are drawn to considering Linux as an option for their systems.

The issue I find of interest here is the politics involved in these decisions. I know from experience that cost is not always a prime issue, and often stability is seen as merely a benefit, not a requirement. It is clear that there is also a level of snugness in choosing a popular commercial OS from Redmond, due to it's reputed ease of use.

Often the thickness of red tape and the complacency of vendor reliance are detrimental factors in choosing the solution. The problem here is that in many cases the best solution is not used, and instead an expensive comfort blanket is re-implemented. The challenge for our community is set...

SPECIAL REPORT

Friday's keynote was given by Elaine Coleman of Sun Microsystems.

# Community under siege

Sam Williams reports from San Diego on the O'Reilly Open Source Conference.

**H**ard to feel gloomy in a city as sunny as San Diego, but attendees of the fifth annual O'Reilly Open Source Conference, (OSCON) certainly made a good show of it last month. "I'd be willing to pay money to anybody who can show me where to get a good beer in this town," lamented Ewan Birney, team leader of Genomic Annotation at the European Bioinformatics Institute and an OSCON keynote speaker.

Birney wasn't the only OSCON person sounding a concerned note. What with two aircraft carriers parked in the nearby harbour, a swooning stock market and Microsoft reps mingling happily among the assembled Perl and Python developers, it was hard to escape the feeling of a community under siege.

"If you can't fight for your freedom, you don't deserve it," intoned cyberlaw activist Lawrence Lessig, alluding to the spate of Free Software-unfriendly

legislation currently idling in the 'to be approved' queue of the US Congress. "So here's my question to you," added Lessig. "What have you done?"

Lessig's frustration, reinforced by an equally spirited keynote speech from Free Software Foundation president Richard M. Stallman, confirmed the unofficial theme of this year's OSCON: survival. Five years after launching OSCON amid a wave of post-*Mozilla* optimism (and soaring book demand), event host Tim O'Reilly was doing

everything in his power to shore up the spirits of underemployed Perl, Python and XML developers.

"Open source may reduce the revenues of some software vendors," he wrote in a pre-OSCON editorial titled, ominously, *The Strange Case of the Disappearing Open Source Vendors*. But that "does not mean that it reduces economic activity or economic success."

To prove his point, O'Reilly invited Rob Glazer, chief executive officer of RealNetworks, the Seattle-based

## What happens next

### Clocking up the airmiles

Once you got past the politics, the most interesting elements at OSCON were, as usual, the emerging technologies. Conference organisers devoted two keynote speeches and a number of tutorials to bioinformatics, a field University of California, Santa Cruz researcher Jim Kent glibly describes as "system administration for biotechnologists."

European Bioinformatics Institute scientist Ewan Birney, another keynote speaker, said bioinformatics and open

source development go hand in hand and was pleased to dangle job offers. "I need four people on my team," Birney said. "What's more, I can promise you better beer in the UK." Look for the momentum to keep building until February. That's when O'Reilly hosts its second annual Bioinformatics Conference. <http://conferences.oreillynet.com/bio2003/>

In the meantime, the first annual O'Reilly Mac OSX conference will be held September 30 to October 2 of this

year in Santa Clara, California. To build anticipation for the event, Apple Computers donated 50 of its sleekest 'desklamp' iMacs to the OSCON Connectivity Room. Although some griped about the default Internet Explorer icon on the desktop, most users seemed impressed by the OS X interface and the underlying BSD functionality. Former FreeBSD project leader and current Apple employee Jordan Hubbard was on hand showing how to port applications to OS X, while *PostgreSQL*

developer revealed that the Mac OS X port of his team's open source database was already the most requested version. **MacOS X Conference**

<http://conferences.oreillynet.com/macosx2002/>

**Joshua Drake Interview**

<http://linux.oreillynet.com/pub/a/linux/2002/07/16/drake.html>

For more technical information about this year's Open Source Conference, visit O'Reilly editor Andy Oram's weblog at <http://www.oreillynet.com/pub/au/36>



Vendors showing off their wares on the expo floor.



Games stations were set up at the local bar for a *Quake 3* LAN Party.



Jason Hunter, Johan Vromans, David Blank-Edelman and Tim O'Reilly.

streaming media company which in late July released portions of its A/V client under an *OpenOffice*-style dual license setup. Speaking under the mammoth lunch tent and over the din of a nearby airport, Glazer said RealNetworks had no choice but to turn to the open source community after a withering five-year standards battle with Microsoft Corp. "I'm a firm believer in [Sun chief scientist] Bill Joy's law," said Glazer. "I believe that the majority of smart people in the world do not work for my company."

Such comments underlined the peculiar nature of OSCON. Launched in 1998 as an outgrowth of the O'Reilly Perl Conference, OSCON has evolved into a pitched battle for the open

source community's soul. Since inviting Bill Joy to defend the Sun Community Source License in a 1999 keynote, O'Reilly and Associates, a company that sells books on Java and .Net in addition to its vaunted line up of open source software manuals, has shown a willingness to court outsiders eager to challenge or exploit the ideas of the open source community. This year's OSCON continued the trend, reflecting the internal angst of a community struggling to decide its rôle in the world.

"Everybody realises we can write the software," said Toshiba engineer Chad Carr, after the Lessig keynote. "The thing we have trouble with is evangelising the freedom. That's what we really need to work on now." **LXF**



Lord of the Rings – Milton Nganof, Weta Digital's visual effects dept.



Business and developers – Paul Pangaro, Sun Developer Network.

## Stallman sets the agenda

### Putting freedom first

It was a busy week for Richard M. Stallman, founder and president of the Free Software Foundation, chief architect of the GNU project and chief evangelist of the free software movement, at this year's OSCON. Conspicuously absent from past OSCONs, Stallman arrived at this year's event as an invited keynote speaker. No more than 10 seconds into his keynote speech, Stallman offered the reason for his belated arrival.

"Unlike some of you," Stallman said. "I am not an open source developer. I have not been associated with the open source movement. I am associated with the Free Software movement."

True to form, The GNU project leader chided developers and reporters who use the term 'Linux' instead of 'GNU/Linux'. "It's only fair that the people who made the system should receive their fair share of the credit," he said, and performed the usual on-stage antics. Before the final question and answer session, he donned a robe and mock-halo, addressing the audience via his 'St Ignucius' alter-ego. "There is no system but GNU, and Linux is one of its kernels," said Stallman, drawing laughter.

Between the antics and the rhetoric, however, Stallman offered some solid insight. Open source may have won the battle, winning over corporate benefactors such as IBM and Hewlett Packard, but it was losing the war. Stallman said the traditional geek desire to hack the system apolitically ignored the fact that the system, at least in the US, relies on legal reinforcement.

"You can leave politics alone, but politics won't leave you alone," Stallman

advised. "You've got to join groups. You've got to send tens of thousands of postcards to your elected representatives."


Such comments, coming two days before open source advocate Bruce Perens was forced to back down on an earlier stated promise to run a DVD player onstage in violation of the US Digital Millennium Copyright Act, proved prophetic. In explaining his reason for the reversal, Perens said his employer, Hewlett Packard, feared reprisal.

For the rest of the conference, Stallman delivered his observations from the back rows of the audience. The strategy had mixed results. While watching an interview of Rob Glazer, during which the RealNetworks CEO discussed the company's Java-inspired RealNetworks Community Source Licence, Stallman shouted out "Don't use it!" only to be shouted down by surrounding hackers. Two days later, however, during a panel discussion of open source software in government, Stallman dominated the microphone, chiding panelists who seemed equivocal in their advocacy of software freedom.

"Governments provide services, but they also have a rôle to lead," interrupted Stallman at one point. "The same goes for government employees."

Stallman's interjections earned the respect of Peruvian congressman Edgar Villanueva, invited panelist and author of a proposed law making Free Software usage mandatory in the Peruvian government.

"I have long admired Mr. Stallman and his work," said Villanueva, speaking through interpreter, and GNOME Project leader, Miguel de Icaza. "Tell him that there are many people in my country who admire him, too!"



Stallman – "but politics won't leave you alone."



# Mailservver

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

## ★ Letter of the month

This month's winner receives a copy of **Programming PHP** published by O'Reilly



### Business desktops

I am the IT Manager for a toiletries company, and part of my brief is to implement a new ERP system on our hybrid Linux/Win2K network.

In a recent meeting with one such ERP vendor I asked the Sales Rep what relational database his product used – to which he replied *SQL Server*. When I mentioned that I favour Linux based systems he replied that “Linux is OK for home use, but it’s not used for anything major!” Plain ignorance? I suspect not.

This particular Sales Rep then tried to get the contact details of the Directors who actually sign the cheques in order to talk to them directly. In my experience, the majority of Directors are not interested in technical details,

they just want to know what it does and how much it costs.

Many smaller businesses do not have a full time IT representative to make board-level recommendations. How many of these businesses have ended up implementing overpriced IT solutions, based on a limited understanding of what is available, and the “recommendations” of biased commission-hungry sales people? Also, to what extent is this restricting the spread of open source software in the SME sector, and locking companies into expensive licensing scenarios which can ill-afford?

The same applies to outsourced IT support who recommend Windows-based solutions because it is what they are familiar with and able to implement, even

though a Linux-based solution can usually do a better job and is considerably cheaper.

I don’t personally believe that Linux is going to make any significant moves onto business desktops any time soon, but will continue apace on the server side. We need to make the non-technical but cost-aware decision-makers aware that there are alternatives to Microsoft out there, and they don’t take a degree in computer science to implement and maintain.

Kind regards

David Jones, IT Manager,  
*Medical Express (UK) Ltd*

Well, you can’t really expect a sales rep to give you an accurate depiction of a ‘product’ that costs nothing! Though his comments may sound ludicrous to those in the know, you do have to wonder how many people are

taken in by such patter.

In many ways only either very big organisations, who have the resources and skills in house, or very small businesses who are lucky enough to employ a single knowledgeable IT person, are really currently able to implement a company wide Linux deployment. However, I do think that there are many advantages for IT directors considering Linux – the reliability and security issues, as well as being effectively able to lock down systems are just as important as the cost savings. Hopefully we will see some big announcements soon and attitudes will change.

In the meantime console yourself with this month’s star prize is O’Reilly’s *Programming PHP*, by Rasmus Lerdorf and Kevin Tatroe.

### Java prescience

I am getting very disturbed about your Java tutorial. I buy your magazine every month, and enjoy it greatly. However on three occasions now I have found that the Java tutorial is exactly what I needed about one week before I bought the magazine.

Are you somehow tracking my Java activities and basing your tutorials on what I am doing? If not.. why do the tutorials always seem to come out just *after* I needed them? Is this some sort of perverse joke?

Many thanks for a great mag!!!  
Stephen, *via email*  
Yes, sorry about that. Mr Drummond has been spying on you for months

now. You know those funny noises in the night? That was him. And that smoke alarm really does have a hidden camera in it. Check the buttons on your favourite shirt.

Unfortunately, due to the delays and lead times involved in printing the magazine and distributing it, the Java articles are arriving just after you needed them, but we’re working on a way we can beam them directly to your brain, or perhaps whisper them in your ear while you are asleep.

### Why Windows ?

I’m a newbie to Linux, and have read your magazine on a regular basis for the past year now, and find it really helpful and informative. However, I sometimes wonder whether I am reading a Linux or a Windows mag!

After a few months of playing with Linux, I decided to take the plunge and remove Microsoft Windows from my system and see whether I could do all I needed to without it. So far, I have managed. I can access my mail and the Internet, my printer works, I can create CDs and watch DVDs, and I can read and create documents, even those created in MS Word format. I have not been able to get my scanner to work, and still don’t know how to capture video from my DV camcorder via the Firewire port, but I’ll persevere and will, hopefully, resolve these issues. I have to say, in response to the ever-present question: “Is Linux ready for the desktop?”, it certainly has been ready for me.

So why do I sometimes wonder whether your magazine is a Linux or Windows magazine ? Well, I regularly read in your magazine “boot into Windows and...” (for example, “Reiser Resize” in LXF30 in the Answers section).

Also, in LXF30, one of your readers explains how to create a CD from an ISO image, using *Nero*, under Windows ! Surely, as a Linux magazine, it would be more beneficial to your readers to explain how to do these things under Linux? I’ve created CD’s from your ISO images under Linux and I am a novice user. So please, assume we have no Windows OS (like me), keep the magazine Linux-only, and help us do what we need to do in Linux.



Otherwise, the magazine is great, the coverdisc is great, and your syntax is improving all the time!

Wayne Jones, via email

The mystery of the CD burning instructions for Windows is simply that many many people asked for them. Picking up the magazine to try Linux for the first time, they lack a working Linux system and want to use the distros on the DVD. We do give brief instructions for burning CDs under Linux in the coverdisc pages when appropriate, and there is a CD writing tutorial running at the moment which covers all sorts of mastering questions.

Running Windows can be useful for people installing Linux, as it's possible to identify hardware components under an OS that is working, and therefore work around any potential incompatibilities.

Some people still use Linux on a dual boot system and operations involving resizing partitions and the like can effect both operating systems – we don't ignore the presence of their other OS, which was Windows in this particular problem.

But thanks for your letter, I hope we've cleared up some of your

**“I am getting disturbed. On three occasions now I've found that the Java tutorial is exactly what I'd needed one week before I bought the magazine”**

worries. It's good to hear your happy with your desktop experience overall. Check out issue 28 for more on video editing and our review of *Mainactor* last issue.

## Big jobs

In the article on quantum computing in *LXF 30*, the problem of factoring large numbers is given as an example of something such a monumentally powerful machine could be used for. The problem of finding the prime factors of 1,658,805,411,577 is described as 'quite a long job'... factoring 267,023 would, we are told, take an average desktop PC 'a little time'.

Well, I just had to!

1,658,805,411,577 is reduced to prime factors 1,278,787 and



**Primed to break the quantum barrier and factorise that large number.**

1,297,171 in 0.9 seconds on a K6/2 400 running *Maple 7*. That's not even an average desktop these days. Although, given the performance of modern machines, I suppose 0.9 seconds is a veritable eternity :-)

Anthony Winters, via email

Sorry, we have discontinued the smart-arse of the month award. I think the numbers involved were mere examples, though perhaps it was misleading to suggest this particular

Are they so frightened of the competition people like Mandrake, SuSE and the like give them, that they need to secure their place at the top by any means necessary.

Are they trying to achieve a stranglehold on the alternative OS market, do they see themselves as the next Microsoft. Do they not see without the open source movement and the FSF they would not exist. How dare they turn their backs on the people who put them where they are today. Surely there is a better way to achieve this without taking such drastic measures.

I for one would cease to write/update any code for open source if this happens, I read somewhere, and I think it is Mandrake that do this, if you write a piece of code that they find useful then they take it from you, give you a special license that enables you to do what ever you like with that piece of code. Whilst everyone else has to adhere to the GNU licensing agreement. Not a perfect solution by a long shot but more workable than what Red Hat is proposing.

Adrian Stokes,

Redditch



**Red Hat – delving into patents for Free Software.**

I believe the Red Hat position on this is that they felt, seeing as software patents had already been legitimised by law, that they needed to protect various “intellectual properties” for fear of someone else obtaining a similar patent and hitting them for cash (yes, prior art could be claimed, but it seems incredible the number of patents that are granted that clearly do not satisfy this condition).

I'm not aware of any restrictive licensing practices by Mandrake, in fact they seem to be the most liberal of the main distro makers as they positively encourage free/low-cost distribution of Mandrake Linux.

## Abandonware

Thanks to Hoyt Duff, for his comment on the *WordPerfect* dilemma. I couldn't help think this situation and others like it could be easily resolved if the owners of no longer supported software were to release the code to the open source community. Do they really have any reason not to? If Corel don't intend to release an updated, replacement, or potential competitor to *OpenWordPerfect* they will definitely not suffer financially. This would in fact demonstrate to Corel customer base, using their current products that they may even care about them. Marketing & PR machines within all multinational corporations do their best to convince us they care. The consumer usually finds out however, this is definitely not the case. It always comes down to the bottom line. Not only are they unwilling to consider anything with a negative effect on this bottom line, but also anything with zero (or no positive) affect.

The company where I work relies heavily on one product that is no longer supported and another for which the company that owns it no longer exists. Not only can't we find updates, but there are no replacements either.

This is stopping us from upgrading the platforms that these legacy apps are running on. If the code for these applications were to be made open, these apps could be kept alive by someone else. It's not just a matter of money either, as improvements to these systems would allow us to give a better quality of service to our



# Helpdex

BY SHANE COLLINGE

shane\_collinge@yahoo.com



« clients. If they were open we could contribute some of our resources to testing & coding, or even keep it alive ourselves.

Could we have a model, something like that in the world of literature and many other artistic fields, where by ownership & rights to royalties expire after a given amount of time. Say one year after a software product is no longer

the commercial product? Wouldn't that make it easier for competitors to see how to beat *WordPerfect*? And so on. Of course, there may be some valid arguments here, but you can bet that few companies will see the benefits of open sourcing. Usually by the time a product is really dead enough that they don't care about it anymore, the original code can no longer be found.

## “How many of these businesses have ended up implementing overpriced IT solutions based on a limited understanding of what is available?”

supported and that company does not release an update, replacement or potential competitor. The code could then be released to the open source community.

Edd, Australia

Here you have hit on one of the major benefits of open source software – it never becomes abandonware while people are still using it. This is a very significant advantage that few people seem to realise – until they have the sort of problems you do.

I would imagine that Corel would perceive of many problems with Open Sourcing *WordPerfect*, particularly when the *W WordPerfect* brand lives on in the *WordPerfect Office 2002* product for Windows 2000/XP. The first of these would be the loss of intellectual property, once upon a time they paid good money for *WordPerfect*, so why would they want to give it away?

There's also the possibility that some of the code from the abandoned version is still in use in *WPO*. Would that cause problems with

### New convert

I just installed Mandrake 8.1 on my HP e-PC. I about went into shock, and then headed for a slow burn. I am an experienced network techie in Windows and Novell.

A computer that was running XP with 256MB of RAM, and a 1.8GHz processor, took 5 minutes to boot. It took me 45 minutes to just *install* XP, let alone touch it up with drivers and updates.

So explain to me why it only took 15 minutes to install Mandrake, and NO MINUTES to update it? Was this too easy or what? I love it so much, I went to the store looking for info and found your magazine. The store clerk slightly mocked me for spending that much on a magazine; I replied that he needed to join the 21<sup>st</sup> century and get a real operating system. It didn't take long for me to join the club. I have been reading your magazine slowly and going to all the websites.

I am looking for some serious, and not so serious games for Linux. I would eventually like to take my mother and friends over to the other side, but it would break my Mom's heart if she had to give up *Ms PacMan*, *Frogger*, and *Quake*. (I have to admit that I am still running Windows on two other computers for web building and gaming.) I read that Loki is going belly-up in your magazine, so do you have any other suggestions?

Ri del Rio, USA

I'm glad you've had a positive experience with Linux so far. On the games side, Loki have closed down, but there are still some excellent games to be had on Linux. *Quake* for example, already runs on Linux – id released binaries for it. So does *Return to Castle Wolfenstein*.

There are also plenty of great open source games – check out the freshmeat games category ([www.freshmeat.net](http://www.freshmeat.net)) or our CDs.

There is also the *WineX* option – a lot of modern windows games can be run on Linux under emulation – take a look at

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

### Consumers...


First off I'd like to say what a great publication you have! Beats any Linux magazine I've seen so far; my only disappointment is that the bookstore I get it from sometimes receives it 2 months late.

I couldn't help responding to your choice for "Letter of the Month" for *LXF28*, June 2002. In regards to Mr. White's comment, "The consumer doesn't want to how things work, they do not want to know what to do when something goes wrong...", etc, it's been my experience as an IT worker the last 4 years that this viewpoint couldn't be further from the truth! Nearly every PC user I've supported in my career would like to "know more". They don't like feeling helpless when an application crashes or the connection to a network share is persistently lost.

The last thing they want to do, especially with their home PCs, is to contact tech support and either be put on hold for ever or made to feel like a dummy because they didn't know any of their basic system info. I'm always being asked



More games, like *Return To Castle Wolfenstein*, are making it to Linux.



what can be done next time to prevent some error or another, or where can they go to learn more. Of all the clients I've had, 100% of them eventually ask me the same question at some time or another and that is "Where can I find info on how I can troubleshoot or upgrade my system by myself?"

So Mr. White's consumer becomes an "educated consumer". This is a critical turning point. Just about everyone who seriously uses a PC in their daily life or work gets to this point. It was at this point in my PC usage I turned to Linux, and why I

invite my clients to at least try it.

Most of the big software firms would like you to believe that you are helpless without them and their systems, cloaking their deficiencies in FUD and a never ending trail of updates or patches. They use strong arm tactics to pay ever increasing licensing fees (BSA sound familiar?). Because to them it's not about knowledge or innovation, it's about the money. There's not a lot of cash flow to be had by a consumer who says "No thanks, I'll wait till you fix it, however long it takes". So they keep us in the dark.

I only hope that others take the plunge and give Open Source solutions a try. They may actually find practical and quality tools they can depend on.

**Robert Wade**, *Desktop Support Specialist, U.S. Air Force*

I expect a lot of it depends on the people involved, but generally I think you're right. Even my dear grey-haired old mother likes to know a bit about computers, mainly because if she has a problem there's usually nobody around to help her out.

Certainly, the motivations for developers of proprietary software are

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

not exactly convergent with those of the user base. A recent example is the latest 'update' to MacOSX, which Apple believe people are going to want to spend \$129 on. More curious when the major improvements seem to be some tinkering with the mail, web and search applications and all built on the very latest Unix technologies – FreeBSD4.4, which is, er, free. The development environment is based around GCC!

## The Positive Internet Company LTD



From E-Commerce superstores, corporate and small business sites, to personal homepages, our flexible web hosting will meet all your requirements. Our dedicated support team are on hand to help, whatever your depth of knowledge.

**we're good**

[good@positive-internet.com](mailto:good@positive-internet.com)

### Fast, Reliable, Powerful Web Hosting

- + State of the art, reliable, high performance Linux hosting
- + Massive 300 Mb/s backbone for lightning fast data transfer
- + Generous data transfer, excellent connectivity
- + A capacious 120 MB of space for maximum flexibility
- + Multiple POP3 mailboxes with unlimited email aliases
- + Full CGI support plus PHP4, Perl, C/C++, Python and more
- + Complete access to your account with FTP and Telnet
- + MySQL database with each account with full ODBC integration
- + Detailed access logs and analysis with graphical summaries
- + Microsoft Frontpage extensions at no extra charge
- + Icecast streaming, Apache::ASP, Java Servlets & JSP
- + Constantly monitored servers with generator backup and UPS
- + 24 hour unlimited high quality technical support
- + SSL option available for secure E-Commerce solutions
- + Try before you buy, no minimum contract

**£250**  
(ex VAT p.a.)



**Web Hosting**

**Web Design**

**E-Commerce Solutions**

**Dedicated Servers**

**0800 316 1006**

**[www.positive-internet.com](http://www.positive-internet.com)**

# 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...

### Debian >>

Debian might be updated every day, but a new stable release is a cause for celebration, especially for our reviews ed. Discover the benefits of the Freest of distros **p20**

### DHCP Turbo

Does the fast, scalable and easy to set up DHCP server really cut the mustard? Read on to find out **p24**

### BX Pro


A point-and-click GUI builder for the industry standard widget toolkit, *Motif*. Maybe there's no need to switch to *Qt* yet **p26**

### Gentoo >>

The ultimate in configurability — a source-based distro that can give Linux From Scratch a run for its money **p28**

### Books

*JSP Examples and Best Practices, TCP/IP Network Administration and Web Database Applications with PHP and MySQL* **p30**



## COMING UP SOON...

### Lycoris

The newbie friendly Windows clone brings Linux to desktops everywhere (thanks to Walmart)

### Books

More forests come crashing down as we bring you the best in informative Linux literature

### Xandros Linux

The successor to Corel Linux is nearing the end of its beta programme, so expect a full review here soon

### Homebase

The commercial *Mozilla*-based desktop environment

### ELX Linux

India is quickly emerging as a centre of IT excellence. Find out how well they do Linux

### Eden box

We put the miniature marvel of a motherboard in an equally tiny case, and slap Linux on the drive

FREE DISTRO

# Debian GNU/Linux 3.0

After two years of waiting, **Richard Drummond** at last gets to sing the praises of Debian's Woody.

Other general-purpose GPL'd distros include Red Hat 7.3, Mandrake 8.2 and Slackware 8.1.


- **DEVELOPERS** Debian
- **PRICE** free download
- **WEB** [www.debian.org](http://www.debian.org)

**D**ebian isn't just a Linux distribution, it's a way of life. For devotees that have been tracking the development of Debian 3.0 over the last 23 months – regularly downloading the new testing packages as they have become available – the official release of Woody, Debian 3.0, is perhaps no big deal: they've already been using many of its new features for some time. (Although the release of Woody does mean that the speed of development will pick up again as the testing tree moves on to the next Debian version, code-named Sarge). For the rest of the world, Debian 3.0 means, at last, there's a new stable version of Debian with up-to-date software.

It's easy to forget just how far behind Debian 2.2 had fallen. Debian 3.0 is the first Debian release with KDE 2, the first with *XFree86 4*, the first with *Mozilla*. But besides these important new software components and other much-needed version-bumps, Debian 3.0 also brings many new features. New architectures have been added – IA64, HP PA-RISC, MIPS (big and little endian) and S/390, bringing the total supported to eleven. The main distro now contains cryptographic software for this first time, due to the relaxing of US export laws. This includes *OpenSSH*, *GnuPG* and browsers with strong cryptography. Boot floppies have been streamlined, making installation simpler and more flexible, and the installer now offers localisation into French, German, Portuguese, Spanish and Japanese. And the system-wide config system, *debconf*, has been improved, now includes a GNOME front-end and sees more extensive use throughout.

## Boot floppies

Debian has a reputation for being difficult to install, which I find not



The package configuration system, *debconf*, now sports a GNOME front-end.

altogether justified. It is true that the Debian installer doesn't do any hand-holding – there's no hardware detection, no partition resizing, and no whiz-bang GUI – but the text-mode, menu-based installer is simple, flexible, and reliable.

The first time you install Debian, you need two components: the boot floppies and the packages that make up the base system. The boot floppies contain the bootable image, the kernel and kernel modules and the installer proper, and, depending on your architecture and the build you opt for, comes on a minimum of three floppy disks – or on a bootable CD. On x86 architectures, the various flavours of boot floppy differ only in the driver support they provide, so choose one appropriate for your hardware. The 2.2 kernel is still standard on most architectures, but experimental boot floppies are offered with a 2.4.18 kernel, including support for *Reiser* and *ext3* filesystems. Otherwise you are limited to just *ext2*.

The installation process is simple. Boot the installer, set-up your partitions, install the kernel and drivers (you can manually choose which kernel modules you wanted loaded), set up your network card, and do the base install. You can then boot into your installed system to perform any

post-install configuration and install what additional software you need. The install CD will at least contain the packages for the base install, so no problem there. Otherwise, you can specify a remote or local Debian mirror from which to download the base packages. This makes installation

```

dselect - Inspection of package states (avail., priority)
-----
#IDM Pri Section Package Inst.ver Avail.ver Description
-----
Up-to-date installed packages
-----
Up-to-date Required packages
-----
Up-to-date Required packages in section base
-----
*** Req base base-files 3.0.2 3.0.2 Debian base system miscellaneous fi
*** Req base base-passwd 3.4.1 3.4.1 Debian Base System Password/Group F
*** Req base bash 2.05a-11 2.05a-11 The GNU Bourne Again Shell
*** Req base bsdutils 2.11n-4 2.11n-4 Basic utilities from 4.4BSD-Lite,
*** Req base debianutils 1.16 1.16 Miscellaneous utilities specific to
*** Req base diff 2.7-29 2.7-29 File comparison utilities
*** Req base dpkg 1.9.21 1.9.21 Package maintenance system for Debi
*** Req base e2fsprogs 1.27-2 1.27-2 The EXT2 file system utilities and
*** Req base fileutils 4.1-10 4.1-10 GNU file management utilities
*** Req base findutils 4.1.7-2 4.1.7-2 utilities for finding files--find,
*** Req base grep 2.4.2-3 2.4.2-3 GNU grep, egrep and fgrep.
*** Req base gzip 1.3.2-3 1.3.2-3 The GNU compression utility.
-----
#IDM Pri Section Package Inst.ver Avail.ver Description
-----
dpkg - Package maintenance system for Debian
This package contains the programs which handle the installation and removal of
packages on your system.
The primary interface for the dpkg suite is the 'dselect' program; a more low-level
and less user-friendly interface is available in the form of the 'dpkg' command.
In order to unpack and build Debian source packages you will need to install the
developers' package 'dpkg-dev' as well as this one.
-----

```

*dselect* is still the most dependable way of selecting apps for installation.

much more flexible and means you no longer need to download the base images onto a dozen or more floppies to do a floppy-based install.

Compared to most mainstream distros, the Debian boot floppies are desperately primitive. At the moment if you want niceties such as hardware-detection, you can only get that post-install with the discover package.

Likewise, if, on x86, you want *GRUB* rather than default *Lilo* bootloader, you'll have to set this up post install.

When you boot into your new system, you'll be asked to set up passwords, create user accounts, and optionally set-up a dial-up connection. You can then either use the revamped task-selection system or manually select packages with *dselect*. The task system has been improved, and makes it easy to select groups of packages to install. For example, tasks offered include the X window system, Desktop (which includes both KDE and GNOME) and C and C++ development. I find taskset too coarsely-grained and unpolished, and prefer to use the much-maligned *dselect*.

Installing Debian isn't rocket science, but it can be long-winded. The base system really is basic, providing only what you need to get

online and download and install any other software you need. The advantage is that the boot images are small (you can fit everything on a business card-sized CD), and it is easy to create a minimal system for low-resource machines or specific tasks. The corollary is, however, if you want to create a useable desktop, there's a lot of stuff you need to select and install.

## Package power

Debian's package management system is second-to-none. This is partly due to the *dpkg* system itself, which is worlds more powerful and useful than *RPM*. With *dpkg*, for example, package developers can not only specify dependencies – other packages which must be installed for a package to work – but they can also recommend and suggest other packages that can enhance or extend a package. This makes the task of manually selecting packages much easier. With *dselect*, for instance, when you select an app, it will pop up a list of its recommended and suggested packages for you to consider – with the recommended packages selected for installation by default. Another benefit with Debian is that the Debian developers tend to break software distributions into small pieces, which means you get less excess baggage and makes installing the system you want quicker. If you want a *Samba* client, for example, you don't need to install the *Samba* server; if you just want *KMail*, you don't need any other non-essential parts of the KDE desktop.

The other weapon that Debian offers for package management is the oft-praised *apt* or *Advanced Package Tool*. This provides a flexible system for

## Getting Debian 3.0

Where to go to download or buy install discs.

You can find out where and how to download official Debian CDs at [www.debian.org/CD](http://www.debian.org/CD).

The official set contains (for the majority of architectures at least) seven binary discs and seven source discs. This page also tells you where you can get unofficial minimal CD images, which contain just the basic

installation system. For floppy based installs, you can visit your local Debian mirror (see [www.debian.org/mirrors](http://www.debian.org/mirrors)) and download the floppy images you require.

Debian Shop in the UK, do full boxed sets for £35, and binary-only CD sets or DVDs for £15 (see [www.debianshop.com](http://www.debianshop.com))

fetching and installing packages and any dependencies. *apt* can handle multiple sources of packages, including remote ftp or http servers, CDROM or from a local filesystem. *dselect* provides a menu-based front-end to *apt*, but it can also be used from the command line. (If you really can't stand *dselect*, Debian 3.0 includes several other *apt* front-ends such as the curses-based *aptitude* and Conectiva's *synaptic*.) A simple **apt-get install mozilla** will install *Mozilla* and any other packages it requires, for instance. New in Debian 3.0 is the **build-dep apt** command which will fetch and install any packages required to build a particular package from source.

As well as *dpkg* and *apt*, Debian also provides a wealth of other tools to make the life of the system administrator easier. Thus there's the *dhhelp* system, which builds and maintains a repository of all system documentation; there's *defoma* which handles the installation and configuration of fonts; there's the alternatives system which lets you install multiple functionally identical pieces of software and choose one as system default (for example, you can

install several Java runtimes, and choose the default to be launched via `/usr/bin/java`); there's the menu system, which builds an integrated desktop menu which can be used by any window manager or desktop environment; there's the kernel-package system for automates the job of compiling a kernel and creating a Debian package from it. The one tool that stands out, though, is *debconf*.

*Debconf* is used by packages to interact with the user during installation employing a consistent user interface and is much more widely-used by packages in Debian 3.0. Various *debconf* front-ends are offered, including text, *curses*, dialog and, new in 3.0, GNOME-based. The user can select the priority of question that *debconf* will display; opting for silent will do a non-interactive install, and packages will left at their default configurations. Packages typically use *debconf* for configuration or to display important notices during installation, and the user can manually invoke a packages *debconf* interface at a later date to perform reconfiguration with the **dpkg-reconfigure** command. The X windows system on Debian makes

## Information

Kernel 2.2.20 (2.4.18 optional)

Glibc 2.2.5

GCC 2.95

XFree86 4.1.0

KDE 2.2.2

GNOME 1.4.1.4

Mozilla 1.0

Apache 1.3.26


MySQL 3.23.49




The *dhhelp* system builds an integrated repository of system docs.



Debian 3.0 is the first release to include the KDE desktop.



Debian's *debconf*-based X configuration can now do hardware detection.



Debian now supports eleven architectures. This is *MOL* running on a Debian PPC box.

extensive use of *debconf* to provide dialogs to automate the installation and configuration of X; there's no longer any need to hack your X config with an editor. If you have the *discover* package installed, it can also detect your graphics card automatically; with *mdetect* it can detect your mouse, and with *read-edid* it can probe your monitor's capabilities.

All of these kinds of tools that Debian offers means that, while it might be more of a chore to install than other distros, it is certainly much easier to manage once installed.

### Software heaven

If all of the above hasn't convinced you to give Debian a try, then perhaps this will: Debian is big. The official Debian 3.0 distribution contains nearly 9,000 binary packages. All of this officially packaged software means that, typically, whenever you need to install some extra software, it's available via **apt-get**.

The range of software that Debian provides is staggering. With a few notable exceptions (the lack of *OpenOffice.org* being an obvious one),

a Debian build of all the significant open source projects is available. This includes networking and development tools, servers, desktop and multimedia apps, text processing and number-crunching software. There's a dozen window managers to choose from, you can opt for KDE 2.2, GNOME 1.4, *XFce*, or *WindowMaker* desktops; *Mozilla*, *Konqueror*, *Nautilus*, *Galeon*, *lynx*, *w3m* number among the browsers on offer; and you can select between *Apache*, *Roxen*, or *boa* as your web server and *Zend* or *Tomcat* for web applications.

Of course, there are limitations. Debian's strict open source policy is laudable, but it means that there is a lot of critical software that is not officially included in Debian and probably never will be. The main omission is a Java2 run-time – there isn't an open source one – but you can download a JDK1.3 as a Debian package from [www.blackdown.org](http://www.blackdown.org). An oft-used solution in Debian, however, is to provide an installer for a particular piece of commercial or non-free software when the software itself cannot be included. Thus there's an

## Bleeding edge

Where to get unofficial and experimental packages

**KDE 3.0.2 (for x86 and PPC only)**  
<http://mypage.bluewin.ch/kde3-debian/>

**XFree86 4.2.0**  
<http://people.debian.org/~branden/>

**GNOME 2.0**  
<http://ftp.debian.org/project/experimental/main/>

**Unofficial multimedia packages, such as lame, DivX, CSS (x86 only)**  
<http://marillat.free.fr/>

installer to download and build nVIDIA 3D drivers as a Debian package, for example. There's also an installer for downloading Microsoft's web fonts.

The other problem is lack of support from commercial vendors. What do you do when they don't provide Debian packages? Of course, most binary packages provided as tarballs will install just fine on Debian. In cases where RPMs only are provided your friend is *alien*, a tool which can convert tarballs and RPMs into Debian packages. Your mileage may vary, but I've found recent versions of *alien* quite successful, although some post-install magic may be required with some converted packages.

### Worth the wait?

To my mind, Debian is the ultimate Linux distribution, at least for those that are familiar with Linux. It offers a combination of bomb-proof stability, simplicity and power that once you have tried, you'll never want to go back from. It's less attractive to newcomers, only because it takes too many steps to install a useful desktop environment. It would take the addition of hardware detection in the boot floppies and a better task selection system at least before I'd recommend it to complete beginners.

The boot floppies, in fact, are in much need of an overhaul, and this is one project that is already a work in progress for the next Debian version. It is possible to create a friendlier and quicker install procedure for Debian, as Progeny proved with their distro. Personally, I would like to see a 2.4 kernel as standard, support for Reiser and ext3 filesystems, *Parted* included for partitioning (preferably with a useable front-end), *GRUB* as the default boot-loader on x86, and support for devfs.

My only other significant complaint with Debian 3.0 is the fact that it took almost two years to get here. To be fair, there were various unforeseen difficulties, such as the legal hoops

that had to be jumped through to get the cryptography software in the main distro. And it can't be understated that simultaneously releasing a distro for eleven architectures – all built from a common source tree – is an impressive feat. Debian has become the main testing ground for the portability of many open source projects, including *XFree86*. For instance, *XFree86 4.2.0* as it was released doesn't build for many of the architectures that Debian supports, and this is the reason that it never found its way into Woody. (You can get experimental packages, however. See box *Bleeding edge*.)

My main worry is that the time it will take for the next version to arrive. Debian's emphasis on stability means that the software included with Debian 3.0 is already fairly conservative when compared to some other mainstream distros. There are some important transitions that need to be gone through for the next release – not least the change to *GCC 3.2* – which could potentially slow things up. It can only be detrimental if we have to wait another two years for the next stable version. If stability is not your main requirement, though, in the meantime you can always track Sarge as it is developed in the testing tree. It will probably prove more robust than you would think. Certainly while I followed the development of Woody, there were remarkably few show-stopping problems. **LXF**

### LINUX Format VERDICT

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

Software and usability updates should bring this new release of the king of distros to a wider audience.

LINUX Format **RATING**  
**9/10**

## DHCP SERVER

# DHCP Turbo 2.1

**Nick Veitch** examines a DHCP server designed to make network administration just that little bit easier.

**DHCP Turbo** is a server daemon and admin tool designed for complex DHCP setups.


- **PRICE** From \$95, See boxout
- **PUBLISHER** Weird Solutions
- **WEB** [www.weird-solutions.com](http://www.weird-solutions.com)

These days it's hard to imagine any reasonable sized LAN being set up that doesn't use DHCP to assign IP addresses and other server settings – it just makes sense. And as you might expect, Linux has a number of options when it comes to setting up a DHCP server, the most popular of which is probably *DHCPd*.

While this is easy to use and configure, with the time-honoured text based config file, for more complicated networks it can become unwieldy. In situations with multiple DHCP servers and complicated rules it's worth considering other options.

One of these options is Weird Solution's *DHCP Turbo*. This server runs on Linux or Windows machines and options can be set via the command line or via a GUI client, which is definitely recommended for initial setup.

*DHCP Turbo* can manage a number of different servers. Each server can have a number of different scopes (which determine the IP address range, lease length, etc) and each scope can contain a number of different policies. The policies



**GUI client makes it easy to change rules and manage several DHCP servers.**

## Prices

*DHCP Turbo* is sold on a client licence basis – the more clients you want to use, the more it will cost. Here's a breakdown of the price structure.

Clients	Price
5	Free
20	\$95
250	\$295
500	\$495
unlimited	\$795

determine the information that is sent to the DHCP client, and cover all the DHCP protocol tags that you might expect, from setting gateway addresses and DNS information to modifying the hostname, forcing

broadcast addresses and whatever else you might reasonably want to do.

## Expressions

One of the most useful aspects of the software is the *Expression Editor*. This can be applied to any of the standard rule tags and allows the flexibility to define dynamic rules based on a number of Functions. For example, it would be simple to construct a ruleset that applied to all devices whose Mac addresses began with '00-06'. Or for every device that requested an IP after 5pm on a Sunday. Or indeed a combination of the above.

The functions cover date and time, various format conversion options, string manipulation and crucially a function which parses the request packet sent to the server. There is even a function to return a default value should any of the other functions return an error due to unforeseen difficulties.

These dynamic scripts can be combined with any number of logical operators, so the flexibility is there to plan for an almost limitless number of scenarios. The expression syntax itself is no harder to understand than the average spreadsheet formula, so you should be able to construct any expression you think you might need without too much recourse to the docs.


The Qt based GUI client is the recommended option for administering the server. This tool is well structured and will make it easier to quickly find and edit

policies and properties of the server, though the command line version may be better for analysis as you can get detailed reports that show at a glance, for example, the policy rules defined.

Both the command line and GUI client can be run on any box that has a connection to the server. For added security, this uses password access as you might expect, but also encrypts data travelling to and from the server. We did have some problems getting the GUI client to run on different Linux distros, but it seemed to be fine on Red Hat and closely-related Mandrake.

*DHCP Turbo* certainly does allow the construction of very complicated policies, and the management of DHCP leasing across many groups of devices. The rules are fairly straightforward to understand, though I can imagine that it will still take a degree of discipline to create a complicated system that will also be easy to spot faults in. The logging feature does give you a head start here (though the log files themselves can get quite long and complicated!) and at least the admin interface does do some error checking before it saves your profile. Worryingly though there doesn't seem to be any syntax checking of the expressions, which could easily result in problems that are hard to track down. Another useful option would be an easy way to copy policies between servers.

Overall though this can make a complicated task much easier. If you need to manage several DHCP servers, or apply different rules to different groups of clients, this could save some time and effort. **LXF**



**Dynamic rules can be applied with the aid of the *Expression Editor*.**

## LINUX Format VERDICT

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

**A good solution for complicated network requirements, and at a reasonable price.**

**LINUX Format RATING**  
**7/10**

## GUI BUILDER

# BXPro 6

Building a GUI isn't often much fun. But **Richard Drummond** discovers that with the right tools, it can be a positive joy!

**BXPro** is a GUI development tool for the Motif toolkit – no similar tools available.

- **DEVELOPER** ICS
- **WEB** <http://www.ics.com>
- **UK DISTRIBUTOR** Scientific Computers Ltd
- **PRICE** £4062+VAT (Full Commercial License including 1 year's support)
- **WEB** <http://www.scl.com/>
- **TEL** 01293 403 636

**M**otif, I must confess, has never been my favourite GUI toolkit on Unix platforms, but since playing with BXPro, a visual Motif builder, my respect for it has increased

enormously (I'm just still not fond of its default look and feel). There are many sound reasons to develop in Motif. It's mature, flexible, and an industry standard on Unix platforms. Ports for Windows exist, so Motif is also cross-platform. Now that Motif has been open-sourced (at least to a degree) and you can distribute it for free on open-source platforms, developing in Motif is now even more attractive (see <http://www.opengroup.com/openmotif/>).

*BXPro* is a bundle consisting of ICS's Builder Xcessory (the GUI builder proper), ViewKit (a C++ framework for developing Motif applications), and EnhancementPak (a set of 34 advanced Motif components, including business graphs). BX currently



supports Motif 2.1 (and hence OpenMotif 2.1), but a free update with support for Motif 2.2 should be available soon. *BXPro* can generate code in C, C++ (optionally using the ViewKit classes), or in a combination of C and UIL. (UIL stands for User Interface Language, and it allows the separation of user interface specification from code. An interface is described in UIL and then compiled into a binary form which can be loaded at run-time to display that interface to the user. The advantage here is that you can then modify the interface without recompiling the code). *BXPro* also lets you build AWT interfaces and generate Java code.

*BXPro* can integrate with various standard open source tools. It automatically generates makefiles for your projects suitable for use with GNU *make*, you can use CVS for source code management (you can also use RCS and Clear Case, too), and *vi*, *emacs* or *nedit* can be integrated for editing code.


## The language of X

If you've used a GUI builder before, you will find much in BX that is familiar. The main window contains the project browser – which displays the hierarchy of widgets and classes in your interface – and the component palette – which displays the widgets which you may use to construct an interface. It has a form editor, where you can drop components and

manipulate them visually. Then there's the resource editor where you can modify the properties (known as resources in X) of widgets. Other aspects of BX will be less familiar.

Motif is built on top of the Xt library (X Toolkit Intrinsics) and Xlib, and so talks the language of X; unless you've done some X programming before, the terminology will probably be unfamiliar. For example, the basis for a window in Xt – the object which talks to your window manager – is called a shell. Two classes of shell are available in BXPro: the TopLevelShell (your garden-variety window) and the XmDialogShell (which is used for displaying dialogues). To create a new window, you have to create a shell, and this will open up a new form editor window in BXPro for you to design the interface that will be displayed in that shell. The XmMainWindow widget, which you might think corresponds to a window, is actually just a container, albeit a container designed to be the basis of your application's main window. Similarly, the XmFrame widget might not correspond to your notion of a frame – it's actually just a container with a border.

*BXPro* comes with some comprehensive printed manuals, and these plus some excellent tutorials are included as PDFs with the package. However, while these constitute a great guide to BX itself, they do assume a knowledge of X, Xt and Motif. If you



**BX's** main window houses the project browser and component palette.

## Requirements

### What you need to run BXPro

Red Hat 6.2+ and SuSE 7+ are officially supported, but any distro with GCC 2.95 and Glibc 2.1 or 2.2 should work.

Glibc 2.1/2.2  
Motif/OpenMotif 2.1  
32MB RAM  
100MB disk space

don't have it, then the Motif Programmer's Guide, available as a PDF from the OpenMotif site, is a good place to start.


## Interface building

Building interfaces with *BXPro* is mostly intuitive once you know your way around Motif. But there are a couple of things to be aware of, before you dive right in without reading the manual. Firstly, you really need three mouse buttons. Secondly, adding components to an interface doesn't work as you might expect. If you select a widget from the palette and left click on the browser where you want that component added, what will actually happen will be a new *TopLevelShell* will be created with that component in it. (If you pick a component and right-click in the browser, a new *XmDialogShell* will be created). If you want to add a child component to an interface via the browser, you either have to drag-and-drop that component by holding down

the middle mouse-button; or you can select the parent component, put the browser into 'Keep Parent' mode, and then simply left-click a widget in the palette to add it as a child. This last method lets you add multiple widgets quickly. Of course, adding components via the form editor works in the usual manner.

The form editor itself is excellent, and includes a powerful set of tools for managing the alignment and geometries of components. In fact, it is one of the best form editors I've seen and makes laying out an interface a joy. In conjunction with Motif's *XmForm* widget, it is easy to control the relative geometry of a widget. Each of the widget's edge has a handle by which you can manipulate it. Right-clicking lets you change its attachment mode – for example, a widget's border can be attached to the edge of a frame or the edge of another widget. Holding CTRL and right-clicking lets you adjust the spacing for that border.

As well as using the form editor to edit widgets, you can use the resource editor. This shows a list of the properties or resources associated with a widget. Many widgets have many resources, but you can filter the list by type (Visual resources, behavioural resource, callback resources, and so on) and a search function is supplied to aid manageability. Depending on their type, resources can be edited with a complex editor. For example, a pixmap editor is supplied for editing pixmaps, and a call-back editor for



The Form Editor boasts powerful tools for managing geometries.

defining call-back functions. You can choose where the value of each resource is saved. Choose 'Code' and that resource is hard-coded in your interface's source code (or UIL); choose 'App' and it is saved in the resource file (*AppDefaults*) for your application. This is a powerful feature, and facilitates all manner of post-build customization of your interface. For example, you can put text strings into the resource file, and allow your interface to be easily translated.


## Development speed

*BX* include many features to speed interface development. Although *BX* is object-based rather than fully object-oriented, it does allow you to take a tree of widgets and turn them into a custom class (as a compound widget). Such a class can be saved out as UIL file, and really aids code re-use. Another great time saving feature is the style editor. You can group together a bundle of widget resources as a style, perhaps to define a particular look-and-feel. Styles may be 'sub-classed', and you can apply styles to a widget or tree of widgets.

*BXPro* provides powerful and comprehensive tools for Motif development. It's certainly not cheap, but if you are doing commercial Motif development, then the cost of *BXPro* is going to be a lot less than your salary bill. You do actually get a lot for your money. *BX* itself is a very

impressive product. Although the initial learning curve is steep, once you find your way around, it provides tools for creating interfaces with unrivalled precision and speed. In addition, *ViewKit* provides an OO framework for speeding Motif development in C++, and the *EnhancementPak* contains a set of advanced application widgets – again saving you time and money.

Of course, no application is perfect, and there are areas of *BX* that could be improved. I would like to see more integration with the compilation stage of a project, that is, more 'real' IDE features. Secondly, although you can save out text strings to a resource file for later translation, a dedicated tool for managing translations (à la Qt's *Linguist*, for example) would be useful; lastly, a widget capable of rendering HTML would be a valuable addition to the otherwise very comprehensive set included. **LXF**



Widget resources or properties can be altered in the Resource Editor.

## LINUX Format VERDICT

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

Not cheap, but this product provides a powerful set of tools for turbo-charging Motif development.

## LINUX Format RATING

9/10

SOURCE DISTRO

# Gentoo Linux 1.2

Use the source, advises **Jon Kent**, and you'll have a much better tailored GNU/Linux installation.

Source-based Linux distro that competes with LinuxFromScratch

- **PUBLISHER** Gentoo Technologies
- **LICENCE** GNU GPL
- **WEB** [www.gentoo.org](http://www.gentoo.org)

Since the increased popularity of Linux, most distributions now provide binary packages for you, pre-compiled and easy to install. All you have to do is download the package and install it. Good isn't it? Well, yes and no. Like all things that are made easier for you, there are drawbacks. Take for example GUI interfaces to system administrations functions. These front-ends help you to configure the system how you want quickly, but the same task can usually be performed far more quickly at the command line.

So, what are the drawbacks to binary packages? There are a few, the main one being that the package is not compiled to suit your CPU, in fact a lot of distributions are i386 compatible which waste a lot of the power you may have in your i686. In addition, package maintainers have to assume things on your behalf, so you may end up with functionality you do not want, thereby increasing the size of the binary and your overall system. It is this situation that Gentoo Linux aims to resolve.

Gentoo Linux is a source-based distribution, that is instead of downloading a precompiled binary, you download source code and compile it on your system, allowing you to use any additional features that your system has. Gentoo currently supports x86, PowerPC, Sun Sparc and Sun Sparc64. This review covers the x86 version of Gentoo.

## Installation

To install Gentoo onto your system you obviously need an install CD of which Gentoo provides two ISO images that can be downloaded. One is a fairly small 16MB ISO which is enough to get a basic system up and running ready to download further source



The ultra minimal window manager *Fluxbox* flies when run with Gentoo.

packages that you will need to complete the installation. The other ISO kicks in at 100MB, but is only available if you have an i686 machine. This ISO then allows you to decide at which point in the installation process you wish to start to self-compile. A word of warning though. Even if you decide to use all of the binaries provided on the 100MB, you will still have to download sources for XFree86, KDE and so on, so it helps if you have a fast link out to the Internet. Building Gentoo over a 56K modem will be painfully slow and is not recommended.

The installation process lacks any fancy GUI interface, or even a console base installation tool such as Debian's. This is not to say that the installation is difficult, it is just a bare bones method of installing Linux. I actually prefer this method of installation as you have far more control over your system than an overfriendly GUI which assumes too much on your behalf, which most of the GUI installers seem to do.

Once the minimal system is in place, the remaining installation is performed by downloading the latest sources, usually from the Gentoo servers or from the package's home server, and compiling them locally. How many sources are downloaded depends upon which CD ISO you have used, and with which tar ball you have built the base system, of which there are three to choose from on the 100MB ISO. With the 16MB ISO you start at, in Gentoo speak, stage one of the installation process and therefore you have to download a large number of sources and then compile them. Stage one therefore takes the longest to perform as you have to build nearly everything from downloaded source code, stage two cuts the installation time in half (roughly) and stage three is basically a running system, without X.

Stage three is, obviously, the fastest installation type, but you lose Gentoo's flexibility slightly. If you do decide to use the smaller ISO be prepared for a

long installation time, which means literally hours. However, there are many advantages to this installation method, namely packages are built for your machine's CPU type and you can customise it to your requirements. The process itself takes care of dependencies and by the end of the installation you will have the latest snapshot of the Gentoo tree.

Despite the length of time involved, installing Gentoo is surprisingly simple and the installation instructions on the website are extremely good. If you have been using Linux for a while you have little to be afraid of here. Booting your chosen CD brings you into a running kernel and once the boot process has completed you are left at the command prompt.

From here you now begin the tasks of preparing the system by hand – load any kernel modules needed (usually just for your network card), configure and startup networking, partition and format your hard drive, mount the partitions, mount the CD,

and then unpack your chosen stage tar ball. Lastly you'll need to configure your kernel and setup *GRUB*, which is Gentoo's preferred bootloader. It may seem like a lot of work, but it is very simple to achieve.

## Source Management

Gentoo does not use either rpm or deb packages, but uses a source management system called *Portage* which currently consists of around 1500 sets of package build scripts called ebuilds. The *Portage* system is written in Python with the ebuilds being *Bash* scripts. The ebuild scripts provide details to the *Portage* backend about how to compile and install a software package and what dependencies there are, if any. By using profiles and the command line utility *emerge*, you can use *Portage* to install, update, remove and generally

maintain the packages that are installed on your Gentoo system.

In this respect you could consider *emerge* as providing roughly the same functionality as Debian's *apt* suite of utilities, albeit for source packages only. If the *Portage* system itself sounds familiar to you, then this may be because *Portage* has borrowed a lot of ideas from BSD's ports system, although the Gentoo developers maintain that "*Portage* is definitely not just another ports rip off". Fundamentally, however, they are basically the same as they both compile packages from source and allow users to safely install and uninstall software from a system and both automatically handle dependencies.

As touched on previously, *Portage* also supports the concept of profiles. Profiles contain a list of package

names and versions that are used to define a set of default configuration options to be used by *Portage*. A profile tells *Portage* which packages and which specific versions of those packages to allow, disallow, or treat as required. This allows you to switch profiles by changing a single symlink (*/etc/make.profile*) to point to your newly created profile. In reality you would probably only need to modify your original *make.profile* file very infrequently but it is good to have the flexibility available to you.

A useful feature of the profile files is the **USE** variable, which is used in the *make.conf* and other files. This variable defines what should, or should not, be enabled when a source tree is compiled. For example, packages with optional GNOME support can have this support disabled at compile time by disabling the **gnome USE** setting. Enabling the **gnome USE** setting would enable GNOME support in these same packages. Enabling or disabling additional support, or feature, is just a case of modifying the line **USE=""** and adding your requirements.

To enable support you simply add the name of feature you require, or to disable a feature you would use a **-** sign in front of its name. There are default **USE** settings defined in */etc/make.profile/make.defaults* that are configured during installation, which you can refer to to see if you need to modify the **USE** variable within the *make.conf* file. This is a very useful feature to have as this enables you to adapt the system to your exact requirements, therefore providing you with a highly tuned system.

However, there may be a few problems with Gentoo depending upon how you access the internet. To ensure that *Portage* is up to date you need to run an *rsync* to the Gentoo server to download the latest ebuild scripts and directory structure. If you are behind a firewall, it is unlikely that you will be able to run such an *rsync* so this will cause you problems when trying to keep Gentoo up to date. Also, some of the ebuild scripts use ftp to download packages. Again, if you are behind a firewall you will most probably not be able to do this directly, if at all. In the case of some of the ebuilds ftp is a strange choice as there are also http download sites available for the same source code. Mostly this can be overcome by editing the ebuild script and


replacing the ftp site with the http site and then re-running *emerge*.

## Performance boost

Once you have Gentoo installed though, there is a noticeable improvement in speed of applications such as XFree86, *Abiword* and other applications that you have compiled specifically for your system, in comparison to the same applications on Debian or Red Hat for example.

However, although there is a gain in performance, you do lose out on some of the additional features other binary distributions supply, such as additional administration tools that you may be used to using. In some respects this is no bad thing as most of these are aimed at making Linux easier to administrate for new users and are not necessarily required. If you need these tools to be available you will mostly have to download the source and install it by hand, which is not a major task in itself, but something that needs to be considered. There are some strange tools missing, such as *vi*, but most of the lower level utilities are available as ebuilds and can be easily installed.

Gentoo does exactly what it sets out to do, provide a source based distribution that you can tailor exactly for your system and requirements, with an easy to use source management system provided by *Portage*. The installation does take a long time to perform, but the gains achieved outweigh the time taken to install Gentoo. If this matches your requirements you cannot go far wrong with Gentoo. The minor problems that there are with Gentoo are just that, minor, and can mostly be overcome quite easily. If you approach it with the knowledge that this is a system that you customise for your own requirements, Gentoo will provide you with a very good example of a source distro. **LXF**



The usual apps, such as *Abiword*, are available to download and compile.



If GNOME or more to your taste than KDE, that is available too.

## LINUX Format VERDICT

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

A well put together source distribution that allows you to use the full power available from your system.

## LINUX Format RATING

9/10

# JSP Examples and Best Practices

Successful web development with Java requires more than simply a grasp of JSP tags, as **Richard Drummond** find out.

- **PUBLISHER** APress
- **AUTHOR** Andrew Patzer
- **ISBN** 1590590201
- **PRICE** £32.00

**J**ava-based technologies such as servlets and JavaServer Pages (JSP) have revolutionised web application development, but technology by itself isn't much use – you have to know the right way to apply it. Andrew Patzer's book can give you a vital head start. It provides design principles and patterns, methods and frameworks that you can use to design and build maintainable web applications.

The book starts off with the basics and reviews the foundations of programming with JSP. Chapters one and two cover the nuts and bolts,

while the next two chapters begin to focus on design issues. Here the reader will find out how to use JavaBeans and later custom JSP tags to separate an application's

presentation from its business logic – a key to building maintainable apps.

This introductory material sets the stage for the Model-View-Controller (MVC) architecture, introduced in chapter 5. The MVC pattern is the Java programmer's favoured way of attaining that separation of content from presentation, and it underpins the topics discussed in the remainder of the book. The design patterns presented in the following three chapters are all derived from the MVC pattern. In particular, the widely-used Front Controller Pattern and how it may be implemented is thoroughly



explored. (The Front Controller Pattern is where a single Controller, typically a servlet, receives requests, processes and forwards them to the appropriate JSP page to provide the View.)

The author looks next at collecting these principles together to design and build a re-usable application framework that can form the foundation for your future JSP projects. The example framework presented is quite functional, and this is demonstrated with a full example app – an online catalogue and shopping cart.

JSP Examples and Best Practices is clear, concise and adopts a no-

nonsense approach. It covers the larger issues not tackled by lesser volumes – such as design, testing and deployment – and thus will be an indispensable guide for the JSP developer wishing to start developing real world applications.

## Linux Format VERDICT

A practical and readable guide to JSP programming in the real world, with extensive, useful example code.

## LinuxFormat RATING

9/10

# TCP/IP Network Administration

## 3rd edition

**Jon Kent** wonders whether this new edition is still an essential read.


- **PUBLISHER** O'Reilly
- **AUTHOR** Craig Hunt
- **ISBN** 0-596-00297-1
- **PRICE** £31.95

**T**his classic technical book has been in publication since 1992. Along with *Essential System Administration* and *Managing NFS/NIS*, they cemented O'Reilly's reputation as one of the best technical publishers. With its third edition, is this book still an essential purchase?

As per the previous editions the book starts at first base, providing a good introduction to TCP/IP covering topics such as the architecture (OSI model), IP addressing, routing and the

underlying protocols. These first few chapters are well written and probably make the book a good purchase. Then it moves into the actual configuration of a Unix box to TCP/IP connectivity, including any Kernel configuration that is required, configuration of *inetd* and a good overview of routing configuration.

Then we move onto configuring network-based services such as DNS, NIS, *Samba*, *sendmail* and *Apache*. The detail the book goes into with these services leaps from fairly in depth (DNS, *sendmail*, *Apache*) to a quick run through (NIS, *Samba*). This approach means that some useful information is either not given enough space or left out completely, which leaves you wondering if it was worth



having the section in the book at all.

The book is aimed at Unix sysadmins, not just Linux Admins, however most of the examples are for Sun Solaris. Unfortunately, the Linux examples here really mean 'this is how it works for Red Hat'. The author assumes that if it works in Red Hat, it works for all, which is not always true. An examples of this is when he recommends adding a static route command to *rc.local* script, but the script doesn't exist on all Linux distros.

If your knowledge base already

includes TCP/IP network, then you may not gain too much from this book. However, if you are starting to look at TCP/IP this book is a good starting point.

## Linux Format VERDICT

A good book to get to grips with TCP/IP, providing good foundations.

## LinuxFormat RATING

8/10

# Web Database Applications with PHP and MySQL

**Nick Veitch** discovers a useful guide to building a commercial website.


- **AUTHORS** Hugh Williams & David Lane
- **PRICE** £31.95
- **PUBLISHERS** O'Reilly
- **ISBN** 0596-00041-3

Perhaps O'Reilly were trying to convey some sort of message with the choice of a platypus as cover creature, but *MySQL* and PHP are a winning combination for dynamic websites, so this book is a welcome guide to using both together.

At first glance, teaching you all you need to know about PHP, all you need to know about *MySQL* and how to bring them together in real-world

applications seems like a huge task for a book of some 550 pages. The book assumes some general familiarity with coding and HTML, though not any experience with PHP. The main features of the language are sketched out, along with the basics of *MySQL* database in the first 100 or so pages.

Most of the real-world element of the book is based around a fictional wine store website. This scenario is perfect for introducing elements for online ordering systems, such as user authentication, transactions, database design, stateful operations, session management and the major topics you might come across when building a commercial website. All the examples



are fully fleshed out in code, reproduced in the book (and as usual, available for download too) – the only slight problem being that it's all based on PHP 4.0.6, and certain minor things have changed in later releases.

This focus on a commercial shopping site might be a bit narrow for all tastes – it might have been nice to see some examples of storing different types of data in the database, or dynamic graphics (for graphs of data stored in the database or whatever). However, the core things that people might want to do are adequately

covered.

The book manages to deliver on two levels at once – the concepts, planning and design process as well as implementation. It's also surprisingly well written, and manages not to be too smug or patronising. **LXF**

## Linux Format VERDICT

Well written and very informative

LinuxFormat **RATING**

8/10

Making The Right  
**Hosting Choice**  
For Your Business

## SOLUTIONS

**With XCalibre you're not just a number  
Let our customers speak for us!**

We have been a customer of XCalibre for over two years now and have to congratulate them on their fantastic service. We moved to XCalibre because our previous provider was so unreliable, and it was definitely the right move. Servers and connection are fast and reliable, and when a rare problem does occur their technical support is excellent.

Ken Hobbs - Integrated Internet Solutions  
<http://www.iis-uk.net>

### Dedicated Servers

- Flexible hosting
- Secure facilities
- Reliable servers

**From £ 75 per month**

### Managed Services

- Total server management
- Quality service guaranteed
- Complete peace of mind

**Call now for a quote!**

### Co-Location

- World class data centre
- Data transfer from £2.50/GB
- Secure private suites

**From £ 52 per month**

### Web Hosting

- personal customer support
- Full control panel
- Regular backups

**From £ 39.95 per year**

**[www.xcalibre.co.uk](http://www.xcalibre.co.uk)**

**Tel: 0870 050 0080**  
email: [sales@xcalibre.co.uk](mailto:sales@xcalibre.co.uk)

# Roundup

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

## Firewall distros

### Our selection at a glance

- MandrakeSecurity SNF
- Astaro Security Linux
- SuSE Linux Firewall on CD
- XSentry Firewall
- Smoothwall GPL
- Smoothwall Corporate
- Securepoint Firewall Server

**Mark Newby and Richard Drummond** take us on a tour of some of the best known Free and non-Free firewall distributions.

**A**s Internet security threats continue to increase rapidly, so too does the number of different firewall solutions now available.

This roundup focuses primarily on ease of installation, usability, admin facilities and features – rather than penetration testing and firewall design evaluation. Therefore, this article is more easily digested by those people that aren't necessarily network security experts, but who are thinking about implementing or changing over to a new firewall product, or those who are simply tired of getting lost in their home-grown *iptables/ipchains* scripts.

## MandrakeSecurity Single Network Firewall

GPL'd SOHO firewall from the simplicity experts.

- **VERSION** 7.2 ■ **WEB** [www.mandrakesoft.com/products/snf](http://www.mandrakesoft.com/products/snf)
- **PRICE** \$299 (boxed)/ Free (download edition)

**Mandrake's Single Network Firewall (SNF)** is designed for use in SOHO environments. It supports a range of connection types including modems, xDSL, and ISDN, but it is simpler than the other firewalls on test in that it only supports two zones, internal and

external (no DMZ) and there's no VPN service. This simplicity means easier config, so it may be a good choice for those with modest requirements. As it's GPL'd; you can download an ISO and use it for free. The boxed edition comes with two CDs (binary and source), an

installation manual and 60 days of online support via *MandrakeExpert*.

*SNF* is based on a hardened 2.2 kernel, and, as well as the usual packet filtering, NAT and port-forwarding, it offers intrusion detection, a DHCP server, and a web proxy (with URL-banning and content-filtering). Firewall administration is performed remotely via a secure web-based interface.

The *SNF* installer is a cut-down version of the now-familiar Mandrake installer, which is even easier. You get all the usual good stuff from Mandrake such as hardware detection and a partition editor.

The web-based admin interface lets you perform routine maintenance, modify your security policy and monitor your firewall – and is mostly intuitive, although it can be long-winded (as is typical of web-based interfaces). Firewall rules are managed in two pages, 'Office traffic' and 'Internet traffic', which let you specify which protocols are allowed to originate from the internal and external networks,

respectively. The 'Internet traffic' is also used to set-up port-forwarding. The default rules on a newly installed firewall are set up rather unintelligently, but, since *SNF* is simpler than the competition, you'll be less likely to make any mistakes. The section for configuring the web proxy and content filter is fairly comprehensive and even lets you restrict access to websites by time. The monitoring tools provided are fairly basic, but you can set up logging to a remote server.

### Stated requirements

Pentium-class processor  
32MB RAM (64MB recommended)  
250MB hard drive space

### LINUX Format VERDICT

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

A simple but effective solution for the SOHO environment.

LINUX Format **RATING**  
 7/10



Easy configuration makes *SNF* particularly suited for SOHO networks.

# Astaro Security Linux

Professional solution from the new kid on the block.

■ **VERSION** 3.2 ■ **WEB** [www.astaro.com](http://www.astaro.com) ■ **PRICE** see text

## Astaro is a German company

formed just over two years ago to provide security solutions based on open source software. One of these solutions is *Astaro Security Linux (ASL)*, a firewall distribution with a slew of integrated services, aimed at medium-sized companies.

ASL is free for home and non-commercial use, and for businesses the pricing scheme is based on the number of IP addresses in use on all the networks protected by the firewall. Prices range from £480 for 25 IPs up to £3,194 for an unlimited number of nodes (additional services such as redundant fail-over, content filtering and virus-scanning are extra). This may seem steep – but *Astaro* offers an unrivalled level of features, and should prove cost-effective across the board. Various levels of support are available from *Astaro* partners around the world, and a community site and forum at [www.astaro.org](http://www.astaro.org) provides an excellent free resource for support.

ASL is based on a hardened 2.4 kernel and includes a number of advanced networking features such as Quality of Service traffic prioritising. It provides *iptables*-based packet-filtering, NAT (masquerading, SNAT and DNAT) and port-forwarding. It doesn't support dial-up connections on the external interface, but does provide PPPoE support and a DHCP client, so xDSL adaptors or cable modems which hook up via Ethernet should work.

An impressive range of application proxies and services are included in ASL – such as DHCP, web proxying/caching, mail relaying, DNS forwarding, VPN, and SOCKS proxying. The web proxy features URL blocking and optional content filtering, and spam filters and optional virus scanning are provided in the mail relay. ASL also boasts a number of enterprise-level features such as optional fail-over redundancy (to a second machine running an identical ASL configuration) and support for load-balancing to servers in your DMZ.


Some might think that all this extra functionality in ASL overburdens

## Stated requirements

400MHz Pentium II  
128MB RAM  
8GB  
2 or more NICs

and weakens a firewall, but most services are run in a *chroot* jail, so even if a service is compromised, it shouldn't compromise the whole firewall. Moreover, all these application proxies make setting up your packet-filtering rules simpler and less prone to error. It does mean, however, that the minimum hardware requirements are fairly high. If you have low bandwidth requirements and don't need all the extra services, a lower-spec machine will probably suffice.

Installation is straightforward with the text-only, menu-based installer. Hardware detection is performed, so the amount of configuration required is fairly minimal. You have to choose a drive to install to (it needs a dedicated drive), assign network addresses to your Ethernet cards, set up passwords, and so on. The remainder of the



Astaro's web-based administration interface is well designed.


configuration can be done remotely post-install with the web-based interface, *WebAdmin*.

Out of the box, ASL's default configuration is a sensible default deny policy with only a minimum of outgoing connections such as DNS allowed and only HTTPS allowed in. It attempts no automatic configuration, though, so it's up to the administrator to manually configure the network interfaces, routing table, packet filtering, etc. Good firewall design skills are thus needed.

*WebAdmin* provides a well-designed, professional-looking interface, but suffers from the limitations of the medium. It employs a

similar model to *X Sentry's* and *SecurePoint's* administration interfaces, but the since the latter pair are traditional thick-client GUIs, they are quicker and easier to use. Still, *WebAdmin* has the advantage of portability, and in use is mostly straightforward. It does have good navigation cues throughout and the online documentation is helpful. Basic shell administration is also possible, but ASL provides no dedicated tools.

If you're looking for an all-in-one, professionally-supported solution and have your Internet connection coming in over Ethernet, then ASL will surely meet whatever requirements you have – considering its wealth of features, flexibility and modularity. The free license for non-commercial use might make it an attractive proposition for home users wanting to protect their local networks, but a lot of the features provided will be just overkill outside of the enterprise space. Unless you have the skills required to design a secure packet-filtering strategy, then ASL probably isn't for you.



Astaro Security Linux offers an excellent package of proxy services.

## LINUX Format VERDICT

Ease of use	5/10
Features	9/10
Performance	8/10
Value for money	8/10

A professional product, bristling with features, suited more to medium-sized enterprises rather than the SOHO environment.


## LINUX Format RATING

8/10

# SuSE Linux Firewall on CD

GUI admin and non-writable medium from SuSE.

- **WEB** [www.suse.com/us/products/suse\\_business/firewall](http://www.suse.com/us/products/suse_business/firewall)
- **PRICE** \$1,149 (no VPN) \$1,499 (with VPN)



SuSE's *Firewall Administration System* is provided with a complete distro.

SuSE's *Linux Firewall (SLF)* runs from a live CD rather than requiring to be installed to a hard drive. This approach maximises security at the expense of complicating installation and configuration. The advantage is that the system files can't be tampered with during a system compromise (the

effectiveness of this approach should not be overestimated, as much of the OS's run-time data must be stored in RAM, which is still volatile). Also, the ability to recover the machine with a simple one-step reboot could be a real time-saver when recovering from a compromise.

There are therefore two components

to *SLF*: the firewall 'Live CD' running on the firewall host itself and the *Firewall Administration System (FAS)* which installs and runs on a separate box, which SuSE call the 'Adminhost'. The *FAS* is based on a SuSE 7.2 distro with GUI admin system, the *FAS* proper. Before you can boot the firewall you must use the *FAS* to configure your policy and save the config to a floppy – which is then read by the firewall when it boots.

We reviewed *SLF* in *LXF27*, so I won't retrace too much ground here. *SLF* offers an *ipchains*-based firewall with the usual port-forwarding and NAT. It includes a small set of application proxies, such as *squid* for web proxying and *postfix* for mail relaying (for these, your firewall does need to access the hard drive). VPN is an option, but there's no support for URL blocking, content filtering, or virus-scanning.

The Adminhost is fairly painless to install with the usual SuSE graphical installer, although we did find unexpected hardware problems on the first two machines we tried. *FAS* is mostly straightforward to use, but the multi-stage process of implementation is cumbersome and time-consuming.

The supplied docs are excellent. The absence of feature-bloat is good to see in a firewall product and the ability to run from the CD is definitely an asset. However, *SLF* offers too few features to justify the price tag.

## Stated requirements

**Firewall:** Pentium, 128MB RAM, 1GB-10GB disk space (for logs and caches), 2 NICs  
**AdminHost:** Pentium, 256MB RAM, 2GB disk space

## LINUX Format VERDICT

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

A solid and functional product marred by overly complex installation and a high price.

LINUX Format **RATING**  
 /////////////// 7/10


# XSentry Firewall

Multiple zones and a Java-based GUI.

- **VERSION** 1.6 ■ **WEB** [www.trustix.com/products/xsentry](http://www.trustix.com/products/xsentry)
- **PRICE** \$780-\$5490

*XSentry* is a basic firewall product which runs atop Trustix Secure Linux, – a hardened Red Hat derivative. It offers

no proxying, no intrusion detection or advanced monitoring features, just NAT, packet-filtering and port-forwarding.



Drag-and-drop firewall building, courtesy of *XSentry*'s Java-based client.

## Stated requirements

Pentium  
 32 MB RAM  
 2+ NICs  
 4GB hard drive

*XSentry* detects hardware and requires minimal config during install. You just need to set up your NICs, pick a host name, choose a password, etc. Boot up your newly installed firewall and you are presented with a dialog-based menu system at which you can perform rudimentary configuration. The actual task of administering the filter rules is done remotely with a Java-based GUI.

The Java admin GUI offers the combined advantage of administering the system from anywhere (anywhere with a Java 2 runtime, that is) and the usability of a traditional, thick-client GUI.

The *XSentry* GUI is excellent and by far and away the easiest and quickest to use of all the full firewall admin systems on test. It provides two views of your network topology: a familiar, hierarchical list and a visual representation, somewhat reminiscent of diagramming apps such as *Visio*. It is a joy to use, and lets you build your security policy largely by drag-and-drop. Despite being simple to use, the high price and limited feature set mean that *XSentry* will be a less than attractive proposition for most.

## LINUX Format VERDICT

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

Quick to set up, and boasts an excellent admin GUI, but just way too expensive for the features offered.


LINUX Format **RATING**  
 /////////////// 5/10

# Smoothwall GPL

Rescuer of many an old 486...

■ **VERSION** 0.9.9SE (Special Edition)

■ **WEB** [www.smoothwall.org](http://www.smoothwall.org) ■ **PRICE** Free



You can use a variety of dial-ups devices with Smoothwall.

Smoothwall was initially designed as way of turning a redundant PC into a Linux-based firewall with a minimum of effort or configuration. It has gained a few more features since then, but it is still incredibly easy to use and will still run on low resource hardware (a 486 and 16 MB of RAM are ample). Smoothwall requires no complex administration of packet filtering rules (although you can add port-forwarding and external access rules if you wish) and thus allows non-experts to implement a firewall on their home or office network without needing complex networking or security skills.

Smoothwall is optimised for SOHO environments, and so includes support for modems and ISDN dial-up connections (with dial-on-demand) and for ADSL (including Alcatel USB ADSL adaptors) and cable modems. A neat feature is built-in support for various dynamic DNS services such as dyndns.org and no-ip.com – which is great for those on a dynamic IP that want to give name-based access to their network from the Internet. Although easy to use at home, Smoothwall scales

surprisingly well to larger networks, too, offering a good range of services including a web proxy, a DNS proxy, a DHCP server, intrusion detection and VPN.

Installation is straightforward and requires a minimum of input from the user. Smoothwall is derived from Red Hat and employs a similar, but cut-down text-based installer. Smoothwall takes over an entire hard drive, and so the first task is to repartition and reformat that

(only IDE drives are currently supported, not SCSI). You then need to pick the Ethernet adaptor (if you have more than one NIC) that accesses the internal network and give it an IP address. This will be known as the 'Green' (safe) interface. Then the system software is installed. Next you need to configure the external or 'Red' interface and the DMZ or 'Orange' interface (if you have one).

On my test system, I had two network cards and a dial-up modem. The first network card was connected to my private LAN and the second to a separate LAN segment allocated as my DMZ. So, the internal interface gets assigned to the Green interface, the second to the Orange, whereas the Red interface was automatically assigned by the set-up routine to the modem, because I bypassed the ISDN and ADSL configuration steps. With this information, Smoothwall automatically configures the packet filtering rules based on the interfaces you're using, so you end up with a firewall that effectively works out-of-the-box without the need to go through a lengthy process of manual configuration.

With version 0.9.9 they have overhauled Smoothwall's web-based administration interface. This now looks a lot more professional and pleasing to the eye, but remains just as easy as ever to use. Since little administration is actually necessary, most users won't use the interface much – other than for monitoring



Smoothwall is easy to install and requires minimal configuration.

## Smoothwall Corporate Server

Smoothwall Ltd offer a commercially supported version of Smoothwall, called Smoothwall Corporate Server. Version 2.0 has just been released, based on a new modular and extensible architecture, with add-on modules for content filtering (SmoothGuard), virtual hosting environments (SmoothHost) and for editing packet-filter rules (SmoothRule). Expect a full review in an upcoming issue of Linux Format.

## Stated requirements

**i386 processor**  
**2 NICs (or 1 NIC plus modem)**  
**8MB RAM**  
**60MB disk space**

network traffic and logs. The web interface lets you reconfigure your network settings, add port-forwarding or external access rules, and activate services such as SASH, web proxying, and intrusion detection. The Smoothwall documentation is available from web interface should you need it.

Smoothwall is an excellent product, which you can use to implement a functional firewall in minutes. The lack of administration necessary is a big point in its favour, since its default policy is secure but useful out of the box – without requiring any, possibly error-prone, configuration. And, despite its simplicity, Smoothwall does offer some advanced features, such as intrusion detection. Bigger enterprises may require more flexible configuration tools, but since Smoothwall is free, it's hard to fault.

## LINUX Format VERDICT

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

It's free, easy-to-use and secure – it's not the most popular Linux-based firewall solution for nothing.

## LINUX Format RATING

9/10

# Securepoint Firewall Server

Professional security at a professional price.


- **VERSION** 3.0
- **WEB** [www.securepoint.de](http://www.securepoint.de)
- **PRICE** see box, Pricing

Securepoint ship a range of firewall solutions at different price points, which differ in the complexity of network they can support and VPN features. All are based on the same Red Hat-based core, with a hardened 2.4.18 kernel and an *iptables* firewall. The administration system is a secure client-server system, with server implemented in Java and hosted on a *PostgreSQL* back-end for logging. The GUI client is currently available only as a Windows binary. This works satisfactorily under *Wine*, but your mileage may vary.

*Securepoint* supports ISDN and ADSL Internet connections, but there's no support for modems. The basic *Office* edition allows only two network adaptors, the *Business* edition eight, while the *Professional* edition can handle up to 16 devices and can therefore cope with very complex networks. As well as packet-filtering, NAT and port-forwarding, *Securepoint* provides web and optionally mail proxies, with URL blocking and content filtering. Professional virus scanning solutions from Viruswall, Antivir and Sophos may be purchased additionally. Intrusion detection is included as standard, and the *Professional* edition features both IPSec and PPTP VPN services.

Installation is simple. Like the majority of solution on test here, *Securepoint* takes over an entire drive. No hardware detection is performed, but a simple text-based interface lets you assign kernel drivers (and hence network cards) to the network devices eth0, eth1, and so on.

The first device eth0 is the internal interface, eth1 the external interface, and eth2 the DMZ, but it's a bit hit-and-miss getting these in the right order, especially if you have multiple cards of the same type. You can always edit */etc/modules.conf* post-install to fix this, though. You also



The GUI runs on a Windows PC, but will run satisfactorily under Linux with the help of *Wine*.

## Pricing

**Firewall Server Freeware**  
Free for non-commercial use

**Firewall Server 3.0 Business**  
Unlimited users \$1,195

**Firewall & VPN Server 3.0 Professional**  
50 users \$1,595  
100 users \$1,995  
250 users \$2,995  
Unlimited users \$4,995

have to assign IP numbers and netmasks to your network devices here. Next you are asked for routine configuration data, such as the default gateway, nameservers, and passwords; then the software is installed and you can boot up your firewall.

You perform a local shell login on your firewall host to perform basic configuration. A dialog-based menu lets you tweak certain settings. To set up your security policy, you need to use the admin GUI. The default policy installed is deny everything,

which is rather draconian. You can ping from the internal network, but nothing else.

The Windows admin GUI has a rather steep learning curve, but once you get used to it, it is straightforward. Unfortunately the documentation is rather vague and doesn't help much. The first task is to define the objects in your network – the internal network, any significant hosts, servers in the DMZ, etc.

You can then add rules allowing traffic between the various objects you have defined. Protocols can be assigned to rules in groups, thus speeding up configuration. For example, a standard group is web which includes HTTP, HTTPS, and FTP.

*Securepoint* requires the necessary skill to set up a secure firewall policy – it doesn't do any of the hard work for you. On the plus side, the thick-client interface is easier and quicker to use than a web-based solution. A handy feature is that you can get a graphical map of your network topology, and view rules diagrammatically – a real boon to

ensuring you set things up correctly.

*Securepoint* provides some excellent accounting and logging functions, and the web interface provides the tools to query these and to view traffic accounts graphically.

*Securepoint* is a professional product aimed at the enterprise market and doesn't fail to impress. However, when put alongside the comparable *Astaro Security Linux*, the price tag is steep – especially when you factor in VPN and virus-scanning. And *ASL* provides more enterprise-level features, too.

## LINUX Format VERDICT

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

A competent and fully-featured professional firewall product, but looks expensive.

LINUX Format **RATING**  
**7/10**

# FIREWALLS THE VERDICT

**Y**ou may have more work to do in making a choice, but it's great to see that such a large number of Linux-based firewall products are available: it's a sure indication of the success of the operating system as a whole. In the last couple of years there's been a significant increase in the number of firewall solutions – as more and more people come to realise just how well suited Linux can be at turning a PC into a full featured, commercial-grade firewall device.

Here we've seen both free and non-free products that aim themselves at roughly the same thing – protecting your network from the big, bad, cracker-infested Internet.


Some are aimed squarely at the corporate environment – such as *Astaro Security Linux* and *Securepoint Firewall Server* – while others – such as Mandrake's *SNF* – haven't been quite as brave and, instead, have narrowed their scope to purely SOHO networks. Your particular environment will most definitely play an important part in guiding you to the most appropriate product.

Some of these products do an excellent job with a minimum amount of configuration, while others require fine-grained control over the firewall's operational specification and hence need a greater degree of skill to set up and maintain. Perhaps as these products mature, their intelligence will increase, thereby requiring less manual configuration.

In this respect, *Smoothwall* wins hands-down – and without any competition biting at its heels from what I can see. I do believe that if the other products want to capture the interest of a wider audience, they too must be somewhat more self-configuring, or at least have this type of feature as an option.

If you're a SOHO user with some technical skills, it would be hard not to recommend *Smoothwall GPL*. It's incredibly easy to use, and does the job required of it with a minimum of fuss. If commercial support is a requirement, then you might consider the commercial edition of *Smoothwall*.

For larger and more complex networks the flexibility and features



**Astaro – not cheap, but well worth the cost if you need its features.**

offered by the more advanced products may be a better fit. Products such as *Astaro Security Linux* and *SecurePoint Firewall Server* may have some features that you or your company simply can't live without or that you can easily justify in terms of cost/benefit.

Whatever your requirements, remember that a firewall's design and robustness is critical, and it's extremely easy to make a mistake that leaves a hole gaping wide open in your defenses. The smallest hole is all it takes for a cracker to get in and, believe you me, they're very good at finding them! Choose wisely, and good luck! [LXF](#)

## Comparison of features

Package	Astaro	Mandrake	Securepoint	Smoothwall	SuSE	X Sentry
<b>Packet filter</b>	2.4/iptables	2.2/ipchains/2.4/iptables	2.2/ipchains	2.2/ipchains	2.2/ipchains	
<b>NAT</b>	Yes	Yes	Yes	Yes	Yes	Yes
<b>Port forwarding</b>	Yes	Yes	Yes	Yes	Yes	Yes
<b>No. NICs supported 2-20</b>	1-2	2-16 (1)	1-3	2-10	2-4	
<b>EXTERNAL CONNECTION</b>						
<b>Modem</b>	No	Yes	No	Yes	No	No
<b>ADSL</b>	Yes	Yes	Yes	Yes	No	No
<b>ISDN</b>	No	Yes	Yes	Yes	No	No
<b>Ethernet</b>	Yes	Yes	Yes	Yes	Yes	Yes
<b>SERVICES</b>						
<b>SSH</b>	Yes	Yes	No	Yes	Yes	Yes
<b>DHCP</b>	Yes	Yes	No	Yes	No	No
<b>Web cache/proxy</b>	Yes	Yes	Yes	Yes	Yes	No
<b>Mail relay</b>	Yes	No	Yes	No	Yes	No
<b>DNS cache/forwarding</b>	Yes	No	No	Yes	Yes	No
<b>VPN IPsec</b>	Yes	No	Yes	Yes	Yes	Yes
<b>VPN PPTP</b>	Yes	No	Yes	No	No	No
<b>SECURITY FEATURES</b>						
<b>Intrusion detection</b>	Yes	Yes	Yes	Yes	No	No
<b>Content filter</b>	Yes	Yes	Yes	No	No	No
<b>Virus scanning</b>	Yes	No	Yes (2)	No	No	No
<b>MONITORING</b>						
<b>Traffic accounting</b>	Yes	Yes	Yes	Yes	Yes	No
<b>Remote logging</b>	Yes	Yes	No	No	Yes	No
<b>Email alarms</b>	Yes	No	Yes	No	No	No
<b>MISC</b>						
<b>Admin system</b>	Web-based	Web-based	Windows client	Web-based	Linux client (3)	Java client
<b>Online updates</b>	Yes	Yes	Yes	Yes	No	Yes

### NOTES

1 Office edition supports 2, Business 8, and Professional 16

2 Optionally supports Viruswall, Antivir of Sophos anti-virus products

3 Only SuSE Linux supported

# 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

ElectricSheep	38
Fluxbox	39
Nano	40
ntop	41
GRUB	42

## 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!



## ENTERTAINING TIME WASTER

# ElectricSheep

■ VERSION 2.1 ■ WEB [www.electricsheep.org](http://www.electricsheep.org)



**Do penguins dream of electric sheep? (With apologies to Philip K Dick).**

There are times when you need to run something just because it looks good on your desktop for its eye candy appeal, and *electricsheep* definitely fits into this category. This software owes its name to Philip K Dick's novel *Do Androids Dream of Electric Sheep*, which was the basis for the *Bladerunner* film. In the words of the creator of *electric sheep* "It realises the collective dream of sleeping computers from all over the Internet." Bizarre but true. *Electric sheep* can be run in two modes, either as an animated background for your desktop or as a screensaver. It can be integrated into xscreensaver quite easily and when the screensaver is activated, the screen goes black and an animated 'sheep' appears. "What's a sheep?" you may well ask, well a sheep is basically an animated fractal

flame for your to gaze at in awe. But what happens behind this image is quick interesting. When the *electric sheep* is activated it contacts an internet server and joins the parallel computation of a new sheep as every fifteen minutes, 24/7, a new sheep is produced and distributed to all clients for display. Therefore your computer aids this computation process.

## Demanding

It has to be said that *electricsheep* is really only recommended if you have a high bandwidth, always-on connection to the Internet such as broadband, cable modem or an Internet connected network. This is because this resulting fractal flame files can be very big, usually running at 2.5MB in size, and obviously you need to have the bandwidth to transfer such

files. Also as a client of the *electric sheep* network your computer renders JPEG frames and then uploads them to the server. Once all the frames are ready the server compresses them into MPEG. This is the MPEG sheep that is then downloaded and displayed in all its glory. The MPEG format used is low bit-rate MPEG-1 video codec, displayed at a resolution of between 256 by 256 pixels and 320 by 320 pixels. There are plans to use a DVD-quality MPEG-2 codec and to increase the resolution to 640 by 480 pixels.

Obviously even at this relatively low bit-rate you would not want to store too many sheep on your computer. You can store up to 100 sheep, but this is a lot of disk space (about 200MB). Luckily, you can restrict the amount of disk space that your sheep will use and once that limit is reached, the oldest sheep files will be removed.

The project itself is designed to create sheep that are attention vortexes. What's that then? Well apparently the sheep created illustrate the process by which the longer and closer one studies something, the more detail and structure appears. So in a nutshell, the more you look the more you should see. The shape of a sheep is specified by a string of 65 numbers which is chosen at random by the server which uses some simple heuristics (a form of trial and error) to avoid malformed sheep. If you want just to view the images created, the website allows you to download the current sheep as well as monitor the rendering of new ones.

Another one of the ideas of this project is to "investigate the rôle of experiencers in creating the experience. If nobody ran the client, there would be nothing to see." In other words as more clients join, more computational muscle becomes available, and the resolution of the graphics may be increased, either by making the sheep longer, larger, or sharper. So the more clients running, the better the graphics look.

So slightly pointless, but fun nevertheless, and sure to get people asking you "what the hell is that??"

## WINDOW MANAGER

## Fluxbox

■ VERSION 1.10 ■ WEB <http://fluxbox.sourceforge.net/>

Many window managers have I used in the past, in my quest to find one with which I felt most at home. I've tried the fully functional environments such as KDE and GNOME, to the bare bones, stripped out managers such as *Blackbox* and *IceWM*. I settled with *WindowMaker* for a while and then someone mentioned *Fluxbox* to me. *Fluxbox*, I was told, is a fork of the *Blackbox* window manager, but with a few bits added to make life slightly easier. I tried *Blackbox* before but it gave too much away in functionality to be useful for me, but overall I quite like its approach. So *Fluxbox* interested me enough to give it a go.

The differences between *Fluxbox* and *Blackbox* are not immediately apparent as visually they look the same when you first run *Fluxbox* – however there are many, both visual and under the bonnet. However, when you start an application such as *xterm* the first visual difference is the window tab which, by default, is located in the top left. The tab is not there just to look pretty but adds a useful function that allows you to group windows together and to switch between them using the relevant tab, much in the same way as you can run multiple shells within one window under KDE. With *Fluxbox* tab function you can group any type of window together that you require. To group two or more windows together simply drag the tab from one window onto another.

Window tabs are not the only difference, as *Fluxbox* provides support for KDE, GNOME and *Windowmaker* dock applications. Support for docked applications is well implemented, and works just as you would expect. *Fluxbox* also adds a native *keygrabber* to the base functionality it inherits from *Blackbox*. The *keygrabber* gives the ability to not only handle keyboard shortcuts for common window manager commands, but it also allows sophisticated multi-key

sequences to be assigned in a manner similar to *Emacs*. On *Fluxbox*'s website there is a conversion utility for users of *Blackbox* who use *BbKeys* to help ease the migration to *Fluxbox*.

## Key grabber

I find the *keygrabber* function very useful as it allows me to configure key combinations to start up applications I use all the time, such as a terminal or *Mozilla*. As I like to have a clean desktop, without icons all over the place, the *keygrabber* function provides an easier method of activating application and switching between desktop and windows than a mouse would, at least for me. Unlike some other window managers the *keygrabber* is very easy to configure. You can edit manually the keys file that is used by the *keygrabber*, which is located in your `.fluxbox` directory, however there is also a gui based configuration tool available for


*Fluxbox* called *Fluxconf*. This additional utility, which is not provided with the *Fluxbox* source, allows you to edit your *keygrabber* preferences and also to change some of the behaviour of *Fluxbox*.

Although it is just as easy to configure the behaviour of *Fluxbox* via the desktop menu, the ability to configure the *keygrabber* function via the gui is useful. However, unlike some other gui window manager configuration tools, you have to restart *Fluxbox* to get *Fluxbox* to activate the changes you have made via *Fluxconf*, which makes this tool less useful than it first appears. But as you make this type of configuration change very infrequently you may find it useful to download and install.

One of the main reasons I did not use *Blackbox* is because of the lack of any iconification which drove me nuts. With *Blackbox* when you iconify a window you can only restore it by using the middle button and selecting the window you wish to restore, which is a bit cumbersome. However, with *Fluxbox* when you iconify a window it is displayed on the taskbar, and to restore it you just use a single click. Much easier and far more efficient.

Obviously there is still a lot that has not changed, and all *Blackbox* themes are 100% compatible with *Fluxbox*. If you are a current *Blackbox* user, the change to *Fluxbox* is almost invisible. You will, obviously, have to add your collected *Blackbox* themes to your new *Fluxbox* configuration and then select the theme you wish to use. This compatibility gives *Fluxbox* a large collection of themes, even though it's a relatively new project.

Overall, *Fluxbox* makes a very nice set of enhancements to the "no-nonsense" style of *Blackbox*, and I like it so much it is now my default window manager. If you are a current *Blackbox* user, *Fluxbox* is worth looking at as it keeps the speed and stability of *Blackbox* and the additional features do make it an easy window manager to use. If you are a user of another window manager who feels that you could get more out of a simpler window manager, but were put off by *Blackbox*, *Fluxbox* is worth looking at as it does provide a more nimble environment without giving up some of the features you are used to in other window managers.



Fast, light, efficient – and you can use *Blackbox* themes – what more do you need from a window manager?

## TEXT EDITOR

# Nano

■ **VERSION** 1.0.9 ■ **WEB** [www.nano-editor.org](http://www.nano-editor.org)

**V**i and I go back a long way, I can use *vi* with my eyes shut, it's all automatic to me now.

Even this very article is written using *vi*, honestly. However, recently I was playing around with the Gentoo Linux distribution and its default editor is *nano*, which I'd never heard of. In fact, until the system is completely installed that is all you have, no *vi*. So I had to use it, but was pleasantly surprised by it. *Nano* is, of course, just another editor but with some differences to the normal method of working. Besides basic text editing, *nano* offers many extra features like an interactive search and replace, goto line number, auto-indentation, feature toggles, internationalisation support, and filename tab completion. However the first thing that strikes you is that, unlike most other editors, *nano* is a modeless editor. This means that you can start typing immediately to insert text, with no need to enter an insert command before you do so.

As per other editors you also have many other functions such as search, cut and paste, replace and so on. The most common functions are available via Control key sequences which are

entered by holding down the **Ctrl** key and pressing the required letter. In addition you also have **meta** key sequences which can be entered in a number of ways, depending upon the type of keyboard you have. You can either use the **Esc** key, releasing it and pressing the desired key, or holding down **Alt** while pressing the desired key. **Meta** key sequences provide a lot of additional functionality such as enabling or disabling auto indentation and enabling or disabling regular expression support.

## Titlebar

*Nano* has a titlebar, displayed at the top of the screen. This is broken down into three sections, left, middle and right. The left displays the version of *nano* being used, not exactly terribly useful, the centre section displays the current file name, or "New Buffer" if the file has not yet been saved. The section on the right will display "Modified" if the file has been modified since it was last saved or opened which serves as a useful reminder.

The shortcut lists are the two lines at the bottom of the screen which show some of the more commonly

## Other console text editors

### More editors to look at

**Vim (Vi improved)**

<http://www.vim.org/>

One of a couple of *vi* clones that have added additional functions to the already endless list, but probably the best one out there.

**Dav (Dav Ain't Vi)**

[dav-text.sourceforge.net](http://dav-text.sourceforge.net)

Aims to be more intuitive than *vi*, if you find the *vi* approach odd this might be worth looking at.

**Minimum Profit**

[www.triptico.com/software/mp.html](http://www.triptico.com/software/mp.html)

Designed with the programmer in mind,

with syntax highlighting, context-sensitive help and multiple simultaneous file editing.

**EDI**

[puma.dkz.de/puma/download/edi/](http://puma.dkz.de/puma/download/edi/)  
Syntax highlighting for Fortran and C and *Wordstar*-like key binding.

**Zed**

[zed.c3po.it](http://zed.c3po.it)

Another editor with text highlighting, multi window and multi buffer and macro functions. Apparently "If you cannot stand *vi*, this is the editor for you", so there you go.

used functions in the editor. The statusbar sits above the this and shows important and information messages, which is quite useful. Any error messages that occur from using the editor will appear on the statusbar.

Something a little different, when compared to other editors is the built in file browser. When you are reading or writing files, you can start the file browser by pressing **Ctrl T**. From the browser you can then navigate directories much like a graphical file browser enabling you to find the file you want. To move with the file browser you use the arrow keys and page up/down. The behaviour of the **Enter** key varies depending upon what

you have currently selected. If the currently selected object is a directory, the file browser will enter the directory and display the contents. If a file is selected then this filename and path are copied to the statusbar and the file browser is exited. Very easy to use and navigate.

Cutting and pasting is no more complex than the *vi* method. To cut a single line you use **Ctrl K** and the line will be removed. To paste it you simply move the cursor to where you want to paste and use **Ctrl U** and the line is pasted in at that point. To move a whole paragraph is a case of simply cutting each line one at a time, and then using paste to insert the whole paragraph elsewhere in the file. Nice and simple. For more control you can also mark areas within the file much in the same way as you would do with *vi*. Searching for a string is straight forward enough. Simply use **Ctrl W** and type in your search string. It is all fairly logical and well thought out. Another nice function is the online help which is useful when you first start using *nano* to help you get your head around the various key sequences. After awhile you find yourself refer to it less and less, but its a nice function to have available.

*Nano* is a very functional editor that deserves more attention. With most of the abilities of other editors, and then some useful additions, such as the file browser, *Nano* stands out from the crowd of *vi* clones. Will *nano* replace *vi* as my favourite editor? Probably not, but nevertheless it is still a capable editor and worth looking into if *vi* or *emacs* is not your cup of tea.

```

GNU nano 1.0.8           File: hints_and_tips           Modified
|
find . -exec grep -il ps- "{}" \; to grep for something (ie ps-) from current d
find / -type f -print | args grep -il planet
find /disk2/tmp -type f -atime +10 -exec rm {} \; - to remove files that have r
find /home/ssg -name ,rhost -exec rm {} \; - remove ,rhosts from dir under ssg
find / -mtime <# of days> - find files older than days specified

Debian/Linux specific commands
=====


dpkg -S <path-to-command> to see which package it belongs to
apt-cache search [package name] - search for a package
discover - Progeny's hardware detection system
kudzu - another (Redhat/Mandrake) hardware dection system ported to Debian
procinfo - displays system status gather from /proc, can be timed with -n flag

Controlling the console
=====

^G Get Help  ^O WriteOut  ^\ Replace   ^Y Prev Page ^K Cut Text   ^C Cur Pos
^X Exit      ^R Read File ^W Where Is ^V Next Page ^U UnCut Txt ^T To Spell

```

The world of text editors is not a glamorous one, but they are our main interaction with \*nix.



**ntop can give you network information in a colourful, easy-to-understand form, through your browser.**

## NETWORK TRAFFIC TOOL

# ntop

■ VERSION 2.1 ■ WEB [www.ntop.org](http://www.ntop.org)



**N**top is a network traffic probe that shows network usage, mimicking the *top* command but for network traffic instead. There are two ways to run *ntop*, either as an embedded web server or as a network shell called *intop* (interactive *ntop*). *ntop*, like most other network utilities, uses the popular *libpcap* packet capture library so you need to ensure this is on your system before installing *ntop*. This will usually be taken care of by your distributions packaging system, but if you intend to build from source you will need to install this first.

Once installed, running *ntop* is very simple, with only a few flags to worry about depending upon your systems configuration. If you only have one interface you can simply run *ntop*, but if you have more than one interface you'll want to use the **-i** flag to indicate which interface is to be used to monitor from. Also useful is the **-d** flag

which informs *ntop* to run as a daemon. Bear in mind that as *ntop* runs the interface in promiscuous mode so it need to be run by root, therefore you should be careful where you run *ntop* from a security point of view.

Once *ntop* is running you can now start looking at the statistics it is gathering. To access *ntop* from the command line you type the command **intop -i eth0** (you must specify the interface to use) which will take you to the prompt : **intop@eth0>**

There is a basic command online help available by entering the question mark which displays the various commands that you can use. For example, if you enter **top** you are presented with a listing of the top talkers on your network, displayed much in the same way as the *top* command does. The numbers that are displayed are dynamic and

connections will disappear from the list as different system stop communicating. In addition to the command listing you can also get keystroke help by pressing the **h** key. For instance, you can switch from local to remote traffic modes by pressing **l**. There are several keystroke commands available to you so it is worth investigating them further.

## Static statistics

The only disadvantage with the web interface to *ntop* is the fact that it does not update, you have to refresh the page to retrieve new statistics. This means that you cannot use the web interface as a monitor you leave

```

ntop 0.0.1 (Aug 2 2002) listening on [eth0]
1680 Pkts/309.5 Kb [IP: 217.5 Kb/Other: 92.0 Kb] Thpt: 12.3 Kbps/33.2 Kbps

```

Host	Act	-Rcvd-	Sent	TCP	UDP	ICMP
10.5.98.180	B	746	652	2.0 Kb	866	0
10.5.98.106	R	728	0	0	589	0
10.5.98.122	B	92	86.0 Kb	3.5 Kb	10.3 Kb	0
10.5.98.2	B	56	3.0 Kb	0	0	0
10.5.98.106	S	0	4.9 Kb	0	0	0
10.5.98.52	S	0	460	122	0	0
10.5.98.2	S	0	9.6 Kb	0	0	0
10.5.98.1	S	0	10.0 Kb	0	0	0
10.5.98.131	S	0	6.0 Kb	0	0	0
10.5.98.102	S	0	89.0 Kb	0	0	0
10.5.98.65	S	0	10.4 Kb	0	0	0

**Dynamically updated statistics are also available from the command line.**

running all day to alert you to any networking issues. *intop*, the command line based interface, does update, but it is not as useful, or as pretty, as the web interface. If all you are looking out for are hosts that are using excessive bandwidth then the *intop* interface will provide you with this ability.

In addition to the command line interface you can also access *ntop* using the embedded web server. By entering the name of the system running *ntop* in your web browser followed by **:3000** you can access this interface. Port 3000 is the default port used by *ntop* but you can easily change this by adding the **-w** flag following by a port number that you prefer to use. The web based interface is well laid out, with data displayed in a variety of ways, some tabular, some graphic, including useful charts and graphs. The web interface is perhaps easier to use as it is all point and click based and provides you with all the info that the command line interface does and then some. Additions include the ability to click on a host within the displayed list to get access to more in depth info about that host, such as the OS it is running, whether it uses DHCP to get its IP address and so on.

*Ntop* is a very useful network analytic tool and both its command line and web based interfaces are very well implemented. It is very easy to gain a good understanding of your network's signature from *ntop* and allows you to track individual systems should you require this. An important note is that if your network uses a switch to provide connectivity you will have to arrange for the port your *ntop* box is connected to to be set up in mirroring mode to ensure that you can see the whole network and not just your individual segment. This is not an issue with *ntop*, but a general note for any network monitoring. *Ntop* is a very impressive tool and worth reviewing to see if it meets your requirements.

## BOOT LOADER

# GRUB

■ **VERSION** 0.91 ■ **WEB** [www.gnu.org/software/grub](http://www.gnu.org/software/grub)

**G**RUB is a very powerful boot loader that can load a variety of operating systems such as Windows, DOS, Linux, GNU/Hurd, \*BSD and so on. A boot loader is a program that resides in the the MBR (Master Boot Record) of your hard disk. After the usual self tests that your system performs the BIOS transfers control to the MBR. Then the program residing in MBR gets executed. It is this program that is called the boot loader. Its main function is to transfer control to the required operating system, which will then proceed with the boot process. Currently *Lilo* is the most popular boot loader but *GRUB* is catching up fast.

*GRUB* currently supports *ext2*, *ext3*, *FAT16* and *FAT32*, *FFS*, *ReiserFS*, *XFS*, *JFS* and *minix*. As *GRUB* natively supports these filesystem it has the ability to look at a filesystem without the need to boot into that operating system. To access this function you need to be in the *GRUB* shell, which you access by typing **grub** at the command line :

```
jkent@smokey:~$ grub
```

```
grub> cat (partition
number)/home/jkent/hotpicks.txt
```

```
grub> quit
```

You will then be presented with the contents of the file.

## Flexible

As you can already see *GRUB* is far more flexible and powerful than *Lilo*. Another good example of the flexibility of *GRUB* is that you can load any kernel on any partition. Usually if you forget to add a newly compiled kernel and you are using *Lilo* that would mean that you would need to boot, add it to the list, run *Lilo* and then reboot to access it. But with *GRUB*, you can simply use the shell, which is available to you on bootup, and load the desired kernel image.

The most critical part of understanding *GRUB* is getting comfortable with how *GRUB* refers to hard drives and partitions. Let's say that your Linux root partition is */dev/hda1*. Under *GRUB* this is referred

to as (hd0,0). Notice the parenthesis around the (hd0,0) these are important and required. The hard drives start from zero rather than "a", and partitions start at zero rather than one. So */dev/hdb3* would get translated to (hd1,2). Once you've got your head around that you are ready to install *GRUB*.

## Installation

As per *Lilo*, *GRUB* is installed into the */boot* partition, so you need to ensure that you backup any existing files in your */boot* directory. The easiest way is to rename the */boot* directory to something like */boot.ORIG* and then create another directory. In fact there is no need, in day to day use, for the */boot* directory to be readily available, so you can use another partition to contain you *GRUB* information and mount it onto */boot* whenever you need to add a new kernel image or change the *GRUB* boot menu. Having a separate partition for */boot*, when using *GRUB*, is a good idea as it will protect it from any accidental damage. You can quite happily use *ext3* on this partition as well to add an extra degree of protection.

Firstly you need to download and install the relevant *GRUB* package for your distribution and follow any relevant instructions for that package. You need to ensure that the following files, at minimum, are installed in your */boot/grub* directory:

```
stage1
stage2
*_stage1_5
```

To start the installation of *GRUB* you will need to start the *GRUB* shell. From within the *GRUB* shell you will install the *GRUB* boot record onto your hard drive. For this example lets assume that you want to replace you MBR (located on */dev/hda1*) with *GRUB*'s. To action this you would type in the following commands:

```
jkent@smokey:~$ grub
```

```
grub> root (hd0,0)
```

```
grub> setup (hd0)
```

```
grub> quit
```

The first command informs *GRUB* where the boot partition is, so for the

example of */dev/hda1*, this translates to (hd0,0) with *GRUB*. The setup command informs *GRUB* where to install the boot record, in this example the boot record will get installed on *hda*, which is *hd0* in *GRUB*. Once these commands have been completed you simply type **quit** to exit the *GRUB* shell.

Now you need to copy your kernel into the */boot* directory and then create a file in */boot/grub/* called *menu.lst* to provide boot menu when the system reboots. The *menu.lst* file is very simple, for example:

```
default 0
```

```
timeout 20
```

```
# Pretty colours
```

```
color cyan/blue white/blue
```

```
#
```

```
title Windows 2000
```

```
root (hd0,0)
```

```
makeactive
```

```
chainloader +1
```

```
#
```

```
title Debian 3.0
```

```
root (hd0,3)
```

```
kernel /boot/bzImage
```

```
root=/dev/hda3 ro
```

If you are using a separate partition and using the example *menu.lst* above you will also need to create a symbolic link in the root of the partition to link boot to the root of the partition, or remove the trailing */boot*.

If you are dual booting into Windows 2000 and Linux the **chainloader +1**

is necessary for *GRUB* to go into Win2K's loader. Be careful with the spacing – it's chainloader space + no space 1 otherwise it will not boot into Windows 2000. If you are feeling paranoid and want to check that your partition entries are correct you can use the *GRUB* shell to verify this.

```
jkent@smokey:~$ grub
```

```
grub> root (hd0, [press tab])
```

```
Possible partitions are:
```

```
Partition num: 0, Filesystem type
is fat, partition type 0xc
```

```
Partition num: 1, Filesystem type
unknown, partition type 0x82
```

```
Partition num: 2, Filesystem type
is ext2fs, partition type 0x83
```

```
Partition num: 3, Filesystem type
is ext2fs, partition type 0x83
```

```
grub> quit
```

Depending upon you distribution packaging of *GRUB* you may find that you have a file called *splashimage.xpm.gz* located in your */boot/grub/* directory. This can be used to provide a splash image (in case you hadn't guessed) on the *GRUB* menu screen. If you do have this file and want to see what this looks like you simply as the following line to your *menu.lst* file usually below the timeout line.

```
splashimage=(hd0,3)/boot/
grub/splash.xpm.gz
```

If you do use a splash image, you will probably find that you do not need the **color** parameter defined, but if you do not use a splash image the **color** parameter makes the *GRUB* menu look a little less bland. To find colours that suit your taste simply play around with the settings until you are happy with the results.

## Options


The default setting within the *menu.lst* indicated which operating system you wish to set as the default to use once the timeout limit has expired.

In the example above the default operating system would be Windows 2000. To change this to Linux you would change the default to 1, or change the positions of the two operating system around if you prefer. The timeout parameter uses seconds to define how long to wait, in this case 20 seconds, which is probably more than enough time to decide which operating system to boot.

To pass any options to the kernel, simply add them to the end of the kernel command. In the example above one option (*root=/dev/hda3*), is being passed to the kernel. Another common option is if you set up scsi emulation for an IDE cd burner, in which case you would pass **hd[x]=ide-scsi**, replacing x which the device for you cd burner.

Another difference between *Lilo* and *GRUB* is that when you edit the *menu.lst* you do not need to re-run *GRUB*, as this file is referred to automatically. With ease of installation, powerful command functions from the boot menu and from the command line *GRUB* really shows how restrictive *Lilo* is these days. If you do not use *GRUB* it is a must have installation. Install carefully and you will be amazed that you managed without it. [LXF](#)

cover feature



# Building better databases

Get some tips from a pro – **Jeremy Cole** tells us how to get the most out of MySQL.



**S**o you have *MySQL* installed on your servers, your company has been using it for quite a while now, and everyone is happy. How would you know if there are some features of *MySQL* that you could use but you don't yet know about? How can you tell if you have some deep lurking problem that will pop out at you later?

In this article, we will cover everything from the basics to the advanced, and hopefully tell you what it takes to make effective use of the best and most widely used Open Source database out there, *MySQL*!

## What can MySQL do?

It is a common misconception that *MySQL* is only for the small stuff. In reality, the truth is quite the contrary. *MySQL* was designed from the start to be the best at handling mission-critical applications, huge amounts of data, to be extremely reliable and just plain blazingly fast. Here are some common categories of *MySQL* users, and some example applications:

**WEB** Everything from shopping carts to catalogues, sessions to customisable webpages. *MySQL* is used every day in most of the popular websites out there. Some of the most common uses of *MySQL* in this area would be to store user information, and to store data for building webpages.

**LOGGING** So you need to keep stock quote tickers polled once every few seconds, for hours a day, indefinitely. That would be an enormous amount of data, but *MySQL* is up to the task. What other database could be better for speedy analysis of all this data you've spent so much time collecting?

**BACKEND** Inventory Control. *MySQL* is quite commonly used for keeping track of warehouse inventory, as well as in various backend product fulfilment systems. The reliability and speed of *MySQL* make it a perfect database for any sort of inventory control.

**FRONTEND** Point of Sale. OK, so you don't see too many *MySQL* command line clients running on cash registers, but perhaps the next time you buy a book or a pair of shoes, that transaction will be back-ended in

*MySQL*. There are already several companies out there running real-time point of sale applications on *MySQL*.

So you get the idea, right? The applications for a general-purpose database like *MySQL* are practically endless. *MySQL* really is good for just about anything.

## Popular MySQL tools

The *mysql* command line client: simple, yet powerful.

We'll take a look at the *mysql* command line client first, mostly because it comes with the *MySQL* product itself. The command line client is the only utility included with the *MySQL* distribution that allows you to issue arbitrary queries. Let us take a look at how it works and what you can do with it.

The *mysql* client is similar to a command shell. It gives you a prompt, you enter a query terminated by a special character (;) and it sends that query to the *MySQL* server, *mysqld*. Upon completing the query, *mysqld* returns the result to *mysql*, and *mysql* displays that result on the console. Seems simple, right? Even with such a simple client, things do not have to be so simple. Most SQL command line clients are rough patched together things that seems to barely work. *MySQL*s, however, is a bit different.

*MySQL*s command line client uses the *readline* library from GNU. Thus, command editing is easy, consistent, and works really well. To make your life even easier, there are numerous time-saving features tucked away that even most experts might not know about. We'll go over a few of them here:

**VERTICAL DISPLAY** One of the most useful features of the *mysql* command line client is the \G option. If you've ever issued a **SELECT** \* query on a table with more than ten or so columns, you might know how ugly the output can get.

If you use \G (remember the capital 'G') to end the query instead of ;, *mysql* will show the results in a more convenient vertical format, like so:

```
mysql> select * from airports where
code="SAN" \G
***** 1. row *****
code: SAN
```

```
fac_type: AIRPORT
fac_use: PU
faa_region: AWP
faa_dist: NONE
city: SAN DIEGO
county: SAN DIEGO
state: CA
full_name:
SAN DIEGO INTL-LINDBERGH FLD
own_type: PU
longitude: -117.189657
latitude: 32.733556
elevation: 14
aero_cht: LOS ANGELES
cbd_dist: 2
cbd_dir: W
act_date:
cert: DS 05/1973
fed_agree: NGSY
cust_intl: N
c_ldg_rts: Y
joint_use: Y
mil_rts: Y
cntl_twr: Y
1 row in set (0.00 sec)
```

Wow! Obviously, vertical output is most useful if you are selecting many columns, but few rows. In this case, the output's clarity is unbeatable. **EDIT BUFFER** You can use the \e command to edit the contents of your current buffer. The \e command is useful in many scenarios. The most common would probably be to allow you to correct any mistakes you've made when typing multi-line queries. However, it is also useful if you'd just like to use your favourite editor to write the query in the first place.

**“You have MySQL installed. How can you tell if you have some deep lurking problem that will pop out at you later?”**

Just type \e at an empty prompt, and you will get a blank editor to type away in. Keep in mind that the \e feature uses your environment variable **EDITOR**, so you can easily customise the editor that you get to be your favourite editor. If it is not set, you will probably get the system default (on most Linuxes that would be vi). You can choose your editor quite easily:

```
bash$ export
EDITOR="xemacs -nw"
```


## Database makers

### What is MySQL AB?

- A virtual company
  - The creators of *MySQL*
  - Owners of *MySQL* code and *MySQL* trademark
  - 50+ employees in 14 countries
  - Committed to GPL
  - Dual licensing: Commercial licenses avoid GPL restrictions
- [www.mysql.com/company](http://www.mysql.com/company)

# Databases

## The *mysql* Command Line Client.



```

mugatu.spaceapes.com - mugatu* - SSH Secure Shell
(jcoles@mugatu) [-] $ mysql
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 36 to server version: 4.0.2-alpha-max

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> prompt \u@h [\d]>
PROMPT set to '\u@h [\d]> '
jcoles@localhost [(none)]> use world;
Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A

Database changed
jcoles@localhost [world]> select Country.name, City.name from Country left join C
ity on Country.capital=City.id where Country.region like "Nordic?";
+-----+-----+
| name          | name          |
+-----+-----+
| Faroe Islands | Tórshavn      |
| Svalbard and Jan Mayen | Longyearbyen |
| Iceland       | Reykjavik     |
| Norway        | Oslo          |
| Sweden        | Stockholm     |
| Finland       | Helsinki [Helsingfors] |
| Denmark       | København     |
+-----+-----+
7 rows in set (0.00 sec)

jcoles@localhost [world]>

```



**RESULTS PAGER** It can happen that you would like to select a several hundred, or even several thousand rows. You issue the query, and of course the result will come flooding back to your console. You will not get a chance to look over more than the last few records, other than by using your scrollbar buffer.

If instead, you issued the command **pager less** before issuing the query, the results of the proceeding queries would be sent through the filter that you've defined, the *less* program. You can turn it

current command shell is using. As of *MySQL* version 4.0.2, you can change the **mysql>** prompt to be whatever you would like it to be. You will have a choice of a few dozen different specifiers in order to build your own custom prompt. A good choice would be **\u@h [\d]>**, which would make your prompt something like:

```

mysql> prompt \u@h [\d]>
PROMPT set to '\u@h [\d]> '
jcoles@localhost [world]>

```

## Perfect for Web

PHP is a relative newcomer to the world of programming. Don't get me wrong, though, that doesn't mean it is not important. PHP has taken the world by storm, and the storm isn't about to let up. There are many advantages to using PHP over others.

There are already many tutorials about how to use PHP with *MySQL*, so I won't bore everyone by repeating the same things here. There are some finer points of optimisation with PHP and *MySQL*, but they are outside the scope of this article. For more information on PHP, please visit the PHP website at [www.php.net](http://www.php.net).

## Swiss army knife

Perl has been around for a long time, and is used all around the world for just about every imaginable task. Some of the more common uses for Perl are web scripting, parsing and transforming text, which is where it fits into the picture here.

One of the great strengths of Perl is the Perl motto: "There's more than one way to do it" When learning Perl, though, that very same strength can become a huge weakness. The fact that there are multiple ways to do absolutely everything in Perl makes it hard for anyone learning the language to know which one is the "right" way.

In particular, when working with Perl and *MySQL*, most people don't have any idea where to start, or what would be considered good programming practice. Hopefully we can make a few things a bit more clear.

## Retrieving results

**fetchrow\_hashref** is one of the most useful methods for retrieving a result set in Perl. For some reason, however, it is often overlooked, and not widely known. To understand the whole situation, you have to take a look at the methods available for retrieving data:

### fetchrow\_array

The **fetchrow\_array** will retrieve a result row and store it into an array, returning that array. That seems nice, but in reality, you will have to copy the result into another array to work with it, so you could waste a bit of memory as well as CPU. There is a backwards-compatible alias for **fetchrow\_array** called just **fetchrow**. Your code would look something like this:

```

while(my @row =
    $sth->fetchrow_array) {
    print "$row[0] is $row[1] years
    old.\n";
}

```

### fetchrow\_arrayref

The **fetchrow\_arrayref** method is similar to **fetchrow\_array**, but instead of returning an array (which would then have to be copied), it will return a reference to the array that it has already created, thus saving the step of copying the data into a local array. The real disadvantage to using **fetchrow\_array** or **fetchrow\_arrayref** is that you will have to reference the columns in the result by their number, rather than their name.

```

while(my $row =
    $sth->fetchrow_arrayref) {
    print "$row->[0] is $row->[1]
    years old.\n";
}

```

**“PHP, a relative newcomer to the world of programming, has taken the world by storm and the storm isn't about to let up.”**

back off again by issuing the **nopager** command.

**DYNAMIC PROMPT** Most database administrators will at some point have to be connected to multiple *MySQL* servers or will use different databases on the same server at the same time. When doing so, it can be very hard to tell which window is which, and you can easily get confused and issue commands on the wrong server. That might be harmless, or it might be really harmful.

Using the **prompt** command is an easy way for you to always know which *MySQL* server or database your

## More information

### Where to go next

#### BOOKS

*MySQL Reference Manual* by Widenius, Axmark, and MySQL AB Published by O'Reilly Community Press

*MySQL* by Paul DuBois Published by New Riders

#### WEBSITES

[www.mysql.com](http://www.mysql.com) – The MySQL AB web site. Nearly anything you could want to find out about MySQL you can find on MySQL.com. Most popular: Downloads and Documentation.

[www.mysqldeveloper.com](http://www.mysqldeveloper.com) – MySQLDeveloper.com has quite a few useful features, including an FAQ, Live Java Chat with the #mysql IRC channel on EFnet, and many other things.

[www.mysqlfaqs.com](http://www.mysqlfaqs.com) – A quick way to get to the Frequently Asked Questions section on MySQLDeveloper.com.

[www.sqlcourse.com](http://www.sqlcourse.com) – SQLCourse.com is a web site devoted to teaching the basics of SQL. It isn't exactly MySQL-specific, but the majority of the concepts you will learn there apply to MySQL.

[www.php.net](http://www.php.net) – PHP.net is the place to find anything you could possibly want to know about PHP. The PHP documentation is some of the best documentation for any software out there.

An extra tip for using PHP.net is: if you want the documentation section for e.g. `mysql_connect`, try [http://php.net/mysql\\_connect](http://php.net/mysql_connect).

[www.perl.com](http://www.perl.com) – The Perl community doesn't really have a single comprehensive source for everything you need to know about Perl (there's more than one way to do it, remember), but Perl.com comes pretty close.

```
}
Using fetchrow_arrayref will save some memory and a little CPU time, but the code certainly will not get any easier to read.
```

#### `fetchrow_hashref`

The `fetchrow_hashref` method will retrieve a result row and store it into a hash, where the hash keys are the names of the columns in the result. Instead of returning that hash, which would require quite a bit of copying, it will return a reference to the hash. For example:

```
while(my $row =
    $sth->fetchrow_hashref) {
    print "$row->{name} is
    $row->{age} years old.\n";
}
```

So, as you can see, using `fetchrow_hashref` makes the resulting code much easier to read, without any added complexity. Since it's using a reference to a hash, there is no additional copying involved, either.

## List assignment

Most queries that you will find yourself issuing to MySQL will likely only access a very limited set of columns. When selecting a small number of columns, you can make your code a bit simpler and easier to read by using a list assignment

construct like this, in Perl:

```
while(my ($name, $age) =
    $sth->fetchrow_array) {
    print "$name is $age years old.\n";
}
```

Or, the same thing in PHP, using the `list()` function:

```
while(list($name, $age) =
    mysql_fetch_array($result)) {
    print "$name is $age years old.\n";
}
```

## Abstraction

After you've written several applications using any language, you will realise that you find yourself

writing the same code to do the same thing, over and over. Creating some functions to do those simple, redundant tasks for you, will save you a great deal of time in the long run. For example, you could create a `db_connect()` subroutine to connect to the database (if you're not already connected).

Another good example is creating a `db_query()` function as a wrapper to the normal Perl routine of prepare and execute, and return a statement handle.


It's still unclear at this point whether any of the prebuilt database abstraction layers are really a Good Thing™. In the words of Rasmus Lerdorf, creator of PHP, "lowest common denominator"-style database abstraction layers are bad, because "the lowest common denominator in SQL is pretty damn low".

You can try some of them, if you like them use them. If you don't, forget about it, no harm done either way.

## Choose tools wisely

One of the most important things, that I can't stress enough, is: Don't be afraid to mix technologies! Use the best tool for the job, don't write something in PHP just because your web site uses PHP, if Perl is better suited, use Perl!

If you could do the same thing very easily just by creating a small SQL script and running it through the `mysql` client, do so. Personally, I find it very comfortable to use PHP for any web programming, but Perl for any



**MyCC**, a GUI client for **MySQL**, produced by **MySQL AB**.

# Databases



system scripts and backend, non-Web things. You must find what is most comfortable for you.

## Configuration is key

The default configuration of *MySQL* is quite good for anyone just getting started, but once you reach some point, you will have to create your own configuration file. Configuration of *MySQL* is somewhat mysterious, but once you get the hang of it, and get a good idea what is important and what is not so important, it doesn't have to be a daunting task. A good first start

how large the key buffer is on any *MySQL* server by issuing the following query:

```
mysql> SHOW VARIABLES LIKE
'key_buffer_size';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| key_buffer_size | 67104768 |
+-----+-----+
1 row in set (0.02 sec)
```

The default for **key\_buffer\_size** is 8MB, but *MySQL* AB generally recommends that it should be set to approximately 25% of available system memory for a dedicated *MySQL* server. In the above example, it is set to 64MB. It is possible to tell how well the key buffer is being utilised by a very simple method. You can find the variables you will need to know with the following command:

```
mysql> SHOW STATUS LIKE
'Key_%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Key_blocks_used | 45232 |
| Key_read_requests | 782625974 |
| Key_reads | 43874 |
| Key_write_requests | 1733 |
| Key_writes | 43 |
+-----+-----+
5 rows in set (0.00 sec)
```

The interesting numbers in the above are:

**Key\_blocks\_used** The number of “blocks” of the key buffer currently in use. One block is 1024 bytes.

**Key\_read\_requests** The number of times a key was requested to be read from the index. If a key was requested to be read, but it wasn't actually read (**Key\_reads**), then it must have been found in the buffer.

**Key\_reads** The number of keys that were actually read from the index.

The principal behind finding out how efficiently *MySQL* is using the key buffer is that *MySQL* should find a key in the key buffer as often as possible. You can determine this ratio as a useful number using the following formula to get the cache hit rate, in percent:

```
cache hit rate = 100 - ((Key_reads /
Key_read_requests) * 100)
```

You should aim for 97% key efficiency (cache hit rate) or higher. Typically it isn't too difficult to get the key efficiency to 99.5% or higher. If you find that your key efficiency is low (less than 70% especially) it normally means that the **key\_buffer\_size** should be increased.

You can check how “full” the key buffer is by examining the **Key\_blocks\_used** variable from **SHOW STATUS**. If **Key\_blocks\_used** indicates that the key buffer is nearly full, you may try increasing it as well. If you get a key efficiency of 99.5% or higher, but **Key\_blocks\_used** says that you aren't using anywhere near the entire key buffer, feel free to make it smaller.

## Indexes

Indexes are magical. At least that's what most people would have you believe. In reality, indexing is a very poorly understood concept which really isn't all that complicated. Once you learn a few simple rules about when and why indexes can and cannot be used, the rest is easy.

*MySQL* currently supports B-Tree and Full-Text indexes on MyISAM tables, and Hash indexes on HEAP tables. Typically most users will only ever use the default B-Tree indexes, and for most practical applications, they are exactly the type you should be using. It is, however, important to know the limitations of such an index, and how they work.

Indexes are used internally by *MySQL* in order to be able to retrieve data quickly. If you add an index to a table, that index *might* be able to be used to retrieve data more quickly, but it will always slow down all deletes, all inserts, and some updates. You should choose which columns you index carefully, only creating indexes which can be used.

“Indexes are magical – at least that's what people would have you believe. In reality indexing is poorly understood.”

would be to ask *MySQL* about its current configuration:

```
mysql> SHOW VARIABLES;
```


*MySQL* will reply with a set of variable names and their current values. You should use the *MySQL* manual to determine what should be changed, based on your application and your use of *MySQL*. Some good starting points for configuration files are included in the *MySQL* distributions as the files *my-huge.cnf*, *my-large.cnf*, *my-medium.cnf*, and *my-small.cnf*. You should find those files in the *support-files* directory of a source or binary distribution from the *MySQL.com* website.

**KEY BUFFER** One of the most important variables in the *MySQL* configuration is called **key\_buffer\_size**. You can find out

Like the UNIX *top* program, *mytop* is a useful tool for analysing *MySQL* usage in real time, written by Jeremy Zawodny.

```
MySQL on db.finance (3.23,47-max-log) up 79+08:23:44 [15:37:21]
Queries Total: 2,068,312,373 Avg/Sec: 301.69 Now/Sec: 242.04 Slow: 798
Threads Total: 45 Active: 5 Cached: 0
Key Efficiency: 99.85% Bytes in: 3,713,492,213 Bytes out: 1,006,022,300

  Id      User      Host      DB      Time      Cmd Query or State
---
5827598  yahoo    proc2     IDB      0      Sleep
5779323  fred     feed1     Fred     0      Sleep
5779350  fred     feed1     Fred     0      Query select a.id as id, a.feed_
5829250  yahoo    museful   mysql    0      Query show full processlist
5825442  yahoo    proc2     IDB      0      Sleep
5826226  yahoo    proc3     IDB      0      Query SELECT * FROM Headlines H,
5825441  yahoo    proc4     IDB      1      Sleep
5829234  root     localhost MySQL_Admin 1      Sleep
5779354  fred     feed1     Fred     9      Sleep
4303469  yahoo    proc2     IDB     11      Sleep
4286987  yahoo    biz5     IDB     14      Sleep
5829245  yahoo    proc3     IDB     19      Sleep
5829242  locker   proc4     Finance 19      Sleep
5829246  yahoo    proc4     IDB     19      Sleep
5829244  locker   proc3     Finance 19      Sleep
5829240  locker   proc4     Finance 20      Sleep
-- paused, press any key to resume --
```



Using *mysql-graph*, you can graphically monitor the performance of *MySQL*.

A B-Tree Index stores data in a “Binary Tree” structure. It’s called a binary tree because at each “node” of the tree, there are only two possible choices to continue on. Since “binary” means “having two parts” I’m sure you can make the connection.

Comprehensive information about how binary trees work is easily available online, so I won’t go into that here. It’s only important for us to know what applies to *MySQL*’s usage of B-Trees.

I find it easiest to think of a binary tree dealing with an index of, for instance, numbers 1 through 100. If you put those numbers on a number line, you would find that 50 falls right in the middle. In a binary tree, the median value of a set will be the “top” node. So the tree starts with value 50. Since it’s a binary tree we know that we have two choices to go onwards from there.

Using the same concept of taking the median value and using it for the top node, the left and right branches of the tree would be the values 25 and 75, respectively. The values 25 and 75 are chosen, because they halve the set to the lesser and greater side of 50. We can continue on down the left hand side, the lesser side of the tree to find value 12 (or 13) and down the right hand side, the greater side to find the value 37 (or 38).

The same concept would continue until all values in the set are represented somewhere in the tree. This means that if we were looking for the value 14 in the index, we would start at the top node, 50, and compare 14 to 50, 14 is smaller than

50, so we would take the left branch of the tree. If 14 were larger than the node we were comparing to, we would take the right branch of the tree. We must compare at each node until we reach 14 itself, or we run out of choices, which would mean that the number we are looking for doesn’t exist in the index.

Now that you know how indexes work, the rules for when an index can or cannot be used might even be somewhat obvious. An index can only be used if you know the start of the index value. Some common types of queries are these:

1. **SELECT ... WHERE name LIKE 'Jeremy%'**
2. **SELECT ... WHERE name LIKE '%Cole';**
3. **SELECT ... WHERE name LIKE '%MySQL%';**

In case 1, an index can be used, as the beginning of the indexed value is known. *MySQL* can search to easily find values that start with “Jeremy”, or, literally, are greater than or equal to “Jeremy”, but less than “Jeremz”.

In case 2, an index cannot be used. Since the beginning of the indexed value is not known, the entire index, or worse, the entire table, must be searched. This is known as a “full index scan” if only the index tree is searched, or a “full table scan” in the most common case, when the entire table must be searched.

Lastly, in case 3, the same rule as case 2 will apply. Since the beginning of the indexed value isn’t known, the index cannot be used. For queries of this type, a Full-Text index, which is an index on whole words in a string, could be used. Full-Text indexes use a

special syntax, so the query would still have to be changed.

## Monitoring tools

There are many useful tools out there to monitor your *MySQL* servers. Some of the things that should be considered to be monitored would be: number of queries per second, read/write throughput per second, network traffic, ratio of select, update, and delete, and many others.

**MYTOP** *mytop* is a program that allows you to monitor a *MySQL* server in real time, very much like the Unix *top* program. It is written in Perl, so that makes it rather portable, and easy to install. *mytop* was written by Jeremy Zawodny from Yahoo! Finance, and you can find *mytop*, as well as any other nice tools that Jeremy Zawodny has created at <http://jeremy.zawodny.com/mysql/>. **mysql-graph** *mysql-graph* allows you to graphically keep an eye on a *MySQL* server.

Based on a very easy to use configuration file, *mysql-graph* will generate PNG or JPEG images at given intervals. The graphs can contain any number of variables graphed, and the data to be graphed can come from any source, as *mysql-graph* allows you to define Perl subroutines and custom *MySQL* queries, as well as easily allowing you to graph the values returned from the **SHOW STATUS** command. The *mysql-graph* package is available from the Software Portals section of the *MySQL.com* website. You can download it at: [www.mysql.com/portal/software/html/software\\_comments-36.html](http://www.mysql.com/portal/software/html/software_comments-36.html)

## About the Author

Jeremy Cole has been an employee of *MySQL AB* for two years, during which time he has worked on the documentation team, the core development team, and the training team. He has contributed to most areas of *MySQL AB*, and is currently teaching training classes. You can email him at [jeremy@mysql.com](mailto:jeremy@mysql.com).

## USB2.0

# USB 2.0

Ready for a high bandwidth future, **Richard Drummond** investigates the next generation of USB, and how it works under Linux.

**T**he latest stable Linux kernel, version 2.4.19, was released at the beginning of August, and despite the surprisingly little fanfare made about the fact, it is the first stable kernel to have support for USB 2.0, the next generation version of the Universal Serial Bus standard.

What is USB 2.0? What does it mean for the user? And how can you use it under Linux? Read on for the answers.

### A little history

The Universal Serial Bus or USB was announced in 1994 and eventually revolutionised the way we connect devices to our PCs. Until USB came along, there were a host of different incompatible bus specifications, each

with several different, sometimes proprietary, protocols for hooking up keyboards, mice, modems, scanners, cameras and printers. You had the decades-old RS232 serial port, the parallel (or Centronics) port, PS/2 ports for keyboards and mice, and others – which together formed an I/O system that was slow, inflexible and burdened by legacy compatibility. You were limited to the number of devices you could connect, and configuration was a chore.

USB solved all these problems. It's an open standard, developed by a consortium of companies including Compaq, HP, Intel, NEC and Philips. It's hot-pluggable and truly plug-and-play. USB connectors are designed so that USB devices can be attached and removed with the power on, and a USB device identifies itself to the system when it is plugged in.

It's the job of the controlling software – the USB stack – to find and associate a driver with that device, so no user configuration should be required. USB is easily extensible by devices called hubs,

which multiplex a single USB port and let you connect four or more devices to a single port. Hubs may be cascaded together to form a tree of up to 127 devices on a single hub. Finally, USB offers better performance than standard serial or parallel ports, with a raw transfer rate of 12 Mbit/s (or 1.5 MB/s).

### The need for speed

While USB offered higher speeds than the legacy interfaces it was designed to replace, the computer industry's ever-increasing demand for bandwidth and today's multimedia applications have outgrown USB 1.1. Using USB 1.1, it would take over seven minutes to transmit the contents of a single CD (even ignoring the overhead of the USB protocol itself), while USB 1.1 is quite unsuitable for streaming full motion video. This need for speed has been addressed by the release of USB 2.0.

USB 2.0 is an evolution of USB 1.1 that introduces a new, high speed transfer rate. USB 1.1 defines 'low speed' as 1.5Mbit/s (suitable for keyboards, mice, etc.) and 'full speed' as 12 Mbit/s. USB 2.0, also called Hi-Speed USB, adds to this a 'high speed' rate of 480Mbit/s

**“The computer industry’s ever-increasing demand for bandwidth and today’s multimedia apps have outgrown USB 1.1.”**



(60MB/s), thus potentially offering 40 times the bandwidth. However, USB 2.0 remains completely forwardly and backwardly compatible with USB 1.1. You can connect USB 1.1 devices to a USB 2.0 adapter or hub and they will continue to work at their designated speeds.

This works because a USB 2.0-capable controller actually has both USB 2.0 and USB 1.1 controllers built-in. When USB 2.0 devices are connected, they are assigned to the Enhanced Host Controller Interface (EHCI), but when USB 1.1 devices are connected, they are assigned to a USB 1.1 "companion" controller, either an OHCI (Open Host Controller Interface) or UHCI (Universal Host Controller Interface) depending on the vendor. Additionally, most USB 2.0 devices will be able to fall back to a full-speed rate and so work,

albeit at a slower speed, when connected to a USB 1.1 bus.

### USB 2.0 hardware

USB 2.0 has been on the drawing board since 1999, and although the adoption rate is not yet high, there are a substantial number of USB devices on the market. NEC were the first to market with a USB 2.0-capable controller, and this is widely used in add-on PCI cards from manufacturers such as Trust and Adaptec. Newer and better controllers are beginning to appear from vendors such as VIA and Intel, however, and both these semiconductor giants are now integrating USB 2.0 into their motherboard chipsets. Once motherboards with USB 2.0 on-board become common-place, expect the demand and hence availability of USB 2.0 peripherals to increase.

The biggest market for USB 2.0 devices at the moment is

### USB 2.0 vs Firewire

VHS and Betamax – or will they co-exist?

Many see USB 2.0 as unnecessary, since the Firewire or IEEE1394 bus is already established as a means for connecting high speed devices such as external drives and digital video cameras. However, despite the existing competition, USB 2.0 is virtually guaranteed to succeed. For one, USB 2.0 is cheaper, both to license and to manufacturer. Now that motherboard chipsets are beginning to appear with USB 2.0 controllers, USB 2.0 will appear on-board with no or little extra cost to the consumer. Secondly, manufacturers can add USB 2.0 support to their products, and retain compatibility with USB 1.1 and a huge install base of applications.

Many argue that USB 2.0 is inferior to Firewire, pointing out that Firewire offers guaranteed bandwidth, essential for multimedia applications such as

streaming video. Although, less well advertised, USB also offers guaranteed bandwidth. The USB specification defines several transfer modes. Devices such as external storage devices typically use the "bulk transfer" mode, but USB also offers an "isochronous transfer" mode which guarantees bandwidth at a specified rate. Few devices use this currently, but expect upcoming USB 2.0 webcams to do so. The principle difference between Firewire and USB 2.0 is that Firewire creates a peer-to-peer connection, while USB is a host-controlled interface.

Firewire's early adoption in the audio and video market and the fact that Firewire is seen as the "professional" solution will be difficult for USB 2.0 to dislodge. However, there is plenty of room for both technologies in the market.

external storage. A range of USB 2.0 hard-drives, CDROM drives and CD writers are available from many manufacturers.


These all tend to be ATA or ATAPI devices hooked up to the



# USB2.0



« USB bus by means of an off-the-shelf ATAPI bridge. The increased bandwidth offered by USB 2.0 is obviously a huge benefit here, and finally makes USB mass storage practical. The prospect of backing up a 40 GB drive at USB 1.1 speeds is not a pleasant one.



**For USB 2.0 support under Linux, you just need to build a kernel with the *ehci-hcd* driver.**

Another useful application of USB 2.0 is networking, and several USB 2.0 Ethernet adapters have now been announced, although we were unable to get one in time to review for this article. Here, again, bandwidth is key. USB 2.0 transfer rates now makes 100 Mbit Ethernet viable over the USB bus.

The other area where USB 2.0 will make an impact is imaging and video applications. We are currently unaware of any products on the market, but USB 2.0 webcams are expected for release soon. Another distinct possibility are USB 2.0-capable digital still and digital video cameras, and USB 2.0 scanners. ALI, for instance, have just

announced a USB 2.0 scanner controller and devices employing this chipset are slated for the end of Q3 this year.

## Linux and USB 2.0

USB support in Linux came of age with the release of kernel 2.4.0 at the beginning of 2001, although the USB sub-system had been back-ported for 2.2 kernels and shipped with most of the major distributions previously to that. Support for the existing USB 1.1 standard under Linux is generally excellent, and it is under Linux that USB really lives up to its plug-and-play nature.

As long as your kernel has adequate drivers and your system includes a hot-plug daemon (as all the major distros do), you can just buy a device, plug it in and expect to work. So how about USB 2.0?

USB 2.0 support first appeared in the revamped USB sub-system in the current development branch of the kernel, in kernel 2.5.2. This isn't recommended for the average user, however, and USB 2.0 support has been back-ported to the 2.4 tree.

The first official, stable kernel to include USB 2.0 was 2.4.19. The kernel *ehci-hcd* driver should support any USB 2.0 chipset and is known to work with silicon from NEC, VIA and Intel.

At the time of writing, only Red Hat have shipped a distribution with support for USB 2.0: this was included in Red Hat 7.3. Expect the other major distro vendors to follow suit with their upcoming releases. However, it is not difficult to build USB 2.0 support yourself. It's a simple task of getting the latest kernel source (available on our coverdisc each month) and recompiling it with the *ehci-hcd* driver built as a module by enabling the

**CONFIG\_USB\_EHCI\_HCD** option (you will need to retain the *usb-ohci/usb-uhci* drivers for your interface's USB 1.1 companion controller).

When you boot your new kernel, you then simply need to ensure that the *ehci-hcd* driver gets loaded. In many cases, if your distro has an up-to-date hotplug system or hardware detection system, this will be performed automatically; otherwise you can add this modules to your `/etc/modules` file to make sure it gets loaded during boot up. That really is all there is to it. Plug in your devices and have fun!

USB 2.0 is very much a work in progress in Linux, though. Currently mass storage devices such as external hard drives and CD-ROM drives will work without any fuss, but the lack of availability of any other kind of USB 2.0 peripheral means support will be sketchy for other types of device. USB 2.0 hubs have just reached the market, but we experienced difficulties with the device we tested under kernel 2.4.19 (although it worked under 2.5.31). As more Linux users starting buying USB 2.0 devices, using them under Linux, and filling in bug reports, the Linux's USB 2.0 stack will mature and stabilise.

## Web resources

### USB.org

[www.usb.org](http://www.usb.org)

USB resources and information sponsored by the USB Implementors Forum, Inc. Here you can download the USB specification, get information on newly-released USB devices, and more.

### Linux USB

[www.linux-usb.org](http://www.linux-usb.org)

Home of the Linux USB project. Visit here for information and support on the

Linux USB stack, and what devices are currently supported under Linux.

### Linux Hotplugging


<http://linux-hotplug.sourceforge.net/>  
Home of the Linux hot-plug system, currently supporting USB, PCI (and CardBus) and Firewire interfaces.

### Everything USB

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

A portal site with USB-related news, information and forums.

## Trust 5 Port USB 2.0 PCI card



■ **MANUFACTURER** Trust ■ **WEB** [www.trust.com](http://www.trust.com) ■ **PRICE** £34.99

### Although motherboards are

beginning to appear with on-board USB 2.0 support, you can add USB 2.0 compatibility to any PCI motherboard with an add-on card such as this device from Trust.

This card is based on the NEC chipset, the first USB 2.0 chipset to market, and offers five USB 2.0 ports, four on the backplane and one internal port – great for hooking up internal USB drives or hubs.

We gave this card a thorough testing – on x86 and PowerPC hardware – and it performed well throughout. The kernel EHCI driver under 2.4.19 and 2.5.x kernels supports this card's chipset, and we had no problems with any USB 2.0 or USB 1.1 devices we tried – including a selection of drives, a camera, a scanner, and a USB trackball. You just plug it in and it works. What more can you say?

The only real complaint I have with this Trust card is not the fault of the card at all but of the NEC chipset that drives it: this chipset doesn't implement the full

USB 2.0 bandwidth.

David Brownwell, one of the USB kernel developers, has stated that the NEC chipset has a bottleneck of around 28 MB/s due to the fact that the controller cannot access the PCI and USB busses concurrently.

This is not much problem if you only want to connect one high-speed device to the card, since most will be incapable of saturating that bandwidth. But start adding multiple high-speed devices, and performance will suffer. This was borne out in our speed tests.

For those with modest requirements, however, this card will be more than sufficient – although Trust themselves have recently announced a card based on the newer VIA USB 2.0 chipset which claims improved speed.

### Linux Format VERDICT

Offers poorer performance than cards based on more modern chipsets, but is cheap and reliable.

### LinuxFormat RATING

7/10

## Amacom Flip2disk USB 2.0 hard drive



■ **MANUFACTURER** Amacom ■ **WEB** [www.amacom-tech.com](http://www.amacom-tech.com)  
■ **PRICE** £289 (30 GB model, including USB 2.0 cable)

### The Flip2disk is a flexible, portable

storage solution. It is basically a 2.5" ATA hard-drive enclosed in a rugged, shock-absorbent case. Capacities from 10–60 GB are offered, and a plethora of interface cables are available to hook it up to just about any host – such as USB 2.0, Firewire, PCMCIA/CardBus, and parallel cables. With the Firewire and USB options the drive can take its power directory from the bus.

We tested a 30 GB model – based on a Toshiba mechanism – and it worked flawlessly under 2.4.19 and 2.5.31 kernels. It literally is plug and play. The drive is supported by the kernel's USB mass-storage driver,

and appears to the system simply as an extra SCSI disk. The Flip2disk comes pre-formatted with a single FAT32 filesystem and no partition-table, but you can re-partition it and use any Linux filesystem you choose, just as you would a regular hard drive.

The performance of the flip2disk impresses. Some quick and dirty benchmarks revealed that sustained

write speeds of 13 MB/s and read speeds of 16 MB/s during filesystem access. In comparison, my aging 20GB ATA/100 drive only managed figures of 21 MB/s and 16 MB/s on the same tests, respectively. (Of course, these figures are when connected via USB 2.0. On a USB 1.1 interface, the drive mustered a creditable 780 KB/s and 990 KB/s, respectively).

The Flip2disk offers high performance storage in a pocket-sized drive, and the range of interface options simply add to its portability. It isn't cheap, though, and additional cables are extra. Similar capacity Firewire or USB 2.0 only drives from companies such as Maxtor cost around half the price, but these lack the portability and flexibility of the Flip2disk.

### Linux Format VERDICT

The ultimate in portable storage, but at a premium price.

### LinuxFormat RATING

8/10

## USB2.0

# Predator 24x USB2.0 CD-R/W



■ **MANUFACTURER** iomega ■ **WEB** [www.iomega-europe.com](http://www.iomega-europe.com)  
 ■ **PRICE** 249.00 (about £181.40, inc VAT)

**Being able to buy a portable CD** writer that burns discs in a respectable time will be the reason that many will go for USB 2.0. Iomega, quick to realise this, have updated their range of external CD-R/Ws accordingly. The Predator 24x shares the same sculpted, two-tone case of its USB 1.1 predecessor, but this now houses a mechanism that's capable of 24/10/40 write/rewrite/read speeds. When connected to a USB 1.1 bus, it can fall back to a 4/4/6 mode.

Like other USB storage devices, the Predator appears to your Linux system as a SCSI device via the kernel USB mass storage driver. Simply plug it and it will work as a SCSI CDROM drive. To actually write discs – just like with any CD writer under Linux – you'll need to use the *cdrtools* package and the kernel generic SCSI driver (and GUI front-ends such as *XCDRoast* can ease things dramatically for you). The Predator is unsupported by the latest stable release of *cdrtools*, version 1.10, however, and the drive refused to recognise or write CD-Rs. Upgrading to the latest beta of 1.11 solved this

problem, and I could then burn discs without a hitch.

The Predator's performance is good, but not amazing. Iomega claim write speeds of 24x, but this is unrealistic in the real world. In tests, I found my 800 MHz PIII system unable to shunt the 3.5 MB/s required for 24x speeds down the USB 2.0 bus. The Predator's buffer-underrun protection handled this admirably, and I was able to write discs in 24x speed mode in just over four minutes rather than the expected three minutes. The drive even coped with old media that was not rated for burning at this speed.

My main worry is Predator's case, which sacrifices ruggedness for style. The flip-open lid, in particular, looks flimsy. Those who require portability above all of their CD-R/W drive, would probably be best looking elsewhere, especially considering there are much cheaper drives on the market.

### Linux Format VERDICT

Looks cool and performs well, but that case isn't designed with portability in mind.

### LinuxFormat RATING

7/10

# Trust 4 Port USB 2.0 Hub



■ **MANUFACTURER** Trust ■ **WEB** [www.trust.com](http://www.trust.com) ■ **PRICE** £39.99

The USB specification lets you attach up to 127 devices to a USB host, but, unsurprisingly, you won't find a USB adapter with 127 ports. The

solution is the USB hub, a device which lets you connect multiple USB devices to a single USB port, and provides power for the extra devices.

You might ask why bother with a USB 2.0 hub, when the chances are you won't have that many USB 2.0 devices, but a USB 2.0 hub lets you scale bandwidth to USB 1.1 devices hooked up to it, enabling full-speed transfers to each. This Hub is quite compact, with a footprint of a mere four inches by two-and-a-half inches. It's smaller even than the PSU brick that powers it.

USB hubs should be the most transparent of USB devices, but we found that the nascent USB 2.0 stack in the 2.4.19 kernel unable to cope. Devices plugged into the hub weren't recognised, and hot-plugging the device caused the USB stack to hang completely. The hub worked perfectly under 2.4.19 when attached to a USB 1.1 adapter, however. The development branch of the kernel, 2.5, is in a constant state of flux, but surprisingly we found the USB stack in 2.5.31

(plus some extra patches) worked perfectly (once I had managed to get the kernel to build and boot adequately) and the hub functioned as expected. A range of USB 2.0 and USB 1.1 devices were tested without problems. Currently there's no support for split transaction scheduling for interrupt transfers, so you cannot at the moment use devices such as keyboard, mice, and joysticks with a USB 2.0 hub under Linux.


We couldn't adequately test the performance of the hub, because we were limited by the bottleneck of the NEC USB 2.0 chipset on our adapter cards. However, connecting devices via the hub appeared not to cause any detriment to performance. **LXF**

### Linux Format VERDICT

Compact and functional, but you might want to let the Linux USB stack mature before buying.

### LinuxFormat RATING

8/10



## What on Earth is...

# PALLADIUM?

**Ross Anderson**, chair of the Foundation for Information Policy Research, gives his views on the real story behind Palladium, TCPA and the Fritz chip.

### »» So, what is the Palladium stuff, and TCPA?


TCPA stands for the Trusted Computing Platform Alliance, an initiative led by Intel. Their stated goal is "a new computing platform for the next century that will provide for improved trust in the PC platform." Palladium (named for a Trojan statue of the goddess Athena, upon whose safety the city depended) is software that Microsoft plans to incorporate in future versions of Windows; it will build on the TCPA hardware, and will add some extra features.

### »» What does TCPA/Palladium do? In ordinary English.

It provides a computing platform on which you can't tamper with the applications, and where these applications can communicate securely with the vendor. The obvious application is digital rights management (DRM): Disney will be able to sell you DVDs that will decrypt and run on a Palladium platform, but which you won't be able to copy. The music industry will be able to sell you music downloads that you won't be able to swap. They will

be able to sell you CDs that you'll only be able to play three times, or only on your birthday. All sorts of new marketing possibilities will open up.

TCPA/Palladium will also make it much harder for you to run unlicensed software. Pirate software can be detected and deleted remotely. It will also make it easier for people to rent software rather than buying it; and if you stop paying the rent, then not only does the software stop working but so may the files it created. For years, Bill Gates has dreamed of finding a way to make the Chinese



pay for software. Palladium could be the answer to his prayer.

There are many other possibilities. Governments will be able to arrange things so that all *Word* documents created on civil servants' PCs are 'born classified' and can't be leaked electronically to journalists. Auction sites might insist that you use trusted proxy software for bidding, so that you can't bid tactically at the auction. Cheating at computer games could be made more difficult.

There is a downside too. There will be remote censorship. The mechanisms designed to delete pirated music under remote control may be used to delete documents that a court (or a software company) has decided are offensive – this could be anything from pornography to writings that criticise political leaders. Software companies can also make it harder for you to switch to their competitors' products; for example, *Word* could encrypt all your documents using keys that only Microsoft products have access to.

### »» So I won't be able to play MP3s on my PC any more?

With existing MP3s, you may be all right for some time. Microsoft says that Palladium won't make anything suddenly stop working. But a recent software update for Windows Media Player has caused controversy by insisting that users agree to future anti-piracy measures, which may include measures that delete pirated content found on your computer. Also, some programs that give people more control over their PCs, such as VMware and Total Recorder, are unlikely to work under TCPA.

It is up to an application to set the security policy for its files, using an online policy server. So, for example, *Media Player* will determine what sort of conditions get attached to protected titles, and I expect Microsoft will do all sorts of deals with the content providers, who will experiment with all sorts of business models. You might get CDs that are a third of the price but which you can only play three times; if you pay the other two-thirds, you'd get full rights. You might be allowed to lend your copy of some digital music to a friend, but then your own backup copy won't be playable until your friend gives you the main copy back. More likely, you will not be able to lend music at all.

This could all be done today – Microsoft would just have to download a patch into your player – but once TCPA/Palladium makes it hard for people to tamper with the player software, and easier for Microsoft to control upgrades and patches, it will be harder for you to escape, and will therefore be a more attractive way of doing business.

### »» How does it work?

TCPA provides for a monitoring and reporting component to be mounted in future PCs. The preferred implementation in the first phase of TCPA is a 'Fritz' chip – a smartcard chip or dongle soldered to the motherboard.

When you boot up your PC, Fritz takes charge. He checks that the boot ROM is as expected, executes it, measures the state of the machine; then checks the first part of the operating system, loads and executes it, checks the state of the machine; and so on. The trust boundary, of hardware and software considered to be known and verified, is

steadily expanded. A table is maintained of the hardware (audio card, video card etc) and the software (O/S, drivers, etc); Fritz checks that the hardware components are on the TCPA approved list, that the software components have been signed, and that none of them has a serial number that has been revoked. If there are significant changes to the PC's configuration, the machine must go online to be re-certified. The result is a PC booted into a known state with an approved combination of hardware and software (whose licences have not expired). Control is then handed over to enforcement software in the operating system – this will be Palladium if your operating system is Windows.

Once the machine is in this state, Fritz can certify it to third parties. For example, he will do an authentication protocol with Disney to prove that his machine is a suitable recipient of *Snow White*. This will mean certifying that the PC is currently running an authorised application program – *MediaPlayer*, *DisneyPlayer*, whatever. The Disney server then sends encrypted data, with a key that Fritz will use to unseal it. Fritz makes the key available only to the authorised application and only so long as the environment remains 'trustworthy'. For this purpose, 'trustworthy' is defined by the security policy downloaded from a server under the control of the application owner. This means that Disney can decide to release its premium content to a given media player application in return for a contract that the application will not make any unauthorised copies of content, will impose a certain set of conditions (including what level of security has to be set in TCPA). This can involve payment: Disney might insist, for example, that the application

# WhatOnEarthPalladium

« collect a dollar every time you view the movie. In fact, the application itself can be rented too, and this is of great interest to software companies. The possibilities seem to be limited only by the marketers' imagination.

## »» What else can TCPA and Palladium be used for?

TCPA can also be used to implement much stronger access controls on confidential documents. For example, an army might arrange that its soldiers can only create *Word* documents marked at 'Confidential' or above, and that only a TCPA PC with a certificate issued by its own security agency can read such a document. This is called 'mandatory access control', and governments are keen on it. The Palladium announcement implies that the Microsoft product will support this. You will be able to configure *Word* so that it will encrypt all documents generated in a given compartment on your machine, and share it only with other users in a defined group.

Corporations will be able to do this too, to make life harder for whistleblowers. They can arrange that company documents can only be read on company PCs, unless a suitably authorised person clears them for export. They can also implement timelocks: they can arrange, for example, that all emails evaporate after 90 days unless someone makes a positive effort to preserve them. Think of how useful that would have been for Enron, or Arthur Andersen – or for Microsoft itself during the antitrust case.

The Mafia might use the same facilities. They could arrange that the spreadsheet with the latest drug shipments can only be read on accredited Mafia PCs, and will vanish at the end of the month. This might make life harder for the FBI – though Microsoft is in discussions with governments about whether policemen and spies will get some kind of access to master keys. But, in any case, a whistleblower who emails a document to a journalist will achieve little, as the journalist's Fritz chip won't give him the key to decipher it.

TCPA / Palladium also seems destined for use in electronic payment systems. One of the Microsoft visions appears to be that much of the functionality now built on top of bank cards may move into software once the applications can be made tamper-resistant. This is needed if we are to have a future in which we pay for books that we read, and music we listen to, at the rate of so many pennies per page or per minute. Even if this doesn't work out as a business model – and there are good arguments why it won't – there is clearly a competitive issue for a number of online payment systems, and there may be spillover effects for the user.

If, in ten years' time, it's inconvenient to shop online with a credit card unless you use a TCPA or Palladium platform, then this could move a lot of people over to the system.

## »» How do Microsoft and Intel propose to make money out of it?

My spies at Intel tell me that it is a defensive play. Because they make most of their money from PC microprocessors, and have most of the market, they can only grow their company by increasing the size of the market. They are determined that the PC will be the hub of the future home network. If entertainment is the killer application, and DRM is going to be the critical enabling technology, then the PC has to do DRM or risk being displaced in the home market.

Microsoft are also motivated by the desire to bring all of entertainment within their empire. But they also stand to win big if either TCPA or Palladium becomes widespread, as they will be able to use it to cut down dramatically on software copying. 'Making the Chinese pay for software' has been a big thing for Bill. With Palladium, he can tie each PC to its individual licenced copy of *Office*, and with TCPA he can tie each motherboard to its individual licenced copy of Windows. TCPA will also have a worldwide blacklist for the serial numbers of any copies of *Office* that get pirated.

Finally, Microsoft would like to make it more expensive for people to switch away from their products (such as *Office*) to rival products (such as *OpenOffice.org*). This will enable them to charge more for upgrades without making their users jump ship.

## »» Where did the idea come from?

It first appeared in a paper by Bill Arbaugh, Dave Farber and Jonathan Smith, *A Secure and Reliable Bootstrap Architecture*, in the proceedings of the IEEE Symposium on Security and Privacy (1997) pp 65-71. It led to a US patent: 'Secure and Reliable Bootstrap Architecture', U.S. Patent No. 6,185,678, February 6th, 2001. Bill's thinking developed from work he did while working for the NSA on code signing in 1994. The Microsoft folk have also applied for patent protection on the operating system aspects.

Markus Kuhn wrote about the TrustNo1 Processor years ago, and the basic idea – a specially trusted 'reference monitor' that supervises a computer's access control functions – goes back at least to a paper written by James Anderson for the USAF in 1972. It has been a feature of US military secure systems thinking since then.

## »» How is this related to the Pentium III serial number?

Intel started an earlier program in the mid-1990s that would have put the functionality of the Fritz chip inside the main PC processor, or the cache controller chip, by 2000. The Pentium serial number was a first step on the way. The adverse public reaction seems to have caused them to pause, set up a consortium with Microsoft and others, and seek safety in numbers.

## »» Why call the monitor chip a Fritz chip?

In honour of Senator Fritz Hollings of South Carolina, who is working tirelessly in Congress to make TCPA a mandatory part of all consumer electronics.

## »» Apart from pirates, industrial spies and activists, who has a problem with all this?


A lot of companies stand to lose out. For example, the European smartcard industry looks likely to be hurt, as the functions now provided by their products migrate into the Fritz chips in people's laptops, PDAs and third generation mobile phones. In fact, much of the information security industry may be upset if TCPA takes off. Microsoft claims that Palladium will stop spam, viruses and just about every other bad thing in cyberspace – if so, then the antivirus companies, the spammers, the spam-filter vendors, the firewall firms and the intrusion detection folk could all have their lunch stolen.

But there are much deeper problems. The fundamental issue is that whoever controls the Fritz chips will acquire a huge amount of power. Having this single point of control is like making everyone use the same bank, or the same accountant, or the same lawyer. There are many ways in which this power could be abused.

## »» How can TCPA be abused?

One of the worries is censorship. TCPA was designed from the start to support the centralised revocation of pirate bits. Pirate software will be spotted and disabled by Fritz when you try to load it, but what about pirated songs or videos? And how could you transfer a song or video that you own from one PC to another, unless you can revoke it on the first machine? The proposed solution is that an application enabled for TCPA, such as a media player or word processor, will have its security policy administered remotely by a server, which will maintain a hot list of bad files. This will be downloaded from time to time and used to screen all files that the app opens. Files can be revoked by content, by the serial number of the application that created them, and by a number of other criteria. The proposed use for this is that if everyone in China uses the same copy of *Office*, you do not just stop this copy running on any machine that is TCPA-compliant; that would just motivate the Chinese to use normal PCs instead of TCPA PCs in order to escape revocation. So you also cause every TCPA-compliant PC in the world to refuse to read files that have been created using this pirate program.

This is bad enough, but the potential for abuse extends far beyond commercial bullying and economic warfare into political censorship. I expect that it will proceed a step at a time. First, some well-intentioned police force will get an order against some offensive pornography, or a manual on how to sabotage railroad signals. All TCPA-compliant PCs



will delete, or perhaps report, these bad docs. Then a litigant in a libel or copyright case will get a civil court order against an offending document; perhaps the Scientologists will seek to blacklist the famous Fishman Affidavit. Once lawyers and government censors realise the potential, the trickle will become a flood.

Now the modern age only started when Gutenberg invented movable type printing in Europe, which enabled information to be preserved and disseminated even if princes and bishops wanted to ban it. For example, when Wycliffe translated the Bible into English in 1380, the Lollard movement he started was suppressed easily. But when Tyndale translated the New Testament in 1524, he was able to print over 50,000 copies before they caught him and burned him at the stake. The old order in Europe collapsed, and the modern age began. Societies that tried to control information became uncompetitive, and with the collapse of the Soviet Union it seemed that democratic liberal capitalism had won. But now, TCPA and Palladium have placed at risk the priceless inheritance that Gutenberg left us. Electronic books, once published, will be vulnerable; the courts can order them to be unpublished and the TCPA infrastructure will do the dirty work.

So after the Soviet Union's attempts to register and control all typewriters and fax machines, TCPA attempts to register and control all computers. The implications for liberty, democracy and justice are worrying.



#### But can't you just turn it off?

Sure – unless your system administrator

configures your machine in such a way that TCPA is mandatory, you can always turn it off. You can then run your PC with administrator privileges, and use insecure applications.

There is one respect, though, in which you can't turn Fritz off. You can't make him ignore pirated software. Even if he's been informed that the PC is booting in untrusted mode, he still checks that the operating system isn't on the serial number revocation list. This has implications for national sovereignty. If Saddam is stupid enough to upgrade his PCs to use TCPA, then the American government will be able to hot-list his Windows licences, and thus shut down his PCs, next time there's a war. Booting in untrusted mode won't help. He'd have to dig out old copies of Windows 2000, change to GNU/Linux, or find a way to isolate the Fritz chips from his motherboards without breaking them.

If you aren't someone the US President hates personally, this may not be an issue. But if you turn TCPA off, then your TCPA-enabled applications won't work, or won't work as well. If the applications that use TCPA/Palladium are more attractive to the majority of people, you may end up simply having to use it.



#### So economics are going to be significant here?

Exactly. The biggest profits in IT goods and services

markets tend to go to companies that can establish platforms (such as Windows, or *Word*) and control compatibility with them, so as to manage the markets in complementary products.

For example, some mobile phone vendors use challenge-response authentication to check that the phone battery is a genuine part rather than a clone – in which case, the phone will refuse to recharge it, and may even drain it as quickly as possible.

TCPA appears designed to maximise the effect, and thus the economic power, of such behaviour. Given Microsoft's record of competitive strategic plays, I expect that Palladium will support them. So if you control a TCPA-enabled application, then your policy server can enforce your choice of rules about which other applications will be allowed to use the files your code creates. These files can be protected using strong cryptography, with keys controlled by the Fritz chips on everybody's machines. What this means is that a successful TCPA-enabled application will be worth much more money to the software company that controls it, as they can rent out access to their interfaces for whatever the market will bear.

So there will be huge pressures on software developers to enable their applications for TCPA; and if Palladium is the first operating system to support TCPA, this will give it a competitive advantage over GNU/Linux and MacOS with the proprietary developer community.



# WhatOnEarthPalladium

## But hang on, doesn't the law give people a right to reverse engineer interfaces for compatibility?

Yes, and this is very important to the functioning of IT goods and services markets; see Samuelson and Scotchmer, *The Law and Economics of Reverse Engineering*, Yale Law Journal, May 2002, 1575–1663. But the law in most cases just gives you the right to try, not to succeed. Back when compatibility meant messing around with file formats, there was a real contest – when *Word* and *Word Perfect* were fighting for dominance, each tried to read the other's files and make it hard for the other to read its own. However, with TCPA that game is over. Without access to the keys, or some means of breaking into the chips, you've had it. And the Microsoft folk say they want Palladium everywhere, even in your watch. The economic consequences for independent businesses everywhere could be significant.

## Can't TCPA be broken?

The early versions will be vulnerable to anyone with the tools and patience to crack the hardware (e.g., get clear data on the bus between the CPU and the Fritz chip). However, from phase 2, the Fritz chip will disappear inside the main processor – let's call it the 'Hexium' – and things will get a lot harder. Really serious, well funded opponents will still be able to crack it. However, it's likely to go on getting more difficult and expensive.

Also, in many countries, cracking Fritz will be illegal. In the USA the Digital Millennium Copyright Act already does this, while in the EU the situation may vary from one country to another, depending on the way national regulations implement the EU Copyright Directive.

Also, in many products, compatibility control is already being mixed quite deliberately with copyright control. The Sony Playstation's authentication chips also contain the encryption algorithm for DVD, so that reverse engineers can be accused of circumventing a copyright protection mechanism and hounded under the Digital Millennium Copyright Act. The situation is likely to be messy – and that will favour large firms with big legal budgets.

## What's the overall economic effect likely to be?

I expect the most significant economic effect will be to strengthen the position of incumbents in information goods and services markets at the expense of new entrants. This may mean a rise in the market cap of firms like Intel, Microsoft and IBM – but at the expense of innovation and growth generally. Eric von Hippel documents how most of the innovations that spur economic growth are not anticipated by the manufacturers of the platforms on which they are based; and technological change in the IT goods and services markets is usually cumulative. Giving incumbents new ways to make life

harder for people trying to develop novel uses for their products will create all sorts of traps and perverse incentives.

## Who else will lose?

One well-known UK lawyer said that copyright law is only tolerated because it is not enforced against the vast majority of petty infringers. And there will be some particularly high-profile hard-luck cases. I understand that copyright regulations due out later this year in Britain will deprive the blind of the fair-use right to use their screen scraper software to read e-books. Normally, a bureaucratic stupidity like this might not matter much, as people would just ignore it, and the police would not be idiotic enough to prosecute anybody. But if the copyright regulations are enforced by hardware protection mechanisms that are impractical to break, then the blind may lose out seriously. There are many other marginal groups under similar threat.

## Ugh. What else?

TCPA will undermine the General Public License (GPL), under which many free and open source software products are distributed. The GPL is designed to prevent the fruits of communal voluntary labour being hijacked by private companies for profit. Anyone can use and modify software distributed under this licence, but if you distribute a modified copy, you must make it available to the world, together with the source code so that other people can make subsequent modifications of their own.

At least two companies have started work on a TCPA-enhanced version of GNU/Linux. This will involve tidying up the code and removing a number of features. To get a certificate from the TCPA consortium, the sponsor will then have to submit the pruned code to an evaluation lab, together with a mass of documentation showing why various known attacks on the code don't work. The evaluation is at level E3 – expensive enough to keep out the free software community, yet lax enough for most commercial software vendors to have a chance to get their lousy code through. Although the modified program will be covered by the GPL, and the source code will be free to everyone, it will not make full use of the TCPA features unless you have a certificate for it that is specific to the Fritz chip on your own machine. That is what will cost you money (if not at first, then eventually).

You will still be free to make modifications to the modified code, but you won't be able to get a certificate that gets you into the TCPA system. Even if a philanthropist does a not-for-profit secure GNU/Linux, the resulting product would not really be a GPL version of a TCPA operating system, but a proprietary operating system that the philanthropist could give away free. There is still the question of who would pay for the user certificates.

People believed that the GPL made it impossible for a company to come along and steal code that was the result of community effort. This helped make people willing to give up their spare time to write free software for the communal benefit. But TCPA changes that. Once the majority of PCs on the market are TCPA-enabled, the GPL won't work as intended. The benefit for Microsoft is not that this will destroy free software directly. The point is this: once people realise that even GPL'ed software can be hijacked for commercial purposes, idealistic young programmers will be much less motivated to write free software.

## I can see that some people will get upset about this.

Yes. And there are many other political issues – the transparency of processing of personal data enshrined in the EU data protection directive; the sovereignty issue, of whether copyright regulations will be written by national governments, as at present, or an application developer in Portland or Redmond; whether TCPA will be used by Microsoft as a means of killing off *Apache*; and whether people will be comfortable about the idea of having their PCs operated, in effect, under remote control – control that could be usurped by courts or government agencies without their knowledge.

## But hang on, isn't TCPA illegal under antitrust law?

Intel has honed a 'platform leadership' strategy, in which they lead industry efforts to develop technologies that will make the PC more useful, such as the PCI bus and USB. Their *modus operandi* is described in a book by Gawer and Cusumano. Intel sets up a consortium to share the development of the technology, has the founder members put some patents into the pot, publishes a standard, gets some momentum behind it, then licenses it to the industry on the condition that licensees in turn cross-license any interfering patents of their own, at zero cost, to all consortium members.

The consortium approach has evolved into a highly effective way of skirting antitrust law. So far, the authorities do not seem to have been worried about such consortia – so long as the standards are open and accessible to all companies. They may need to become slightly more sophisticated.

Of course, if Fritz Hollings manages to get his bill through Congress, then TCPA will become compulsory and the antitrust issue will fall away, at least in America. One may hope that European regulators will have more backbone.

## When is this going to hit the streets?

It has. The specification was published in 2000. Atmel is already selling a Fritz chip, and although you need to sign a non-disclosure agreement to get a data sheet, you have been able to buy it installed



in the IBM Thinkpad series of laptops since May 2002. Some of the existing features in Windows XP and the X-Box are TCPA features. For example, if you change your PC configuration more than a little, you have to reregister all your software with Redmond. Also, since Windows 2000, Microsoft has been working on certifying all device drivers: if you try to load an unsigned driver, XP will complain. There is also growing US government interest in the technical standardisation process. The train is rolling.

The timing of Palladium is less certain. There appears to be a power struggle going on between Microsoft and Intel. Palladium will also run on competing hardware from suppliers such as Wave Systems, and applications written to run on top of vanilla TCPA will need to be rewritten to run on Palladium. This seems a play to ensure that the secure computing platform of the future is controlled by Microsoft alone. It might also be a tactic to deter other companies from trying to develop software platforms based on TCPA. Intel and AMD appear to plan for the second generation of TCPA functionality to be provided in the main processor for free. This might provide higher security, but would enable them to control developments rather than Microsoft.

I do know that the Palladium announcement was brought forward by over a month after I presented a paper at a conference on Open Source Software Economics on the 20th June. This paper criticised TCPA as anticompetitive, as amply confirmed by new revelations since.

### »» But isn't PC security a good thing?

The question is: security for whom? You might prefer not to have to worry about viruses, but neither TCPA nor Palladium will fix that. Viruses exploit the way software applications (such as Microsoft *Office* and *Outlook*) use scripting. You might get annoyed by spam, but that won't get fixed either. (Microsoft implies that it will be fixed, by filtering out all unsigned messages – but the spammers will just buy TCPA PCs.

You might be worried about privacy, but neither TCPA nor Palladium will fix that; almost all privacy violations result from the abuse of authorised access, often obtained by coercing consent. The medical insurance company that requires you to consent to your data being shared with your employer and with anyone else they can sell it to, isn't going to stop just because their PCs are now 'secure'. On the contrary, they are likely to sell it even more widely, because computers are now 'trusted'.

The most charitable view of TCPA is put forward by a Microsoft researcher: there are some applications in which you want to constrain the user's actions. For example, you want to stop people fiddling with the odometer on a car before they sell it. Similarly, if you want to do DRM on a PC then you need to treat the user as the enemy.

Seen in these terms, TCPA and Palladium do not so much provide security for the user as for the PC vendor, the software supplier, and the content industry. They do not add value for the user, but destroy it. They constrain what you can do with your

PC in order to enable application and service vendors to extract more money from you. This is the classic definition of an exploitative cartel – an industry agreement that changes the terms of trade so as to diminish consumer surplus.

No doubt Palladium will be bundled with new features so that the package as a whole appears to add value in the short term, but the long-term economic, social and legal implications require serious thought.

### »» So why is this called Trusted Computing? I don't see why I should trust it at all!


It's almost an in-joke. In the US Department of Defense, a 'trusted system or component' is defined as 'one which can break the security policy'. This might seem counter-intuitive at first, but just stop to think about it. The mail guard or firewall that stands between a Secret and a Top Secret system can – if it fails – break the security policy that mail should only ever flow from Secret to Top Secret, but never in the other direction. It is therefore trusted to enforce the information flow policy.

### »» So a Trusted Computer is one that can break my security?

Now you've got it. [LXF](#)

You can find out more about Ross Anderson, the FIPR and public policy issues relating to IT at [www.cl.cam.ac.uk/users/rja14/](http://www.cl.cam.ac.uk/users/rja14/)

## Emulators



# AMIGA VARIATIONS

**Simon Goodwin** accepts a tough challenge to emulate the ultimate home computer.



**A**tari maestro Jay Miner, Commodore released the Amiga in 1985. It eclipsed all previous micros, and the design elegance remains unmatched. Three revolutionary custom chips offered 4096 colours, four-channel sampled sound, plus a 16/32-bit Motorola 68000 CPU, serial, parallel and two mouse ports and an expansion bus called Zorro.

The way it fitted together was more important than raw specs, or the sub-systems. Amiga architecture transcended the sum of the parts, augmented by a ROM-based multitasking device-independent WIMP OS. Commodore squandered profits from millions of Amiga 500 sales trying to buy into the commodity PC market. 32-bit Amiga production shifted to Europe in the 1990s, after a buy-out. When Escom followed Commodore down the pan, Gateway 2000 took a turn; now a fifth Amiga company continues to develop hardware and software, but it's hard to imagine any new Amiga matching the impact of the original A1000.

The subtlety, performance and tight integration of the Amiga system makes it exceptionally hard to emulate. Modern PCs are very complicated, but their parts don't mesh seamlessly like

Amiga ones, struggling to do the same even at 100 times the clock rate. Amiga emulation was a myth for more than decade – the acronym *UAE* once stood for *Unusable Amiga Emulator* – but now the *U'nix Amiga Emulator* really works. It spawned variants like *XFellow* and *Amithlon*. We'll focus on the main version.

### Organisation

*UAE* comes with well-organised C source, and substantial technical docs. The mantra is 'you'll probably have to experiment a little'; a friend with Amiga experience can help a lot if you're a newcomer.

Compiled *UAE* versions are on the coverdisc but, for best compatibility, rebuild from source. Start with `./configure`, adding `--with-svgalib` for optional full screen display and `--enable-sound` for OSS audio. 'docs/unix/README' notes other options.

A default build took my K6-2/500 eight minutes, making a 1.6MB executable. There is no `make install` but you can run the executable from the directory you configured in, after listing Amiga resources in the `'-/.uaerc'` file, as explored in the 'Amiga fuel' section. The `jit-0.8.21.tar.gz` package gives an `--enable-jit` option for experimental code translation, adding a control panel and boosting CPU emulation at the expense of compatibility. An MMU emulation patch reduces speed but boosts stability, especially when writing new code. PC hosts have more blocking issues than a real Amiga, as PC hardware was never designed for multi-tasking. Linux inherits PC limitations, so an emulated Amiga cannot be as smooth as the real thing – most obviously in mouse movement and screen updates – but may be faster in other respects, especially if you're used to 16-bit Amigas from the 1980s rather than later 32-bit A1200, A3000 or A4000s.

### UAE Versions

The current stable release *0.76* of *UAE* emulates 16-bit Amigas reasonably (A1000, A500, A600, A1500 and A2000 without PC

**Deluxe Paint** was the Amiga's first killer app – former Amiga artists still wonder why there's nothing else like it.



co-processor) but lacks support for the 32-bit AGA chip set and later CPUs. 0.8.x releases are notionally betas, but more capable and quite reliable, so some recent distros include 0.8.x versions; 0.8.21 was the latest when this article was finished in August. UAE 0.7.6 can emulate a 32-bit 68020 processor with optional 68881 floating point unit, and retargeted graphics via *Picasso96* which lets system-friendly Amiga programs drive PC-style frame buffers. Few games support *Picasso96*, but it offers extra resolution and colour depth for later productivity applications.

A 15 or 16-bit screen depth is recommended. Extra bits per pixel benefit *Picasso96*, but hamper original Amiga graphics emulation. An 800 by 600 pixel window works well, as the original Amiga TV output was limited to 724 by 576 plus a border. Hires *Picasso96* requires the Kickstart 3 ROM, introduced in 1992, and 68020 or later 32-bit CPU emulation. These options confound some older Amiga software. Config can use a bewildering variety of switches on the command line, or a `~/uaerc` file, *GTK+* configuration window, or front-ends like *MyMiggy*.

Printer output defaults to *lpr*. *Ghostscript* filters give quality printout via the AmigaOS Postscript driver. BSD-style network sockets can link UAE to host networks. Serial port access from the emulator is diverted to `/dev/tty` but shaky, especially on input to the emulator. Mouse tracking is problematic unless you config the Linux pointer to disappear or run an Amiga program called *Mousehack* to tie it to the Amiga one.

## Compatibility

UAE runs most of the Amiga software from the 1980s and many later titles, but it's far from perfect even on ostensibly 'fast' PCs. Most classic games and demos give recognisable if not identical results, but I found *Octamed6* sequenced sound was unlistenable at any quality setting on my machine, stuttering even on simple four-channel modules. *Deluxe Paint IV AGA* ran OK, if slowly. *Vista* got stuck while 'Initialising' and didn't render anything. I got the French A1200 demo to play without obvious sound and graphics glitches by reducing the render rate to one frame in four in the default window on my Debian 2.2 Linux box with SoundBlaster Live and Matrox G400. Attempts to render frames faster led to jerky graphics and sound, yet this demo runs smoothly on a basic A1200 with a 14MHz 68020 sharing access to 2MB slow memory.

Jerky sound can be tuned up from the command line. These parameters allow double the buffer size and the lowest common sampling rate you can use without shredding samples. The size and rate parameters are not on the setup tabs but once set from



The new look for AmigaOS on UAE and real Amigas with the latest 3.5 or 3.9 system release, Kickstart 3.1 and 'newicons'.



Some of the open sourcerers who develop UAE.

the command line you can save them with the default config from the GUI. Editing the `~/uaerc` file gives the finest control.

## Bottlenecks

UAE is an interpreting emulator; overhead of instruction decoding is high, especially on PCs, architecturally limited compared with UAE's original RISC workstations. A 65,536-entry despatch table chews up data cache, and 70 Pentium instruction dispatch slots are consumed in decoding before starting to interpret each 68K instruction, which can do a lot more per cycle than an x86. UAE also vets every memory access, using another cache-churning table of 65536 pointers, to mediate memory-mapped ports.

'Just In Time compilation' of code blocks drastically reduces such overhead. Extra memory holds translated blocks and an index. After each block JIT emulators check to see if the next is pre-translated. If so, it can be run without decoding; otherwise it is translated and buffered. JIT requires at least 6MB for its tables, and prefers 12MB. Some programs benefit from even more.

The 64K word table is used a bit more cleverly by the JIT engine, and 2.4 Linux kernels let the host Memory Management Unit help with the sifting. Another speedup comes from using the CMOV Pentium Pro instruction. Despite these complications part-translated combination code can run programs several times faster than the interpreting code. Register-bound loops may outrun the fastest 68060s on a fast PC.

## JIT variations

*uae-jit* is a pre-beta compiling version of the emulator. It only emulates 32-bit 68Ks and relies on defensive programming to flag the need for cache flushes, upsetting many old Amiga apps. The JIT code only works on x86 hosts, though the technique may be easier to implement on other processors, and cannot handle the PowerPC code in the fastest apps written for dual-processor Amigas. The JIT releases, named after UK football teams, lack AGA chipset emulation and access to host TCP/IP, but can compile FPU code as well as that for 68020, 68030 and 68040 integer units. They even support a second PCI graphics card, dedicated to the emulation, using Amiga Virge 3D drivers and XFree86 v4 tweaks. The JIT patch has excellent docs.

The author now develops a commercial emulator, *Amithlon*. This only runs on Pentium Pro or later PCs, using a custom Linux kernel, emulating AmigaOS but little classic Amiga hardware – just enough to put up a grainy boot menu, after that you are dependent on emulation of third-party Amiga add-ons for peripheral access. However you do get fast nVIDIA and Matrox graphics, plus easy access to PCI and USB devices rarely available on real Amigas.

*Amithlon* only runs a small minority of Amiga games, but storms through power applications from the mid-Nineties like *ImageFX* and *Audio Evolution* that leverage add-on sound and



Amiga 1500 – the A500 is reborn with revolutionary "two floppy drives as standard" format.



# Emulators

« graphics. *AHI* and *Picasso96* extensions divert pure software requests from the Amiga hardware to PC cards. Licensing arguments have stifled *Amithlon* sales. We've waited a year for a review copy, or even demos.

*Amithlon* has split the Amiga scene; for some AmigaOS is all that matters now, and a fast CPU emulation on new hardware is more important than compatibility with the old Amiga; others argue it's not an Amiga without the custom chips, or the PowerPC chosen for future Amigas. While *Amithlon* only runs a late subset of Amiga programs it does so fast and expandably, and shows a neat way to embed Linux, reminiscent of *EC64 (LXF15)* although CD rather than floppy-booted.

*Amithlon* emulates *Intuition*, the Amiga GUI, rather than the Amiga as most people remember it. JIT compilation speeds up the CPU dramatically, though not entirely reliably, but does nothing for the bigger bottleneck, custom chip emulation. I had to drop the redraw rate to one frame in two to run the A500 game *Turrican 3* on *UAE*. Sound and scrolling were still grainy by real Amiga standards, though the game was playable. Graphics diversions like *DGA* and *SVGALib* avoid the substantial overhead of X, but custom chip emulation, for sound as well as graphics, still burns far more *UAE* cycles than CPU when running typical Amiga programs.

The interpreting *UAE* will seem sluggish to anyone who has used an accelerated 32-bit Amiga, but the custom kernel and direct drivers in *Amithlon* mean it feels like a very fast Amiga with 32-bit peripherals.



Amiga emulation JIT maestro Bernie Meyer.

## Amigaesque

There are ways to get an Amiga-like environment on Linux without running a full emulation. *Amiwm* is an X window manager modelled on the Amiga *Workbench*. It can move, position and resize windows and supports multiple screens with backdrops, which can be dragged up and down as on the Amiga. *Worker* and *Gentoo* mimic Amiga *Workbench* alternatives like *SID*, *DirWork* and *Directory Opus*. These two-window file managers are ideal for computer housekeeping if you grew up with the Amiga equivalents, and they are configurable to launch your favourite Linux tools.

## Latest UAEs

*UAE1.0* hovers on the horizon. Current 0.8.x betas add support for AGA playfields and sprites, though collision detection remains a



This *AmiWM* desktop brings Amiga sliding screens to X.

weakness. Audio is better, though even more CPU intensive, and direct SCSI access allows CD burning from the AmigaOS emulation, via a virtual 'uaescsi.device' that supports direct SCSI commands but not (yet) filesystem access. This should suit CD burners and scanners, but not disks, which are better accessed through the Linux file system.

Experimental snapshot code can save and restore the machine state; this is good for floppy-based games but little use for applications because it can't save or restore the hard disk emulation state. Emulated drives are now called *DHn* rather than *UAEn*, matching Commodore defaults, and work more like the real thing. Copper co-processor emulation is more accurate and the 68040 CPU is emulated, though the choice is academic as the instruction set and speed are the same.

*UAE* is a very accomplished emulator and if you've never used an Amiga you'll find it an impressive system. It won't put my A4000/060 into retirement yet, but the gap is closing. Your coverdisc holds *UAE* versions, patches, demos, games and utilities – why not track down a *Kickstart*, *Amithlon* or *Amiga Forever CD*, and try for yourself?

## The end is nigh

Next month we get into *MESS*, the *Multi Emulation System*, in the final column of this series, plus a round-up of a dozen or so other emulators that might otherwise slip the net.

## XFellow 0.03

Ported from DOS

*XFellow* is a port based on *WinFellow 0.4.2*. *Fellow* was designed for MSDOS so the port only runs on x86-compatibles. *Fellow* was derived from portable *UAE* source, adapted because that was sluggish on PCs because it used C exclusively and did not take short-cuts to the PC hardware. Chunks of *UAE* were re-written in assembler, simplified and bodged for mid-Nineties commodity hardware, yielding an emulator which was less stable and compatible but could run many Amiga programs on old PCs that would only crawl if asked to run them in *UAE*.

*XFellow* requires *NASM 0.98* for the assembler parts, as well as *GCC* for filesystem support, and *SDL1.2*. Nine *GTK+* tabs configure aspects of the system. It emulates 68000 to 68030 processors, including cut-down *EC020* and *EC030* versions, at the standard A500 clock rate or half, double or quadruple that.

*XFellow* can play 8 or 16-bit mono or stereo sound at rates from 15650 Hertz, one sample per TV line, up to the CD rate of 44.1KHz, optionally diverting output to a *WAV* file. The default writes to '/dev/dsp' via *OSS*. Two gameports accept Linux key, mouse or joystick input, vital for many Amiga programs. Up to four floppy disk images are configurable. I couldn't get the hardfile and host file system controls to do anything, and some configuration table options are ignored.

*XFellow* is slow and prone to graphics corruption, hard to exit without closing the launch shell or bombing the emulator. 'unhandled event messages' fill the shell as it runs; it tends to make temporary files with MSDOS paths line C:\ prefixed, and segfaults if the ROM is not found. The X version lacks docs. I did get the *Kickstart* animation and a *Workbench*, but the Amiga screen occupied just the middle quarter of a big window with



Original Amiga *Workbench 1*, with graphical glitches caused by *XFellow* infelicities.

only 320 low-resolution pixels per line and the emulated mouse tagging the Linux one in the wide borders. This format is adequate for old games, and contributes to the speed of *Fellow*, but unless you have a sluggish Intel-based Linux box, *UAE* is likely to be more useful to you than *XFellow*.

## Amiga fuel

The Amiga OS is built around a kernel ROM called *Kickstart*, 256KB in the early *Kickstart 1* machines. 512KB *Kickstart 2* and 3 versions were more capable and configurable, though most *Kickstart 3* images require a 32-bit processor.

*Kickstart*, like AmigaOS, is still on sale and you need a copy from an Amiga or the officially-licensed *Amiga Forever* emulation pack to emulate an Amiga legally. *UAE* includes skeletal *Kickstart* emulation, enough to run a few simple demos, but you need the real thing to launch a desktop or shell commands. Once you've got your *Kickstart*, *UAE* requests a bootable disk by showing a picture of a floppy. The next hurdle is that the Amiga disks process a track at a time, rather than dividing tracks into small sectors with gaps. Amigas automatically synchronise to read one track every rotation, whatever the initial position of the disk.

PC disk controllers can't work this way, so you can't insert an Amiga floppy and use it like one in a FAT or HFS format. Amiga hardware can read and write those formats, and any other, but you need a real Amiga.

### Floppy tools

*UAE*, like other emulators, can use images of the contents of a disk held as Linux files. *transdisk* and *trackwizard* read and write such 'ADF' files on real Amigas. You need to compress the result to fit it on a 720KB disk for transfer to Linux, using an AmigaOS port of *ZIP*, *GZip*, *BZ2* or Amiga standards *LHA* and *LZX*.

*CrossDOS*, bundled with AmigaOS 2 and 3, and the free *FAT95* or *Messydos* from Aminet, will format, read and write PC-compatible floppies. You can pack files onto those and transfer them either way via *mtools* or the Linux FAT mounter. If your Amiga has an HD floppy there's no need to compress the 900, 120byte ADF file for 1.44MB media. Amiga-format high density floppies hold 1760KB so they will need compression or splitting.

The Amiga's own compressed disk image format, DMS (short for Disk Masher), can run in the emulator, reading and writing ADFs to any emulated floppy drive. The Linux tool *readdisk* moves files from an ADF to your local filesystem, but only understands the old Amiga *Kickstart 1* file system, not the later Fast File System the custom systems on commercial releases. To read FFS disks:

```
mount df0.adf /df0 -t affs -o loop
```

This mounts disk image 'df0.adf' at the mount point '/df0', created as a subdirectory by a super-user.

The *adskutil* Perl script gives FTP-style command line access to the contents of ADF files. It supports **LS**, **CD** and **GET** but lacks **PUT**, **DELETE**, **MKDIR** and other commands to change the disk, so it's really only useful for looking inside ADF files without loading the emulator though it is a good guide to Amiga disk internals.

*UAE's* **-m** option allows native file directories to be accessed from within the emulator:

```
uae -m Work:/home/simon/amigawork
```

Put Aminet 'lha' archives there and unpack them within *UAE* to preserve file attributes. Unpacked files can be accessed more quickly, as there's no need to emulate floppy rotation, but some loaders use those timings and insist on reading an ADF image.

There are two hardware ways round the disk incompatibility. The *Disk2fdi* hack may let you read 880K Amiga floppies from the command line or by patching later *UAEs*. It switches a PC disk controller abruptly between a second drive of any type, with an



arbitrary blank 'scratch' disk in, and a 3.5" disk with an Amiga disk in it. Running the disk controller at double speed, to pick up clock as well as data bits, with luck and a following wind this snarfs a snatch of Amiga data in a format that the PC can unravel later.

*Disk2fdi* is worth a try, but won't write Amiga disks and can't read some at all. There is a reliable answer for those prepared to buy a special controller. *Catweasel* from Individual Computers lets you read and write Amiga floppies directly in a Linux shell or *UAE*. This ISA card and its open-source drivers can handle old FM and MFM (1771, 1793, etc.) disk formats as well as the peculiar Amiga ones, in single, double or quad (1760K) density; the hardware supports old Commodore and Apple GCR formats too.

### Wired links

Another way to transfer files is via the serial ports – an original Amiga can manage rates up to 19,200 baud with a properly wired RS232 cable, and accelerated 32-bit models can drive the motherboard port reliably at up to 57,600 baud. *Term* for AmigaOS and *Minicom* for Linux work well together.

Power Computing and Eyeteck sell Ethernet kits with *Samba*, and there are NFS ports too. Since Linux can read and write Amiga hard drives, the quickest way to move a lot of data may be to move a drive or ZIP or similar disk between systems. Linux can mount Amiga hard disks using the **-t affs** switch as long as you have 'Amiga Fast File System' support in your kernel:

```
mount /dev/sda1 -t affs -o rw /amiboot
```

```
uae -m Workbench:/amiboot
```

Once you have a bootable system you'll find mounds of free software for your emulated Amiga. Gigabytes of well-organised Amiga files, rarely crippled or undocumented, appear online in the Aminet archive, or on about 50 CDs. Aminet trounces any other platform's single-source site. [LXF](#)

### Links

Amiga accessories: [www.eyeteck.co.uk](http://www.eyeteck.co.uk)

Amiga CDs: [www.schatztruhe.de](http://www.schatztruhe.de)

Amiga Forever: [www.cloanto.com/amiga](http://www.cloanto.com/amiga)

Amiga hardware: [www.powerc.com](http://www.powerc.com)

Amiga software: [www.aminet.net](http://www.aminet.net)

AmiWM:

[www.lysator.liu.se/~marcus/amiwm.html](http://www.lysator.liu.se/~marcus/amiwm.html)

Catweasel Linux: [www.soundtracker.org/raw](http://www.soundtracker.org/raw)

Disk2FDI: <http://fast.emuunlim.com/disk2fdi>

LXF Emus online:

<http://simon.mooli.org.uk/LXF>

MyMiggy launcher:


<http://mymiggy.sourceforge.net>

UAE home: [www.freiburg.linux.de/~uae/](http://www.freiburg.linux.de/~uae/)

UAE-JIT: <http://byron.csse.monash.edu.au>

XFellow: <http://xfellow.sourceforge.net>

**Amithlon boots a tiny Linux kernel and JIT 68K compiler from CD into a PentiumPro, turning it into a fast new-generation Amiga which lacks Classic Amiga compatibility.**



**Amiga 600 – bring it down from the attic, and relive the glory days of Amiga's (and your) youth.**

# 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

## How code is represented

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap
between lines:
begin
mniWordWrap.Checked := not
end;
Usually, you'll find the code on
our CD/DVD too.
```

## THIS MONTH...

### Low resource >>

Dig out that redundant old laptop or PC as we learn how to optimise standard GNU/Linux distros for low memory and small disks **p68**

### Perl


Our look at the newly-released Perl 5.8 – learn about the benefits as Perl beefs up its signal handling and goes all unicode **p72**

### Java

Time for a spot of housekeeping as we tidy up our build processes with *Ant*, the Java build tool with XML build files **p74**

### Kylix

In our penultimate *Kylix* tutorial we use what we've learned to build a Linux process viewer **p76**



### CD writing

This month we get GUI with *Gnome Toaster* and *KreateCD*. Then we take a look at writing mixed mode CDs **p82**

### PHP

Introducing the first of a two part look at *MySQL* and *PHP*. Find out what you can achieve with this dynamic combination **p86**

## TIP OF THE MONTH!

KDE may be a graphical environment, but it can often be useful to know how to invoke KDE applications from the command line, especially when customising your desktop or writing scripts. One essential KDE tool to become familiar with is *kfmclient* which invokes *Konqueror*.

You call *kfmclient* with a command as the first argument (run **kfmclient --commands** for a list of supported commands). For example, the **openURL** command will open a new *Konqueror* window to display the specified URL (this can be a relative path or a fully-qualified URL).

## Command-line KDE

For instance

```
kfmclient openURL .
will open the current directory in the
file manager, while
```


```
kfmclient openURL
http://www.kde.org/
will open the KDE home page in the
web browser.
```

The command **openProfile** opens a new *Konqueror* window using the specified 'View profile'. Thus, **kfmclient openProfile midnightcommander ~/Desktop** will open a two-pane window in the *Commander*-like mode (look in `$KDEDIR/share/apps/`

`konqueror/profiles` for the names of defined profiles).

Also useful is the command **openProperties**, which opens a properties dialog for the specified URL to let you edit its properties

Another handy KDE tool is *kprinter*, which is used to talk to the KDE printing system. Running it with no arguments will open up a print dialog, or you can specify the filename to print or use the switch **--stdin** to spool from the standard input. This command is a life-saver when you want to print from a non-KDE application.



## EFFICIENT DISTRIBUTIONS

# Low resource Linux

**Marco Fioretti** illustrates some methods to trim modern distributions to make them run faster, or on limited computers. Rejuvenate that old laptop or PC.

**L**inux is getting friendlier: installation is a breeze (well, most of the time), configuration is point-and-click, documentation improves slowly but surely.

This is fun, and brings more users to Free Software, but there is a problem: most default installations swallow hard disk space like crazy, and many applications appropriate almost as many CPU cycles and RAM chips as their closed source counterparts we used to call bloated. Sysadmins who only need a low traffic firewall end up buying a box which will sleep 99% of the time just because a “base” install needs three times more megabytes than a five years old computer will hold. Desktop users must upgrade their hardware just to keep doing the same things they were doing before, at the same speed. A lot of perfectly good computers, which certainly were not ‘free’ (as in beer) are made worthless by ‘free’ software. Laptops, carrying

smaller disks and less RAM, become obsolete even faster.

This article is about fighting this tide, or, more exactly, about how to get as much functionality and performance from modern Free Software without paying its default price, which is the the cost of the hardware it seems to need.

Since there is not a unique solution to this problem, we will present several areas where performance can be improved, focusing on solutions and methods accessible even without high hacking skills.

## Why bother?

This is the default reaction, the “hardware is so cheap today” school of thought. Bad school. Wasted money is wasted money, and ‘cheap’ is a terribly relative concept: computers are still very expensive for the majority of the world population.

At the same time, there are countless schools and non profit institutions that must compute and access the Internet to function. However, they must make better use of every single penny, when they have one, and their computing needs are really basic. If they could afford it, often a bigger monitor would make

them much more productive than two gigabytes of RAM.

On a different level, less is safer: any unused package or feature that can be removed is one possible security hole closed. Last but not least, if a 2GHz processor can go 30% faster with a few tweaks, why not do it?

## Possible solutions

The software overload can be reduced in several ways: the hardest and purest is certainly doing it from scratch. Choosing every single application and library, and compiling it optimised for a specific CPU certainly gives the most performing systems. Projects like RockLinux ([www.rocklinux.org](http://www.rocklinux.org)) or Linux From Scratch ([www.linuxfromscratch.org](http://www.linuxfromscratch.org)) do exactly this. However, they require a lot of previous knowledge, and quite some time to install and, later, to keep all the software current.

The second most popular suggestion is to use older distributions: if five years ago one could already do all really useful things with much less MHz and megabytes, why update? Many (like Tiny Linux, [tiny.seul.org](http://tiny.seul.org)) choose this path, and are happy with it. Sincerely, this sounds more like an admission of defeat than a solution. Why give up on every modern feature (from iptables to IMAP to multi-language support in Emacs), and maybe live with well known security holes?

In many cases, especially for servers or network administration appliances, the perfect solution is a specialised distributions. There are several built just to be firewalls or other kind of servers. Of course, they do not help on the desktop side.

For those desktop users who cannot be full time hackers, but enjoy plenty of modern and well documented applications, the best solution turns out to be some (heavily) tweaked mainstream distributions, those with the greatest amount of ready packages, documentation, and fellow users.

It is possible to squeeze a lot of functionality in a little space. The rest of this article suggests some methods to achieve this result with the most popular distributions. The only requirement is that the target PC can bear the standard base install of the chosen distribution. Overcoming this limitation is a task for specialised projects like RULE for Red Hat - see box.

## Pre installation tasks

**PACKAGE SELECTION** The first step to maximise performance is to have very clear what the computer will do. This is easy for servers, a bit less so for desktops. Some people will only need a browsing/ email terminal, others a text processing system, and so on.

Once the usage is clear, the proper applications can be chosen: when shopping for a very limited machine some particular selection criteria must be applied.

First of all, the Unix tradition is good for interactive programs: those that do *one* thing well and can easily cooperate are better, regardless of all those fancy screenshots around. Real people who need to get the job done quickly, (almost) never simultaneously chat, play DVDs, read email, browse and write their diary. The big, super integrated, do-it-all mammoths and desktops environments can be ignored, as the average brain will never simultaneously stimulate all the RAM that they require.

Dependencies are the second thing to check, and the hardest to control. Many programs claiming to be very little and fast but sit on a full blown KDE or GNOME installation: just avoid them. Another problem with dependencies is that: "...Most [packagers] are exceedingly careless. They require all

sorts of other packages - most of which are not required at all!" (from page 9 of *Multitool Linux*, Addison Wesley 2002)

The author is talking of RPM, but this happens also in other package systems, sometimes even in raw makefiles. For example, some recent *GhostScript* RPMs mandate Asian fonts even if the chosen system language is European, and the *Xemacs* binary for Red Hat 7.2 from the [www.rpmfind.net](http://www.rpmfind.net) repository will abort if the Canna library for Japanese input is not aboard.

## Toolkits

Even toolkits can create dependency problems. As far as we are concerned here, a toolkit is "simply" the set of graphic libraries used to draw the graphic widgets, talk to other programs, drag, drop, etc... The two most recent and popular toolkits today are Qt (for KDE) and GTK (for GNOME), but there are many others still around.

Each toolkit needs RAM and disk space: for this reason, it is important to choose applications that collectively require as little of them as possible: single KDE or GNOME programs might still find a place on a very lean and mean system, but it is better to not install from both families, or at least to check that only one version of each toolkit is required.

In general, do not trust "mean and lean" programs without checking that they will not fill the drive with friends: for RPM packages, the command to check dependencies is:

```
rpm -qp --requires package.rpm
```


**POST INSTALL CONFIGURATION** Once the program list is ready, there are three major phases to go through. The first is to perform the very basic base install, and remove as much as possible from it. The second is to add all (and only) the other programs which are needed, and the last to do some system wide tuning to squeeze the maximum performance out of your box. Installation and package management are covered in many other places, so let us discuss the system tuning part (without forgetting that some of the actions below may have to be repeated whenever packages are added, even during the initial configuration!). Again, your mileage may vary wildly depending upon your use and hardware: experiment, and let us know!

**TUNING THE HARD DISK** Often, the older machines at which

## TinyX

### An alternative to X11

When a real graphic server is needed, but XFree86 is too heavy, why not try something like *TinyX*? This server is available at [www.handhelds.org](http://www.handhelds.org), and already used on several PDAs.



A small hard disk, ancient CPU and a collection of tiny SIMMs are no barrier to Linux.

◀◀ this article is targeted may not have a lot to tune, as far as hard disks and their controllers are concerned. It does not hurt to check, however: the man page of the `hdparm` command lists several ways to test the hard disks and access them more quickly, if possible.

**COMPILING THE KERNEL** Should it be done? Maybe, maybe not. If RAM is \*really\* scarce it might be worth it. Otherwise, modern kernels being modular, there might be almost nothing to gain.

**DAEMONS** By default, install procedures set up several daemons to start at boot time. All the daemons not really needed should be turned off: the `ntsysv` command provides a nice console based interface for this. Right after that, the packages could be removed.

**VCONSOLES** Multiple virtual consoles are great, if not for the reason that they let you still enter to clean up after something froze on the first one. However, they should be limited to the smallest possible number (see the Small Memory HOWTO for details). Vconsoles are started in the `/etc/inittab` file – with lines like this for Red Hat 7.3:

```
6:2345:respawn:/sbin/mingetty tty6
```


Just comment out those you do not need, starting from the higher numbers.

**UNNEEDED IMAGES AND DOCUMENTATION** Several distributions install a lot of documentation and images no matter what one tells them. If disk space is scarce, consider removing documentation altogether, or not installing it when adding new packages (`--excludedocs` option in RPM). The places where to look for documentation are the directories listed in the `/etc/man.config` file, and `/usr/share/doc`. Another space saver is looking for graphic files in the `/etc` and `/usr/share` directories, and remove all icons and backgrounds that will not be used.

**VIDEO SERVER** On a server, or a terminal to be used only to run a few text based applications, X should be excluded (but see also the next paragraph). If looking at more windows simultaneously, and at high resolution, is actually needed, there are still a few tricks to make a slow machine faster.

The most obvious is choosing the lowest colour depth that still make the machine usable. Unless heavy graphics are involved (and in that case more powerful hardware should really be used) 8 bits per pixel are enough. The second step is to avoid

**Not just old junk – with careful tuning it can do all the tasks required.**



**Small distros are not the only solution.**

background images and graphical screen-savers. Many other tips to make is faster can be found at:

[ftp://rtfm.mit.edu/pub/usenet-by-group/news.answers/x-faq/speedups](http://rtfm.mit.edu/pub/usenet-by-group/news.answers/x-faq/speedups)

**MULTIPLE WINDOWS WITHOUT X** When only console applications are used, the only real gain of a windowing interface is the visibility of more windows simultaneously, at the maximum possible resolution. In such scenarios, depending again from the typical usage pattern, several solutions are possible. The simpler one is to just leave the system as is, and open each text-based program in a different virtual console: for example, log in at boot and start the `mc` file manager, then hit **Alt+F2**, log in again and start `mutt` to read email, and then jump at will from program to program.

The next step is to get more screen lines: this is possible with the VGA option to the boot loader (`man lilo.conf`) or using the `SVGAText` package, depending from our needs, willingness to hack, and hardware. At this point, to have multiple text terminals on the same text root window, try `TWIN` (<http://linuz.sns.it/~max/twin/>), which does exactly this.

**LEAN WINDOW MANAGER AND MENUS** With a real graphic server, the most important trick is a correct choice of the window manager. There are a lot available (for more details see [www.plig.org/xwinman/comparisons.html](http://www.plig.org/xwinman/comparisons.html)), and the same general criteria already mentioned should be applied: avoid excessive dependencies and extra tasks, just look for windows and menus management. The fact that older computers and laptops usually can display less colors, and come with small monitors, should not be overlooked: "light" managers eating screen space with colourful icons do not really help from that point of view.

Speaking of menus, a terrific way to get almost all the functionality and friendliness of the big leagues at a much lower price is to massively bind shell scripts (to dial up, fax a document from the scanner, etc...) to entries in the window manager root menu. Using tools like `xask` ([www.hdk-berlin.de/~rasca/xask](http://www.hdk-berlin.de/~rasca/xask)), this approach can be extended to interactive scripts.

**FONTS** Several popular distributions use a dedicated daemon, `xf86`, to serve fonts to all applications. This is shown by the presence, in the `XF86config` file, of one font line in this format:

```
FontPath "unix:/7100" # port where the server is listening
```

```
FontPath "/usr/lib/X11/fonts/first_font/" FontPath "/usr/lib/X11/fonts/second_font/"
```

Limiting the number of available fonts and hard code them instead of using a server can often be a good way to speed up

## The RULE Project

Giving new life to old computers

The RULE project aims to modify the Red Hat installer and to create an installation option for Red Hat that will work in less than 32MB of RAM, and consume very little disk space, while remaining fully RH compatible. The project home page is [www.rule-project.org](http://www.rule-project.org).

both the boot sequence and normal operations. To do this, the system administrator should choose the smallest possible set of fonts which will still keep users happy, avoiding (almost always) TrueType and other scalable ones, list them as shown above, and do not start *xfs* at boot anymore. It should be kept in mind that this is one of the areas where the optimal solution varies wildly according to which applications are run most frequently, the users' tastes, and similar issues. Sometimes, turning off the font server may overload X, and TrueType fonts, containing in one set multiple size and encoding information, could save CPU cycles. More details on this issue can be found at [www.dnallounge.com/backstage/src/kiosk/laptops](http://www.dnallounge.com/backstage/src/kiosk/laptops).

**TERMINALS AND SHELLS** Terminal emulators allow us to run text-based program even inside a windowing system, and many people do not need anything more. Sticking to terminals, however, is not by itself enough to decrease RAM and CPU demands. The most modern terminals, like *gnome-terminal*, require a lot of libraries and third party packages to load, and quite some RAM even before running anything inside them: adding transparency or custom backgrounds just makes things worse.

Even the "efficient" ones, however, like *rxvt*, must be configured properly, and be closed as soon as they are not needed anymore: even a lean terminal like *rxvt* set at 80 columns reserves 80 or 160 bytes, depending from the font used, for each lines it is asked to remember. If several thousand lines must be kept available (as happens with most default settings), several megabytes of RAM may disappear just for that. For this reason, whatever terminal is chosen, it should be set to store the minimum amount of lines found usable. In *rxvt*, this setting is controlled by the `-sl <number>` options, or by the `saveLines` resource in `~/.Xdefaults`.

Another useful suggestion on the same line, but valid even in console mode, is to check and reduce the number of commands saved in the shell history. In *Bash*, this number is contained in the `HISTFILESIZE` variable, and its default value is 500.

**UNNEEDED LANGUAGES** Linux can be localised for hundreds of languages and dialects, even those that a commercial company would never bother to maintain. This is excellent for cultural diversity, and must be maintained at all costs. When it comes to working on very limited disk space, however, it must be kept under control. As a matter of fact, several binary packages (*gkrellm* and *CUPS*, for example) install all the locale files they contain, regardless of the language(s) chosen at install time. Removing all the locale settings that will never be used may free space for some more (little) packages.

Therefore, even if at install time only one or two languages have been selected, we recommend looking inside the directories `/usr/share/locale` and `/usr/X11R6/lib/X11/locale`, and, if space is "really" scarce, inside `/lib/kbd` too (on Red Hat: other distributions may have different paths) and remove what is possible.

**CONFIGURATION TOOLS** Though Penguin users stick to editing



text files by hand for any configuration (we too, almost always), and nothing saves more space and CPU than that.

That said, there have never been so many (apparently) user friendly graphical front-ends for configuration as there are today – should they be all ignored?

Maybe not. Our main concern about fancy, GUI based tools is that on most systems configuration happens very seldom: why waste disk space permanently to set up a firewall once and never touch it anymore?

Again, there no single ideal solution. A good strategy may be to learn to manage by hand only those settings that need to be changed often, and use front-ends only for one time tasks. Even when using front-ends, prefer console-based to graphical ones, if available. Speaking of console driven configuration, do not forget the *webmin* plus *links* pairing. The former is a multipurpose system configuration tool which can be accessed through a web interface: if that happens to be *links* or another text browser, we achieve maybe the best compromise between user friendliness and low resources!

Of course, cheating is allowed, if not outright encouraged here. If a certain GUI tool is really much easier to use, install it, configure, and then remove it! Even better, when possible, run it on some other, more powerful, computer, and then copy the resulting files on the older machine (applying the proper dose of common sense!!)

**PACKAGE SPECIFIC OPTIMISATION** The final thing one should do to make one's system faster is to make the most used programs lighter. Since this can mean anything, we'll use *Emacs* as a test case to make our point clear. On Red Hat, the *Emacs* RPM requires about 31Mb of disk space. This include *PostScript* and *TeX* reference cards, tutorials in seven or eight languages, ancient calendars and some games. Removing all this stuff would save more than thirteen megabytes. The *lisp/language* and *lisp/progmodes* directories contain almost every human and programming language known on earth: removing those that will never be used will give up some more megs.

## Conclusion

Linux can still be small and efficient. Apart from this or that trick, we hope to have demonstrated that there are plenty of reasons to keep it that way, and that it does not take an expert programmer to do it: just start to experiment, and enjoy. 

**Just take care with the packages that you install and old need not be slow.**

## NEXT MONTH

A lean, well tuned base system is mandatory to give new life to an old computer, but how does one make it useful too? In the next issue, we will present a collection of Free programs giving the maximum real functionality at the minimum price.

UNICODE, IO AND THREADS

# A quick tour of Perl 5.8



**Charlie Stross** gives a whistle-stop tour of what's new in the latest version of Perl, released last month.

**P**erl was first released to the public around 1987, and evolved rapidly into Perl 4.0; this was the standard version until the first release of Perl 5, in 1994. Perl 5 has mushroomed in popularity and is the standard flavour of Perl; work has been underway since 2000 on developing a radical successor (Perl 6, which we will cover in great detail in the next Perl tutorial in *Linux Format*), but for the time being Perl 5 has progressed slowly, with an emphasis on bug fixes and stability improvements rather than changes to the core language.

Since 2000, we've been running Perl 5.6 (actually 5.6.1 for the latest patched release); this is the stable branch of the Perl development tree, and unless your Linux system is more than two years old or you like installing bleeding-edge development releases, it's the version on your computer right now. Development of the Perl 5 tree since 5.005 (released in 1998) has followed the naming convention of the Linux kernel; that is, there's an even-numbered stable version, and an odd-numbered development tree. Around April 2002, the Perl 5.7 development branch was considered stable enough to start building release candidates of Perl 5.8: Perl 5.8 was officially released in July 2002.

What has this got to do with Perl 6.0?

The answer is: very little. Perl 6 is a complete redesign of the core language, from the ground up. When it surfaces, it will probably bear a slightly closer relationship to Perl 5.x than Java does to C++ – it will be recognisably of the same family, and most Perl 5.x code will actually compile under Perl 6, but it will fundamentally be a new language, at least as different as Perl 5 was from Perl 4. (Perl 5 added references, object-orientation, and modules – not exactly minor changes.) But Perl 6 is still some way off, and before it arrives there will be a Perl 5.10 release. For now we working stiffs are stuck with Perl 5.8. So what's changed?

## Read me first

Perl 5.8 is a maintenance release, but one with an eye on Perl 6. We know – from the list of RFCs and Larry Wall's *Perl Apocalypse* papers – a little bit about what features to expect in 6.0, so it's no surprise to see funny stuff happening around the I/O side of things. There's a full list of changes in <http://dev.perl.org/perl5/news/2002/07/18/580ann/perldelta.pod> – but here's an overview of the gotchas you'll run up against.

There are three major aspects to Perl 5.8. Firstly, it's not binary-compatible with existing XS (extension system) modules – the whole input/output system has been ripped out from under the hood and replaced. Secondly, Unicode support has been

beefed up considerably, with several side-effects. And finally, the old multi-threading model has been tossed on the scrapheap and replaced. Most existing Perl 5.6 code will run happily enough on Perl 5.8, but there are some constructs that will fail as a result of these changes – we'll tackle them in turn.

Binary incompatibility can be a major gotcha when upgrading Perl versions. Because some Perl modules include extensions written in C and compiled to shared libraries (XS modules), you will need to reinstall all your existing modules (see the box, *Installing Perl 5.8*). More importantly, you must ensure that old binary modules don't exist in the **@INC** search path of your new Perl, otherwise you may experience erratic segmentation faults. (This is a particular problem on MacOS X, and may affect you if you installed Perl in some non-standard location, but if your Linux installation uses the default settings you should be alright.)

Perl traditionally provided file handles as a user-level abstraction for dealing with input and output. Perl 5.8 still uses them, but the underlying C library Perl relies on – *stdio* – has been replaced by the *PerlIO* framework. *PerlIO* relies on a lower level library to handle direct input/output to files or operating system devices. As a result, it allows layers to be added that do **\n** to **CRLF** translation, or some other useful task, or to talk to different types of file store. Layers can use different buffering schemes, and extra layers can be inserted under Perl – for example, to translate between Perl's native character encoding (Unicode UTF-8) and whatever native format is used by the operating system.

This is an important move for the future, but has several side-effects. Firstly, any modules that use XS need to be recompiled when you switch to Perl 5.8 from 5.6.1. Secondly, XS modules that aren't *PerlIO*-aware may be unsupported in future – this probably won't affect you immediately, because the *PerlIO* system is designed to look identical to the older *stdio*-based interface, but it may have effects on modules that try to do odd things to file handles. Globbing on filehandles is deprecated – we're supposed to use IO objects instead, when passing references to data sources around. And there are changes to the way layers are handled: the **:raw** layer (aka "discipline") is now formally defined as equivalent to **binmode()**.

There are some other fun effects. For example, the old **IO::Stringy** module is now obsolete: it's legal to open a file handle on a variable:

```
open($fh, ">", \$_output)
```

This directs the output of writes to **\$fh** into the scalar **\$\_output**. And you can create anonymous temporary files: **open(\$tmpfile, "+>", undef)**

Unicode support was added to Perl in 5.6; in a nutshell, Unicode is a character set (and encoding scheme) that is intended to supplant the old ASCII character set by providing

support for just about any writing system, including the largest Chinese, Japanese and Korean dictionaries. Unicode uses a number of encoding schemes, including UTF-8, a transitional 8-bit scheme roughly equivalent to the traditional Latin-1 character set, but Unicode characters aren't bound to any integer width. Unicode characters consist of a "code point" (an entity, such as **LATIN CAPITAL LETTER A**) and various modifiers (such as **COMBINING ACUTE ACCENT**). Code points also have properties (**uppercase, lowercase, punctuation**) and collating sequences. The combination of a code point and its modifiers and properties is called a "combining character sequence".

Perl 5.8 is the first fully unicode-compliant release of Perl. Normally, if all code points in a string are of value **0xFF** or less, Perl treats the string as being of the native 8-bit character set; otherwise it assumes that the string is UTF-8 encoded. If you specifically want to output UTF-8, you can use the **:utf8** output layer in *PerlIO* by explicitly attaching it to a filehandle with **binmode()**:

```
binmode(STDOUT, ":utf8");
```

You can use other output layers too:

```
open($fh, ">:crlf:utf8", "myfile.$$")
```

Applies a **CR->CRLF** filter layer and the UTF-8 translation layer to **myfile.\$\$** when it is opened for output.

You can create Unicode characters in string literals in Perl by using the **\x{}** notation in double-quoted strings or regular expressions, or **chr()** to return a unicode character at runtime:

```
my $smiley = "\x{263a}";
```

```
# or
```

```
print "Smiley detected!\n" if $string =~ /\x{263a}/;
```

The basics of Unicode handling are explained in the POD documentation "perluniintro" – if you're likely to have to handle Unicode text you really need to read this, because it explains how to apply Perl's text mangling capabilities to these character sets.

A few related things have happened to string handling in the migration to unicode. *E.g.*, the string relational operators **ge**, **lt**, **eq**, and so on used to have uppercase aliases (**GE**, **LT**, **EQ** ...). These have now been dropped. A couple of unimplemented

POSIX regular expression features that formerly failed silently now cause fatal errors, and so on.

Threading is a hairy subject; essentially, when you spawn a thread you tell your program that execution can proceed in parallel instances of the same program, with some access to shared data. The new *ithreads* implementation forces data sharing to be explicit, rather than implicit – it's explained in the "perlthrtut" POD file. *Ithreads* is now considered stable. (I'm not going to go into it here – threading with *ithreads* will be covered in a future tutorial.)

## Signal handling

Maybe as a side-effect of the multithreading work, Perl 5.8 has considerably beefed up its signal handling capability. Signal handling is not handled robustly – signals are deferred until Perl finished processing the current opcode, in order to prevent them from corrupting Perl's internal state. However, use of signals to break out of potentially blocking operations is still possible.

On top of these three significant changes (Unicode, *PerlIO*, and *ithreads*), a whole load of new modules have found their way into the core Perl distribution. For example, there are now switch and case constructs in Perl – just use the **Switch** module:

```
use Switch;
```

```
switch($key) {
  case "a" {print "you pressed 'a'\n"}
  case "b" {print "you pressed 'b'\n"}
  case "q" {print "quitting"; last;}
  else {
    # do something here
  }
}
```

There's no substitute for reading the *perldelta* pod document; a whole lot has changed in 5.8. However, for the most part it will be pleasant experience (unless you rely on **taintperl** or on globbing filehandles, both features that have died or are on the way out). In particular, most of the changes make life easier – for example, the new *PerlIO* layers make a bunch of IO modules obsolete and unnecessary. [LXF](#)

## Installing Perl 5.8

### Safe upgrading

Installing Perl 5.8 goes pretty much the same as for any previous version of Perl. If you don't want to use pre-packaged RPMs from your Linux distributor, go to a mirror of CPAN – the combined Perl archive network – such as <ftp://ftp.demon.co.uk/pub/perl/> CPAN. Look in the "src" subdirectory and grab the file `perl-5.8.0.tar.gz`. Then become root and type the following magical incantation (bearing in mind that it'll take some time to run):

```
tar xvf perl-5.8.0.tar.gz
cd perl-5.8.0
./Configure -des
make
make test
make install
cd /usr/include && h2ph *.h sys/*.h linux/*.h
cd -
```

### installhtml --help

This should – if nothing blows up – tell the Perl distribution to autoconfigure itself, compile and test itself, install the results, then create the Perl header files and install the help text as HTML.

The place where Perl installs itself is usually `/usr/local`; it's dictated by the file in the "hints/" subdirectory of the Perl source tree that corresponds to your operating system. (You can live dangerously and tell Perl 5.8 to install in `/usr` by either running `Configure` interactively, without the **-des** arguments, or by editing `hints/linux.sh` or the `config.sh` file that `Configure` generates.)

This installs a new copy of Perl, but it doesn't convert all your old modules over. To do that, *before* you install your new Perl you should do the following with your old Perl:

### perl -MCPAN -e autobundle

"Autobundle" generates a special bundle file – a listing of all the modules installed under your current Perl's library tree. The bundles are written into your `.cpan/Bundle` subdirectory (with a name beginning "Snapshot" followed by the current date – such as **Snapshot\_2002\_07\_22\_00.pm**).

If you generate a bundle file, you can make your freshly installed Perl reload all the modules listed in it by first configuring the CPAN module (type `perl -MCPAN -e shell` and answer the questions), then telling CPAN to install the bundle:

```
perl -MCPAN -e install
Snapshot_2002_07_22_00
```

As long as the bundle is in your **@INC** search path, Perl will find it and reinstall each module listed in it.



## BUILDING WITH ANT

# Speaking Java

**Richard Drummond** does a little housekeeping and tidies up his build process with Ant, the Java-based make equivalent.



You can find the source code for this month's tutorial and the latest version of *Ant* on the coverdisc in the drawer Magazine/Java

**W**hen writing maintainable software, programmers often forget to document the build process. You may write clear and well-documented code, but if nobody can compile it, it's not much good.

The easiest solution to this problem is self-documentation – that is, to automate the build process using make or some similar tool.

The problem with make when building Java projects is that make is designed to build C programs: it's not a good fit for Java at all. Makefiles can be unwieldy at the best of times, and when you try to apply make to the problems of creating a cross-platform build process for Java, it suddenly gets a lot more hairy. This prompted James Duncan Davidson to write *Ant*. *Ant* is a build tool designed from the outset for Java. And in case you're interested, the name *Ant* is because *Ant* is a little thing which can be used for building big things. (*Ant* is part of the *Apache Jakarta* project. See <http://jakarta.apache.org/ant/> for more information and complete documentation.)

*Ant* works much like *make* in certain respects. A project has targets, which are stages of the build process, and rules (called tasks in *Ant*), which accomplish those targets. In other respects, *Ant* is very different. *Ant* is written in pure Java and *Ant* build files are written as XML documents, they are not scripts. Instead of writing shell commands to build a target, a target is accomplished by built-in tasks specified as XML elements and attributes. Thus *Ant* has a **javac** task to compile java source code, a **copy** task to copy files, a **jar** task to build a jar file, and so on. Not only are *Ant* build files simpler to use and read than traditional makefiles, using XML means that the hierarchical structure of the build process is readily apparent. It's not any less flexible than make, either, since anybody can extend *Ant* by writing custom tasks.

## Show, don't tell

While *Ant* is conceptually straightforward, it would take more space than we have here to discuss it any great detail (for that you should definitely read the *Ant* manual). In this intro I'm just going to dump in you in at the deep-end with an example:

```
class Hello
{
    public static void main( String args[] )
    {
        System.out.println( "Hello World" );
    }
}
```

The following file is an *Ant* build file which will make the above class. To use it, enter it and save it as a file called 'build.xml' in the same directory as the above file. To build, (assuming *Ant* is in your command path) simply invoke *Ant* from that directory with **ant**

You can specify a build file with the switch **-buildfile** or let *Ant* default to the file called 'build.xml' in the current directory.

```
<?xml version="1.0"?>
```

```
<project name="HelloWorld" default="compile" basedir=".">
  <property name="src.dir" value=""/>
  <property name="build.dir" value="build"/>

  <target name="prepare">
    <mkdir dir="${build.dir}"/>
  </target>

  <target name="clean">
    <delete dir="${build.dir}"/>
  </target>

  <target name="compile" depends="prepare">
    <javac srcdir="${src.dir}" destdir="${build.dir}"/>
  </target>
</project>
```

The root element of an *Ant* build file is the **project** element. The only mandatory attribute here is **default** which specifies the default target to build when none is passed to *Ant* on the command-line. The **name** attribute assigns the project an ID, while **basedir** sets the directory against which all other paths specified in the project will be relative.

The build targets are defined by **target** elements, one for each target. Each target is named with the **name** attribute, and targets may depend on other targets with **depends** and a comma-separated list of targets. In the above, the target **compile** depends on **prepare**; that is, the **prepare** target must be fulfilled before **compile** can be. To invoke a target other than the default target, you simply pass that target's name as a parameter to *Ant*. Thus to execute the **clean** target, you would call *Ant* with **ant clean**.

Each target contains the tasks (equivalent to make's rules) which will build that target. *Ant* has a number of core tasks built in for most of the jobs that need to be performed to build a typical Java project. This includes general file manipulation, as we can see in the **prepare** and **clean** targets above. The **mkdir** task creates a directory, and **delete** deletes files or directories.

You may have noticed the constructs in the various tasks above – such as **\${build.dir}** – which look like variable substitution. You're half right. *Ant* build files are not programs, and so there are no variables, but symbolic substitution is supported. The **property** element associates a string value with a key; the value of the key **foobar** can be substituted in a string literal with **\${foobar}**. The properties we have defined above, which specify source and build directories, simply serve to make our build file more maintainable.

Properties may also be used to implement conditional compilation. If you specify the **if="<property>"** as an attribute of a target, that target will be executed only if that property is defined; conversely, if you specify **unless="<property>"**, the target will execute only if the **property** is undefined.

*Ant* properties map on to Java properties. Thus any of the standard Java properties are accessible from *Ant*, such as **os.name** (the host operating system's name) and **usr.home** (the user's home directory). *Ant* also has a number of built-in properties such as **basedir** which gives the base directory of the project and **ant.file** which gives the path of the current build file. Properties can be specified on the command-line to *Ant* with the **-D** switch as is usual when invoking Java programs.

## Building code

The meat of the example build file given is the target **compile**. This contains the *Ant* task **javac** which invokes the Java compiler to compile Java source code into Java byte code. By default, *Ant* will use *javac*, from Sun's JDK, but other compiler's are supported with the attribute **compiler** or by setting the **build.compiler** property. To use *Jikes*, for example, set either of these to **jikes**.

The root of the directory tree containing source files to build is specified with the attribute **srcdir** and compiled bytecode is placed in the tree rooted at **destdir**. The *javac* task scans **srcdir** for Java source files and compiles those that have no corresponding class file in **destdir** or are newer than the corresponding class file. Thus with the **javac** task you can simply build a tree of Java source code. You don't need to name individual files and you don't need to specify dependencies. (*Ant* itself doesn't calculate dependencies: this is up to the Java compiler.) You can add or remove source files to your project and the build file does not need to be modified, making maintenance simpler.

You can achieve finer control over which Java source files *Ant* will build with some attributes of the **javac** task. The **includes** attribute lets you specify a coma-separated list of files (including shell-style wildcard patterns) which *Ant* will build. Similarly, the **excludes** attribute lets you list files to omit. More flexibly perhaps you can also specify files to build and omit by using nested **include** and **exclude** elements in the body of the **javac** task. For example:

```
<javac srcdir="{src.dir}" destdir="{build.dir}">
  <include name="**/*.java">
  <exclude name="**/Test*">
</javac>
```

will build all Java files except those beginning with the word **Test**. The pattern **\*\*** matches zero or more path elements.

The classpath that the Java compiler will use may be declared with the **classpath** attribute. Other attributes which control the compiler's behaviour include **optimize** which indicates whether compiler optimisation should be turned on (it defaults to off), and **target** which lets you specify the target virtual machine to compile for. The default value depends on which compiler you use, but for instance, if use JDK 1.4, then you must specify **target="1.1"** if you wish your compiled code to run on a 1.1 JVM.

## Back to the index

The following is a simple build file for our Index-searching project. If you take a look at the code on the coverdisc, you'll find this is a re-organised version of last month's code. All the source code is now in the directory 'src' and I've arranged the source files into packages, rather than simply letting them be in the default package. Thus the code for building constraint expressions is in the package *com.linuxformat.expression*. It is a good idea to put your code into packages, since it avoids namespace collisions with other people's code. (The standard practice is to use your

## Building Ant

### A straightforward guide

Unpack the *Ant* source tarball and enter its folder with:

```
tar jxvf jakarta-ant-1.5-src.tar.bz2
cd jakarta-ant-1.5/
```

Before you can start, you first need to specify where your *JDK* is installed. Do this by specifying it's location in the environment variable

**JAVA\_HOME**. For example,

```
export JAVA_HOME=/usr/lib/j2sdk1.4/
```

Now you can build with

```
./build.sh -Ddist.dir=/usr/local/ant jars
```

The property **dist.dir** specifies where you

want *Ant* to be installed on your system. We've chosen **/usr/local/ant**, so we need to gain root access before we can install proper. Do this with **su** or similar. Now enter

```
./build.sh -Ddist.dir=/usr/local/ant dist-lite
```

to complete the install. Alternatively, pass the target **dist** instead of **dist-lite** to build and install the Javadoc documentation as well. Once installed, *Ant* can be invoked with

```
/usr/local/ant/bin/ant
```

(Or you might want to add the directory **/usr/local/ant/bin** to your command **PATH**.)

domain name in reverse as a basis of your package names.)

```
<?xml version="1.0"?>
<project name="LXF32" default="compile" basedir=".">
  <property name="src.dir" value="src"/>
  <property name="build.dir" value="build"/>
  <property name="jar.name" value="indextest.jar"/>
  <property name="jar.mainclass"
value="com.linuxformat.index.TestQuery"/>

  <target name="prepare">
    <mkdir dir="{build.dir}"/>
  </target>

  <target name="clean">
    <delete dir="{build.dir}"/>
    <delete file="{jar.name}"/>
  </target>

  <target name="compile" depends="prepare">
    <javac srcdir="{src.dir}" destdir="{build.dir}"/>
  </target>

  <target name="jar" depends="compile">
    <jar jarfile="{jar.name}" basedir="{build.dir}">
      <manifest>
        <attribute name="Main-Class" value="{jar.mainclass}"/>
      </manifest>
    </jar>
  </target>
</project>
```

This simple build file takes the source code in the directory 'src' and puts the compiled byte code in a separate directory 'build'. The only new feature here is the extra target **jar** which builds the compiled code into a jar file. The **jar** task should be self explanatory. This task creates a default manifest for a jar, and the nested **manifest** element lets you include additional attributes to add to the manifest. Here we've specified the **Main-Class** attribute so that we can execute the jar.

You can build the project simply with **ant**, and build the jar file with **ant jar**. You can then run the example with, for example,

```
java -jar indextest.jar LXFCD23.xml.gz apache
```

This will search in the supplied test index file for packages contained the keyword 'apache' [LXF](#)

## NEXT MONTH

We'll get back to our CD index-searching project proper next time, when we'll have a look at creating a user interface. I had intended to do that this issue, but I ran out of time to do the topic justice.



## BUILDING A PROCESS VIEWER

# Systems analysis

**PART 14** This month **Brian Long** builds a Linux process viewer with Kylix.

**T**he goal this month is to return to more traditional Linux programming tasks. It is very common for Linux coders to build utilities that display information about the system and so this month we will build a graphical process viewer.

First things first, let's peruse some existing process viewers and get an idea of what the app should display. Linux may come with a variety of process viewers, some of which are console apps (such as *ps*), and some of which are written for the GNOME desktop (such as *gtop*) or KDE (such as *ktop* and *kpm*).

Some of the common aspects of these programs are that they display a list of all running processes (although this list can often be filtered down to a more specific list of processes). For each process a certain amount of info is listed. The specific info may vary from one utility to another, but it usually includes the process name, the process identifier (PID), the amount of RAM it is consuming, its status (running or sleeping) the name of the user that started it and so on. Typically you are also able to terminate a process from these tools by sending it an appropriate signal. We'll take these ideas on board in our own Kylix-built process viewer.

### The outline

The application will make use of actions to implement the user-invokable behaviour. These actions can then be connected to menu items or toolbar buttons as required by the UI. If you are unfamiliar with actions, dig out your back issues of *Linux Format*,

we looked at them in *LXF 22* (Christmas 2001). If you don't have that issue you can read the article online at: [www.blong.com/Articles/Kylx%20Tutorial/Part4/Tutorial4.htm](http://www.blong.com/Articles/Kylx%20Tutorial/Part4/Tutorial4.htm)

In this case we will have an action list, main menu and popup menu connected to an image list (so they all have access to some icons). We also require a status bar to display menu item hints and a list view control to display the process information. The popup menu is associated with the list view through the latter's **PopupMenu** property and the status bar will display hints if you set both **SimplePanel** and **AutoHint** to **True**.

The list view needs some columns defined and this can be done using the **Columns** property, as we did back in *Issue 25* (March 2002). You can find that article online by changing the two occurrences of **4** to **7** in the URL above.

In order to allow the process list to update periodically we need a timer on the form. The **Interval** property will need to be set, for example to **2500** (for it to tick roughly every 2.5 seconds), but it would be easy enough to allow the user to choose their own preferred update interval (that's left as an exercise for the reader). The design-time view of the application form can be seen in **Figure 1** with all the columns set up, components renamed, menus designed and actions set up.


We have yet to see what actions will be used and reflected through the menu tree. However you can see some of the menu items in **Figure 2**, which shows the program running and displaying all processes running in a sample system. This finished version of the project can be found on this month's disc as *pslist.dpr*. The rest of this text will look at how the app was built.

### The actions

Every menu item in the application is implemented via an action component. The **Options** menu has a **Close** item that terminates the program by calling the form's **Close** method. It also has a **Kill Selected Process** item that does exactly what its caption says. We'll look into this once we've seen how the process list is built.

The **View** menu is visible in **Figure 2** and has a **refresh** item that updates the process list in its event handler. This same event handler is triggered whenever the timer ticks – it clears down the

Figure 1: **The main form at design time.**



list view then re-populates it with process info. The other four menu items allow you to filter the processes listed in some way. Whichever one is selected has a checkmark placed next to it and the process list is refreshed. Each menu item (or more correctly the action behind it) shares the same undemanding event handler:

```
procedure TfrmMain.actViewProcessesExecute(Sender: TObject);
begin
  actAllProcesses.Checked := False;
  actNonRootProcesses.Checked := False;
  actYourProcesses.Checked := False;
  actRunningProcesses.Checked := False;
  (Sender as TAction).Checked := True;
  actRefresh.Execute
end;
```

As you can see the four actions have their **Checked** property cleared, then the one that was clicked is checked before the main refresh action is invoked.

The refresh action is obviously the where the main functionality resides and is where we will start our proper look at the code, but we should also note the last action, used to display an **About** box through a call to **MessageDlg** (see **Figure 3**):

Now we can delve into the guts of the program..

## Linux processes

Getting info about Linux processes is straightforward as the kernel maintains a memory-mapped file system in the `/proc` directory. No file in that directory tree is a physical file on the drive; instead they all map onto various kernel data structures. The files in `/proc` give information about the system, for example **Figure 4** shows what you find in the `meminfo` file. You also find a number of directories there, most of which have numeric names (see **Figure 5**).

There is one directory for each running process (the directory name matches the process identifier or PID) and the files inside each directory give plenty of information about that process.

**Figure 6** shows an example process directory where you can see all the files are listed as being 0 bytes (that's how much disk space these memory-mapped files consume).

Of potential interest to us are the following:

- **cmdline** – which contains the command-line used to invoke the process
- **cwd** – a symlink to its current working directory
- **environ** – the process environment variables
- **exe** – a symlink to the process executable
- **status** – status info, such as the process name, its memory usage, its UID, its parent process PID and state (see **Figure 7**)

## Quick start system programming

To talk to the Linux API you need to use the *Libc* unit, which contains type definitions, constants and import declarations for the *glibc* routines. For example, whilst UID and PID values are simply integers, *Libc* defines the `__pid_t` and `__uid_t` types that are clearer in their purpose.

You can get help on these *Libc* routines either in the *Kylux* editor (just click on a reference to the routine in the editor and press **F1**) or using the *man* command at the Linux prompt. Both approaches show you a Linux manual page and the routine will be described in C terms (you can see the Pascal declaration by locating the routine in the *Libc.pas* source file – click on **Libc** in your uses clause and press **Ctrl+Enter** to open it).

You can see a PID variable declared as a field of the form class below, as well as declarations of the various helper routines involved in this application.




Figure 2: The process viewer in all its glory.




Figure 3: The process viewer About box.




Figure 4: System information found in `/proc`.




Figure 5: A listing of `/proc`.

```

Konsole - blong@Cube-Mandrake81.Polyhedra: /proc - Konsol
[blong@Cube-Mandrake81 proc]$ ls -l /proc/2306
total 0
-r--r--r-- 1 blong blong 0 Jul 5 15:
lrwxrwxrwx 1 blong blong 0 Jul 5 15:
-r--r--r-- 1 blong blong 0 Jul 5 15:
lrwxrwxrwx 1 blong blong 0 Jul 5 15:
dr-x--x--x 2 blong blong 0 Jul 5 15:
-r--r--r-- 1 blong blong 0 Jul 5 15:
-rw-r--r-- 1 blong blong 0 Jul 5 15:
lrwxrwxrwx 1 blong blong 0 Jul 5 15:
-r--r--r-- 1 blong blong 0 Jul 5 15:
-r--r--r-- 1 blong blong 0 Jul 5 15:
-r--r--r-- 1 blong blong 0 Jul 5 15:
[blong@Cube-Mandrake81 proc]$

```

Figure 6: A subdirectory of /proc.

```

<< uses
    Libc, ...

type
    TfrmMain = class(TForm)
    ...
private
    //Some convenient user name strings
    CurrentUser, Root: String;
    //Directory of process being worked on (/proc/XXX)
    PIDDir: String;
    //PID of process being worked on
    PID: __pid_t;
    //Used to keep same item selected
    LastSelectedPID: String;
    //An easy to read version of /proc/XXX/status
    ProcessStatus: TStringList;
    procedure LoadTextFile(const Filename: String; List: TStringList);
    procedure GetProcessInfo;
    function GetProcessStatusLine(const Prefix: String): String;
    function GetProcessName: String;
    function GetProcessExe: String;
    function GetProcessUser: String;
    function GetProcessStatus: String;
    function GetProcessMemVal(const Prefix: String): String;
    function GetProcessCWD: String;
    function GetProcessCmdLine: String;
end;

```

The following listing shows the form's **OnCreate** and **OnDestroy** event handlers. Apart from ensuring that any subsequent file access is read-only, setting up a string list for later use and kicking off the program with all processes listed, the code calls a couple Linux APIs to find the name of the root user (which will almost certainly be root, but we make sure) and of the current user.

```

procedure TfrmMain.FormCreate(Sender: TObject);
begin
    //Default to safe file access mode
    FileMode := fmOpenRead or fmShareDenyNone;
    //Set up root & current user name strings
    Root := getpwuid(0)^.pw_name;
    CurrentUser := getpwuid(getuid)^.pw_name;
    //Set up string list for status files to be read into
    ProcessStatus := TStringList.Create;
    //Do first read of processes
    actAllProcesses.Checked := True;
    actAllProcesses.Execute;
end;
procedure TfrmMain.FormDestroy(Sender: TObject);
begin
    ProcessStatus.Free;
end;

```

```

Konsole - blong@Cube-Mandrake81.Polyhedra: /pro
[blong@Cube-Mandrake81 2306]$ cat status
Name: Kylix
State: S (sleeping)
Pid: 2306
PPid: 2175
TracerPid: 0
Uid: 502 502 502 502
Gid: 502 502 502 502
FDSize: 256
Groups: 502 504 81 80 43 22
VmSize: 91532 kB
VmLck: 0 kB
VmRSS: 36852 kB
VmData: 60416 kB
VmStk: 24 kB
VmExe: 504 kB
VmLib: 25772 kB
SigPnd: 0000000000000000
SigBlk: 0000000000000000
SigIgn: 8000000000011000
SigCgt: 00000003800004da
CapInh: 0000000000000000
CapPrm: 0000000000000000
CapEff: 0000000000000000
[blong@Cube-Mandrake81 2306]$

```

Figure 7: A process status file.

**getpwuid** takes a user's UID and returns a pointer to a record of useful information about that user including their name in the **pw\_name** field. This field is defined as a **PChar** (a C-compatible string) but assigning it to a **String** will successfully copy its content to our variable.

Note the ^ symbol used to de-reference the pointer returned from **getpwuid**. Also note that the man page does not advise that we free the memory pointed to by the returned pointer, so presumably the function returns the address of a record maintained by the Linux libraries. As you can tell from the code, root's UID is zero and the UID of the current user is returned by calling the trivial **getuid** function.

## Listing Processes

The main refresh action has the job of iterating across each of those process directories, reading the required pieces of information and adding them to a new row in the list view. A list view row is represented by a **TListItem** that is managed by the list view's **Items** collection property. The first column is represented by its **Caption** property and the additional columns are all stored in the **SubItems** property, a **TStrings** object.

The first thing the routine does is wrap the whole block between calls to methods of the list view's **Items** property. **BeginUpdate** and **EndUpdate** prevent the list view from updating its appearance whilst we clear and refill it. Note the call to **EndUpdate** is protected through a **try..finally** statement.

```

procedure TfrmMain.actRefreshExecute(Sender: TObject);
var
    SearchRec: TSearchRec;
begin
    IstvProcessInfo.Items.BeginUpdate;
    try
        if Assigned(IstvProcessInfo.Selected) then
            LastSelectedPID := IstvProcessInfo.Selected.SubItems[1]
        else
            LastSelectedPID := "";
    IstvProcessInfo.Items.Clear;
    if FindFirst('/proc/', faAnyFile, SearchRec) = 0 then
        try
            repeat

```

```

if SearchRec.Attr and faDirectory > 0 then
begin
  PID := StrToIntDef(SearchRec.Name, 0);
  if PID <> 0 then
  begin
    PIDDir := Format('/proc/%s/', [SearchRec.Name]);
    GetProcessInfo;
  end
end
until FindNext(SearchRec) <> 0
finally
  FindClose(SearchRec)
end;
end;
finally
  IstvProcessInfo.Items.EndUpdate
end
end;

```

The next thing it does is remember which process was selected in the list view so it can be re-selected (if still running) when we've finished refreshing the list. Remember the list gets refreshed every 2.5 seconds and it could be irritating if you select a process (e.g. to kill it) and it deselects itself a second or two

later. The info it requires is in the 3rd column, so can be found in the second string in the list view item's **SubItems** property.

After clearing the list the code uses **FindFirst**, **FindNext** and **FindClose** to iterate through the contents of /proc. Each item is checked to see if it is a directory and if so, whether its name can be treated as an integer (i.e. a process directory). For each one that meets the criteria, the form's PID data field is set to the process PID and PIDDir is assigned the full path of that directory.

Next a helper routine, **GetProcessInfo**, takes over to add the details into the list view:

```

procedure TFormMain.GetProcessInfo;
var
  Item: TListItem;
  State, User: String;
begin
  LoadTextFile(PIDDir + 'status', ProcessStatus);
  //If user only wants running processes, leave for any other state
  State := GetProcessStatusLine('State:#9);
  if actRunningProcesses.Checked and ((Length(State) = 0) or
  (State[1] <> 'R')) then
  Exit;

```

```

  //If user only wants non-root processes, leave if this is root process >>>

```

## Code Templates

Save on typing




Figure 8: Selecting a Code Template from the list.

Note that when implementing classes, or entering any other common code snippets, you can save on the typing by using *Kylux* 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




Figure 10: Customising the editor productivity options.

that can be entered by pressing **Ctrl+J** and selecting from the list (see **Figure 8**). 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 9**).

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 10**. 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



Figure 11: Adding a new Code Template.




Figure 9: Selecting a Code Template from a smaller list.

box and enter a name and description of your new Code Template (see **Figure 11**). Then you can enter the code snippet you want to be represented by the Code Template. An example is shown in **Figure 12**. 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 12: A simple Code Template for a call to MessageDlg.

```

<< User := GetProcessUser;
if actNonRootProcesses.Checked and (User = Root) then
  Exit;
//If user only wants user processes,
// leave for other users' processes
if actYourProcesses.Checked and (User <> CurrentUser) then
  Exit;
//Add row to list view and fill in the columns.
//Also try and re-select last selected process
Item := lstvProcessInfo.Items.Add;
Item.Caption := GetProcessStatusLine('Name:#9);
Item.SubItems.Add(GetSymLinkPath(PIDDir + 'exe'));
Item.SubItems.Add(Format('%d (%:0x)', [PID]));
if Item.SubItems[colPID] = LastSelectedPID then
  Item.Selected := True;
Item.SubItems.Add(User);
Item.SubItems.Add(State);
Item.SubItems.Add(GetProcessStatusLine('VmSize:#9));
Item.SubItems.Add(GetProcessStatusLine('VmRSS:#9));
Item.SubItems.Add(GetSymLinkPath(PIDDir + 'cwd'));
Item.SubItems.Add(GetProcessCmdLine);
end;

```

Since most of the information we require comes from the process status file this is read straight away. However due to the file reporting its size as 0, various common *Kylix* file-reading techniques don't cut the mustard. For instance it would be convenient to read this text file into a **TStringList** via its

**LoadFromFile** method, but this gets us nowhere fast.

Another helper routine uses basic Pascal text file reading code to keep reading lines from the file until either a blank line is returned or an I/O error occurs. Normally I/O errors cause exceptions but the **IOChecks** compiler directive in the listing stops this happening; instead the **IOError** variable is set to a non-zero value if something goes awry. In normal situations you'd use the **Eof** function to see when you reach the end of the file, but this is not a normal file.

```

{SIOCHECKS OFF}
procedure TfrmMain.LoadTextFile(const Filename: String; List:
TStrings);
var
  F: TextFile;
  S: String;
  Stop: Boolean;
begin
  AssignFile(F, FileName);
  Reset(F);
  try
    List.Clear;
    repeat
      ReadLn(F, S);
      Stop := (IOResult <> 0) or (S = '');
      if not Stop then
        List.Add(S)
    until Stop
  finally
    CloseFile(F)
  end;
end;

```

Going back to **GetProcessInfo**, the code loads the status file into a **TStringList** that you can see created in the form's **OnCreate** event handler and freed in the form's **OnDestroy** event handler above. This string list is examined by additional routines that look

for particular pieces of information about the process under scrutiny.

Before committing to adding a new row to the list view, the code examines the actions behind the process filter menu items to see whether the selected process should be listed or not. If not, the code exits from the routine to cycle round to the next process.

If the **Running Processes** menu item is checked then we should only deal with processes whose state indicates they are running. You can see an example state string in **Fig 7**; it starts with **S** for a sleeping process and **R** for a running process so anything other than an **R** means we should skip out. If the **Non-Root Processes** item is checked and the process owner is found to be root we also skip out and the same is true if the **Your** processes item is checked and the process owner is not the current user.

To get the state and process owner (and other information) from the status file a couple more helper routines are employed.

The common one is **GetProcessStatusLine**, which takes some target string at the beginning of the sought line in the file and returns the remainder of the line, if found. You can see above that **GetProcessInfo** calls **GetProcessStatusLine** to get the process name, status, virtual memory size (**VmSize**) and resident memory size (**VmRSS**). Notice that it is careful to pass the correct suffix, for example Name: followed by a tab (character 9) to get just the important part of the line back.

```

//Locate a particular line from the status file, given a prefix string
function TfrmMain.GetProcessStatusLine(const Prefix: String):
String;
var
  I: Integer;
  Line: String;
begin
  Result := '';
  for I := 0 to ProcessStatus.Count - 1 do
    begin
      Line := ProcessStatus[I];
      if Pos(Prefix, Line) > 0 then
        begin
          Result := Copy(Line, Succ(Length(Prefix)), Length(Line));
          Break;
        end
      end
    end;
end;

function TfrmMain.GetProcessUser: String;
var
  UID: String;
  TabPos: Integer;
begin
  UID := GetProcessStatusLine('Uid:#9);
  TabPos := Pos(#9, UID);
  //Typically there are 4 UID values separated by tabs, but check
  if TabPos > 0 then
    Delete(UID, TabPos, Length(UID))
  else
    UID := Trim(UID);
  Result := getpwuid(StrToInt(UID))^ .pw_name
end;

```

**GetProcessUser** has a bit of extra work to do since the **Uid:** line in the status file contains four variants on the UID (which include the effective UID and sticky UID); the real one is the first one so it looks for the tab that follows the first number and trims the rest of the string away. The resultant string is turned into the

integer UID and the user name attained using `getpwuid` again.

The remaining helper routines called from `GetProcessInfo` are used to find the underlying process executable, the current working directory and command line used to initiate the process.

If you recall from **Figure 6**, the executable and current working directory are both symlinks and they can be followed by the `realpath` *libc* routine. Note that if you don't have permissions to see where the link points, `realpath` gives you the same path you gave as input so the code checks. We'll see how to ensure the app runs with root privileges later for those who don't know.

```
function TfrmMain.GetSymLinkPath(const SymLink: String): String;
```

```
var
```

```
RealPathBuf: array[0.. POSIX_PATH_MAX] of Char;
```

```
begin
```

```
realpath(PChar(SymLink), RealPathBuf);
```

```
if SymLink <> String(RealPathBuf) then
```

```
Result := RealPathBuf //if nothing doing, return N/A
```

```
else
```

```
Result := 'N/A'
```

```
end;
```

Getting the command line is a bit fiddly as the file is a series of zero terminated strings (rather than line feed separated). Also it seems there may or may not be an extra zero terminator after the last string so the helper routine has to keep reading past any zero terminator it finds. If it gets a second one or is told it has read past the end of the file then that's it, otherwise it has started reading another part of the command line. Again, this code could be much simplified if this were a regular disk file.

```
function TfrmMain.GetProcessCmdLine: String;
```

```
var
```

```
F: file of Char;
```

```
Ch: Char;
```

```
begin
```

```
Result := '';
```

```
//The file /proc/XXX/cmdline is a series of zero-terminated strings,
```

```
//possibly followed by an extra zero terminator
```

```
AssignFile(F, PIDDir + 'cmdline');
```

```
Reset(F);
```

```
try
```

```
repeat
```

```
Read(F, Ch);
```

```
if Ch <> #0 then //read till a 0 char
```

```
Result := Result + Ch
```

```
else
```

```
begin
```

```
//We will find an extra #0 at EOF in some cases here
```

```
//otherwise we'll trigger a read past EOF IOError
```

```
Read(F, Ch);
```

```
if (IOResult = 0) and (Ch <> #0) then
```

```
Result := Result + #32 + Ch;
```

```
end
```

```
until (Ch = #0) or (IOResult <> 0)
```

```
finally
```

```
CloseFile(F)
```

```
end
```

```
end;
```

The last thing to look at is how to kill the selected process. The action's event handlers are shown below. The first one uses the simple Linux routine `kill`, which takes a PID and a signal to kill the process with. The `OnUpdate` event handler ensures the action is only available when an item is selected in the list view.




Figure 13: Using the sticky bit to read root's symlinks.

```
procedure TfrmMain.actKillProcessExecute(Sender: TObject);
```

```
var
```

```
PIDStr: String; //e.g. '1955 (7A3)'
```

```
PID: Integer; //e.g. 1955
```

```
begin
```

```
PIDStr := IstvProcessInfo.Selected.SubItems[colPID];
```

```
PID := StrToInt(Copy(PIDStr, 1, Pos('#32', PIDStr) - 1));
```

```
kill(PID, SIGKILL)
```

```
end;
```

```
procedure TfrmMain.actKillProcessUpdate(Sender: TObject);
```

```
begin
```

```
(Sender as TAction).Enabled := Assigned(IstvProcessInfo.Selected)
```

```
end;
```

## Setting up the process viewer

Now that the utility is complete it is important to set it up correctly. If you log in as a non-root user when developing, which is advisable, then the final executable file will be owned by your UID. The upshot of this is that the program will not be able to see the executable or working directory of any root processes as your UID doesn't have permission to read symlinks owned by root. You can prove this by getting a long listing of the process directory of a root process (the `init` process is owned by root and has a PID of 1):

```
ls -l /proc/1
```

The symlinks will not be followed.

You can circumvent this limitation of the program by changing its owner to be root and setting the sticky UID (SUID) bit in its permissions (you'll need to be logged in as root). The sticky bit ensures that no matter who runs the application, it will run with owner's permissions (*i.e.* those of root). **Figure 13** shows the commands needed to change the program's owner to root and set the SUID permission. You can see in the long format listing that the sticky bit is represented by a letter `s` in the owner executable position (it changes from the normal `x` to an `s`).

Also in the screenshot you can see the `kylixpath` script being called to ensure that the `PATH` is set up correctly for *Kylix* applications to be run from the command prompt (as opposed to from the *Kylix* IDE). The process viewer can then be seen successfully reading the executable and working directory symlinks of root's `init` process.

## Summary

We looked at building an app to list details of running processes in the system. This demonstrates that the business of building system level utilities is not only within the realm of `gcc` coders; *Kylix* is just as capable of lower-level systems programming. 

## NEXT MONTH

Next month's will be the last instalment in this tutorial series and we'll round off by looking at building console applications. Until then, happy coding.

## 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).

## GUI BURNING

# Burning with Linux

**PART 2** Robert Smith explains burning for the point'n'click generation, then delves into mixed mode CDs.

Last month I showed you how easy and powerful it was to burn a CD under Linux under the command line. By using a few short commands, an ISO image could be created and then that and audio could be quickly burned to a CD. This month we will look at the various options available for burning CDs under an X interface, and how to use both effectively.

In this article we will look at the following software.

- *Gnome Toaster*
- *KreateCD*

## Gnome Toaster

*Gnome Toaster* is one of the better CD GUI burning tools under Linux. Although the interface may be confusing to begin with, if you have ever used *wincd* in Windows, you should be able to master the program quickly and easily. If not, the following guide will get you started with the basics.

The interface is divided into three main parts, and a toolbar and menubar along the top. The left hand panel (long and thin) allows you to access your filesystem, along with any CD devices or SCSI data devices you may own. The main large panel shows you the files and folders in a certain part of your filesystem, and the tracks on the other storage devices. The bottom panel provides you with your CD layout options, along with the main burning options to create your end CD.

The bottom panel is further subdivided via side graphic tabs to provide these options. Tab 1 which is shown by a folder symbol

is the ISO image creation part of the software. This is where you drag and drop files from your filesystem to produce a data CD. The next tab is shown by a T symbol, which stands for Track, this is where you add audio tracks and existing data tracks, and then reorganise them so that they are in the order you require. The third tab is indicated by a CD symbol. This panel opens up and gives you all the CD burning options that are available under *Gnome Toaster*.

## Configuration

The majority of the settings are already configured when you start *Gnome Toaster* and unless you know exactly how the majority of the commands work, it is not recommended to change most of the options. The main configuration options that you need to make sure are working correctly are the configuration of your CDROM and CDRW drives, and precaching for audio and data. To configure these options click on the preferences button on the toolbar. This should open up a new window.

**Configuring CDROM + CDRW devices** – Click on the “CDROM and Recorder Setup” tab and check to see that your devices have been recognised. If they have not you can add them using the ADD button at the bottom of the page, or choose for *Gnome Toaster* to scan your scsi bus again. If you have a IDE CDROM you can use this by adding a new device and pointing all the options to its normal mount point and its /dev/hd\* device instead of a /dev/scd\* device file. Make sure that it is not selected as a writer, as these need to be found by scsi emulation to work. Click on your CDRW and then click the edit button. Make sure that the “This drive is a CD writer” option is checked.


**Data Caching** – To setup data caching you have to select the following options. Under the “Common” tab click the “Precache Data tracks” option. Then click on the “ISO9660” tab and select the “Precache before writing” tab. This will ensure that caching has been activated. Caching is useful as it means that your system requires less resources when it finally burns and there is less chance of a buffer underun.

**Audio Caching** – This is extremely useful when recording audio off multiple CDs as it records the audio data straight onto your PC. It also reduces the chance of problems, if the audio CD has any physical problems such as scratches which stop consistent reading of the audio.

## Creating a data CD with Toaster.


This process is extremely simple and can be very quick and easy to do. Select the part of the file system that you are wishing to put on the CD using the left hand side panel. Use the + and - signs next to the roots to expand and contract the trees respectively so that you can select the parts of the filesystem. Then click on the folder that you wish to burn and drag it into the Filename/Filesize window in the bottom panel of the program (set to the file tab). The data should then be added to the ISO filesystem. If you wish to place the contents of the folder into the root of the filesystem, select the folder you want in the left hand

**Burn your way through it, with the huge number of burn options in *Gnome Toaster*.**





# LinuxFormatTutorialCDwriting



## Ordering audio around with KreateCD.



```
for ogg in *.ogg; do ogg123 --cdr - "$ogg" |
cdrecord -audio -pad -nofix
done
cdrecord -fix
```

## The Perl way

This way will just convert the files to wav files. These can then be burnt using the

```
cdrecord -audio <file 1> <file 2> etc.
```

command in the order required.


```
#!/usr/bin/perl -w
use strict;

foreach (@ARGV){
my @arr = split(' ', $_);
$arr[0] =~ s/ \\/ /g;
system("mpg123 -w $arr[0].wav $arr[0].mp3");
}
```

Save this in a file and run it with the following command in the mp3 directory. `./<programe> *.mp3`

To do the same for ogg files replace the system line with `system("ogg123 -w $arr[0].wav $arr[0].ogg");` and then run the command `./<programe> *.ogg` This can be repeated in multiple directories and the results burned from

## Checking your options before the burn.



the same directories. Just remember to delete the WAVs afterwards to free up space.

## KreateCD

This is the best Qt based KDE offering for CD burning. It has simple wizards and a great design interface, which means that it is very simple to use. When first started a wizard interface comes up asking you what sort of CD you wish to burn. Choose the type of CD which you want to burn, or choose classic mode. The advanced mode is similar to that of *Gnome Toaster*, so if you are happy with this style use this.

When you first start *KreateCD* you have to configure the program once so that it knows how to burn a CD using your CDRW. To do this click on the "Settings" menu and then the "Configure KreateCD" option. A new window will appear. Click on the "SCSI" tab and select, using the radial buttons, which of your devices is the primary writing device and the primary reading device. To configure the speed at which your machine burns click on the "Burner" tab and configure speed and post burn options. The "ISO Defaults" tab will allow you to configure the sort of ISO9660 filesystem that is created, such as a Rockridge compatible ISO file.

## Audio CD mode

To add audio tracks using this interface, click the "Add Track" button. Then click on the "Track Source" drop down menu. This will bring up 3 options. CD Track, File or Filesystem. To add a track from a CD use the CD Track option and it will scan your reader device and give you a list of the tracks.

Select the one you want and then click ok to add the track. To add a MP3, WAV or OGG file to the audio CD choose the File option, navigate to the file, select it, and then click OK twice to add it to the line up. Do this for all your files, and then use the "Track Up" and "Track Down" buttons to order the tracks. When you have finished click the "Write CD" button at the bottom. Change any final burning options and then click "Write".

## Data CD only Mode

Navigate through the tree on the left panel, and then drag across files and directories to the right panel. Click on the "Calculate ISO-image size" button to check how full the CD is. Once you are happy click the "Write CD" button at the bottom. Creating a data CD really is that easy.

## CD-Copy Mode

This is one of the best features of the package. Put the CD you want to copy in your CDRW device and a blank CD in the CDRW device. Click "Write CD" and the CD will be copied after you have chosen the final options.

## Classic Mode

This allows you to easily create multimode CDs. To change the type of track that you are inserting, click on the "Add Track" button. Under the "Track Type" drop down menu select "Audio" for an audio track, or "Mode1" for a Data CD.

If you choose Audio follow the same instructions as above in the Audio CD tutorial to add the track. If you chose Mode1, and selected filesystem, then a window such as the Data CD only interface will come up. Use this to drag and drop the files into place. Click OK to save the track ready for burning. If you have a data track on a CD that you wish to copy onto a new CD, then click the CDRW option under "Track source" then select the data track on the CD after the TOC has been read.

## Types of Burning

All at once, or a bit at a time

There are many ways in which your CD writing device can write the information to the CD. Disc-at-Once, Track-at-Once and Packet writing.

### Disc-At-Once (DAO)

In this method the CDRW drive writes all the data to the CD without turning off the writing laser. This causes the disk to be finalised after it has been written. In this way, it requires a blank CDR and cannot support multisession.

The advantages of this method, is that there is no gap in between the tracks burnt on the CD, in comparison to Track-at-Once which leaves a 2 second gap between tracks. With an audio CD this is the difference between the tracks appearing continuous or with large gaps in between.

### Track-At-Once (TAO)

With this burning style the CDRW device's laser is turned off after each track has been written even if more than one track has to be burned in a session. This gives more flexibility as you can use multisession CDs, partially used CDs and you can leave the CD open.

Unfortunately you get a 2 second gap in between each track, which may not sound as good on an audio CD.

Track-At-Once can be used with variable-gap; this means that you can control how large the gap is in between the tracks burnt. The gap can be lowered to as little as  $\frac{2}{75}$  of a second.


### Packet Writing (PW)

This is where you can format a CD much like a hard disk, and then write information onto it in small

increments. This is very useful when copying small amounts of data from one computer and taking it elsewhere. Unfortunately this too has drawbacks; the computer requires certain software to access the data and only one track is burnt to the CD, which is data.

All CDRW devices will support TAO mode, and the more modern ones will support TAO with variable-gap.

Unfortunately only a few CDRW devices support DAO, but this can be found out by experimenting with the settings in such programs such as XCDRoast and Gnome Toaster. Packet Writing is still in its infant stage under Linux. More information about this can be found at <http://packet-cd.sourceforge.net> and you will need an experimental kernel module to get this working. You will also have to enable UDF whilst building the kernel.



**KreateCD takes away the hassle of burning complex audio files with its inbuilt decoders.**


There is a large number of X-based CD authoring tools available now for Linux. If you are not happy with the way either *Gnome Toaster* or *KreateCD* then try another, with the tens of programs appearing day by day there must be surely one that you can become acquainted with. Have a look for the following packages: *gcombust*, *cdrtoaster*, *cdbakeoven*, *XCDRoast*. Or if you feel at home just stick to the command line.

Unfortunately the *KreateCD* package is still very much in development. I personally found that the Copy CD option under the package didn't work correctly with audio CDs on my Linux setup but worked fine with data only CDs. The same thing occurred when using the audio track from CD under the other modes.

## User interface

Audio grabbing works from all other X-based CD writing packages, so this leaves me to believe it is a bug with the code in *KreateCD*. The method for adding tracks can be effective when working with small numbers, but when adding similar files this can become a bit time consuming. The great user interface provided by *KreateCD* though, more than makes up for the speed difference and for a beginner this is a good option.

The simplicity of *KreateCD* means that almost anyone can create CDs under linux. In many ways the process is actually easier than using programs under windows such as *Nero* and *Create-A-CD*. Coupled with the command line, you will find that authoring professional quality CD layouts can be easily accomplished. Creating a backup CD can be so quick, that



**KreateCD happily burning a CD.**

you have no excuse not to do so, and this will mean that your data will be well protected from any crashes or hardware failures that may occur.

*KreateCD* is easy and fast, with its drag and drop system creating a easily usable interface but its stability may be questionable and *Gnome Toaster* has very powerful ISO layout tools, and audio conversion tools. This choice will mean that no matter what you want to do, how you want to do it and how quickly, you will be able to accomplish the task under Linux. [LXF](#)

## Further Information

### useful links

<http://gnometoaster.rulez.org>

– Gnome toaster home page

[www.kreatecd.de](http://www.kreatecd.de)

– KreateCD home page

[www.xcdrtoast.org](http://www.xcdrtoast.org)

– XCDRoast home page

[www.fokus.gmd.de/research/cc/glone/employees/joerg.schilling/private/cdrecord.html](http://www.fokus.gmd.de/research/cc/glone/employees/joerg.schilling/private/cdrecord.html)

– CDRecord home page

[www.tidp.org/HOWTO/](http://www.tidp.org/HOWTO/)

[CD-Writing-HOWTO.html](http://www.tidp.org/HOWTO/CD-Writing-HOWTO.html)

– CDwriting HOWTO

[www.linuxlinks.com/Software/CD\\_Writing](http://www.linuxlinks.com/Software/CD_Writing)

– CD writing software links

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

[modules.php?op=modload&name=phpBB\\_14&file=index&action=viewforum&forum=1&6794](http://www.linuxformat.co.uk/modules.php?op=modload&name=phpBB_14&file=index&action=viewforum&forum=1&6794)

– The *Linux Format* help forum (always useful)



## STRUCTURED QUERY LANGUAGE

# Practical PHP programming

Having given teasers for the first two articles in the series, Paul Hudson now starts a two-part SQL primer using PHP.

Many people believe database access in PHP is its most important feature, and the PHP team have indeed made it extremely easy to interact with databases using the language. It's fair to say that a PHP developer who has yet to come into contact with SQL really has only touched the tip of the PHP iceberg.

In this article I will be introducing you to Structured Query Language (SQL) using *MySQL*. Don't expect PHP just yet – that's next month...

## What does it do?

Structured Query Language is a standalone language designed to allow users to manipulate databases. If at this point you're not sure what a database is, see the box *Database jargon made easy*, where we explain some of the jargon.

Being a fourth-generation language, you use SQL to say what you want out of the system – the database manager does all the hard work to provide the requested functionality. So for example, you would use SQL to add information to a database simply by specifying the data you want to add and the place you wish to put it – you wouldn't need to write any sort of code to actually write the data to the database, as this is done for you.

Sound easy yet? It should – simple SQL is very easy to pick up, and, as the language is very small, you tend to re-use your code a lot.

## PHP and SQL?

Through the use of a database management system, PHP allows you to do all sorts with database information – from querying information to extracting it in a structured manner. PHP supports many database systems, but for the purposes of this article (and

indeed the entire series) I will be working with *MySQL*. *MySQL* is fast, simple, open source, has good online documentation, and as such makes a great choice for beginners and experts alike. Not to worry though, much of the information presented here will work in other database systems – see the box *DBMS Choices* for a list of PHP-supported alternatives to *MySQL*.

Note that because SQL really is a vast topic, and that even though I have tried to stick to the core elements, there just isn't enough space this month for me to go into using *MySQL* with PHP. Next month, I'll be spending a lot more time showing you how PHP interacts with SQL. If you can't wait, I suggest you read over the past two practical PHP tutorials for information.

## Installing MySQL

At the time of writing, *MySQL 3.23.51* was the latest stable release available for download from the *MySQL* website ([www.mysql.com](http://www.mysql.com)). While v4 has proved in my own tests to be perfectly stable, I would still not recommend it for use on production servers – it's just not worth the risk.

The *MySQL* team strongly recommend using official (as in, downloaded direct from their web site) binaries. This might seem alien to many of us, but my own personal experience has proved the *MySQL* developers quite correct in this recommendation. If, like me, you're using Debian, you can install official *MySQL* binaries simply by typing

```
apt-get install mysql-server
```

RPM users are provided for on the *MySQL* web site itself – they keep a binary RPM of the latest *MySQL* server ready for download. There are only two real pitfalls you are likely to encounter while installing a *MySQL* binary. Firstly, the RPMs require that you have *glibc2.2* or above installed. You can correct this problem merely by upgrading your *glibc* libraries. The other problem, however, is a little trickier – users running on kernel 2.2.14 may experience corrupted *MySQL* data as a result of an I/O bug in that release of the kernel. The only solution here is to upgrade to a more recent version of the kernel – anything newer should work fine.

For those die-hard source compilers out there and those experiencing problems using a pre-packaged distribution, *MySQL* offers the complete source code for download. They provide a very strict warning for those of you running *GCC 2.96* as it is known to create bad binaries. Furthermore, I ask you to make quite sure you want to compile your own binary – a good *MySQL* binary is very difficult to create unless you have read the manual thoroughly.

Once you have your *MySQL* binary created, go ahead and run the server with the command **safe mysqld**. Your installation may allow you to run *MySQL* using */etc/init.d/mysql start*, which is a

## SQL books

The most popular book on SQL, at least amongst the SQL developers that I know, is *Database Systems* by Thomas Connolly and Carolyn Begg (Addison-Wesley). It provides information on every aspect of databases, but might prove a little too thorough for those wishing merely to dip their toe into the SQL pool.

If you're in that category, try *SQL in a Nutshell*

by Kline and Kline (O'Reilly) – be warned, though, it really is little more than a reference guide, as you would expect from an O'Reilly Nutshell book. For a book specific to *MySQL*, go no further than *MySQL* by Paul DuBois (New Riders) – the book is considered by many to be *the book on MySQL*, and, although it is just a tiny bit dated now, I would still agree.

## Database jargon made aasy

### Putting it into plain English

There are lots of terms used in SQL and database systems, several of which mean the same thing. Here's our quick guide to some of the most popular, and some of their synonyms.

#### DATABASE

An arbitrary group of sets of information called tables.

#### TABLE, RELATION, FILE

An individual set of rows in a database. Each table has its own set of attributes, and each row in the table has a value for each attribute.

#### ROW, RECORD, TUPLE

One entry in a table. This could be a person in a Staff table, a software item in a product table, etc.

#### COLUMN, FIELD, ATTRIBUTE

Each table has many columns, and they describe a certain property of each row. "Age", "Name", "Sex", etc, are all examples of attributes. Make sure your column names don't contain spaces, and, more importantly, aren't keywords like **INT**, **DATE**, etc.

#### PRIMARY KEY

A unique, not null column in a table. This is indexed for you by *MySQL*.

#### STORED PROCEDURE

A stored procedure is like a user function, kinda like a batch file. Many database managers (not *MySQL* yet, sadly) allow you store SQL commands up on the server, pre-compiled for extra-fast performance, and execute them all at once with one call.

#### INDEXES

Indexes are used to find rows with a specific value as quickly as possible. If you choose not to use an index, *MySQL* must sequentially read through every row in your table until it finds the matching data. Sometimes reading data sequentially is not a problem – for example, if you want to read the entire table. However, more often than not, indexes are better – they allow *MySQL* to jump almost immediately to the data you request.

#### NORMALISATION

A normalised table is one that has no data redundancy and satisfies a variety of very strict requirements for table design. We'll be looking at normalisation, which is the process of producing normalised tables, next month.

little easier. *Warning:* By default, *MySQL* installs with no password for the root user. So, the first two things we are going to do is update your root *MySQL* password to something other than blank, then add a user for day-to-day usage. *MySQL* works like Linux in that you should generally not use the root account for day-to-day work.

With the *MySQL* server started, enter **mysql -u root** at the command prompt. The **mysql** command starts the *MySQL* client software with default parameters. If you see

```
Welcome to the MySQL monitor
```

then you have successfully installed *MySQL* and we're ready to configure. Enter the following into the *MySQL* prompt, separating each command by pressing **Enter**.

```
use mysql;
update user set password = PASSWORD('<somepassword>')
WHERE user = 'root';
insert into user (host, user, password) VALUES ('localhost', 'lxf',
PASSWORD('<differentpassword>'));
create database lxfdb;
grant all privileges on lxfdb.* to 'lxf';
flush privileges;
exit;
```

The above commands update the root user password to whatever you entered in place of **<somepassword>**. Note that it is surrounded in a **PASSWORD()** *MySQL* function which converts your chosen password into an encrypted format. We also add a new user, lxf, with a password of its own. A database called lxfdb is created, and we give the user lxf full permissions to the lxfdb database – this allows the user to read data, make changes, etc. Finally, we call flush privileges. This forces *MySQL* to update all users on the system immediately, and you need to remember to do it whenever you make a user/access change on your system.

## First steps

At this point, you should now have *MySQL* installed and configured for your system. Before we proceed, it's important to have a clear understanding of how databases work, so prepare yourself for a little bit of theory...

Firstly, databases are most often likened to filing cabinets – they can hold lots of information in an easily retrievable format. Of course, filing cabinets can be well organised or they can be poorly organised, and this works the same for databases. Like a filing cabinet with many drawers, each database can have many tables. The similarity goes on further, but I think you get the point.

While some people like to have a database for each set of tables (ones that relate to each other), others, especially those virtually-hosted by a large hosting company, often only get one database to themselves, so they put all their tables in one place. In reality, there is little difference between the two – it's all horses for courses.

Each table can have many rows (records), and each row can have many columns (attributes). All rows in a given table share the same column names, but each can have different values. Confused yet? Here's an example of how it might work.

We could have a database "gc\_timbuktu" which contains all the information about a golf course in Timbuktu. Inside this database, we could create several tables, for example "Staff", "Members", and "Equipment". Each table is independent of the other tables in the database, and also independent of any other tables in other databases. Our Staff table could be defined as having four columns: Name, Age, Job, and Pay. This would mean that any rows in our table, which would presumably be staff members, could have information stored about their name, age, job, and pay. We would then add a row to this table for each staff member. Our Members table might be defined as having three columns: Name, DateJoined, and Handicap. Notice that this table has different columns from the staff table, and this allows us to store different, more relevant data about club members. The equipment could contain different columns again, specifically to keep information specific to golfing equipment stored at the club.

If you're not quite sure how it all comes together, don't worry – it will all come clear once you see it in action. From your shell prompt, enter the following commands to login to *MySQL* and create the staff table in your lxfdb database.

```
mysql -u lxf -p<yourpasswordhere>
```



```

<< mysql -u lxf -p<yourpasswordhere>
USE lxfdb;
CREATE TABLE staff (Name VARCHAR(255), Age INT, Job
VARCHAR(255), Pay INT);

```

The **mysql** command, as explained already, logs you into the *MySQL* monitor. This time, we're using the **-u** parameter to specify your username, and the **-p** parameter to specify your password. Note that there should be no space between **-p** and your password!

Once logged into the *MySQL* monitor, we execute two commands (note that SQL is *not* case sensitive on the whole, but I use upper case words to point out commands and functions) – **USE** and **CREATE TABLE**. **USE** takes just one parameter – the name of the database you wish to work with. All subsequent operations (until you quit, or specify otherwise) are considered to be for this database. **CREATE TABLE** lets you, unsurprisingly, create tables in the working database. Note that table names are one of the few case sensitive things in *MySQL*! Each column in the table you are creating is specified by giving it a name and a data type, and also optional modifiers that we'll look at shortly.

Data types are what kind of data a column should hold. In the example above we use two different data types – **VARCHAR** and **INT**. These are short for variable-length character string and integer. A variable-length character string allows you to specify how much data you want to set aside for a value in that column – in our example, we use 255 characters.

Note that **VARCHAR**, being variable in length, only uses as much space as is required. You can specify a column as **CHAR(255)**, which means *MySQL* will set aside 255 characters of space for the value, whether it's used or not. Using **CHAR** is faster, but takes up more room – I'd recommend you stick with **VARCHAR** unless you really want to eke out every last bit of speed.

Other column types include **BIGINT** for very large integers, **TINYINT** for very small integers, **FLOAT(x,y)** for floating point numbers (allows you to create an x length number with y number of decimal places), and **TEXT** for large amounts of text. The

## DBMS Choices

PHP supports a multitude of choices, including (but not limited to) the following: *MySQL*; *MS SQL Server*; *PostgreSQL*; *Oracle*; *IBM DB/2*; *mSQL*; *SESAM*; *Sybase*; *InterBase*.

## IS NULL?

### To be or to be null

A value of **NULL** in your row doesn't mean no text, or 0 like it would for character fields and integers. Instead, it means "Nothing here", and is treated specially in SQL.

For example, you cannot test against **NULL** like you normally would, and this can cause problems. Try these commands which help illustrate the problem:

```

SELECT NULL;
SELECT NULL + 1;
SELECT NULL * NULL;
SELECT NULL = NULL;
SELECT NULL > NULL;
SELECT NULL = 0;
SELECT NULL > 0;
SELECT NULL < 0;

```

See the problem? You simply cannot compare null to any other value without receiving the return **NULL** unless you use special syntax. This can make your tables act in strange and unexpected ways, which is why it is often helpful to use **NOT NULL**.

The correct way to handle **NULL** values is to use the two special **NULL** operators: **IS NULL** and **IS NOT NULL**. For example:

```
SELECT * FROM staff WHERE Age IS NULL;
```

*MySQL* manual has a complete list online of all the different column types available for *MySQL* – it's generally best to be as specific as possible with your data type to ensure maximum speed with the minimum of space requirement.

To make sure your table has been created correctly, try executing **DESC staff**; in the *MySQL* monitor. It returns the list of columns in that table, and also the modifiers on those columns.

Modifiers can include **NOT NULL** to force a value for that column, **DEFAULT <x>** to assign a default value to that column, or **PRIMARY KEY** to make this field special in the table. I'll explain about keys shortly, but for now I want to explain **NOT NULL** and **DEFAULT**.

Defining a column as **NOT NULL** means that when it comes time for you or your users to enter data into the table, a value *must* be provided for this column. Failing to provide a parameter for a **NOT NULL** column results in *MySQL* attempting to provide a value for you – this is most often **0** for integer columns, or an empty string for **VARCHAR** (variable-length character) columns. The reason *MySQL* provides this special functionality is because **NULL** columns are quite tricky to deal with – have a look at the box **IS NULL?** to see why.

The **DEFAULT** modifier allows you to specify a default value for your fields. This can often be helpful – for example, in our golf club scenario, we could set the default handicap score for new users to 30, which means we wouldn't have to keep re-entering the handicap each time we added a new member.

Here's how a **CREATE TABLE** would look for our Members table – note that we're using **NOT NULL** and **DEFAULT** for the Handicap attribute.

```
CREATE TABLE members (Name VARCHAR(255), DateJoined
INT, Handicap INT NOT NULL DEFAULT 30);
```

Try creating the equipment table yourself with appropriate columns and column types.

## Adding data

With our tables created, let's take a look at adding information to them. *MySQL* has two special commands to add and edit data inside a table – **INSERT INTO** and **UPDATE**, with the difference being that **INSERT INTO** adds a wholly new row, whereas **UPDATE** merely amends an existing row.

Using the table definition for Staff given above, here are a few examples of these two commands:

```

INSERT INTO staff VALUES ('Joseph Smith', 28, 'Security Man',
29000);
INSERT INTO staff VALUES ('Harold Barnes', 32, 'Security Man',
29000);
INSERT INTO staff VALUES ('Carmen Major', 30, 'Security Man',
29000);
UPDATE staff SET Age = 29 WHERE Name = 'Joseph Smith';
UPDATE staff SET Name = 'Carmen Hobbes', Pay = 30000
WHERE Name = 'Carmen Major';

```

As you can see, **INSERT INTO** requires a table name, followed by the keyword **VALUES**, and then the values you wish to enter into the table. Note that, unless you specifically say otherwise, you need to insert values in the order they come in the table.

**UPDATE** requires a table name, followed by the keyword **SET**, then the amendments you wish make, and finally qualifiers, if any. In our example above, we change the name of all rows to have the name Carmen Hobbes, and move their pay up slightly as long as the Name column of matches "Carmen Major". Given the

## Tackling versioning problems

### Plus PHP support for Apache 2

Several readers contacted me by email citing problems with the code printed with my prior articles. Putting typographical and printing errors aside (they do happen from time to time, sorry! :), the main problem was that users found they didn't have the `$ _POST` super-global variable available to them, and instead had to rely on using `$HTTP_POST_VARS`.

If you are experiencing this kind of problem, it's best that you upgrade – `$ _POST` (and its sibling super-global variables) is only available in newer versions of PHP, which is why I recommend you use the latest release when trying out

scripts from this series.

Using `$HTTP_POST_VARS` is, on the whole, a fairly good substitute for those unwilling (or unable) to upgrade, except for the fact that `$ _POST` is automatically global across your entire script – even inside functions – whereas `$HTTP_POST_VARS` isn't quite as flexible. The PHP development team has made a very strong push towards the use of the new super-globals, so it's probably best to start getting used to them now!

Also, just as a side note, I've had various queries about PHP support for Apache 2. At the time of writing, with PHP at version

4.2.2, Apache 2 is \*not\* supported in PHP beyond being tagged as experimental. Various people have tried and succeeded in getting PHP and Apache 2 to communicate, but it really isn't reliable, and certainly not recommended – just as many people, if not more, have tried and failed to get the two working together!

Better Apache 2 support is pencilled in for PHP 4.3, and I'd strongly recommend you stick with 1.3.x at least until PHP 4.3 is released, unless you really fancy a challenge :)

Thanks for the questions – please, keep them coming!

three rows we entered above, this will match only one person.

There are several problems with the above layout – firstly, our qualifier in our second **UPDATE...SET** command is matched by name. What would happen if we had two Carmen Majors? Or ten? In that case, we'd need to have:

```
WHERE Name = 'Carmen Major' AND Age = 30
```

But what if we had two (or ten!) 30-year old Carmen Majors? The best way to solve this problem, and I strongly recommend this solution for all tables you create, is to use an extra field: **ID INT**. By giving a unique ID number to every row, you can simply say:

```
WHERE ID = 492" rather than "WHERE NAME = 'foo' AND Age = 30 AND... AND... AND ...".
```

If you're wondering how you would assign a unique number to each ID, that's not a problem – MySQL allows you to specify one auto incrementing column for every table with the **AUTO\_INCREMENT** keyword. Used with the **PRIMARY KEY** keyword, which ensures that you cannot have duplicate values in a column and that the given column is indexed for speed, ID fields become very powerful indeed. See the box *Database jargon made easy* for further information about primary keys and indexes.

The second problem with our layout above is that there is duplication of job title, "Security Man". If we wanted to be politically correct and use the title "Security Officer" rather than "Security Man", we would have to update every row in the table. The process of normalisation (again, see the box *Database jargon made easy*) rids your tables of duplicate data like this, which improves performance, and cuts down on wasted space in your database. We'll be looking at this next month.

The third and final problem with our table definition for staff is that we have no index. Indexing your data allows you to find data much more quickly, and you should always have at least one index in your tables – generally on the field you search by most. In my experience, people using ID fields to distinguish between rows tend to search mostly by ID key, e.g. "Find all information on member #49834", and luckily enough setting your ID field as **PRIMARY KEY** also has the added side effect that MySQL indexes it for fast searching.

## Deleting data

You can delete data easily using SQL by making use of the **DELETE FROM** command. You need to specify the name of the

table you wish to delete from and also any qualifiers. You can delete all rows in a table by simply specifying no qualifiers, e.g.: **DELETE FROM staff;** To delete with qualifiers works just as with **UPDATE**:

```
DELETE FROM staff WHERE Name = 'Harold Barnes';
```

There really isn't much more to **DELETE**, so let's move on to the main event – querying data.

## Reading data from your tables

The real power of SQL is seen in its comprehensive querying options – you can pick out very precise pieces of data using the same qualifiers seen in **UPDATE** and **DELETE**. The command to use here is **SELECT**, and it has many, many options. Let's take a look at a few, starting with the most simple. I've commented the code throughout, using the **#** symbol:

```
# a simple calculation
SELECT 1 + 1;
# * = "all columns"
SELECT * FROM staff;
# here we just select the Name column
SELECT Name FROM staff;
# our qualifier limits the results
SELECT * FROM staff WHERE Age = 30;
# LIKE allows you to do fuzzy searches on strings -
# % means "anything", like *, so this will match
# "Bart", "Marge", "Carmen", and "Harold"
SELECT * FROM staff WHERE Name LIKE '%ar%';
# How to use AND and OR
SELECT * FROM staff WHERE Name LIKE 'H%' AND Age > 30;
```

As you can see, **SELECT** can be put to a lot of uses, and it's really not hard to pull out the exact information you want.

## Conclusion

You should now be fairly competent at manipulating data inside your database using SQL. Try various combinations of modifiers to get different results, although remember to be careful with **NULL** values.

Next month the series is continued with a second article on SQL. This time we'll be looking at how PHP makes it easy for you to integrate SQL with your pages, and also, if there's space in the magazine, we'll be looking at some of MySQL's built-in functions, and how you can normalise your tables. As always, email the author with any requests for coverage of particular topics. [LXF](mailto:lxft@linuxformat.co.uk)

## 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

We'll be focusing on integrating your new-found SQL skills with PHP. 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*...



## Where's the icon?

**Q** I have recently installed Red Hat Linux 7.0. One of the problems that I'm having is that I can't find the CDROM icon to enable me install a Linux version of JDK 1.3.1. The other problem is that I have not been able to find the appropriate menus to get me connected to the web. I would be glad to receive suggestions on how to solve these problems.

**James**

**A** You should have a CDROM icon on your GNOME desktop, although you can still mount it manually with;

```
mount /mnt/cdrom
```

You can create a symlink from this to your GNOME desktop with GNOME, and change the icon and filename to suit. Once you have this icon, you can mount and unmount the device using the right-button menu.

GNOME uses *gnome-ppp* to connect to the Internet, which should be within the Networking menu. Alternatively you can run **gnome-ppp** from the command line. If it is not installed, you may want to check the Red Hat CD to ensure that all of the appropriate GNOME RPMs are installed on your system.

## Mutt concern

**Q** Thanks for including my recent email in your Answers column. The only problem is that your advice didn't work. I don't know what version of *Mutt* that you used in reference to answering my query, but when I use the command line

```
command:
mutt -f pop://<account>@<pop server>.<ISP>.<whatever>
```

and then press the return key while highlighting an email, I get the following error message:

```
sh: mutt.view: command not found
```

There is as far as I can see, no possibility for perusing the content of an email with *Mutt* version 1.3.22.1i while on the POP server, before deciding to forward, bounce, delete, reply, etc. What I am looking for is a way to do just that.

**Brian Durant**

**A** We were using 1.4i, which was released back in May, so you may want to upgrade. 1.3.22i is a development release, and there have been many updates since then, so you may be experiencing a known bug within *Mutt*. More importantly, the release of 1.3.25 back in January fixed a remote security hole in *Mutt*, so you're probably wanting to upgrade to avoid falling victim to that. Head on over to **www.mutt.org**, find your closest mirror and build it from the source tree, or download 1.4 from your distribution's FTP server and install a package.

Alternatively, you may want to check your `~/muttrc`, to see if it is calling an external program with the **pager=** option. We certainly don't have a *mutt.view* binary installed, and *Mutt* is configured to use its own pager. Upgrading to 1.4 should fix your problem, although if you do continue

to fail in having *Mutt* access your POP3 account, then you may want to check out the *LXF* forums, or post a message on the mutt-users mailing list, which is a great source of information and knowledge for *Mutt* users.

## Is it secure?

**Q** I've recently got myself an ADSL connection and would now like to do some work from home. The company network has VPN access using PPTP & MS CHAP V2 which I've connected to from a home WinXP machine, and all I had to supply the VPN network wizard was an IP address and login/password.

I'd like to know how to configure the same PPTP connection on a Mandrake 8.2 machine to the company network. I also have to \*ensure\* it is secure, well, as secure as the winXP connection. I can't tell the difference between some insecure 1&0s on a wire and some secure 1&0s on a wire. I wouldn't want to inadvertently open up the company network to the Internet with my Linux connection, or something equally as bad. I'd probably get fired. Or should I not

```
david@niamh:~ (pts/21)
1:Exit --PrevPg <Space>:NextPg v:View Attachm. d:Del r:Reply j:Next ?:Help
 1 07/29 davidapia@attb (1.2K) Re: David - Tux shopping
 2 07/30 Spencer Tarring (8.9K) Re: Dedicated Servers
-> 3 07/30 Wayne Walker (1.6K) Re: [uml-user] usermodelinux.org logo
 4 01/07 Yuwraj Kailash (2.4K)
 5 07/31 Lars Marowsky-B (8.9K) Re: [uml-user] usermodelinux.org logo
 6 07/31 Melanie J. Goet (3.0K) Re: The Story of Dave
 7 07/31 Lunchtime Fool (3.5K) Big Brands And Market Bottoms
---Mutt: pop://david@10.1.1.5/ [Msgs: 7 26K]---(date-received/date)-----(all)-----
Date: Tue, 30 Jul 2002 21:59:25 -0500
From: Wayne Walker <wwalker@ubunt.com>
Subject: Re: [uml-user] usermodelinux.org logo
To: David Coulson <david@davidcoulson.net>
Cc: UML user <user-node-linux-user@lists.sourceforge.net>

Cool, I've no artistic ability - I can't draw stick figures even with a
ruler and a quarter :)

But, I envision Tux with flippers held out "pain up" with a smaller Tux in each
flipper.

I know there are win32 and FreeBSD ports under development, so maybe a
- 3/7: Wayne Walker Re: [uml-user] usermodelinux.org log -- (26/4)
```

**Mutt can do POP3 natively, but you do need to have an up-to-date copy.**

### do this and stick to my WinXP connection?

John Smith


**A** The *PPTP-Client* project should handle all of your requirements, as it supports both MPPE and MS-CHAP-v2. The use of MPPE, which is used for link encryption, requires a kernel patch and a patch to *pppd*, but if you only have encrypted authentication via MS-CHAP, you can use the standard kernel and *pppd* distribution. Generally, with a point-to-point tunnel such as this, it will either permit encryption, or it won't, so you won't end up with an authenticated tunnel if you don't tell the server that you support encryption.

There is a great deal of documentation, as well as the required downloads, at <http://pptpclient.sourceforge.net/>. They also have a few mailing lists which may come in useful should you get completely stuck.

### GNOME or KDE?

**Q** I'm trying to find out the pros & cons to using GNOME or KDE.

Most articles written about these seem to say that it is a matter of personal preference and leave it at that. There must be more to it than picking whichever one that you think looks prettiest.



If you don't have very much RAM, a lighter window manager, such as *Sawfish*, will run better than GNOME or KDE.

In particular, is one hungrier on resources than the other? I am looking at the possibility of installing on an older PC with only 1MB of video memory (Pentium 166 MMX with 64MB RAM). Is it easy enough to disable the 3D

capabilities on either? I know there are other window managers that are better for slower systems, but I'd rather stick to one of these if possible. Is there anything else to consider?

Eddie

**A** Neither GNOME nor KDE use any 3D capabilities, so a low-end video card is not a problem. However, both GNOME and KDE are fairly graphically intensive, so you may choose to install KDE, so you have access to the applications, >>

## A QUICK REFERENCE TO: chroot

It would be nice to think that our Linux machine was impervious to an outside attacker, through the use of firewalls, up to date services and limiting access to the system.

Unfortunately, life is not that simple, and every once in a while, we are faced with a security exploit which may enable someone to gain administrative access, or root, our poor like Linux box.

Often we can't firewall all services, as some need to be publicly accessible, in the case of mail servers or other similar servers. There isn't much which can be done on the network side of things, since packets need to

get to the server in order for it to justify its own existence.

Ensuring that the server is up to date with all security patches, and that there are no obvious configuration issues which could lead to a security compromise helps, but sometimes things happen.

While we can't protect our machine 100%, we can limit the damage an attacker can create once they have access through an insecure service. A popular method is to create a fake/directory, so the service only sees that directory, and its contents, rather than the entire system. As one would expect, this can prevent an attacker being able to damage the entire machine.

Forcing a service into a specific new/directory is done with the **chroot** command, which changes the root directory for any commands run through it. Having a service run within a *chroot* is a non-trivial task, but it is well worth doing if the service is something with a chequered security history, such as *BIND*. A quick way to install something into a *chroot* is with the **--prefix** switch for *.configure* scripts, so the prefix for installation can be pointed to `/home/chroot/bind/usr`, rather than `/usr/`. As well as the main software, any libraries which are linked to a library, such as *libc.6.so*, need to be copied into the appropriate location within the *chroot* directory. We can

check with libraries link to a binary with *ldd*:

```
$ ldd /usr/local/sbin/named
libnsl.so.1 => /lib/libnsl.so.1 (0x4002e000)
libc.so.6 => /lib/i686/libc.so.6 (0x40044000)
/lib/ld-linux.so.2 => /lib/ld-linux.so.2 (0x40000000)
```

To get the service up and running, we just need to run *chroot* as root:

```
chroot /home/chroot/bind
/usr/sbin/named
```

In the case of *named*, it will attempt to change its uid to that of the 'named' user, so including `/etc/passwd` and `/etc/group` with only the required users present will enable it to do so.

## FREQUENTLY ASKED QUESTIONS: NETWORKING

### FAQ Where do I find a list of supported NICs?

If you're considering purchasing a NIC, and want to find out which are supported, you can check `linux/Documentation/networking`, which contains information on the major network card manufacturers, including 3com and D-Link. However, if you know the chipset a card uses, you can usually quickly find the appropriate kernel module within the 'Network device support' section of the kernel configuration. As always, [groups.google.com](http://groups.google.com) can be a great assistance if you're unsure as to which kernel module to use with a specific card.

### FAQ How do I give my NIC an IP?

Interface configuration is done with the `ifconfig` command, and you need to be the root user in order to change the settings of network interfaces. If our NIC is to have the IP '10.1.1.2' and a netmask of '255.255.255.0', then we do;

```
ifconfig eth0 10.1.1.2 netmask
255.255.255.0 up
```

You will then be able to `ping` 10.1.1.2 along with other systems within that IP range.

Of course, unless it's your own network, IPs shouldn't be picked out of thin air, and your network administrator will be able to supply you with the IP and netmask for the network to which you are connected.

### FAQ Help! eth0 doesn't exist. Now what?

If `eth0` does not exist, then either the kernel module has not been loaded, or it is unable to find the network card. Checking the kernel history is done with `dmesg`, and you should be able to spot what kernel modules are being loaded and any errors they have produced. If the module is not loaded, you can use `modprobe` to load it. Once the correct module is loaded, and `eth0` exists, you can setup an alias in `/etc/modules.conf` so that when anything tries to access `eth0`, it will load the module for you without causing any errors. The appropriate

line for `modules.conf` is as follows, although the Ethernet module needs to be change for your own configuration:

```
alias eth0 eeepro100
```

### FAQ Can I setup eth0 when I start my machine up?

While this depends upon your distribution, many store their network configuration in `/etc/sysconfig/networks`, although Debian uses `/etc/network/interfaces`. If a manual configuration is not desired, the distribution's GUI configuration utility can be employed in order to setup the NICs appropriately.

### FAQ How do I tell Linux to get Internet access through the LAN?

Assuming you know what the gateway IP for the network is, you can setup Linux such that all packets not for the local LAN will be sent to this machine, and from there, off into the Internet. This is known as a 'default route', and is setup using `route`. As with `ifconfig`,

you need to be root to change the routing tables:

```
route add default gw 10.1.1.1
```

### FAQ I can access things by IP, but not by hostname. Can I fix that?

Accessing machines by hostname relies on a DNS resolver, which uses both `/etc/hosts` and a remote DNS server, as specified in `/etc/resolv.conf`. `/etc/hosts` is the easiest way to add DNS entries on a single machine, as they do not affect the rest of the network, but unless you want to list the hostname for everything, a DNS server is needed. You generally have two IPs for DNS servers for redundancy, which are supplied by your ISP. A typical `/etc/resolv.conf` looks like:

```
nameserver 10.1.1.4
```

```
nameserver 10.1.1.5
```

### FAQ Is a single NIC able to have more than one IP?

A NIC can have as many IPs as you like, on as many networks as you

but run a much smaller Window Manager, such as *Sawfish* or *Blackbox* to save on memory.


It is possible to run GNOME and KDE on low-end boxes, by disabling many of the features within the desktop to save on memory, but it really depends on the applications you want, as some people prefer the KDE apps to GNOME. Both are going to run fairly slowly, so it's purely your own preference which, if either, to run. My personal opinion is that you should just run something simple on the machine, then launch *kicker* or *gnome-panel* in order to access the apps for the various suites.

## CD solutions

The advantage of buying the DVD version of *Linux Format* is that there are fewer media to keep track of and, as in the case of the Mandrake 8.2, (May 2002), the extra software. However, if one has multiple PCs, then it's unlikely, as in my case,

you'll have have a DVD reader for each machine, as they are still using the CDROM reader. My question is simply this (and you may want to post this for other users on your web site with a similar need);

How can I split up the DVD into multiple CDs so that it will install on PCs with only CD readers so that they can work similarly to the CD versions that you also issue?  
*Nick*



Making ISOs and burning them to CD is easily done with *gtoaster*.

When we include distros on the DVD they are often as ISO images, so you can simply burn these to a CD, using the method outlined in the back of each issue of *LXF*. However, in the case of Mandrake we were able to create a booting DVD. There is no easy way to extract this back to a set of CDs, but if your computers are networked you can create a network install boot floppy to install via ftp etc. The network install image is in the Mandrake/images directory. You can find more info on this in the onlune Mandrake documentation ([www.linux-mandrake.com/en/fdoc.php3](http://www.linux-mandrake.com/en/fdoc.php3)) For other software, you can simply copy what you want into a directory until you hit 640MB, then use *mkisofs* to build the ISO image, before burning this to CD. Generally, an ISO is generated using:

```
mkisofs -o disk1.iso my_dir
```

There are numerous ways to create the ISO, with various options if Windows machines will be reading the

choose. Adding extra IPv4 IPs to a NIC is done with IP Aliasing, which needs support in the kernel. We can add an IP by creating an IP alias for a specific interface. If we want to add '10.1.2.4' to eth0, we could create a 'eth0:0' interface, and add the IP:

```
ifconfig eth0:0 10.1.2.4 netmask
255.255.255.0 up
```

The part after the `:` can be anything you choose, for easy reference.

### **FAQ** I can't reach many hosts on the Internet. Is there a way to find out where the problem is?

A utility such as *traceroute* or *tracert* can be employed to establish which routers your traffic is being handled by, and makes it easy to find a failure on a network.

A more informative packet, known as *mtr*, is also available, which both provides a traceroute through all the routers as well as a list of *ping* times and packet loss information for each of the routers.

CDROM, so you may want to check our CD-Writing tutorials (p82). Once the ISO has been generated, it can be burned using *cdrecord* just like any other ISO from the DVD. An application, such as *gtoaster*, can help you do this with a GUI.

### Upgrade woes

**Q** I have recently upgraded my Mandrake to Mandrake 8.2, from 8.1. My machine is a

Medion with a 1800MHz processor with 40GB hard drive and 256MB RAM. I have four problems:

- 1) My machine does not recognise the modem, it says no modem is connected. I have a Creatix V.90 Ham Data fax modem. Does this work with Linux?
- 2) It is very slow. Do you know I could speed it up?
- 3) The CD Player does not work.
- 4) The CD and CDROM player keep on becoming unmounted.

I am quite new to Linux although I do use Unix a bit at

work. One another question are there any programs for Linux that can be used for editing photos?

David

**A** 1) The Creatix modem is a winmodem, so does not work happily with Linux. Looking on google.com, it appears that Creatix have a driver for it, but it is very unstable and is not all that reliable. You may wish to invest in a hardware modem with Linux

2) It depends what exactly you mean by 'speed'. It could be that your kernel is not compiled for the specific CPU architecture you are using, so is unable to use many CPU optimisations. Another possibility is that you are not using the correct X server for your video card, so that it is not taking advantage of the capabilities of your adaptor. Of course, it does entirely depend upon what exactly you are using to measure the 'speed' of the system, and if it occurred under 8.1. You may also want to check with *top* that there are no processes eating up CPU time.

3) We're assuming you mean an audio CD player here, rather than the actual CD drive itself. Does it produce any useful error messages which may help you debug the problem?

You may want to check that it is pointing to the correct /dev device for your CDROM drive, and that the permissions are such that a regular user can read the CD drive, which can be done with *chmod*.

4) Again, we're going to need more information. How do you know that it is being unmounted? What are you doing in between it being mounted and unmounted?

To answer your final question, the *GIMP* is a great graphics package. You can find it over at [www.gimp.org](http://www.gimp.org), or you can find it on Mandrake's FTP server for an easy install.

### Apache poser

**Q** The Linux in my system is Mandrake 8.0 and I was following your recent tutorial on CGI programming article in which you had a HOWTO configuration of Apache. The Apache server is running OK on my system as I get the web server default page. But I am unable to obtain the list of environment variables when I look for <http://localhost/cgi-bin/printenv>. I have tried to locate the cgi-bin file

```
roozbeh@bamdad.sharif.ac.ir: /home/roozbeh/temp
[roozbeh@bamdad temp]$ cat hafez.utf8
اا با ايها الساقى ادر كاسا و ناولها
كه عشق آسان نمود اول ولي افتاد مشكلها
به بوى نافه‌اى كاخز صبا زان طره بگشايد
ز تاب جعد مشكيش چه خون افتاد در دلها
مرا در منزل جانان چه امن عيش چون هر دم
جرس فرياد مي‌دارد كه برينديد محملها
به مي سجاده رنگين كن گرت پير مغان گويد
كه سالك بي‌خبر نبود ز راه و رسم منزلها
شب تاريك و بيم موج و گردابي چنين هائل
كجا دانند حال ما سبكاران ساحلها
همه كارم ز خود كامي به بدنامي كشيد آخر
نهان كي ماند آن رازي كز او سازند محفلها
حضورى گر همي‌خواهي از او غايب مشو حافظ
متي ما تلق من تهوى دع الدنيا و اهملها
[roozbeh@bamdad temp]$
```

Linux supports many international languages, but you need the fonts to make it happen.

as suggested but was unable to find it on Mandrake. I have found the *httpd.conf* but in a different location from the one you suggested in the article. However, I am not able to find *cgi-bin* and therefore not able to configure *Apache* as recommended.

Also I would like to know if there are any distros that support the Iranian/Persian language. In particular, I would like to know if a program or web browser exists that would translate English web sites into Persian. Any suggestions on where to look of information on this subject would be appreciated.

I have been a casual Linux follower since 1996 and to date I find your magazine the best informative publication. Keep up the good work.

Anvar Alizadeh

**A** The *cgi-bin* directory generally lives one folder up from the 'htdocs' directory. You could use `rpm -ql apache` to list all files *Apache* installed on your system, and grep for *cgi-bin* within there, to ensure that is has actually installed the *printenv* CGI. Have you checked the default configuration and error logs to see where *Apache* is trying to source the CGIs from? It is, of course, possible that certain distributions have modified the *Apache* install such that *printenv* is no longer present (Mandrake isn't one of these though), although it is a simple task to obtain the *printenv* script from the Internet. You can still test the *cgi-bin* directory by creating a simple script which can

be executed by the web server.

The use of the Persian language with Linux is covered in some depth over at [www.farsiweb.info](http://www.farsiweb.info), which looks at the various programs which work with the language, as well as supplying a selection of fonts for use with the system. We couldn't find any translation service for the Persian language though.

### Windows hell

**Q** So, I'm happily playing on my Linux-only box, but then I had the crazy idea of putting Windows on it.

I really have no idea what I was thinking because I knew good and well that it would overwrite my MBR. In any case now I can't boot into Linux because my floppy drive bit the dust. I was wondering if there was some way I could re-write my MBR so I could get back into Linux.

Zane, from the LXF forums.

**A** If you still have the original installation CD, you could boot from that in rescue mode, pointing it to your root file system with the `root=` option, then proceed to re-run `/sbin/lilo` to install your MBR. If you don't have the CD, you can use *loadlin* within Windows to boot a Linux kernel without rebooting, which should get you far enough to actually reinstall the MBR.

### Tune in, turn off

**Q** I've set up a firewall/gateway/router thing using an old Pentium >>

« kicking about – it works great, but I don't want to leave it on all the time and I don't want to attach a monitor and keyboard to it, so how do I turn it off? I've been telneting to it and issuing a **shutdown -h now**, and obviously I could just push the power button – but that isn't satisfactory. Has any one done this or have a simpler solution?

From the LXF forums.

**A** Doing a remote login to shut it down is probably the cleanest option, as you could use *SSH* with keys to automate the shutdown, so you can write a simple script which does:

```
ssh root@gw /sbin/shutdown -hnf 0
```

This script could be run from another machine on the network, and will cleanly turn the gateway box off. If you can mount the root file system read-only, depending on the services running on it and the methods of logging, you could turn it off with the button without corrupting the file systems.

## Secure boot

**Q** I've installed Mandrake 8.2 on a PC, to use as a test server. When installing, I chose the system to be paranoid. However, now when it boots, it boots into Linux-secure, which promptly restarts the machine. The machine starts fine with Linux, so I edited the *lilo.conf* file, so this was the default, but it still uses Linux-secure.

I chose *Lilo* as the default boot loader, so I know that the change should have affected the system, but it hasn't. Any clues?

From the LXF forums.

**A** Once you edit */etc/lilo.conf*, you must run */sbin/lilo* as root for the changes to take effect. Until you do this, it will continue to use the existing configuration, which is why it's continuing to use Linux-secure.

## Print problem

**Q** If you want to stop printing a document before it has finished, is there a way to empty the spool? I have found that when I stop printing a document, after the printer is switched on again the printer starts printing gobbledygook for page after page when it is switched back on – sometimes there is only one line or word to a page. Even closing the

```
david@niamh:~ (pts/21)
lp: \
:sd=/var/spool/lpd/lp:\
:nh#0:\
:sh:\
:lp=/dev/lp0:\
:if=/var/spool/lpd/lp/filter:
photo:\
:sd=/var/spool/lpd/photo:\
:nh#0:\
:sh:\
:lp=/dev/lp0:\
:if=/var/spool/lpd/photo/filter:
mono:\
:sd=/var/spool/lpd/mono:\
:nh#0:\
:sh:\
:lp=/dev/lp0:\
:if=/var/spool/lpd/mono/filter:
rau:\
:sd=/var/spool/lpd/rau:\
:nh#0:\
:sh:\
:lp=/dev/lp0:
```

The */etc/printcap* file controls where printer spool files are held.

editor down, and even rebooting, does not stop it from continuing printing rubbish, when switched on again, until it gets to the end of whatever is in the spool.

From the LXF forums.

**A** Printing is handled by the *lpd* daemon, which stores its spool within a directory configured by */etc/printcap*, which is usually */var/spool/lp/<printer-name>*, where *<printer-name>* is often 'lp' as default. If you shutdown *lpd* with *'/etc/rc.d/init.d/lpd stop'*, then delete the spool files, which are the ones which start with **df** and **cf**, you can restart the *lpd* service and not have it send more to the printer. Of course, what is already buffered within the printer will be printed, but if you power cycle the printer, it will flush the buffer.

You may also want to look at the **lpc** command, which controls the *lpd* daemon.

## Ping failure

**Q** I'm using *BIND9* on my Red Hat 7.2 setup at home and when I ping the name of one of my machine the lookup works fine and it resolves to the correct address that I specified in my zone file. It also correctly passes on requests when I'm connected to the web so I can use it as my primary DNS. The problem is that the local machine, although resolving fine, fails pinging most of the time.

The reply time is ridiculous and when I start the ping we get about 85% packet loss which gets worse. I left it running for a while and when I came back the reply time was steadily increasing and as a result the packet loss was also

increasing. It is really annoying, because the other machine on the network can use the DNS server without problems, just localhost cannot ping itself. More an annoyance than a problem but please help anyway.

From the LXF forums.

**A** Pings are not always a reliable way to test a network, as firewalls and kernel options can affect the throughput of the ICMP replies. Your machine will have two IPs, the 127.0.0.1 localhost IP, as well as its external eth0 IP, so you can test pinging either of those. It may be that a incorrect firewall rule is limiting the number of ping replies you are getting. Have you tried to ping the machine from another network node? You can then find out if the problem is local, or on the network. You may also want to do some testing with other IP services, such as *telnet* or *SSH*, and see if they experience any problems, as you may only be seeing issues with pings, rather than with all IP traffic.

## The path to Linux

**Q** We currently have a Windows 2000 SBS server with 25 workstations. We intend to migrate this entirely to Linux based systems. However the first hurdle is to add the Red Hat 7.2 server to the network. Because the Windows 2000 server is acting as DHCP how do we get the Linux server to request an IP address from the network. Documentation for this situation seems to be very thin on the ground, yet it must be one of the most common things to do when starting on the path to Linux. All help greatly appreciated.

Andrew Ballentine

**A** Configuration information for DHCP is within the DHCP mini-HOWTO at [www.tldp.org](http://www.tldp.org). Having Linux act as a DHCP client under Red Hat is done by editing the */etc/sysconfig/network-scripts/ifcfg-eth0* script to read;

```
DEVICE=eth0
```

```
BOOTPROTO=dhcp
```

```
ONBOOT=yes
```

This will then use *pump* to send a DHCP request out to the network.

## How to delete

**Q** Should be a quickie, I don't know why this is happening.

How do I delete a file named "#bind#". Using the *rm* command doesn't work, I get a **too few arguments error**. Trying to force it doesn't get rid of it, and I can't even open it.

From the LXF forums.

**A** There are two ways to do this, depending upon your tastes.

Type either of the following:

```
rm "#bind#"
```

```
rm \#bind#*
```

You can get the latter by typing **rm #** then by pressing the **Tab** button to complete the filename. [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 etc. Try reading the documentation!

We will try and answer all questions. If we don't answer yours specifically, you'll probably find we've answered one just like it. We can't really give personal replies to all your questions.

WRITE TO US AT:

Linux Format, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or email: [lxf.answers@futurenet.co.uk](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 overleaf for more details).



**September 2002**

Product code:  
**LXFB0031** (cd)  
**LXFD0031** (dvd)

**DVD HIGHLIGHTS:**  
Slackware 8.1, Cinelerra, Ogle & Zine, Gnome2, MultiCD, Phobia III, Grip, Zinf, TKVoice

**MAGAZINE FEATURING:**  
Linux goes to Hollywood, HTML Editors roundup, Internet security special, Ruby scripting language

**CDs HIGHLIGHTS:**  
Slackware 8.1, Fluxbox, Kallers, Torcs, SaveMyModem



**August 2002**

Product code:  
**LXFB0030** (cd)  
**LXFD0030** (dvd)

**DVD HIGHLIGHTS:**  
Intel's C++ & Fortean compilers, CPAN, Gnome2, Gentool Linux, Boson, Mozilla, Normalize, EarCandy

**MAGAZINE FEATURING:**  
Ultimate Office: Every current office solution on test, plus Quantum computing, Red Hat 7.3, PHP tips, Eden motherboard test

**CDs HIGHLIGHTS:**  
Highlights: Intel's C++ & Fortean compilers, Gnome2, Transcode, Album Cover Grabber, Office Suites



**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

**MAGAZINE FEATURING:**  
Customise your kernel, WineX latest, Inside IPv6, Astronomy applications roundup

**CDs HIGHLIGHTS:**  
Highlights: OpenOffice.org 1.0, Evolution, Omnis Studio 3.01, Clam AntiVirus, Python 2.21



**June 2002**

Product code:  
**LXFB0028** (cd)  
**LXFD0028** (dvd)

**DVD HIGHLIGHTS:**  
KDE3, Mozilla 1.0, Gentoo, Beehive, PixiePlus, Funk You, NetComics, Fluxbox

**MAGAZINE FEATURING:**  
Mozilla special issue, Video editing, Networking with Macs, KDE3 uncovered, Crossover Office

**CDs HIGHLIGHTS:**  
Highlights: KDE3, Mozilla 1.0, Linux from Scratch, Tomsrtbt, GKrellm, evoBB




**May 2002**

Product code:  
**LXFB0027** (cd)  
**LXFD0027** (dvd)

**DVD HIGHLIGHTS:**  
Mandrake 8.2, knoda, Wine, Gspy, Morphon CSS-Editor, Robotcop, WebCalendar

**MAGAZINE FEATURING:**  
HP's vision of a Linux future, Linux Format Awards, Security update, Web know-how tutorials, Making Linux accessible

**CDs HIGHLIGHTS:**  
Mandrake 8.2 (featuring Kernel 2.4.18, XFree86 4.2, KDE 2.2.2, KOffice 1.11, USB2 support)



**April 2002**

Product code:  
**LXFB0026** (cd)  
**LXFD0026** (dvd)

**DVD HIGHLIGHTS:**  
Kylix 2, Trustix Secure Linux 1.5, Sorcerer Linux

**MAGAZINE FEATURING:**  
Network special - features and tutorials, Multimedia players round-up, Intermezzo, Linux Expo NY, Samba special

**CDs HIGHLIGHTS:**  
Trustix Secure Linux 1.5, KAlarm, Spectix, Adonthell, Avidemux, Animap, MuSE

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.

# LINUX FORMAT

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

# Benefits – why subscribe?

- **Subscribe by Direct Debit – it's easy! Fill in the form below or call 0870 444 8645\* and have your bank details ready.**
- **Huge savings on a year's supply of *Linux Format*.**
- **Subscription copies are sent hot off the press – be one of the first to get yours!**
- **Guarantee your copy – never miss an issue.**
- **All you need to do is fill in the form. If you live in the UK, you don't even need a stamp!**

**RISK-FREE GUARANTEE!**  
 You can cancel your subscription at any time and receive a full refund on all unmailed issues, no questions asked

\* or call the overseas number shown at the bottom of the coupon, below

Title \_\_\_\_\_ Initials \_\_\_\_\_ Surname \_\_\_\_\_  
 Address \_\_\_\_\_  
 \_\_\_\_\_ Postcode \_\_\_\_\_ Country \_\_\_\_\_  
 Email address \_\_\_\_\_  
 Daytime telephone number \_\_\_\_\_



Please choose your method of payment **1** or **2**  
**1.**  Direct Debit (UK only – double CD Edition) – I wish to pay £27.45 every 6 months  
 or  Direct Debit (UK only – DVD Edition) – I wish to pay £29.95 every 6 months  
**Please complete the form on the left**  
 I understand that I will receive 13 issues during the next 12 months.

Instruction to your Bank or Building Society to pay Direct Debits




Originator's Identification Number **768195**  
 Please send this form to address shown

Media with passion  
 Future Publishing Ltd, Cary Court,  
 Somerton, Somerset, TA11 6BR

**2. Cheque/credit card 13 issues**

<b>UK</b>	<input type="checkbox"/> £54.90	<input type="checkbox"/> £59.90
<b>Europe (including Eire)</b>	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
<b>North America</b>	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
<b>Rest of the World</b>	<input type="checkbox"/> £91.90	<input type="checkbox"/> £96.90

**1.** Name and full postal address of your Bank or Building Society branch  
 To: The Manager \_\_\_\_\_ Bank/Building Society  
 Address \_\_\_\_\_  
 \_\_\_\_\_ Postcode \_\_\_\_\_

**2.** Name(s) of account holder(s) \_\_\_\_\_

**3.** Branch sort code \_\_\_\_\_  
 (from the top right-hand corner of your cheque)

**4.** Bank or Building Society account number \_\_\_\_\_

**5.** Instruction to your Bank or Building Society  
 Please pay Future Publishing Direct Debits from the account detailed on this instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Future Publishing and if so, details will be passed electronically to my bank or building society.  
 Signature(s) \_\_\_\_\_ Date \_\_\_\_\_

Ref No (Office use only) \_\_\_\_\_  
Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

Cheque (payable to Linux Format. Sterling cheques drawn on a UK account.)  
 Visa  MasterCard  American Express  
 Switch Issue No/Valid date \_\_\_\_\_

Card No \_\_\_\_\_ Expires \_\_\_\_\_  
 Signature \_\_\_\_\_ Date \_\_\_\_\_

**UK READERS RETURN THIS COUPON TO:**  
 Linux Format, Subscriptions, Future Publishing Ltd, Freepost BS4900, Somerton, Somerset, TA11 6BR.

**OVERSEAS READERS RETURN (POSTAGE PAYABLE) TO:**  
 Linux Format, Subscriptions, Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6TB, UK

**WEB:** [www.futurenet.com/secure/subs1.asp?ID=92](http://www.futurenet.com/secure/subs1.asp?ID=92)  
 Please tick this box if you would prefer not to receive information on other offers.

**ORDER HOTLINE ☎ 0870 444 8645**  
**OR FROM OVERSEAS ON +44 1458 271 100**

QUOTE ORDER No. **LXFP32**

**OFFER ENDS: 8.10.02**

# EVERY ISSUE IS A SOFTWARE BONANZA!

Every issue *Linux Format* comes with either 2 CDs or a DVD crammed full of the latest Linux software. Get all this delivered direct to your door every month with a subscription costing from just £27.45 for six months!

In the last six months you would have received on DVD:

The latest **FULL LINUX DISTROS** including ● Mandrake 8.2  
● YellowDog ● Trustix Secure Linux ● Sorcerer Linux ● Libranet  
● Red Hat 7.2 ● Slackware 8.1  
**PLUS** Over **750** of the best Linux apps and utilities.

Don't miss these great software packages – save yourself hours of download time with every issue of *Linux Format*!

« **Subscribe from just**  
**£27.45**

**every 6 months by Direct Debit**  
**saving £22.97 a year (CD issue)**

DVD issue from £29.95 for 6 months



# Essential disc info

Read this important information before you use your Linux Format coverdisc – CD or DVD.  
We've collated some helpful info to help you get the most from these jewels of data!

## Finding the essentials

### Missing something?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases the latest libraries and

other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

### Package formats

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

### Documentation

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

## What are all these files?

If you are new to Linux, you may find the profusion of different files and extensions confusing. As we try to give as many packages as possible for compatibility, there will often be two or three files in a directory covering different types of Linux, different architectures and usually source and binary versions – so which do you install? They can be identified by their filenames, and usually just by the file extensions.

Someap-1.0.1.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.0.1.i386.deb – The same, but a debian package.

Someap-1.0.1.tar.gz – This is usually source code.

Someap-1.0.1.tgz – Same as the above, tgz is abbreviated form of tar.gz

Someap-1.0.1.tar.bz2 – Same, but uses bzip2 compression instead of zip

Someap-1.0.1.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.0.1.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.0.1.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7.x PPC Linux.

Someap-devel-1.0.1.i386.rpm – A development version.

## Installing from tarballs

A tar ball is a two stage archive. First the files are archived into a single file with tar and then compressed with Gzip or Bzip2. To unpack, cd to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
```

```
tar xvf -bzip /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.**

## 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.

# Coverdisc



**Neil Bothwick** is your guide through the wonders of this month's jam-packed *Linux Format* CD. You'll almost believe a penguin could fly...

## On the CD



Wherever you see this logo it means there's related stuff on the CD

### Essential info

On page 107 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

The British "summer" is drawing to a close, the rain is getting colder and the nights are closing in. Never mind though, because we've

got plenty to keep you occupied on this month's coverdiscs. The whole of the first CD is given over to *Flightgear*, complete with maps covering much of Europe. DVD

users also get maps for almost all of North America. *Flightgear* is covered in depth on pages 102–103, so we will concentrate on the second CD here.

### Desktop/Epsutil


I have an Epson Stylus Photo printer. Until recently, I used to plug this into a Windows laptop whenever I needed to clean the ink nozzles or check the level. Then I found *esputil*, a Linux command line program for doing the same thing. The only trouble with it is I forgot the arguments in the weeks or months since last using it. Now it gets easier with the arrival of *epsutil*. This is a GUI for *esputil*, which is part of the *Gimp-Print* package. *Gimp-Print* is on this month's DVD and has been on previous CDs, as well as being standard in some distros, so it's most likely you will already have *esputil*.

### Desktop/Netclipboard

The clipboard is a great way to transfer data between apps, but only if those apps are on the same computer. When you need to transfer the data between computers on the same network, you either have to retype it or email it. The Network Clipboard is a solution. It lets you copy your clipboard from one computer to another, making the transfer of data almost as simple as between programs on the same box.

### Distros/UnofficialSuSEFAQ

There are no new distros on the CD this month. Sometimes it's too easy to



**epsutil brings a graphical interface to Epson printer management.**

keep installing the next latest version, or jump from one distro to another without spending time getting to know one so you can get the most from it. This month we have tried to address that issue by including the Unofficial SuSE FAQ. As the name suggests, this

is not sanctioned or provided by SuSE (the name also suggests this is a set of frequently asked questions about SuSE, but you'd probably already worked that one out for yourself). This FAQ is compiled from postings to the SuSE English Linux mailing list.

### Games/Mah-Jong

This is actually a suite of three programs for playing Mah-Jong against the computer and other players. At the core is a server, so the game can be played over a network. Then there is a client, used by each player to connect to the server. Finally, there is a computer player. Up to three copies of the computer player may be run, allowing any combination from one real player against three computer players to four people playing over a network, using the computer only for communication and game display. There are several versions of Mah-Jong. This program currently supports "Chinese Classical", although the author may add Japanese and Hong Kong variants in future releases.

### Mobile/AccessPointUtilites

Networking has seen enormous growth over the past couple of years. This is partly because of the downward spiralling of the cost of networking hardware and partly because it is now quite common for households to have



### Transfer clipboard contents between computers on a network.

more than one computer (giving your partner your old computer is a great excuse for buying yourself a new one). Following the trend of lower prices for wired networking, the hardware for the catchily named 802.11b wireless networking is now becoming affordable to more users.

Most access points simply plug into your network, so they work equally well with Linux as any other OS. However, most of them are only supplied with config and monitoring software for Windows. *Access Point Utilities* is a set of programs to work with access points using the most common chipsets. These include well known names like D-Link, Linksys, Netgear and SMC as well as several other manufacturers. With these program you can configure most settings of the access points and view status and traffic information.

### Internet/WebSuck

Downloading two or three files from a web page is as easy as clicking on the links, but what happens when the page contains dozens, or even hundreds, of links? Image galleries often consists of a large number of links, as do file listings from FTP sites. You could use something like *wget* to grab the page and all files linked from it, but this is more suited to mirroring a site or page.

### Odd files?

#### Security checks

You may have noticed some files with extensions like *md5*, *asc* or *sig* and wondered what they were for. They are concerned with the integrity of files. The *.md5* files, sometimes called **MD5SUMS**, are created to ensure that files downloaded from the Internet are not corrupted in transfer. The program *md5sum* generates a checksum for a file, something like "7a07b59ade5c137bc493fe8ec1703697". While that may look like a portion of a *sendmail* config file, it's actually the checksum for this text so far. any change to the file would give a completely different checksum, so

*WebSuck* is a Java program that will search for particular types of files, like images, and build a list of all the URLs. This list can be in one of two formats, to suit either *wget* or *GetRight*. Now you can point *WebSuck* at the page or site in question and have all the files queued for download later, instead of sucking up all your bandwidth on downloads while trying to browse the web or load the next list of images.

### Server/HTML-Mason

I reviewed *HTML-Mason* a year and a half ago. At the time I thought it was a good choice for dynamic web scripting for those who preferred Perl over languages like PHP and ASP. Eighteen months on, a new version has been released with several enhancements. *Mason* allows you to mix Perl and HTML within a page, much like you do with PHP, with the obvious advantage that it saves learning another language if you already know Perl well.

*Mason* requires a few Perl modules to work, most of which are part of the standard Perl distribution nowadays. Those that aren't are included on the CD, to save CD owners a trip to CPAN, and to save DVD owners rummaging in that pile of magazines and discs in the corner for issue 30's DVD.

### Development/Perl

Speaking of Perl, the language has been updated with version 5.8.0 released recently. Naturally, we would not ignore a significant update to such an important part of most Linux setups. You probably use Perl every day, although you may not realise it. So many useful programs are written

it can immediately tell you if the file has become corrupted. To compare the file with its original checksum, type `md5sum --check MD5SUMS` replacing **MD5SUMS** with whatever the checksum file is called. It will check the files you have against checksums in the file and let you know if any are corrupt.

The *.sig* and *.asc* files are for a different kind of integrity check. They contain a PGP or GnuPG signature to verify that the files are genuine. By downloading a copy of the author's public key file from a reputable source, such as a public keyserver, you can

### README before INSTALL


#### Info made easy

Most tarballs of Linux software include a file called **README**, containing important information about the program. This file may also contain installation instructions, or they may be in a separate file called **INSTALL**.

From this month, we have extracted any **README** and **INSTALL** files from the main archive in each directory so that you may read more about the program

and what is needed to install it without unpacking the tarball.

If there is no **README** or **INSTALL** file in the directory, none of the tarballs in that directory contained a file with either of these names. Where the program was packed into a Zip archive instead of a tarball, this was checked instead, although *.jar* archives aren't included yet.



### Everything you wanted to know about SuSE but were afraid to ask. The Unofficial SuSE FAQ contains the answer to everything SuSE.

in Perl, from small utilities to large programs like *Webmin*. It is even used for complete distribution installers, so it's well worth making sure you have the latest version installed.

### Server/HTMLHelp.comValidator

Writing HTML or XHTML that conforms to the standards won't

necessarily ensure that it will work as you expect in all browsers, but it is the best starting point. Far too many web sites are written so that they look fine on the author's browser, with no consideration for sticking to the rules that the browsers use to interpret the HTML. There are various online validation services, including one provided by the WWW Consortium, but they all have one disadvantage.

You need to either put the page online and give the URL to the validator, paste the HTML into a web form or upload the page to the validator. Then you make any correction it suggests and repeat the process. Once that page has passed, you have to go through the whole thing again for every other page.

It would be so much easier to have a command line program to make the checks locally. This program is just that – the same as used by the validator at [www.htmlhelp.com/tools/validator](http://www.htmlhelp.com/tools/validator). There are two packages to install, *lq-sp* is a validation lib while *wdg-sgml-lib* contains the various DTDs (Document Type Definitions). These describe the rules for each version of HTML and

# LinuxFormatCoverdiscCD

XHTML. Then you copy the validate program itself to somewhere in your path and set the execute permissions.

You can check a single page with `validate /var/www/htdocs/index.html` or check your whole site in one go:

```
find /var/www/htdocs/ -name "*.html"
-exec validate {} \;
```

## Server/QuantaPlus

While hard core HTML geeks insist on coding everything by hand, for many

people an HTML editor is a great time saver. HTML editors fall into two main categories. The first is more like a DTP program, presenting a finished doc and shielding the user from the actual HTML used to generate it. While these may seem easier to work with at first, being able to work with an understand the HTML really makes an editor more effective and easier to get good results with in the long run. The second type of HTML editor lets the user work with the HTML, speeding up the process by

providing shortcuts and templates, as well as preview facilities. *Quanta Plus* belongs to the latter category: It is the GPLed version of *Quanta*. This is a release candidate for *Quanta Plus 3.0*.

## System/ACPMoDemDriver

The general advice when choosing a modem for use with Linux is "avoid winmodems at all costs". In fact, this is also good advice when choosing a

modem for Windows. However, there isn't always a choice. Many notebook computers come with integrated modems, which are often winmodems. Some manufacturers are more Linux-friendly than others and some even provide Linux drivers for their winmodems. IBM is one such company, they released drivers for the ACP modems in their Thinkpad notebooks quite a while ago. They work well, I have a Thinkpad that has stayed online for hours on end using

## »» CD CONTENTS AT A GLANCE

### Disc B

#### Magazine

Emulators  
HotPicks  
Intel  
Java  
Kylix  
PHP

All the files mentioned in this month's Emulation article.  
All the programs covered in this month's HotPicks section.  
Evaluation versions of Intel Software Development Products  
The files to go with this month's Java tutorial  
Example files from the Kylix tutorials  
PHP source code and example scripts

#### Desktop

Calcoo  
CDox  
CRecord  
di  
DialogCDWriter  
DirectVNC  
Dnotify  
Epsutil  
FeaturificFileManager  
FoldingAtHomeGrapher  
GKrellIM  
GKrellMBUPS  
GNOMEClipboardManager  
Grocgret  
KEGS  
Kile  
Klicklack  
Kover  
Netclipboard  
PackMan  
PipeMeter  
PipeViewer  
ProcViewer  
QtMyAdmin

Scientific calculator designed to provide maximum usability  
Generates, edits, and prints CD documents like CD covers  
Creates home-burned CDs with a CD-R/CD-RW recorder  
Disk information utility  
Dialog front end to burn CDs under GNU/Linux and Unix  
VNC client using hardware accelerated DirectFB library  
Execute a command when the contents of a directory change  
Graphical Epson inkjet printer toolbox  
Filter, sort, view, and use files easily and quickly  
Graph and analyze Folding@Home team statistics  
Monitor program that charts SMP CPUs, disk, load, network  
Plugin for GKrellIM to display the status of a Belkin UPS  
Manage selections and clipboards  
Shop for groceries faster and more reliably  
Apple II emulator, originally written for HP workstations  
LaTeX source editor, TeX shell and gnuplot front end  
Plugin-based dictionary application for KDE  
WYSIWYG CD cover printer  
Clipboard synchronization over a local network  
Packman is an easy-to-use graphical front end to apt.  
Displays the speed and progress, of piped data  
Monitors the progress of data through a pipeline  
GUI interface for the Linux /proc file system  
Handle the administration of MySQL, like PHPMyAdmin does

#### SignLanguageLearningAid

Set of programs help the beginner student learn ASL  
Handyman's text editor, laden with tools  
Small VNC server that can be run on 8-bit microcontrollers  
Implementation of the Windows 3.x and Win32 APIs  
Network interface monitor dockapp for Window Maker  
Powerful, extensible text editor with full GUI support  
Assign an action to any key that is otherwise unused in X

#### Development

Alzabo  
Avifile  
BFCD3UnixForthCompiler  
Bugzilla  
Data modelling and reverse engineering tool  
Library to read and write compressed AVI files  
Forth compiler for Linux, FreeBSD and OpenBSD  
Leading open-source/free software bug tracking system

Fastdep  
GTK+  
JDebugTool  
Jmark  
Jonpy  
liba52  
libcfg+  
libfs++  
Libss  
LibVNCServer  
Mail-Freshmeat  
Meld  
Netclasses  
Perl  
PythonCard  
Qt  
Readline  
SSL++  
Xiquall

Preprocessor to generate makefile dependency information  
Library for creating graphical user interfaces  
Standalone Java debugger  
Encode and decode digital watermarks in Java class files  
Python modules with multi-threaded, OO facilities  
Library for decoding ATSC A/52 streams  
Command line and configuration file parsing  
Small C++ library that provides a file\_iterator class  
Cross-platform library for controlling screensavers  
Library for creating VNC servers  
Perl module parse the daily Freshmeat email newsletters  
GNOME 2 diff and merge tool  
GNUstep interface to socket programming in Objective-C  
GUI construction kit for cross-platform applications  
GUI toolkit for software developers  
Library that provides functions for command line editing  
OpenSSL functions in an easy-to-use C++ class hierarchy  
Frequently used functions for games and related utilities

#### Distros

UnofficialSuSEFAQ  
Compiled from the SuSE English Linux mailing list

#### Games

Mah-Jong  
Monopd  
NewAdventureShell  
QStat  
Raptor  
RotateMania  
Stoned  
TetriStation  
Networked Mah-Jong game  
Game server daemon for Monopoly-like board games  
Turns your shell into a text adventure game  
Displays the status of Internet game servers  
2D vertically scrolling shoot-em-up game  
An addictive mix of elements of arcade and puzzle games  
Simple but fully functional curling simulation  
A score of tetris-like games

#### Graphics

AVIPlayerForXMMS  
Digikam  
DivXCalc  
Drip  
Earthview  
FilmGimp  
Grace  
Gv4l  
KMencoder  
Kooka  
myDVDs  
Ogmttools  
oKle  
PHPaint  
ScreenSniper  
Sketch  
Tubesock  
Vobcopy  
XMMS plugin to play Windows AVI, DivX and ASF files  
Frontend to gphoto2 to communicate with a digital camera  
Simple DivX calculator  
DVD to divx conversion tool  
Produces a view of an arbitrary region of the Earth  
Film Gimp is the 16-bit per channel variant of Gimp  
WYSIWYG 2D plotting tool for X Windows  
Transcode front end to watch or record from a v4l device  
KMencoder is a front end for Mplayer/Mencoder.  
Raster image scan program that supports OCR  
Web-based graphical DVD inventory database  
Tools for working with Ogg Vorbis streams  
oKle is a KDE front end to the Ogle DVD player.  
Quick, easy, mass-production of Web page "buttons"  
Graphical front end to the import screenshot program  
Vector drawing program similar to CorelDraw or Illustrator  
GTK/GNOME Shockwave file player  
Copies DVD .vob files to disk, decrypting them on the way

#### Internet

Acidblood  
Aku  
IRC robot featuring channel and user management  
Download queue that handles HTTP and FTP downloads

the ACP driver. Now they have released an updated driver, so here it is for all you IBM notebook users, as well as users of other brands of notebook with internal modems based on the same chips.

## Sound/Aglaophone

*Aglaophone* is a program for recording, processing and playback of real-time audio. The package is entirely modular, so any new effect can be added by

creating a new module. This means the program is not reliant on only the author adding new bits and pieces, anyone can expand its capabilities. That said, *Aglaophone* already has a useful range of filters and functions. If you need to convert, process mix or combine audio, this is certainly worth a look. As for the name, according to the website *Aglaophone* was one of the Sirens of Greek mythology. It means "Bright voice". It doesn't say how it should be pronounced. [LXF](#)



**Quanta Plus is the GPLed version of the Quanta HTML editor.**

<b>Amethyst</b>	A bot for IRC and other protocols
<b>Balsa</b>	Gnome e-mail client
<b>Bocazas</b>	Connect to IRC and use IRC commands from a web page
<b>FFT</b>	Alternative traceroute program
<b>GnomeICU</b>	Communicate with ICQ users and other GnomeICU users
<b>Gnuzza</b>	Peer to peer encrypted chat client for Linux and Windows
<b>Grepmail</b>	Searches normal or compressed mailboxes
<b>HepMessagingServer</b>	Transfers messages between different messaging systems
<b>HTun</b>	Create a fully-bidirectional IP VPN over an HTTP proxy
<b>KMyIRC</b>	IRC client for KDE
<b>Licq</b>	Licq includes all the basic features of ICQ
<b>MozillaStarterScript</b>	Startup wrapper for Mozilla
<b>TkRat</b>	Graphical Mail User Agent (MUA) that handles MIME
<b>Tuxaator</b>	IRC bot that can keep a database of keywords or definitions
<b>Vamp</b>	An automatic URL downloader
<b>WebJob</b>	Downloads a program over HTTP/HTTPS and executes it
<b>WebSuck</b>	WebSuck goes through a webpage making a list of data files

## Mobile

<b>AccessPointUtilites</b>	Configure and monitor Wireless Access Points
<b>CountDown</b>	Count down from now until some event
<b>Hanoi</b>	Classic Tower of Hanoi puzzle
<b>Kismet</b>	802.11b network sniffer and network dissector
<b>MyEricsson</b>	Administration tool for the Ericsson T39m mobile phone
<b>P2Z</b>	Converts Palm Memo/ToDo/Address/Datebook to Zaurus
<b>Perl</b>	ARM port of perl 5.6.x
<b>Troll-ftpd</b>	A port of Trolltech's ftpd to the Zaurus
<b>Winzig</b>	Small Python/GTK applications suitable for PDAs
<b>WirelessToolsHowto</b>	Generic tools to set up and monitor Wireless LAN devices

## Office

<b>Chronos</b>	Web agenda calendar for intranets
<b>cWriter</b>	Web-based collaborative writer
<b>Dataxi</b>	Database system with single form to query and modify data
<b>Jeremi</b>	Groupware software for real time document sharing
<b>KOffice</b>	Release candidate of the new version of this office suite
<b>LABE</b>	Web application to administer a centralized LDAP directory
<b>OpenOfficeIndexer</b>	Generates an index for OpenOffice.org documents
<b>ToDoManager</b>	A simple task manager
<b>TOEJAM</b>	Telephone answering machine using a voice modem
<b>TUTOS</b>	The Ultimate Team Organization Software
<b>TuxMonkeyIssueTracker</b>	TuxMonkey Issue Tracker is a helpdesk application

## Server

<b>AlberT-EasySite</b>	PHP-based site generation system
<b>Ampoliros</b>	Advanced and easy to use PHP Web applications platform
<b>Anna</b>	Set of AIML files that constitute a chat bot
<b>Banner</b>	Sends commands to daemons and displays the replies
<b>BlackHoleSpam_VirusFilter</b>	Spam and virus filtering for email
<b>Boa</b>	Single-tasking HTTP server
<b>Camorama</b>	GNOME 2 Webcam application featuring various image filters
<b>Dyn2Stat</b>	Web application that keeps your IP in a central database
<b>HardwareCompatibilityDB</b>	Web-based database for hardware compatibility information
<b>HTML-Mason</b>	Full-featured web site development and delivery system
<b>HTMLHelp.comValidator</b>	Offline version of the WDG HTML/XHTML Validator

<b>Peacock</b>	HTML Editor for GTK+/GNOME
<b>PHP-PaiShoppingSystem</b>	Shopping cart system for users of the PayPal Online Store
<b>PHPgraph</b>	Simple tool for creating PNG images containing line graphs
<b>phpMyAdmin</b>	Administer MySQL over the WWW
<b>POP-before-SMTP</b>	POP-before-SMTP daemon for use with Postfix
<b>ProMA</b>	Web interface for managing users in a ProFTPd/MySQL setup
<b>QmailAdmin</b>	Web control panel for qmail/vpopmail-based mail accounts
<b>QtMyAdmin</b>	Handles the administration of MySQL
<b>QuantaPlus</b>	Web editor for KDE supporting HTML and more
<b>ScrollKeeper</b>	Cataloging system for documentation on open systems
<b>Vertigo</b>	Makes virtual hosts management easy
<b>Vsftpd</b>	FTP server written to be free of security holes
<b>WebAPP</b>	Web Automated Perl Portal
<b>Zorbstats</b>	Web statistics generator using PHP and MySQL

## Sound

<b>Aglaophone</b>	Recording, processing, and playback of real-time audio
<b>BeatForce</b>	DJing system, with two players
<b>BurnCDDA</b>	Console frontend to cdrdao, cdrecord, mpg123, oggdec etc.
<b>Cjukebox</b>	Management system for audio files and play lists
<b>Fultron2002</b>	Sample-based drum machine
<b>Gogg</b>	Ogg player featuring a nice tree-structured playlist.
<b>JukeX</b>	Multi-user web Jukebox application implemented in Java
<b>MakeAudio</b>	Graphical tool to create audio CDs from MP3 or wav files
<b>MP3AlbumOrganizer</b>	Organise your large collection of MP3s in a database
<b>mp3isofs</b>	Patch to mkisofs aimed at creating CDs of MP3 files
<b>NRip</b>	Frontend to CDDA rippers and MP3 encoders
<b>RhythmBox</b>	Music management and playback software
<b>Songwrite</b>	Music tablature editor
<b>Vorbix</b>	Graphical front-end for the Ogg Vorbis tools
<b>XMMSdWebControlPlug-in</b>	Allows XMMS to be controlled with a Web browser
<b>Zina</b>	Graphical interface to your MP3 collection

## System

<b>ACPModemDriver</b>	Drivers for the ACP modem (Mwave), a WinModem
<b>ArrayInfo</b>	Retrieves RAID information and logical drives status
<b>BackerUpper</b>	Simple, easy to configure backup utility
<b>ClamAntivirus</b>	Powerful anti-virus scanner for Unix
<b>E2tools</b>	Utilities for manipulating files in an ext2/ext3 filesystem
<b>HPOfficeJetDriver</b>	Rrinting, scanning and photo-card support for HP OfficeJets
<b>LifebookTouchscreenDriver</b>	Touchscreen driver for Fujitsu-Siemens Lifebook B-Series
<b>Im_sensors</b>	Essential drivers for motherboard hardware monitors
<b>MS-sys</b>	Linux program for writing Microsoft compatible boot records
<b>NetMon</b>	Lightweight, flexible Web-based system to manage networks
<b>Netstat-nat</b>	Displays NAT connections created by iptables conntrack
<b>Nmap</b>	Utility for network exploration or security auditing
<b>PCMCIA-CS</b>	Complete PCMCIA or PC Card support package
<b>PEANUTS</b>	Console-based Internet user management system
<b>RPMgraph</b>	Generates a graph of installed RPMs and their dependencies
<b>Schedtool</b>	Alter or query a process' scheduling policy
<b>Serel-FastBootSoftware</b>	Boot faster by allowing services to be started in parallel
<b>Sonar</b>	Network reconnaissance utility
<b>UserFriendlyIptables</b>	Generate optimized iptables packet filter rules
<b>ValuHack</b>	Aims to teach that hackers are not criminals or vandals

# Coverdisc



**Neil Bothwick** is your guide through the wonders of this month's jam-packed *Linux Format DVD*, which this month doubles as a demo and rescue disc.

**W**e have something unusual for DVD users this month. Normally the DVD contains a complete distribution or two, either to be installed direct from DVD or after burning to CDRs. In either case, you still need to reformat your hard drive to try out Linux. If you are an existing Linux user who wants to try out a different distribution, this probably won't be much of a problem.

What if you currently use Windows and want to see what this Linux lark is all about? Your hard drive is probably set up in a single large partition, so it's quite a job to rearrange this just to take Linux for a spin round the block. Now we have the ideal solution for

you, a complete Linux distribution, based on Debian, ready to run direct from the DVD.

## Install-free Linux


Knoppix is a complete Linux distro using the latest kernel, KDE 3.0.2, full Internet and networking software, *GIMP* for image manipulation, the *X Multimedia System (XMMS)*, *OpenOffice.org* and much more.

In all, there are over nine hundred packages in Knoppix, a total of 1.7GB, compressed to be able to fit on a single 700MB CD. See the box if you want to recreate this CD from the DVD.

Provided your PC is able to boot from CD (this includes booting from DVDs), you only need to reboot with the Linux Format DVD in your drive to start Knoppix. At the splash screen you can enter various options to change the way it starts up, such as using a different screen resolution. Press **F2** to see the options, **Enter** to continue or just wait a few seconds and Knoppix will carry on booting. During startup, it tries to identify your hardware, which worked well on the three different systems I tried it on. Then it loads the relevant drivers and boots into the KDE desktop.

Now you can start experimenting. you don't need to worry about breaking anything, as the entire distribution is on the DVD so you can't overwrite anything important. You'll notice some disk icons down the left hand side of the screen, Knoppix has identified your hard drive partitions, so you can access your existing data. Note that the default behaviour of KDE is to open an icon with a single click, which can be a bit of a surprise when coming from another OS. This can be changed in the KDE control centre, started from the icon in the toolbar at the bottom of the screen.

Most programs are started from the *K-menu*, using the icon at the



**A complete Linux distribution, with a full range of software, bootable from this month's DVD.**

bottom left of the screen. This is also where you logout/shutdown. After logout, Knoppix ejects the DVD and asks you to press **Enter**. This will shut down your computer, if you want to reboot, it's safe to press the reset button at this point.

If you are unable to boot from the DVD, the */KNOPPIX* directory contains boot images to create a bootable floppy, together with full instructions. There is also a FAQ (Frequently Asked Questions) file in here.

## Rescue disc

If you are an experienced Linux user, happy with your current distribution, you may have skipped the section on Knoppix as being irrelevant to you. Think again. As a complete distro on CD/DVD, it makes an ideal rescue tool. In addition to the productivity and games software, there are also the tools needed to recover a broken system, both GUI and command line. And if your Windows using friends keep asking you what is so good about GNU/Linux, put a copy of the Knoppix CD in their drive and reboot.

## Games/Flightgear

DVD users get to spread their wings and fly further afield, with the extra *Flightgear* maps on the DVD. In addition to the European maps on the CD, we have maps covering almost all of North America. Much as we would have liked to bring you the rest of the World too, the full map set occupies more than 4GB, the entire capacity of the DVD. That in itself says a lot about the scope and ambition of this project. See the separate coverage of *Flightgear* for more information on using these maps.

## Sound/OggVorbis

MP3 (more correctly, MPEG-1 Audio Layer-3) is the most popular format for compressed audio, particularly for transporting music over the Internet. However, it is not without its problems, particularly because of the licensing of the compression methods used. In the same way that similar problems with GIF gave rise to the superior PNG image format, *Ogg Vorbis* was

**On the DVD**

Wherever you see this logo it means there's related stuff on the DVD

## Important notice

**Before you even put the DVD in your drive, please make sure you read, understand and agree to the following: The Linux Format DVD is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker on ANY new software. While every care is taken in the selection, testing and installation of DVD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system which may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all important files. Please read individual licences for usage terms.**

## From DVD to CD

Best of both worlds

We can't win. When we include ISO images you say you'd rather have an installer boot from the DVD. So we set up bootable DVDs and you ask us for ISO images to install on non-DVD computers.

In an effort to please everyone, here are instructions to create a Knoppix CD from the data on the DVD.

Create a directory called Knoppix somewhere on your hard drive, say ~/Knoppix

Copy the following directories and files from the DVD to this directory

Demos  
KNOPPIX  
Talks  
autorun.bat  
autorun.inf  
knoppix.html

rename knoppix.html to index.html

Now use your favourite CD writing software to create a bootable CD from this directory. The boot image you'll need is KNOPPIX/boot.img. Command

line fans can create and burn an ISO image with


```
mksifos -R -b KNOPPIX/boot.img -c
KNOPPIX/boot.cat -o knoppix.iso
~/Knoppix
```

```
cdrecord -v dev=0,0,0 knoppix.iso
```


With a fast enough machine, you can skip the image file with

```
mksifos -R -quiet -b KNOPPIX/boot.img
-c KNOPPIX/boot.cat ~/Knoppix |
cdrecord -v dev=0,0,0 -
```

Note the trailing '-'.  
If your style is more point-and-click than hunt-and-peck, you can create the CD with one of the many CD burning GUIs. The screenshot shows how the booting is set up in *X-CD-Roast*. Note that you must give the path to the boot image relative to the CD's source. The Browse button in *X-CD-Roast* gives an absolute path, you need to manually remove the leading part. *X-CD-Roast* is on Knoppix, you can create a new Knoppix CD while running Knoppix.



Set *X-CD-Roast* to create a CD from wherever you copied the DVD files.



Making a bootable CD, note the use of a relative pathname for the boot image file.



*XMMS* and *Zinf*, two of many audio players that work with Ogg Vorbis files.

conceived as a solution. A new audio compression system without MP3's licensing issues, and with slightly better compression and quality into the bargain. One of the reasons it hasn't taken off yet is that until it reached a stable version 1.0, hardware manufacturers were reluctant to commit to a format that may change. That's no longer the situation as *Ogg Vorbis 1.0* is out now, and on the DVD.


Most players, such as *XMMS* and *Zinf*, now support *Ogg Vorbis*, and the DVD contains all the tools you need to create your own *Ogg* files, as well as libs to add support to any app you write.

## Internet/WWWOFFLE

Mention proxy servers and most people think of *Squid*. It is good in the traditional role of a proxy server, but that's not what *WWWOFFLE* is for. This proxy is specifically designed for use on dial up connections, where the PC is not permanently connected. If you try to access a website while offline, *WWWOFFLE* will cache the request and fetch the page the next time you connect. The web interface has options to request pages in advance, or even whole sites. These can be one-off fetches or regular updates.

## Graphics/Drip

We had the *Xine* and *Ogle* DVD players on last month's DVD and make no apologies for including updated versions this month. Most people with a DVD drive in their computer like to be able to watch video DVDs with it. With bigger and cheaper hard drives it makes sense to store films on the hard drive instead of playing direct from DVD, in the same way that many people keep their favourite CD tracks as MP3 or Ogg files. *Drip* is the DVD equivalent of *Grip*. It reads movies from DVD and saves them as compressed DivX files. There are GUI and a command line versions. Both produce the same results, it's simply a question of the environment you prefer to use. [LXF](#)



Convert your DVD movies to DivX files, ready for playing direct from hard drive, with *Drip*.

## LinuxFormatCoverdiscDVD

## » DVD CONTENTS AT A GLANCE

## Desktop

<b>BetabriteMessenger</b>	Commandline utility to program a BetaBrite LED sign
<b>CreamForVim</b>	Easier to use Configuration of the famous Vim text editor
<b>Gnumeric</b>	Powerful and easy to use spreadsheet program
<b>IcuKrell</b>	gKrellm plugin to display the status of GnomeICU
<b>KFish</b>	KDE 3.x panel applet animated with a fish, bubbles etc.
<b>Percy</b>	Talking penguin that also listens to you and responds
<b>PilotLink</b>	Move information between the desktop and a Palm device
<b>SambaLink_Q</b>	Qt smb.conf file editor
<b>SimpleCDR-X</b>	GTK+ based frontend for CD writing and mastering
<b>Tksql</b>	Edit the tables of PostgreSQL databases
<b>WMALMS</b>	Monitors sensor chip: temperature, fan speed and voltage

## Development

<b>Alma</b>	Software workshop for modeling and analyzing
<b>Cgixx</b>	Modern Common Gateway Interface library
<b>Chart2D</b>	Java library for adding 2D charts to Java programs
<b>CVSTrac</b>	Low-ceremony Web-based bug and patch-set tracking system
<b>DynamicProbes</b>	Generic and pervasive system debugging facility
<b>HTTPAgent</b>	PHP library for server-to-server GET and POST requests
<b>libferris</b>	Exposes hierarchical data through a common C++ interface
<b>PHPScrew</b>	PHP script encryption tool
<b>PHPXMLClasses</b>	Classes for XML processing using PHP
<b>Pypect</b>	Make Python a better glue
<b>PySMB</b>	Experimental SMB/CIFS client library
<b>Scriptix</b>	Lightweight scripting language for embedding applications
<b>Sfront</b>	Compiles MPEG4 Structured Audio bitstreams into C programs
<b>SPIKE</b>	API that helps reverse engineer unknown network protocols
<b>Squeak</b>	New OpenSource and portable Smalltalk-80-based language
<b>TinyQ</b>	Core QT functions for embedded applications
<b>WideStudio</b>	Multi-platform integrated development environment

## Distros

<b>ClarkConnect</b>	Transform an old beat up PC into a secure Internet gateway
<b>FAI</b>	Non-interactive installation of Debian on a group of PCs
<b>Knoppix</b>	Complete Linux distribution, bootable from the DVD
<b>PXESLinuxThinClient</b>	Micro Linux distribution for thin clients
<b>RxLinux</b>	Centralised configuration & management of multiple servers

## Games

<b>Flightgear</b>	The Flightgear flight simulator with extra maps
<b>GeeWhiz</b>	Role playing game in the style of Wizardry(tm)
<b>OrbitalEunuchsSniper</b>	Control an orbital laser to protect the VIPs
<b>Project1</b>	Cross-platform space simulation game
<b>RolePlayingGameDaemon</b>	Medieval fantasy role-playing game

## Graphics

<b>CRWInfo</b>	Extracts thumbnails and information from Canon .crw files
<b>EXIFtags</b>	Extracts EXIF information from JPEG files
<b>Gmandel</b>	Explore the intricate details of the Mandelbrot set
<b>jFunction</b>	Function visualizer written in Java
<b>Kallery</b>	Highly configurable image gallery generator
<b>Mplayer</b>	Command line movie player
<b>OgleDVDPlayer</b>	DVD player that supports DVD menus and navigation
<b>Transcode</b>	Text console video-stream processing tool
<b>Xine</b>	Video player for MPEG1/2 video, DVDs, VCDs, and AVI files

## Internet

<b>AutoResponder</b>	A simple email receptionist
<b>DjVuLibre</b>	Web-centric platform for distributing documents and images
<b>Getleft</b>	Given a URL, Getleft will try to download all links
<b>Lftp</b>	Sophisticated command line based FTP client
<b>Liamail</b>	Graphical mail client like Microsoft Outlook Express
<b>Mutella</b>	Gnutella client optimized for a high-bandwidth connection
<b>WWWOFFLE</b>	Proxy server for computers with dialup cconnections

## Mobile

<b>GPSdrive</b>	Map-based navigation system
-----------------	-----------------------------



## IconEdit

Make new icons on the Zaurus.

## Pjptools

Libraries and utilities to communicate with a Psion palmtop

## Zbedic

Dictionary for Zaurus SL-5x00 handhelds

## Office

## CK-Ledger

Double-entry ledger accounting system for PHPGroupWare

## DataVision

Database reporting tool similar to Crystal Reports

## DocumentManager

Document management system

## SMTM

Perl/Tk ticker, profit/loss calculator, and chart tool

## tkFaxSpool

Frontend for mgetty's faxspool/faxq/faxrm send FAX programs

## Server

**BessieTheAnnihilator** Allows teachers to post assignments and grades on the Web  
**BigDumbFingerDaemon**

Replacement fingerd with new features and configurability

## DaDaBIK

Creates a customizable Web interface for a MySQL database

## FortuneWorld

Create and access fortunes from a MySQL database

## jo

Pure Java webserver

## LinkChecker

Check HTML documents for broken links

## Lurker

Mailing list archiver

## MyDNS

DNS server implemented from scratch using MySQL

## Nanoweb

Modular HTTP server written in PHP 4

## Ringlink

Tools to run one or more rings of Web sites

## SaintSQLAdmin

Web-based interface for MySQL administration using PHP

## Zabbix

Monitors your servers and applications

## Sound

## gAlan

Audio-processing tool

## GneleMixer

Yet another OSS mixer with a good looking GUI

## jMusic

Library of classes for generating and manipulating music

## MuSEStreamer

Mixing, encoding and network streaming of sound

## Ogg Vorbis

High-quality lossy audio codec, free of IP restrictions

## ReZound

Graphical audio file editor

## StompBoxes

Realtime audio effects processor, designed for guitar

## System

## Atop

ASCII full-screen performance monitor

## Bash

sh-compatible command language interpreter

## Diskfree

Track changes in your filesystem and report on size changes

## GimpPrint

Collection of very high quality printer drivers

## IPTablesTutorial

How to install and set up iptables and netfilters

## Sourcer

Package installer that uses source tarballs

## Syslog-ng

Syslogd replacement, but with new functionality

## TimosRescueCDSet

Generate bootable CDROMs with a ready-to-use rescue system

## UtilLinux

Essential utilities for any Linux system

## Xinetd

Replacement for inetd, the internet services daemon

## Essentials

## gcc-3

The best compiler in the world?

## LDP

A complete mirror of the Linux Documentation Project

# 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)  
Contact Tony Dyer

## 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.wylug.lug.org.uk](http://www.wylug.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)  
Contact Ken Guest

## 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 [www.midlandsLUG.cjb.net](http://www.midlandsLUG.cjb.net) *WARNING: Popup ads*  
Contact Pete Thompson

## 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 East London**

URL [www.eastlondon.lug.org.uk](http://www.eastlondon.lug.org.uk)  
 Contact Jonathan Spriggs

**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

**All the following are new details:**

**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@dpn.co.im](mailto:john@dpn.co.im)

# LinuxUserGroups

## LUG OF THE MONTH!

### Huddersfield

The Huddersfield Linux User Group was formed in late 2000 with the aim of promoting Linux and open source. Since then it has steadily grown to around 50 members. Abilities range from total beginners to Linux gurus with members coming from wide geographical areas and backgrounds. We have a website, an active mailing list, an IRC channel called #hudlug on Undernet and monthly meetings. Mailing list topics are quite diverse

but are mainly related to computers and Linux.

The meetings are held on the evening of the first Tuesday of every month at The Head of Steam in Huddersfield. Members get the chance to compare experiences and discuss any problems they might have with Linux. The turnout at these meetings varies quite a lot. IRC provides a friendly place to meet between these 'real' meetings.

The website is maintained by

the founding member, Adam Brookes, and is becoming an online resource for the group as tutorials, reviews and other materials are added to the site by various contributors.

In the future we would like to host events such as install fests or Linux demo days.

We look forward to your valued contribution to the LUG.

<http://www.hud.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.

**Hello folks, and welcome to a** new series of articles based around LUGs, advocacy and spreading the good word of Linux. The aim of these articles is not to just inform you about this topic or that topic, but to give you the incentive and knowledge to get out there and tell people about Linux and how it can help them.

The other aspect to this series is that we want to get feedback from you folks. If you have sold someone on Linux, or managed to get Linux in your local school or college, please let us know. Feel free to use the Discussion forum on the *LXF* website at [www.linuxformat.co.uk](http://www.linuxformat.co.uk). The aim is to use this series in conjunction with your feedback to get Linux in more places and used by more people.

So, where do we get started spreading the word about Linux?

The answers to many of your questions lie in the humble Linux User Group (LUG). A LUG is a great place to meet like minded Linux folk, and to share knowledge and

discussion about Linux and where it is going. There are LUGs organised across the country, and every LUG will cater for users of all walks of life and knowledge; you do not need to be a boffin to attend. You can use the lovingly crafted map of LUGs on these pages to find your nearest LUG, and I suggest that you pop along to see what you think.

LUGs are an important part of this series of articles as they represent a nationwide network of Linux enthusiasts. This network is a great resource to the Linux community, and this series will discuss concepts and ideas to stimulate this community. If you don't go to a LUG however, do not feel left out as we will be providing useful tips and advice for individuals to get involved also.

Next month we will take our first steps into the word of Linux advocacy and discuss some of the concepts and resources available. Until then, feel free to post some suggestions and ideas on the *LXF* forums.

## 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)

**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**

Ross Anderson, Jono Bacon, Neil Bothwick, David Cartwright, Andy Channele, Jeremy Cole, David Coulson, Hoyt Duff, Marco Fioretti, Simon Goodwin, Paul Hudson, Jon Kent, Brian Long, Mark Newby, Robert Smith, Charlie Stross, Sam Williams

**COVER CD PRODUCTION****CD Editor** Neil Bothwick**Senior CD Editor** Kate Hadley**ART CONTRIBUTORS****Photography** Louise Parker, Philip Sowels, Getty Images, Powerstock**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 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 96**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 CDROM** demos and reader submissions are supplied to us on the assumption they can be incorporated into a future covermounted CDROM, 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

Media with passion

# NEXT MONTH

## Issue 33 on sale Tuesday 8 October



# ...FROM OUT OF THE SUN

Sun have announced the LX50, backed by their own Linux distro. What does this mean for the average user, what is this Sun ONE platform they harp on about and perhaps most importantly of all, is the LX50 as cool as it looks? Find out the answers, only in Linux Format, next issue...

**PLUS:**

There'll be bagfuls of tutorials, and a hatful of reviews including Kylix3, Lycoris, Homebase and ELX Linux. Find out which mailserver you should be using in our latest roundup.

## DON'T FORGET

We'll also be bringing you another dose of Linux Pro, packed with real world Linux case studies, security updates and sysadmin know-how.



## DON'T MISS YOUR COPY OF LINUX FORMAT!

Get it delivered to your door every month – subscribe on page 96

The exact contents of future issues are subject to change

# LINUXPRO

From the makers of LINUX Format

October 2002

## UNITED LINUX

A great leap forward or a marriage doomed to failure?



**PLUS**

**Bind 9** – What you need to know about the new nameserver  
**Hill House Hammond** – managing the move to Linux

## Contacts

### 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)

#### Contributors

Andy Channelle, Chris Denton,  
 Jon Kaye

#### Photography

Getty Images

#### Publisher

Dave Taylor

### ADVERTISING

#### Advertising manager

Matt Stanley  
[matt.stanley@futurenet.co.uk](mailto:matt.stanley@futurenet.co.uk)

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street,  
 Bath BA1 2BW  
 Tel +44(0) 1225 442244

# Welcome

Twenty-four pages of real-world Linux for IT professionals

One of the biggest stories of recent month's has been the announcement of the United Linux conglomeration – several Linux distro developers hoping to combine their efforts and create a more uniform Linux distribution. The advantages would be obvious for developers – a more stable distribution base with less variance to worry about. But isn't this already being tackled by the Linux Standard Base? And isn't everyone who was around in the Unix space a few years back getting a strange sense of déjà vu? With high profile names left out of the group, there's clearly more questions to be asked, so we charged the irrepressible Andy Channelle with the task of extracting the answers from the participants and non-participants alike – you can check out the results over the page.

Our regular security spot this issue centres around wireless networking. WiFi might be the cool technology of the



moment, but implementing it could leave your network vulnerable, so here's a few ideas on how to stop your secrets being broadcast all over town. There's plenty more in this issue, and, as ever, if you have any comments or suggestions, I'd love to hear them.

**Nick Veitch** Editor  
[nick.veitch@futurenet.co.uk](mailto:nick.veitch@futurenet.co.uk)

## Contents

**United Linux:** A marriage made in heaven? **p4**

**Wireless Security:** Hardening WiFi **p12**

**Bind 9:** Using the latest DNS **p14**

**Case study:** Hill House Hammond **p20**



**FEATURE | UNITED LINUX**

# United Linux



4 **LINUXPRO**

OCTOBER 2002

Some have called it a bold move which secures the future of the operating system; some have called it an act of desperation; others said it is merely an understandable reaction to the success of Red Hat. But what does United Linux mean for business? **Andy Channele** talks to newly appointed project leader Ransom Love.

Over 30 years ago the computer scientists of Bell Labs began their haphazard endeavour to create a brand new operating system. Through their efforts, Unix – an operating system ‘of unusual simplicity, power and elegance’ – and a computer revolution was born. Since then computer prices have plummeted, while at the same time power has soared, and Microsoft has taken over the world leaving Unix banished to the back room. Of course it didn’t help that Unix became a largely incompatible range of Unices, with IBM, Hewlett Packard, Santa Cruz Organisation (SCO) and many others chasing a diminishing market with their bank-busting proprietary solutions.

Twenty years later Finnish student Linus Torvalds began work on his own ‘Unix-like’ operating system (‘just a hobby, won’t be big and professional like GNU’) that, though it’s taking a little longer, is also having a profound effect on the way people develop, use and sell software. Linux – the kernel – was free, meaning software vendors could take that core, package up their own selection of applications, build a nice installer and sell on their distribution to those without the time, skill or inclination to ‘roll their own’ operating system. Of the many distributions available – over 200 discrete editions at the last count – Red Hat has emerged as a leader to such an extent that for many users and buyers in the corporate world, it is Linux.

### Love and unity

In June 2002 four of the biggest Linux distribution vendors joined forces to break down what they had identified as the main barriers to widespread adoption of Linux in enterprise. Caldera, SuSE, Conectiva and

TurboLinux announced that they would collaborate on a common Linux core to create the ‘next generation’ of distributions to make deploying and supporting software easier, and resolve the common problem of binary incompatibility between distributions.

These individual distros and, the partners hope, others will be marketed under their respective brands but will be ‘powered by UnitedLinux’, meaning the user can be confident that software compliant with the published UnitedLinux spec will install and run uniformly across all similarly-branded systems. The UL brand will reduce the number of mainstream distros to a manageable, and well targeted five – Red Hat, Mandrake, United, Debian and Slackware – but most pundits predict that in the enterprise sector Red Hat and UnitedLinux will come to dominate.

**“Software compliant with the UnitedLinux spec will install and run on all UL-branded systems”**

The first part of the plan was for the four partners to pool their development resources to create the UL core. This is a single CD containing the basics of the operating system, drivers and a ‘state of the art’ installer; it is designed to be an enterprise grade solution and is currently being developed around SuSE’s Enterprise Linux Server product. The core specification itself is based on work already completed by the Linux Standard Base (see below) and will be compliant with many



## FEATURE | UNITED LINUX



of the current standards for things like web services, authentication and file system hierarchy. It is intended, in the first instance, to run on x86-32, IA64 and x86-64 platforms and will be aimed squarely at business users.

On top of this UL-base, participants will include supplementary media which will provide opportunities to configure or brand the distribution's 'look and feel' and install applications and value added services. The obvious point of this – apart from branding – is that each partner could produce task or region specific offerings that still adhere to the UL specification. The first products are due to hit the shelves toward the end of this year or, at the latest, the beginning of 2003.

A month after the official launch of the project Ransom Love, former CEO of Caldera, opted to dedicate himself full-time to its success and assumed the leadership. In his previous incarnation Love was responsible for a number of Caldera acquisitions, including the little brother takeover of 'traditional' Unix maker SCO, the launch of *Volition* and also the controversial imposition of 'per seat' fees for the company's business-focused distribution. This latter position – an attempt to give Caldera's product a premium gloss – raised concerns among some in the open source community that UnitedLinux was heading down a similar path. However, Love says both Caldera and UnitedLinux's stand on Free Software has been misinterpreted: "Neither Caldera nor UL has ever advocated per seat licensing for Linux and open source components. The components on Caldera's

CD fall under multiple licenses," he says. "Some of those components have had more restrictive licensing terms than the other Linux or open source components."

And it will be a similar situation with UnitedLinux?

"As far as we're concerned, there are absolutely no plans to put proprietary components in the common CD of United Linux," Love says. To comply with the GPL the source code to all of the open source components will be made freely available, but "to protect the certification with many different ISVs and IHVs, the binary will not be made freely available.

"There will, however, be developer programs where the developers can get easy access to the binaries." The product price will include a 12-month maintenance agreement that will be tied 'per seat' style to just a single CPU.

As UnitedLinux gets closer to launching actual product, the details of the licensing will be more fully explained, but none of the partners relish the idea of calling in the lawyers to test the strength of the GPL or alienating the open source community. "We have absolutely no intention of violating the GPL or any other open source licence," Love told *LinuxPro*.

The response the initiative has inspired so far from potential partners seems to suggest there was a definite desire for standards, regardless of what other distribution builders may think. Love says companies have not just been supportive of the project, but also genuinely excited. "The concept of bringing

## UL genesis The heart of the enterprise

SuSE's Enterprise Linux Server (SELS) forms the basis of the UnitedLinux distributions. It is a business optimised product built to handle a range of services including email, Internet and application services, ERP systems and file/print services on multi-platform networks. Like Red Hat's Active Server, SuSE promise to limit the product to a

one year release cycle and support periods tailored to individual users. It is available for a range of architectures including IA-32, IA-64, S390 and zSeries (31bit), iSeries/pSeries and 64bit zSeries. There is also a maintenance-excluded edition available for Sun's Sparc.

For desktop use, SuSE pointed us

their Professional distribution and Linux Pro-Office solution which will apparently be UL compliant. The Pro-Office package is designed to be installed on top of SuSE 8 and features patches for the core distro and KDE 3.0.1. The main draw though is a full copy of Sun's *StarOffice* productivity suite.

Linux together to simplify certification for independent software vendors (ISV), hardware vendors and users is extremely compelling." At the launch an impressive range of supporters lined up behind the partners including AMD, Borland Software, Computer Associates,

Fujitsu Siemens, Hewlett-Packard, IBM, Intel, NEC and SAP, with many more expected to join as the products become available.


"Vendors spend considerable effort certifying their products and services on individual Linux distributions to ensure compatibility for their customers," Love says. "UnitedLinux will significantly diminish the number of distros that vendors are asked to certify." And the next step, he says, is to give these companies an opportunity to make their voices heard throughout the development of the project, and offer a range of 'membership' levels to encourage the full spectrum of participants.

### Have we been here before?

If all this sounds a little familiar, it's because at the heart of UnitedLinux (and other rival development efforts) is the Linux Standard Base (LSB) which was set up to define a common specification for Linux distributions and applications. The LSB has the support of all the major distribution makers. Its stated aim is to 'develop and promote a set of standards that will increase compatibility among Linux distributions and enable software applications to run on any compliant Linux system'.

Sounds like a pretty big overlap with the UnitedLinux project, but Love says UL goes a lot further: "The LSB specification covers the Linux application programming interface and the Linux application binary interface with a view to allowing an ISV or a developer to create an application that can be deployed across all LSB compliant operating system platforms.

"To do so it specifies a minimum set of functionality and the mode of implementation that needs to be shared in common with all Linux compliant platforms." What the LSB doesn't do, he says, is define "how the semantics of installation, administration, and additional functionality should be delivered as a complete customer-oriented solution." Which



Ransom Love.

**“Most pundits predict that in the enterprise sector UnitedLinux and Red Hat will come to dominate”**

is where UnitedLinux comes in. "We aim to address all of these areas and to provide a mechanism with single point contact by which IHVs and ISVs may get dependent drivers/utilities included into a commercial LSB compliant Linux OS implementation."

According to Mandrake and others, the success of the Linux Standard Base negates the need for additional 'standards' such as UnitedLinux, in fact it just adds an extra layer





of confusion. Love, inevitably, disagrees: "We have been and will continue to be LSB's strongest advocates, but LSB is only a specification. UnitedLinux is a global, united product and service offering." It will succeed, he says, because it offers software makers a cost efficient way of getting their applications

## “Neither Caldera nor UL have ever advocated per seat licensing for Linux or open source components”

onto the widest range of systems.

"ISVs cannot afford to test and certify multiple Linux offerings, even if they are LSB compliant. The differences at the driver level, the installer, and hundreds of other areas force the ISV and/or IHV to test each new product on all of the hardware or software out there.

"Take IBM as an example, they have totally

different hardware platforms and many, many different models or offerings of each. Just certifying one Linux product on all of their products is extremely costly. Now, add to that all of their middleware products... Each Linux product that an ISV or IHV certifies is incredibly expensive." The only reason that these companies have certified multiple products in the past is to get global coverage. "With the current economic climate, most companies cannot always afford to go through certification with every application. United Linux provides global coverage with the major market leaders around the world and it is a single product certification." The other advantage is that individual companies can add or remove aspects of their operating systems without forcing hardware or software developers to reapply for certification.

While LSB is a general standard, UnitedLinux is also committed to the LiN18ux internationalisation standard (administered alongside the LSB by the Free Standards Group) which aims to achieve application portability across international boundaries.

### The unknown quantity

While most *LXF* readers will have used, or at least evaluated, Linux releases from SuSE, Caldera and TurboLinux, Conectiva is something of an unknown quantity for many, yet in South America it is as synonymous with Linux as Red Hat is elsewhere. Based in Brazil, Conectiva is a Linux distro aimed at a Spanish and Portuguese speaking audience often tied to legacy hardware. This, according to some commentators, will both ease the internationalisation effort (The four partners now effectively cover the globe's major language groups) for UL and ensure a commitment to supporting legacy hardware, as many of Conectiva's

potential users have less than cutting edge computers available to them. This isn't to say Conectiva's current product isn't powerful, in fact like SuSE, the Latin American offering uses the Distributed Replicated Block Device (DRBD), a high availability kernel module (it mirrors a whole block device via a dedicated network thus ensuring the data from one node doesn't get lost if that node falls over) which makes it the perfect distro for mission critical apps. The other novelty of Conectiva is that, even though it has the hallmarks of a typical RPM-based distro, it uses *apt*, Debian's package management tool, to update the OS and apps.

### The Mandrake Position

Of course, not everyone is getting behind the UnitedLinux brand. The two highest profile critics are also, as it happens, the project's biggest rivals; the two vendors who have essentially cornered the US Linux market – Red Hat in the expanding enterprise sector and Mandrake on the desktop. While Red Hat rushed out an announcement of a deal with Oracle and Dell to create Unbreakable Linux, Mandrake made a much more direct response and went on the record to tell users why they wouldn't be joining the consortium, why the basic idea is fundamentally flawed and may, ironically, lead to less compatibility.

"A primary concept of United Linux is the false idea that Linux – like Unix in the 1980s – is splintered and diverging. The communication campaign being conducted by UnitedLinux, and even the name itself, implies that this phenomenon is happening within the Linux community! The statement goes on to point out that the Unix industry failed in the fruitless quest for the 'definition of common standards'.

## Opinion The ayes and the nays

**"This organisation should help in making it easier for Borland to offer our technology running on Linux from more vendors and open up new opportunities and channels."** *Simon Thornhill, Borland*

**"Customers now have a true international version of Linux. UnitedLinux enables more rapid adoption of Linux in the enterprise sector which, in turn, allows customers to reap the benefits of**

**Linux with lower risk and cost."**  
*SuSE Linux*

**"Today all the commercial packagers of GNU/Linux add non-free software... Several of them develop non-free software to add to GNU/Linux. Caldera has been one of the worst offenders. UnitedLinux carries this regression one step further with its 'per seat' licensing. Users of that distribution will be as restricted as if they were using Unix or Windows."** *Richard Stallman*

**"We have absolutely no intention of violating the GPL or any other open source licence."** *Ransom Love*

**"Joining UnitedLinux could destroy many of the features that have made Mandrake Linux so popular, such as our 'easy to install, easy to use' approach. It should be noted that several recent polls indicate that the four UnitedLinux companies currently rank lower than Mandrake Linux in market share."** *MandrakeSoft*



**UnitedLinux business structure.**

The Linux industry, Mandrake suggest, is analogous to the car market: cars compete in a number of areas, but must conform to a certain set of standards in order to make them road legal.


"It's the same with Linux. There are several 'mainstream' Linux distributions (Caldera, SuSE, Debian, Mandrake, Red Hat); others are dedicated to specific needs (firewalls, clusters, embedded...), others are adapted to local languages (Chinese, Finnish, Thai, etc.), and so on. Each version addresses a certain demand."

The irony of Mandrake's argument is that, fundamentally, it's the same as that proposed by Ransom Love and UnitedLinux; it's simply the approach that differs. While Mandrake are happy to work to the spirit of the LSB (which

means applications should be broadly compatible) but continue to use its own – extremely capable – installation routine, UnitedLinux hope to stabilise the specification and, more importantly, reduce confusion for potential new users.

Of course, while Mandrake are the vocal critic, Red Hat aren't going away, but Love claims UnitedLinux is nothing to do with Red Hat – even though "they are and will continue to be welcome in the UnitedLinux camp" – and all about making it easier for sellers to sell and for buyers to buy with confidence. "This is all about ISVs and IHVs and being able to deliver quality Linux based solutions to a global market of business customers who need





quality global support options.

“Red Hat’s dominance is more mind share in the US and Europe than market share around the world. They are beginning to gain some acceptance in Japan, but they have almost no market share in China or the rest of Asia. In German speaking countries, SuSE is the number one distribution by far; in Japan and China, Turbo is still the number one distribution by a long way; Caldera has a very strong presence in Korea and Taiwan and other parts of Asia; and Conectiva dominates the market in Latin America. This is not about Red Hat,” he said.

### Breaking down?

It’s a common (and much loved by management guru’s) line that any chain is only as strong as its weakest link. So which of the four partners is most likely to be Anne Robinson bait? It didn’t take long for one candidate to put its head above the parapet. Hot on the heels of the project launch it

**Each UnitedLinux distribution is built on a common core (specifying things such as libraries, file locations and kernel drivers), but the architecture allows for partners to hook their own features into the OS without forcing recertification on the part of software and middleware developers.**

was widely reported that TurboLinux had collapsed.

The story started with a number of anonymous sources claiming the company had ceased operating on Monday 15 July after finally exhausting the goodwill of financiers such as Dell and August Capital. Maureen O’Gara, writing in *LinuxGram*, said the collapse, if true, would be a black eye for the UnitedLinux consortium. This closure, O’Gara wrote, was a salient demonstration of “why the revenue-thin, layoff-prone distros had to band together in the first place” and called into question the commercial viability of open source development.

A day later Turbo responded saying that,

though the US arm of its operation was undergoing restructuring and a 'reduction in force' due to the withdrawal of funds from an unnamed investor, the company was still open for business. In a written statement TurboLinux President, Ly-thong Pham, said that despite being in the early stages of a major reorganisation, the company would be able to satisfy US customers' needs for support and future products. Comprehensive details of the restructuring are due to be announced as we go to press.

In response to the rumours, SuSE Linux issued a brief statement saying that the UnitedLinux project wouldn't be derailed by problems experienced by TurboLinux or any other UnitedLinux partner, as "SuSE's development team is responsible for the quality and schedule" of the core distribution.

At the height of the Linux boom TurboLinux was valued at an impressive \$200 million, and over the past few years the company is said to have worked its way through \$100 million of capital investment from the likes of IBM, Novell, Compaq and SGI. The last round of funding saw the company valued at a far more modest/realistic \$7 million.

### So what does it all mean?

For buyers or IT managers, at its most simple, standardising on UnitedLinux should make it easier to build and maintain a heterogeneous network. Software that installs on SuSE's Professional distribution, for instance, should work identically on products from Caldera, Conectiva or any other UL branded package. And with so many software developers supporting the project it shouldn't be too long before we start to see the orange UL swoosh on application packaging. In the short term this will mean less headaches for the IT department, while the long-term effect should be a drop in the cost of support licences as migrating or upgrading operating systems will be a far less demanding task, stimulating competition. It also means, of course, that software developers and distributors don't have to test and certify their products on every potential Linux flavour out there which, again in the long term, should reduce costs. Users,

## Basic specifications

### MAIN COMPONENTS

- Kernel 2.4.18+
- glibc 2.2.5
- gcc 3.1
- Xfree86 4.2
- KDE 3.0

### STANDARDS COMPLIANCE

- LSB
- LIN18ux
- GB18030

### SUPPORTED LANGUAGES

- English
- Japanese
- Simplified Chinese
- Traditional Chinese
- Korean
- Portuguese
- Spanish
- Italian
- German
- French
- Hungarian

meanwhile, should be able to sit down at a machine and – within reason – find most things where they expect them to be.

Love predicts that in the future there will be just two platforms certified by the major hardware and software sellers – Red Hat and UnitedLinux – and, as both of these entities are publicly committed to the LSB specification, incompatibilities should be pretty much wiped out in the enterprise space.

Despite the occasional contrary voice, most people in the industry have been broadly positive about UnitedLinux, seeing it at the very least as an opportunity to present a united front and reduce the distro jungle for the benefit of potential buyers whose experience of Linux may be limited or non-existent. Whether United Linux can locate a 'holy grail' of compatibility, reliability and lower costs to users is the real test, and we should find out the answers soon. ■

## Web links

**UnitedLinux** – [www.unitedlinux.com](http://www.unitedlinux.com) The latest news from the UL camp.

**Free Standards Group** –

[www.freestandards.org](http://www.freestandards.org). Home of both the Linux Standard Base and Lin18uk.

**SuSE Linux AG** – [www.suse.com](http://www.suse.com)

**Caldera** – [www.caldera.com](http://www.caldera.com)

**Conectiva** – [www.conectiva.com](http://www.conectiva.com)

**TurboLinux** – [www.turbolinux.com](http://www.turbolinux.com)



802.11 has brought greater freedom to laptop users, but **Nick Veitch** has a few words of caution for the WiFi generation.

**A**fter a few false starts, IEEE 802.11b is finally taking off. From Apple Airports on laptops to everything-in-a-box access points, more and more people are turning on to the joys of networking without wires. But the very thing that makes wireless networks desirable, the lack of a physical wire, also brings with it new

security concerns. When you know where all the possible access points are, you at least have some element of control over who is using them.

One might think that office workers would be suspicious if someone came in off the street and plugged themselves in to a network socket. But installing an

unsecured wireless network is like have many invisible access points, even outside your physical premises.

## WEP

The Wired Equivalent Privacy encryption protocol was developed to prevent unauthorised access to a wireless network. Like other encryption techniques it revolves around a key which must be supplied to the access point before connection to the main network is permitted. Unfortunately, early implementations of this protocol only use a 40-bit key – the equivalent of a five character password. In fact, although the provision is there to use the full 40-bits, many people opt for an easy to remember passphrase.

I don't really need to spell out how quickly any combination of five alphanumeric characters could be determined by brute force, and sadly, even the full gamut of 40-bits wouldn't take forever. If 40-bits is all you have, then at least use it well! Later implementations provide 128-bit security. Again many people choose to use a 15 character passphrase which seriously reduces the effectiveness – with a full binary key at least there's a very good chance that the guy standing outside trying to crack the code will run out of patience and laptop batteries before he strikes it lucky.

## Mac address

Most WiFi access points will also allow only specific MAC addresses (the individual 'number' assigned to each network interface) to connect. This adds a significant level of security, though it may require more effort to create a list of the allowed interface addresses. MAC addresses can be 'spoofed' (it's incredibly easy to do this in Linux for example), but the range of potential addresses is huge – the only likely way an unauthorised person could 'guess' such a number is if they already had detailed information on your network. Certainly a combination of defining permissible hardware numbers and WEP encryption would prove an effective barrier against the casual hacker.

However, the access point isn't the only

vulnerable point in your network – so are all your WiFi enabled devices connecting to it.

## Ad hoc

WiFi supports a number of connection modes, one of which is a peer to peer 'Ad Hoc' mode – you may think you have only one access point, but every device with 'Ad Hoc' enabled is an access point for a potential hacker. This is very useful for people sharing data between laptops etc, but really isn't necessary in a more static environment. The safest option is to simply turn off the 'Ad Hoc' capability in any device that doesn't strictly need it, and set up appropriate local security on the machines which do.

## Conclusion

You can't hide your WiFi LAN – you are literally broadcasting its presence all the time it is on! You can't (easily) stop the signals reaching beyond the walls of your offices/home (though it is a good idea to locate the access point somewhere where the signal is less likely to leak), so it's more important than ever to take appropriate precautions. For even more security, it may be appropriate to use VPN from the wireless

## Public WiFi

**OF COURSE, YOU might actually want people to connect to your WiFi network. There are many projects, ranging from commercial initiatives to individuals, who don't mind sharing their network with anyone who happens along. Such free access points have become popular at international airports, certain coffee houses and even just in residential areas. You can find a map of known WiFi access points at [consume.net](http://consume.net)**

**“Every device with ‘Ad Hoc’ enabled is an access point for a potential hacker. Turn it off”**

network to the main LAN, which would add a further level of protection.

Bear in mind that although you might be able to initially prevent someone connecting to your wireless network, there is still the possibility that packets can be captured, particularly from nodes, which can give away a fair amount of information. Armed with a variety of open source tools such as *etherreal*, *kismet* and others, a determined cracker could capture network packets and attempt to crack them at their leisure later. ■

# BIND 9 nameserver



Without it, there wouldn't be an Internet. Chris Denton sorts the masters from the slaves.

**B**IND stands for Berkeley Internet Name Domain. As the name-check for Berkeley suggests, it was originally developed on BSD UNIX. Now, though, it will run on just about anything, up to and including Windows. BIND is by far the most popular Linux implementation of the Domain Name System (DNS). BIND 9 is the current version, although its predecessor, BIND 8, is still in wide usage.

BIND is produced and maintained by a not-for-profit organisation called The Internet Software Consortium (ISC). The ISC is sponsored by a multitude of big corporations, including HP, IBM and Sun. BIND is Open Source software, released under its own modified BSD-like licence. BIND comprises three major components: the server software, the resolver library and some useful tools for verification and management.

## Where to get it

Well, since BIND 9 is included in most distros, odds are you already have it. However, if you are planning to put your nameserver onto the Internet at any stage, it is really important that you make sure you've got the current release installed. Check your vendor's website for updated binaries, and it is a good idea also to review [www.isc.org/products/BIND](http://www.isc.org/products/BIND) for latest version numbers and announcements.

Some would argue that it is good practice not to be at the cutting edge of new software development, as bugs and vulnerabilities then get tested out on you. However, because BIND is a widely used Internet program, you really don't want to be running anything with known weaknesses.

For instance, at the time of writing, the current version of BIND is 9.2.1. There is a security advisory notice covering all previous versions, so running them could be an open invitation to hackers. If there is not an up-to-date binary available for your distro you are probably best compiling it from source. There are detailed installation instructions provided for this, but it should just be a case of the usual **./configure; make; make install** steps. After you've performed the initial installation, clearly you still need to be vigilant. Keep an eye on the relevant websites, as well as the computer press, and be prepared to upgrade quickly if you have to.

## Using a BIND nameserver

To make use of an existing nameserver from a client PC, you need to ensure that it is set up for DNS name resolution. First check that the following line is present in `/etc/nsswitch.conf`:

```
hosts: files dns
```

This tells your computer to first resolve name requests locally, in `/etc/hosts`, and then look to DNS. Lines in `/etc/resolv.conf`, e.g.:

```
nameserver 10.1.1.222
```

tell the system which nameservers to use.

These files are not BIND specific, and can

be used with any type of DNS server, even a Microsoft one. What BIND itself does provide are several excellent querying utilities. These are included when you compile from source, but are also usually provided as a separate 'BIND utilities' package, and likely as not you already have it installed.

The **host** command performs simple DNS lookups. To find out the IP address of a domain, use the syntax **host domain**, e.g.:

```
host yourdomain.com
www.yourdomain.com has address
10.1.1.111
```

If you want more info than this, then use the *dig* (*domain information groper*) command:

```
dig www.yourdomain.com
```

The output will be far more extensive, and might not make too much sense at this stage. Rest assured, though, that all should become clear very shortly. For experienced BIND administrators, *dig* is an extremely versatile and powerful tool, with many options and uses. See the accompanying man pages for a complete run down on all its myriad features.

The *host* and *dig* commands were included in BIND before version 9, although they did not tend to be so widely used. The program they have superseded is *nslookup*. Formerly very popular, this command is still included as part of BIND 9, but only for historical reasons. It's likely to disappear altogether fairly soon.

## Setting up a slave

There are two types of nameserver with BIND, masters and slaves. A master maintains all the necessary domain information, and then this is disseminated automatically to all configured slaves. Strictly speaking this is managed on a domain by domain basis, but usually servers are either masters for all or slaves for all.

So, let's say that we want to set up a slave from a properly configured master. This is a relatively simple three-stage process.

Firstly, we have to obtain a special file called the hint file. This is needed to tell our nameserver about the top of the Internet DNS tree. Without it, you'll only be able to resolve names for which your machine itself is responsible, but with this file in place, the whole World Wide Web opens up.

The hint file is known as *named.root* or *db.cache*. It's available to download from several places, but be warned, sometimes the info included is not 100% accurate. To be on the safe side, get it from <ftp://ftp.rs.internic.net/domain/named.root> and put it in a directory you want to use for DNS, e.g. */usr/local/named*.

The second step is to create the BIND config file. This is called *named.conf* and generally resides in */etc*. The layout of this file is fairly simple. Basically, all we need to do is tell BIND the directory path we're using, what the hint file's called and then what domains we want it to look after. Here's how it's done:

```
options {
    directory "/usr/local/named";
    pid-file "/usr/local/named/named.pid";
};


zone "." in {
    type hint;
    file "named.root";
};

zone "yourdomain.com" in {
    type slave;
    file "db.yourdomain";
    masters { 10.1.1.2; };
};
```

This file is based on keywords followed by settings, so it's important to get the syntax right. Notice that domains are defined by the keyword **zone**, and that where **type slave** is specified, the master(s) are defined by IP



This is where BIND lives.





address. In the case of slaves, the **file** should be the same as that on the master, but does not have to exist at this stage. It will be created upon successful startup of the BIND daemon.

It may seem slightly strange that slaves can have more than one master, but in fact they can happily cope with several or more. The benefits and drawbacks of multiple master nameservers should be carefully considered. More than one master gives you the added piece-of-mind of not relying on just one ultimate source of domain information, but the big drawback is that each master must be managed separately. Multiple slave nameservers are a more straight-forward issue, as they offer redundancy and resilience plus the support overhead is low.

Finally, to complete our slave setup we need to get the BIND daemon running. It's called *named* and to kick it off, just type:

```
named
```

If everything has gone according to plan, an **ls** of the `/usr/local/named` directory should now reveal a file called `db.yourdomain`. The following will now produce a valid output:

```
host www.yourdomain.com
```

```
host www.linuxformat.co.uk
```

Note that the latter will only work providing the server is connected to the Internet. All logging is done to the system messages file, so it's a good idea to do a *tail* of the messages file after starting the daemon, and then on each occasion you restart it after a config change. With this kind of nameserver, any problems are generally going to be caused by either a typo in the `named.conf` file or difficulty in communicating with the master server.

## Setting up a master

Right, now it's time to get down to the nitty-gritty. Up to now, we've assumed that someone else has defined our domain info for us, but the time has come to go it alone and get ourselves a master server.

As before, the hint file needs to be acquired. Then, the `named.conf` file needs to be edited so that this time the `yourdomain.com` entry should look something like this:

```
zone "yourdomain.com" in {
    type master;
```

```
    file "db.yourdomain";
    also-notify { 10.1.1.222; };
};
```

So, this time our zone is a master, and we have the **also-notify** keyword, followed by the slave(s) that need to be kept up to date. Note that the entry for **file** is identical to the previous one. There is, however, one very important difference. The file `/usr/local/named/db.yourdomain` is not created automatically. This time, we have to create it manually.

Zone files are divided into records. These records are usually a single line of text in length. Each zone file begins with a variable declaration that is always the same but is required to avoid an error message:

```
$TTL      86000
```

Then there is a special **SOA** (Start of Authority) record. This can be abbreviated to a single line, but most commonly looks like this:

```
yourdomain.com IN SOA ns1.yourdomain.com. you.yourdomain.com. (
```

```
    200207301      ; serial
```

```
    10800      ; refresh (3 hours)
```

```
    3600      ; retry (1 hour)
```

```
    604800    ; expire (1 week)
```

```
    86400     ; minimum (1 day)
```

```
)
```

The first entry on the top line tells BIND the domain that the file is going to be authoritative for. As this is usually always going to be the same as the zone name specified in `named.conf`, you can simply put `@` instead of the domain name and it will work just as well. It's one less thing to change if you're copying a file to use as the basis of a new domain name.

The second and third entries on the first line tell BIND about the type of record that is being defined. **IN** stands for 'Internet' and is present within all of the records. The record-type is defined by the following sequence of letters, in this case **SOA**. The fourth entry is the name of the master nameserver for this domain. Subsequently, there is an email address for the administrative contact. This can be confusing, as the `@` you would expect to see has been replaced by a dot. It has to be like this, though, because the `@` symbol, as we have seen, is a special character.

After this comes some parentheses, and

## rndc the named controller program

BIND 9 comes with a name daemon controller utility called *rndc*. This can be extremely useful for administering names, especially as this works for daemons on remote hosts as well as local ones. However, the problem is that *rndc* usually doesn't just work upon installation but requires some setting-up.

Firstly, verify that we have some work to do:

### rndc status

If this comes back with something meaningful then all well and good, but the chances are you'll get some 'connection refused' type error instead. The reason it's complaining is that some authentication is required before control commands are accepted. To sort this out, we'll need to generate an 'rndc key':

```
dnssec-keygen -a hmac-md5 -b
256 -n HOST rndc.key
```

This will create the file *rndc.key*, which will look something like this:

```
key "rndckey" {
    algorithm hmac-md5;
```

```
    secret
    "lpHLaBGS8um6zVE+xcR4hA==";
};
```

Note that *rndckey* is just an assigned name, and can be anything as long as it's consistently referred to as the same thing. This file's existence gives us a base to build on, and we do that by editing the */etc/rndc.conf* file so that it looks something like this:

```
options {
    default-server 127.0.0.1;
    default-key "rndckey";
};

server 127.0.0.1 {
    key "rndckey";
};

key "rndckey" {
    algorithm hmac-md5;
    secret
    "lpHLaBGS8um6zVE+xcR4hA==";
};
```

Note that the initial version of this file sets *default-server* to 'localhost' which may seem the same as

127.0.0.1 but can cause problems if IPv6 has not been implemented, and should be amended. The final lines are achieved by simply using:

```
# echo rndc.key >> /etc/rndc.conf
```

Once *rndc.conf* is right, the final step is to create some corresponding entries in the *named.conf* file. Placed before the zone statements like this:

```
controls {
    inet 127.0.0.1 allow { localhost; }
    keys { rndckey; };
};

key "rndckey" {
    algorithm hmac-md5;
    secret
    "lpHLaBGS8um6zVE+xcR4hA==";
};
```

After these steps are completed, you should find that *rndc* now behaves itself. Try it again with the *status* option and this time you should get some interesting output. The other options you will probably find yourself using are *refresh* and *reload*, but see the man page for a full listing.

between them are five numeric fields. The text after each field is simply a comment, and marked as such by the preceding semi-colon. Only the first need really concern you, as the others are in fact perfectly acceptable defaults. The serial number, on the other hand, is an integer that must be incremented after each modification to the zone file. If this is not done, BIND will not re-read the file, and the information it provides will not change.

This being so, you might very well wonder why is the number in the example so big? Surely it would be better to start with 1 and increment from there? Well, there is nothing to stop you from doing it that way, but sooner or later you are likely to then encounter a problem with identifying when a domain was last amended. Because domain information only changes when the serial number is

incremented, the standard Linux last modified field (that you would get from running an *ls -l*) does not necessarily give you an accurate idea of the last effective change.

A good way to get round this is to use a system of serial numbering that also gives you the date last modified. To do this the number must be made up of the year, month and day, as well as containing a final figure, which should be used to indicate the number of changes made on that day. So, in the example above, the zone file was last modified once on 30<sup>th</sup> July 2002.

It is very important, when using this system that you stick rigidly to the syntax of *year/month/day/change*. If you alter this order then you will end up with a number that may actually be lower than the previous one, and BIND will subsequently reject it. For instance,





120820021 is far less than 300720021 and so this serial number would not be accepted. Also, it is a good idea not to add any extra figures to the number as if you do you are stuck with them. Having said this, if you operate within these constraints it's virtually certain that date-based serial numbers will pay off in the long run.

After the **SOA** entry come the rest of the records. There are several different types, but this example includes most of the main ones:

```
IN      NS      ns1.yourdomain.com.
IN      NS      ns2.yourdomain.com.

IN      MX 10  mail.yourotherdomain.com.

IN      A       10.1.1.111
www     IN  CNAME yourdomain.com.
ns1     IN  A     10.1.1.2
ns2     IN  A     10.1.1.222
```

The record type is set in the third column along. The most common one is the **A** record, and this simply maps a hostname to an IP address. The first **A** record here associates 'yourdomain.com' with the address 10.1.1.111. Underneath this is a **CNAME** record, and these are used to set up aliases. In the example, www.yourdomain.com's canonical name is defined as yourdomain.com, and the latter's IP address is used.

What's happened at the top of our zone file is that two nameservers have been specified as authoritative for the domain with the **NS** records. The last two **A** records are needed to ensure the **NS** records are valid, by actually giving addresses for both nameservers.

Notice there's also an **MX** record in here. 'MX' stands for 'Mail Exchanger' and it simply defines where email to this domain should be forwarded. Notice the number **10** has been inserted here. This is a priority rating, and, in the event of other mail destinations being included, tells BIND the order in which they should be tried, lowest first. Since our particular **MX** record points to a machine on another domain, there's no corresponding **A** record here.

Okay, so we're nearly ready to start our master nameserver up now, but there are still a couple of things about the above records

that need explaining. Why, for instance, is there nothing before the **IN** for some of the records? Also, what's the deal with that trailing dot after all the fully-qualified domain names? The answer to both these questions is actually the same. Without the dot, BIND expands to a fully-qualified domain name using its zone name. So, everything to the left of all the **IN**s will be reinterpreted in this way. Where there is nothing, read 'yourdomain.com' Where there is a hostname, read 'hostname.yourdomain.com'

This also applies to the names on the right of the records. So, without the trailing dot, BIND actually thinks 'yourdomain.com' is really 'yourdomain.com.yourdomain.com' and that is palpably no use to anyone.

Right, time to kick-off the name daemon in the same way we did for the slave nameserver, and if all has gone to plan the master should start doing its job.

## Reverse lookups

So far, we've set up a slave and a master nameserver to perform name to IP address translation. This works well up to a point, but to really complete the job properly support needs to be added for IP address to name translation, generally known as reverse lookup.

All the hosts we have been dealing with have been sitting on the 10.11 network. Therefore, to get reverse lookup working, we need to start by setting this up as a zone in named.conf:

```
zone "1.1.10.in-addr.arpa" in {
    type master;
    also-notify { 10.1.1.222; };
    file 'db.10.1.11';
};
```

Note the strange way the zone name is made up: the network address backwards and then **in-addr.arpa**. This slightly unwieldy convention must nevertheless be observed.

Next comes the zone file **db.10.1.1**. The **SOA** record follows the same rules as detailed above, but the rest of the file should only contain records of type **PTR** (for 'pointer'). To match the information we've already set in 'db.yourdomain', we need to input

```
2.1.1.10.in-addr.arpa.
IN PTR ns1.yourdomain.com.
```

```
111.1.10.in-addr.arpa.
IN PTR yourdomain.com.
```

```
222.1.1.10.in-addr.arpa.
IN PTR ns2.yourdomain.com.
```

Note, again, that the IP details are backwards, and always followed by **in-addr.arpa**. Also, don't bother putting aliases in this file. In most cases all but the first defined name will be ignored anyway. When this file is safely saved away, restart the name daemon (you can use **pkill -HUP named**) and then to verify it has worked:

```
host 10.1.1.111
```

If you see the lookup performed successfully, then congratulations, you're done!

## Getting the most from BIND

As well as configuring your architecture on the Internet, BIND makes a fine nameserver on an internal network. Care should be taken with this though, as the two sets of information should really be mutually exclusive. It's clearly a bad idea to put details of your whole network out in the open, and there's not much point in keeping info that should be on the Internet within the confines of your own setup. Therefore, what you really need are separate internal and external nameservers.

There is a helpful-sounding utility called *h2n* that will convert standard hosts files into BIND-style zone files. It's a Perl script that is included with some distributions, or alternatively can be obtained from <http://bullwinkle.deer-run.com/~hal/h2n/h2n>. If you're planning on converting a large host file into a BIND server this may come in handy, but the output generated does not always seem that satisfactory and in most cases it is probably better to create the new zone files from scratch.

One handy tip for using BIND to implement some degree of load sharing is the feature known as Round Robin DNS. Say you have two servers both serving out the same website using *Apache*. Simply specify both of them as **A** records within the zone file and BIND will alternate the order in which it lists them to clients, ensuring that both are used pretty much equally. Of course, if one of the servers goes down half the connections are going to

```
File Sessions Settings Help
[root@xwing root]# dig www.microsoft.com

;<<<> Dig 9.2.1 <<<> www.microsoft.com
;; Global options: Pr:infofd
;; Got answer:
;;->>>HEADER<<< opcode: QUERY, status: NOERROR, id: 5173
;; flags: qr rd ra; QUERY: 1, ANSWER: 7, AUTHORITY: 7, ADDITIONAL: 7

;; QUESTION SECTION:
;www.microsoft.com.      IN      A

;; ANSWER SECTION:
www.microsoft.com.      6200    IN      DNAME   www.microsoft.akadns.net.
www.microsoft.akadns.net. 300    IN      A       207.46.230.248
www.microsoft.akadns.net. 300    IN      A       207.46.230.220
www.microsoft.akadns.net. 300    IN      A       207.46.197.102
www.microsoft.akadns.net. 390    IN      A       207.46.197.100
www.microsoft.akadns.net. 300    IN      A       207.46.230.219
www.microsoft.akadns.net. 300    IN      A       207.46.197.113

;; AUTHORITY SECTION:
akadns.net.             160175  IN      NS       Z4.akadns.net.
akadns.net.             160175  IN      NS       ZC.akadns.net.
akadns.net.             160175  IN      NS       ZD.akadns.net.
akadns.net.             160175  IN      NS       ZE.akadns.net.
akadns.net.             160175  IN      NS       ZF.akadns.net.
akadns.net.             160175  IN      NS       ZG.akadns.net.
akadns.net.             160175  IN      NS       ZH.akadns.net.

;; ADDITIONAL SECTION:
Z4.akadns.net.          160175  IN      A       216.32.65.105
ZC.akadns.net.          160175  IN      A       63.241.199.50
ZD.akadns.net.          160175  IN      A       208.132.160.36
ZE.akadns.net.          160175  IN      A       12.47.217.11
ZF.akadns.net.          160175  IN      A       63.215.198.79
ZG.akadns.net.          160175  IN      A       204.248.36.131
ZH.akadns.net.          160175  IN      A       63.208.48.42

;; Query time: 266 msec
;; SERVER: 172.24.200.31#53(172.24.200.31)
;; WHEN: Fri Jul 5 11:02:48 2002
;; MSG SIZE  rcvd: 417

[root@xwing root]#
```


have problems, but with that in mind, it's still a far sight easier than clustering!

Two of the nice new features of BIND 9 are support for Dynamic DNS, or DDNS and DNSSEC. DDNS works with DHCP to dynamically update all IP address info as soon as it changes. DNSSEC is an encrypted version of DNS. Setting up these features is beyond the scope of this article, but they are fully documented. BIND 9 also comes with support for IPv6, which may prove useful in years to come when the IPv4's dominance wanes.

The above should provide you with a good grounding, but there is a lot more to BIND than has been covered here. Further information is available from the usual suspects, such as man pages, online help resources and the inevitable O'Reilly textbook. It's worth taking the trouble to go down these avenues as BIND 9 is not only an essential part of any self-managed Internet presence, but also a potentially invaluable backbone of any medium to large network. ■

**The output from *dig* just about fits on one page.**

## PROFILE | HILL HOUSE HAMMOND



# Open Source POLICY

When Hill House Hammond needed to push its IT systems and infrastructure into the 21<sup>st</sup> century, it chose to work with Linux. **Jon Kaye** investigates.

**F**ounded in 1959, Bristol-based Hill House Hammond (HHH) is an insurance intermediary with some 240 high-street branches throughout England, Scotland and Wales. The wholly owned Aviva (until recently CGNU) subsidiary

provides over one million customers with core motor and home insurance plus a growing range of commercial and specialist policies through its branches, seven-day call centre, and through its web site.

Growing organically and through a robust broker acquisition programme, the company had by 1998, however, become a sprawling mass of disconnected branches, separate call centres, and 10-year-old computer equipment that was falling apart at the

seams. Broken equipment was sent from branches to the Bristol HQ, and repair effectiveness was very low.

This was a company which, unlike Aviva, had no history of IT. The first branch and commercial systems were acquired without any serious investigation, database sizes mushroomed in an administrative vacuum, and servers ground almost to a standstill trying to process the overblown files.

There was little understanding of testing, integration or even design. Applications were bought by sales teams without recourse to any IT policy, and were delivered without even a nod to enterprise integration or corporate standards and strategies – which was fair comment, since there were none.

And yet the company continued to flourish. The question was, for how much longer? At this point in time, the highly competitive and increasingly squeezed insurance sector was throwing millions at technology just to maintain parity, let alone carve any leadership position. Ever decreasing development times were needed to generate a growing range of products, services and special offers, as newly empowered consumers heartily embraced the new world of choice while throwing off the shackles of life-long brand loyalty.

Hill House Hammond's response was the appointment of a new managing director with a background in IT began a process that would bring technology closer to the business. But it wouldn't happen overnight. While much needed to be improved, the company maintains a policy of 'wood chip, not blue chip' when it comes to in-house investments – delivering the best value to customers by keeping operating costs as low as possible.

The question about where to start proved challenging, but not overly vexing. What was needed above all else was a moratorium on all IT projects, a three-month period to iron out what were predominantly known faults and to stabilise a hitherto reactive environment built around a fire-fighting mentality: system availability outside the branches was often barely more than 80%, the claims system too often was slower than treacle, the financial system was often not

up until lunchtime, and perhaps most critically, the call centre would grind to a halt.

"It was a question of breaking out of a vicious circle," says Operations Manager, Alex Dunn, "and showing how IT can deliver value to the company. It was about delivering our staff from maintenance and allowing them to create."

**“There was a fire-fighting mentality – system availability outside the branches was 80%”**

But several problems needed fixing before any advance could be made. The time to load data from branches for Claims and HHH's insurers' reports needed to be hours not weeks, sales and financial systems were running on separate servers without any integration, management information was virtually non-existent, and there was no electronic communication with branches, which received new business details from call centres through the post on floppy disks!



## PROFILE | HILL HOUSE HAMMOND



In all, some 150 outstanding issues had to be addressed.

### Centralisation

Away from the fire fighting, plans were being drawn up to provide a platform for growth, a technology infrastructure capable of supporting and adding value to any new business initiative.

Working with IBM infrastructure business partner, Anix Group, the decision was made to deploy one large central server in a newly created data centre environment that would run all core business applications, and to consolidate the different call centres into a single site, integrating the disparate telephony functions. "The idea," says Dunn, "was to move away from the leaky old boxes, and to house the finance, claims and other systems on new equipment that would support our current and future needs."

justifying the choice and configuration of the original system, the availability of which remains at well over 99%.

Allied to that centralisation was a programme to replace all branch servers with new IBM xSeries Intel-based machines, configured to run Pick software under Linux, making HHH one of Europe's biggest Linux implementations. Next, the original dumb terminals in the branches would be replaced by desktop PCs, which in turn would act as thin clients running Linux and provide access to the original sales applications via a terminal emulator.

What's more, branch staff could now have access to internal and external mail via Outlook Web access, and there is limited access to common office applications via Citrix. Ultimately, the company wants all applications to be browser-based. In the meantime, the benefits have already kicked in: new branches can be deployed through just a CD that includes all configuration components as well as a desktop operating system; a process that in future will be done electronically.

Likewise, network requirements would be met through IP-based virtual private networks (IP VPNs), furthering the company's shift towards a comprehensive Web-based operating environment.

**“New branches can be deployed through a CD that includes config components and the desktop OS”**

At the heart of the system was an AIX-based IBM pSeries SP supporting financials, sales and claims applications, running Pick software (in which the branch sales applications were written). Neither processors nor storage has required upgrading since the system was installed three years ago, easily

### Hosted web commerce

Anix's role in HHH's transformation has been substantial. Besides configuring and implementing the pSeries, disaster recovery services and the 240 branch-based servers, the company has also provided extensive

## Your case studies Diverse Linux installs

ONE OF THE PLEASURES OF producing this mini magazine has been coming across such a diverse and interesting set of examples of real-world Linux usage in businesses big and small.

Of course we don't know everything that's going on, so we're more than happy to hear from you about your

examples of interesting Linux installations big and small.

Email us now at [linuxformat@futurenet.co.uk](mailto:linuxformat@futurenet.co.uk) using the subject line "case studies."

Or send your letters to:  
Linux Format, Future Publishing,  
30 Monmouth Street, Bath, BA1 2BW



training, software support, ongoing infrastructure support and hosting services to support the company's online presence, which now provides quotes, premium comparisons and the ability to buy online.

"The one thing we didn't have was any particular skills in setting up a Web-type infrastructure," recalls Jon Whittle, e-Commerce Manager responsible for setting up the online service. "And it wasn't possible then to include our own Web presence in the CGNU group infrastructure, which was very much based on a Microsoft platform. We, on the other hand, were using Pick, Linux and AIX, and so were forced to go it alone to a certain extent. What we needed was help to deploy the application we built on line."

Anix built the firewall and the windows-based Web server, and provided the Internet connection, all of which sits in the Anix site as a hosted service, maintained on a 24x7 basis, that supports HHH's proprietary application ([www.hhh.co.uk](http://www.hhh.co.uk)). Following a successful pilot, the system went live and now provides a significant new source of sales, and between 2,500 and 3,500 quotes a week.

Meanwhile, the company has also created a management information environment, with its own dedicated server that produces high-quality reports based on Business Objects applications.

### IT in the ascendant

Four years on, HHH's IT operation is unrecognisable, both physically and in terms of its value to the company. It has enabled the firm to become much more automated and truly multi-channel, without diluting its core emphasis on the value of human interaction. Its message remains a pertinent one: Human, Helpful and Here.

But Anix's work is not yet done. Besides providing ongoing services and skills, the company is involved in the development of a second computer centre in Yate, just outside Bristol, which will support the insurer's centralised and secure data model.

"The value of Anix's contribution is implicit in much of the improvements that have taken



place," says Dunn. "From 80% availability our branches are now running at 99.96%. In the data centre the pSeries SP is very robust and is now running at 99.6%, and an excellent Disaster Recovery policy has been put in place.

"And now that people can see IT delivering real value, its standing within the company has risen greatly. Not long ago it was seen as a backwater having little importance, and whose staff were given little or no respect. People in finance, sales and claims operations can clearly see the improvements, suffering just one or two incidents a month instead of 30 or 40.

"Finally," concludes Dunn, "IT is able to serve the business and help HHH to move forward into the 21<sup>st</sup> Century." ■