

LINUX

www.linuxformat.co.uk

FORMAT

**GET MORE
FROM LINUX!**

Effortless backups p61
Set up web proxies p74
Customising Mozilla p78
Advanced graphics p84
Dynamic Flash animation p88

Is Linux ready for the DESKTOP?

Experts and industry insiders
answer the ultimate question

PLUS

Which Windowmanager?

The best and the worst that Linux has
to offer – rounded up and rated. Find
out which is best for your needs p36

UNREAL TOURNAMENT 2003

The best Linux game of the year
is here already p30

CHARGE YOUR WEAPONS!



WHAT XFT2 MEANS FOR YOU

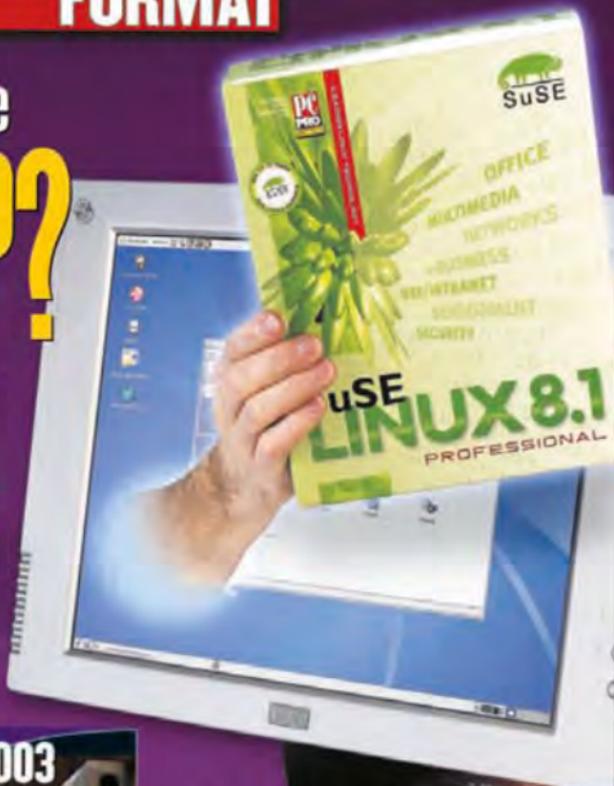
The new X font server has
implications for users and
software developers p56

Serial ATA

What is it? Why is it? And
will it work with Linux? p66

Lindows on test!

First full review of the
controversial distro p22



DVD issue also available Printed in the UK

LXF37 FEBRUARY 2003 \$5.99

02

9 771670 425019

(\$5.99 outside UK and Republic of Ireland)

THE UK'S BEST-SELLING LINUX MAGAZINE

Are you ready?

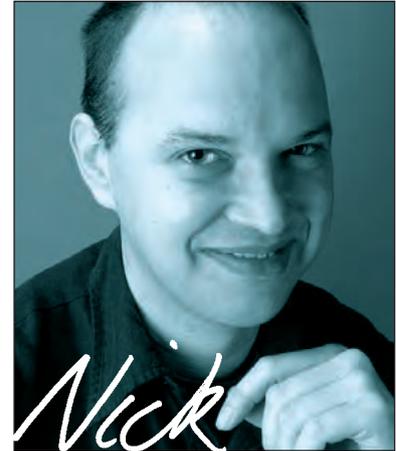
We don't shirk the big challenges here at *Linux Format*, so when we decided to have a stab at the ultimate question – “is Linux ready for the desktop?”, we approached the task with gusto. Five minutes later, in a moment strangely reminiscent of *The Hitchhikers Guide To The Galaxy*, we decided the answer was probably simple, but it's the question that causes the problem. What exactly do we mean by ready? What actually is the 'desktop'? Both of these concepts mean different things to different people and different situations. There is a world of difference between the office desktop and the home desktop for example (or at least there should be). Nevertheless, we have soldiered on to bring you the most up-to-date assessment of Linux's suitability for a number of potential desktops, and dragged in the leading lights of Linux to help us out with their take on what seems to be often an emotive topic.

Moving on to a less contentious, but slightly related and very important topic – fonts. You may

not have noticed yet, but the way fonts are used, and the software that uses them is changing. *Xft2* not only brings new functionality and answers a lot of long-standing problems related to unicode and internationalisation, but it also fundamentally changes the way developers and users will be able to access fonts in the future. Find out exactly what has changed and why, and most importantly what it means to you in our font feature that starts on page 56.

Lindows has probably been one of the most talked-about distros of the last twelve months. Wild claims and cunning marketing have made it very visible, but is it any good? Our reviews section has the answer, and we also examine the benefits of Jool's Kwartz server, the latest *Arkeia* incarnation and a contender for game of the year.

All that should certainly be plenty to keep you busy until the next issue. In the meantime, if you have any unfeasibly hard questions about the Linux market that you need an answer to, or an opinion to share, please just drop us a line...

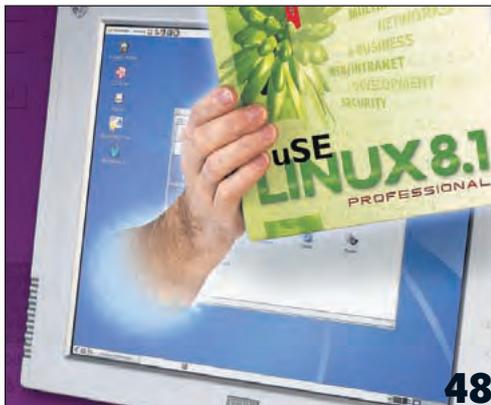


Nick Veitch EDITOR

Find out the answer to the question of the moment, and what else needs doing to attract more Windows refugees? **p48**

The most controversial distro of recent times gets a very thorough going over **p22**

Xft2 changes the way fonts work on your system **p56**



AIMS OF THE MAGAZINE

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

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

MEET LINUX FORMAT'S TEAM OF WRITERS...



Andy Channelle
Regularly writes the news, and co-authors this issue's cover feature – essential reading for all computer users.



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



Rich Drummond
He was off to the USA, but is now in Scotland. Thai readers look out for him next month as he plans going to Paris!



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



And introducing... Matt Nailon
New Production Ed and newbie defector from 'another' OS who enjoys beta-testing MMOGs.

Neil Lucock
Whatever you use your Linux box for, you need *OpenOffice.org*, and Neil's handy tutorials to optimise your abilities.

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

Charlie Stross
Perl guru with a beard that Edward Lear would get all poetic about, Charlie concludes his Perl series this issue.

Mike Saunders
Distros and desktops are Mike's bread and butter, this month he will be mostly eating Window Managers.

Les Hatton
The Prof from the UK Computing Laboratory has been sharing his know-how since God was in short trousers.

CONTACT US

Letters for publication:
lxf.letters@futurenet.co.uk

Subscriptions/back issues:
subs@futurenet.co.uk

Technical help/Ask the Experts:
lxf.answers@futurenet.co.uk

Disc problems:
lxf.support@futurenet.co.uk

General enquiries:
linuxformat@futurenet.co.uk

Website: www.linuxformat.co.uk

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

More contact info on p114

LXF37 February 2003

Contents



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



COVER HIGHLIGHTS

Is Linux ready for the desktop?

Our Linux luminaries muse over the current state of our favourite operating system

48

22 Windows on test!

Whatever happened to the 'Windows compatible' tag?

36 Roundup Window Managers

Some familiar, some new – eleven WMs compared

66 WhatOnEarth is Serial ATA?

If you think you know about IDE, you're in for a surprise

30 Unreal Tournament 2003 review

The Linux port is hidden on the Windows version!



'I have never damaged any electronic components with static, but this is probably because I am a dull computer scientist'

61 Professor Les Hatton tempts fate...

'Let's get rid of snobbery in the Linux community and start getting Linux accepted for what it is – a damn good OS!'

14 Mailserver – Integration is the key

LINUXPRO

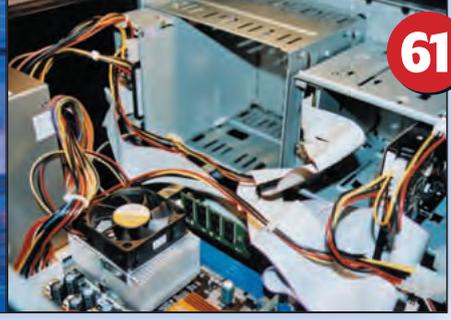
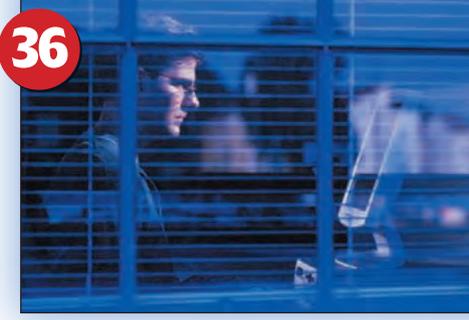
Globecast real-world accountancy solutions
Recycling use Linux to save you money and save the environment
Business Continuity Can your business cope with disaster? **Open MP** computing to extremes





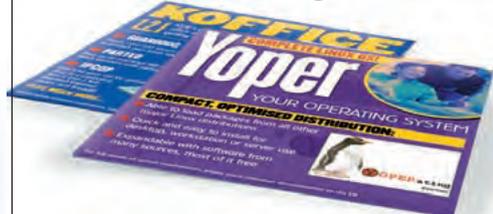
26 Jool Kwartz server – as pretty as this on the inside?

<p>6 News Analysis of the latest Linux stories</p> <p>12 Fosdem 2003 Jon 'Maddog' Hall & others speak</p> <p>14 Mailserv Got an opinion? Have your say</p> <p>20 Arkeia review Making a backup isn't a chore</p> <p>22 Lindows review Contraversial OS on test</p> <p>26 Jool Kwartz review Small office server solution</p> <p>30 Unreal Tournament 2003 An FPS you say? Never heard of it...</p> <p>34 Book reviews A couple of Red Hats and a PHP</p> <p>36 Roundup Window Managers They're all on the coverdisc too!</p>	<p>42 HotPicks Bluefish, Gambas, Phoenix & more</p> <p>56 Xft2 – The Changing face of fonts How fontservers make life easier</p> <p>61 Build your own Linux backup box Prof Les Hatton spends only £267</p> <p>66 WhatOnEarth is serial ATA? It just gets faster and faster</p> <p>72 Perl tutorial The final part of the popular series</p> <p>74 Squid tutorial An Open Source web proxy cache</p> <p>78 Mozilla tutorial New series – Rhapsody In Orange</p> <p>82 Java tutorial Compile Java to native code</p> <p>84 OpenOffice.org tutorial Simple vector graphic effects</p> <p>88 PHP tutorial Flash and Ming add pizzazz to PHP</p> <p>92 Answers Your queries answered</p> <p>110 User Groups Local support, events and fun</p>
---	---



COVERDISCS

A DVD or 2 CDs packed full of the latest Linux goodies **100**



CDS A AND B

Yoper a compact, optimised Linux distro for desktop, workstation or server use; **Guarddog** quick-'n'-easy firewall protection; **parted** work magic with your hard disk; **IPCop** turn your old PC into a full-featured Internet router & firewall; **GTK-Gnutella** share files with thousands of other Internet users; **jEdit** powerful & flexible text editor for programmers ; **KOffice 1.2.1** new version of office suite



DVD

Flightgear Take to the skies in Orville and Wilbur's kite; **MSC Linux** highly optimised distro for extreme performance; **PostNuke** website creation; **NVidia drivers** updates etc

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



SAVE MONEY!

SUBSCRIBE TO LINUX FORMAT

See page 98 or phone 0870 4448645



Newsdesk

Sony and Panasonic collaborate on Linux; Israeli military to use Linux; Red Hat in the black; Turn-Key Linux Audio; IBM helps India try to avoid proprietary trap; MySQL upgrade; New Lycoris Desktop/LX

LES POCHETTES VIDENT?

Mandrake's problems continue

Despite consistently releasing one of the most popular desktop-orientated distros, France's MandrakeSoft are again facing financial problems and have launched an Increase Of Capital drive, hoping to mobilise users and help raise an estimated US\$4 million.

"This level of cash infusion would resolve outstanding debts, cover the expenses needed to become profitable, plus secure an extra amount to satisfy the needs of future growth," the company said.

Ironically, Mandrake have been edging toward profitability and had forecast to break even with their next

major release in 2003. The increase of capital is said to be needed in part to cut a number of 'expensive external contracts' containing break-penalties. In a statement, MandrakeSoft also said the money was needed to bridge a shortfall after a major investor pulled a US\$4 million package in July 2002 – "As a result we have a big short-term cash issue."

Been here before?

MandrakeSoft's efforts to weather the 'technology bubble' are common to many hi-tech outfits after a period of initial profitability – expanding too fast, entering into expensive long-term contracts, diversifying into e-learning and bringing in new staff, including a 'world class management team'. As the tech-boom started to falter in 2001, Mandrake was burning its way through \$1.5 million per month, while the business model failed to deliver the expected returns. After a change at the top, the new management cut expenses and refocused the business back towards Linux applications (including products such as the CLIC cluster management tool and Multi-Network Firewall), support services and their user-friendly distribution. The change in strategy appeared to work,



Make the switch

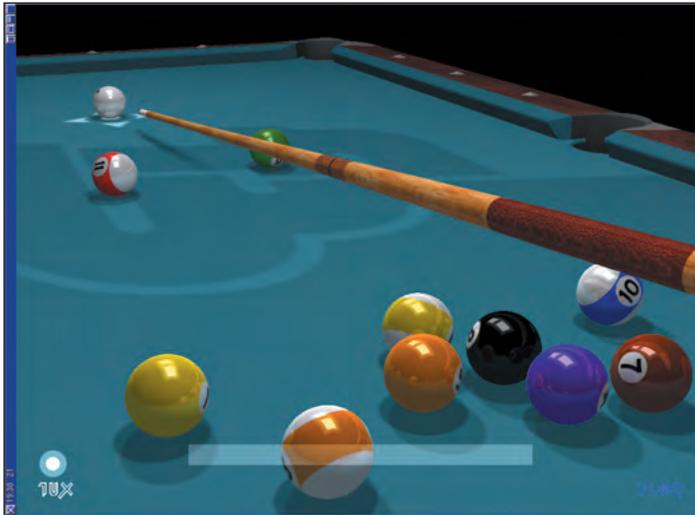
Bring me your Windows, your Windows, your Lycoris

Echoing Apple's Switch campaign (and the embarrassing MS 'spoof'), Mandrake are offering users forced to buy an operating system with their PC the opportunity to 'upgrade' to either the single CD edition of Mandrake 9 or the full version of 8.1/2 for the bargain basement price of 12.50 (or US\$12.50). The packs also include a voucher knocking US\$12.50 from the cost of the full version 9 pack.

with Mandrake posting results edging towards the break-even point. This latest cash-flow problem may hint that more changes are needed, and recent developments in the licensing of the company's Multi-Network Firewall, may point the way to future growth.

Vice President Gael Duval said the dual licensing regime did not mean Mandrake was moving away from Open Source software, but that users would have more choice. He told Newsforge.com that the system was

very similar to the model that Trolltech use to develop Qt. "People who want GPL can use the GPL, and others can use the Commercial License," he said. Under the GPL, developers would still be able to release a commercial version of the firewall, but would also be compelled to distribute the source code; under the MNF Commercial License you can make a commercial distribution without making the source available, thus protecting any alterations or improvements.



The realistic 3D effects with a decent graphics card are quite spectacular.

A LOAD OF BALLS

Anyone Foobilliards?

Foobillard is free, open source billiard/pool game for Linux built with OpenGL and featuring realistic physics and 3D mode that requires users to don fetchingly traditional 1950s-style red and blue paper specs. Author Florian Berger is currently undertaking a diploma in the Physics of Condensed Matter at the Institute for Theoretical Physics, hence his attention to detail when it comes to plotting the actions and reactions of balls on a lovingly rendered baize table. The latest

release is version 2.2; and Berger plans further development.

The game is available in both GLUT and SDL formats – the latter also having the benefit of full sound support – and requires **libpng** and **freetype2**. Players lucky enough to have a Gforce 2/3 or greater can witness the full splendor of reflective balls, real shadows and bumpmapping on the baize, though the game is quite playable on slower cards. More info: <http://foobillard.sunsite.dk/>

NWN Client news

Forgotten Realms not forgotten



The long-delayed native *Neverwinter Nights* client is close to finished according to Bioware. An intended dual-release with the Windows version was put on hold as the development team concentrated on the Linux NWN server. The client was then delayed due to a lack of native Bink and Miles support (both of which were solved thanks to the intervention of Rad Games Tools), the release of the second *Lord Of The Rings* film and of course, Christmas.

Bioware says all outstanding issues have now been resolved and the Linux client is imminent.



One of the reasons why the online fanbase of the Windows version of *Neverwinter Nights* numbers well over half a million? Hmmm...

NEWSBYTES

■ IDC analyst Dan Kusnetzky says his research points to Linux surpassing Apple's OS-X product to become the number two operating system, perhaps as soon as the end of 2003, and definitely by 2005. IDC has based its forecast on the recent developments that have seen Linux becoming available on pre-made systems. Growth, though, will not be as explosive on the desktop as it has been, for instance, in the server market. Kusnetzky claims that the main problem facing Linux as a desktop OS is the lack of well established applications. "Consumers and organisations select applications first and let that choice direct the choice of the operating environment," he told ZDNet.



Michael Robertson

■ Speculation has been rife throughout 2002 over the identity of the **anonymous benefactor** offering US\$200,000 to the individual or team who could get Linux to run on Microsoft's Xbox. US\$100,000 has already been won, but the deadline for the second part of the prize has been extended by a year to December 31st 2003. There's another US\$100,000 up for grabs for the first hacker(s) to produce a version of Linux to run on an Xbox with no hardware modifications. And mysterious Mr Moneybags's identity isn't such a surprise – it's Michael Robertson, CEO of Lindows.com, tipped as prime suspect from the start by those in the know.

■ Linux creator **Linus Torvalds** has been nominated as one of 365 modern saints by Dutch broadcaster KRO. Other beatific nominees included Bruce Springsteen, Tina Turner and Wild Swans author Jung Chang. No St IGNUtious though.

■ **The Eclipse Project** – which aims to create a complete Open Source development environment – has launched a trio of new projects and announced the addition of 13 more members. A donation from Fujitsu brings COBOL tools to the environment, the Eclipse Modelling Framework allows effective management of metadata and Project Hyades provide a suite of automatic testing tools to check the status of Eclipse applications. The new cadre of members include Hewlett Packard, SAP and Oracle.

Jono Bacon

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



COMMENT

The memory remains

“ Some LXF readers may remember with affection a game called *Rise Of The Triad* – a Doom-esque FPS. Like many great games superseded by newer titles, it took on a seemingly obscure cult status – until it was GPLed recently by its developers.

A major game developer has realised that software can have its shelf life significantly extended by making the source free? By no means a new idea, but it has not caught on with major vendors (aside from a few notable cases).

If a vendor stops development of a particular app, it can mean that it will no longer be available. In many cases there are people who will still use the application or would like to rebirth it and hack the source code to keep it going; aside from some possible third-party source code that may be in the archive, I see little practical explanation why vendors cannot give away this older technology.

It strikes me that the Free Software world should seek to promote with vendors a policy of Open Sourcing applications when they are no longer commercially useful to the vendor. Some may see it as a gift, and some may see it as nothing special, but the culture of giving to those who have kept the vendors in business could become another offshoot of the Free Software revolution.

I'm sure that many of us have thought this – but now it's time to act upon the strength of our convictions. Let's get on the phone to vendors and let them know our thoughts. Free Software, its policies and its culture is developed by the community – when we pull together, we can make things happen!

POWERFUL ALLIANCE

Electronic giants embrace embedded Linux

Sony and Matsushita Electric Industrial (better known as Panasonic) have cast aside their traditional enmity to collaborate on a new version of the Linux operating system tailored to digital consumer devices. On top of this ground-breaking partnership, the duo have suggested that they also have a commitment to the project from other electronics giants including Hitachi, IBM, NEC, Koninklijke Philips, Samsung and Sharp.

The project will go beyond mere tinkering, intending to rebuild the OS from the kernel up, and the first release is tentatively scheduled for mid-2003. A Sony spokesman stated that though the operating system is

already well-established, in the embedded space there are still areas, such as start up times and real-time performance, that needed attention; saying "These examples could be said to be shortcomings of Linux so we are going to develop it further."

Both Sony and Masushita have invested in Linux (Sony's Cocoon PVR system uses the OS) and are said to be concerned about the prospect of Microsoft dominating home entertainment in the same way it has with desktop computers. They are also both backers of embedded developer MontaVista; unconfirmed reports suggest that MontaVista's already established OS will be the starting block for the new project.



Everything from DVD players to camcorders could soon be running the Sony/Panasonic embedded version of Linux.

Linux Web Watch/



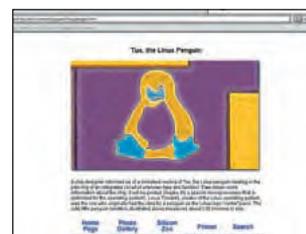
www.linux-on-laptops.com



www.linuxlookup.com



<http://lnx-bbc.org>



The smallest penguin ever?

Just browsing

Linux and music on the move, a very small distro and the world's smallest penguin picture.

The first place you should look when it comes to adding Linux to your laptop is www.linux-on-laptops.com. This is an extensive resource containing information and documentation on the suitability of various laptops to Linux. It's a great site to visit before you actually go out and part with your cash. The main part of the site is a manufacturer-by-manufacturer breakdown on tested hardware. The testing is actually done – and hosted –

by unconnected users. While this allows for comprehensive testing, it does mean that there is no standard layout for the information, and some of the links are dead – but don't be put off, a little trial-and-error can end up being very rewarding indeed.

If it's an iPod rather than a laptop that's giving you grief, LinuxLookup (www.linuxlookup.com) have a very useful primer on making sure Apple's miracle of musical miniaturisation works

as expected with your Linux box. The information on the iPod is brief and to the point, but there's also plenty of other stuff to read on the site.

The LNX-BBC project, a rescue distribution house on a bootable 60MB business card CDROM (see – it is a useful format) is nearing its 2.0 release. The developers suggest LNX-BBC is ideal as a rescue disk for post intrusion examination or as a portable workstation OS. <http://lnx-bbc.org/>

And finally, take a break by wandering around the Silicon Zoo at <http://micro.magnet.fsu.edu/creatures/index.html> a site specialising in extreme microscopic photography, we chanced upon a rendering of Tux, the Penguin, on an unspecified chip. This is the sort of website that can amply kill half an hour per day if you're at all curious about very small things. This penguin, for example, is a mere 130 microns tall. You'll go blind...

NEWSBYTES

■ **IBM** are out of the hard disk business after selling up to Hitachi for \$2.5 billion.



■ Australia's **University of Wollongong** is shifting all of its 1,700 computer studies students from dual boot Windows/Linux systems to Linux only, thanks to the OS's ability to be 'locked down' preventing first year students from breaking things. They will only be let loose on Windows PCs when they've mastered the basics.

■ **NVIDIA** have released Linux drivers (covering both 32 and 64bit versions) for their most recent graphics cards. Get them from <http://www.nvidia.com/content/drivers/drivers.asp>

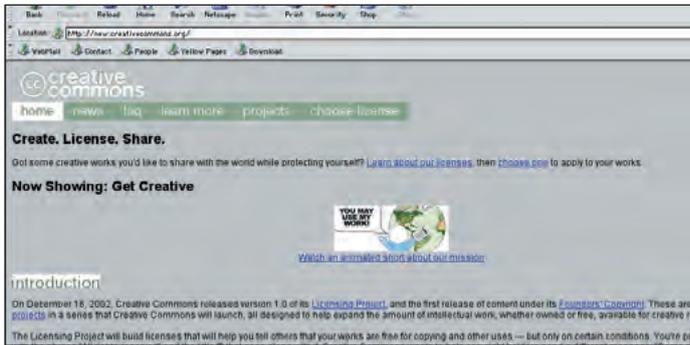
■ The latest **OneStat** website usage statistics suggests that Mozilla has surpassed Opera to become the third most popular browser after Internet Explorer and Netscape 4. OneStat's results have been criticised in the past as they don't factor in users who may use a User Agent String to convince a website they're running IE instead of one of the alternatives. The company have also revealed that most people surf at a resolution of 1024 x 768.

■ China's first **Dragon** chips have rolled off the production lines. The first batch of processors (totalling 10,000 units according to the Chinese Academy of Sciences) are said to be running at around 260MHz and will be making their way to PCs, mobile phones and TVs.

■ **AOL** have sought and been given a patent covering instant messaging systems. There has been speculation that AOL/Time Warner will use the patent to extract royalty payments from its rivals Microsoft and Yahoo! though a spokesman said no decision had been made about the future of the patent.

■ **Bochs 2.0**, the latest version of the open source x86 virtualisation software has been release and is, according to project leader Bruce Denney, about twice as fast as the previous major version and boast support for the MMX instruction set.

■ **Groupe Bull SA** of France have launched their LinuxEXPRESS product, which aims to combine the advantages of industry support with the flexibility of Open Source. For more see www.bull.com



The CC mark lets authors specify how their work can be used and distributed by others in a way that is easily understood and documented.

LICENSING MADE EASY

The end of Copyright confusion in sight?

Creative Commons is an intriguing project to address the shortcomings of the all-encompassing copyright laws as large multinationals encroach on every part of our cultural lives. The **CC mark** covers a series of four licences designed to simplify access to your work in a similar way to the various GPL/BSD/MPL style licenses; but CC licenses can be applied to any body of work in any format – useful for documentation, free music etc.

The licenses provide progressively conservative rights to alter and/or distribute your original work. *Attribute* allows others to copy, distribute or display work (or derivatives) so long as the original author is credited. *Noncommercial* allows copies to be made and distributed only on a non-commercial basis. *No Derivative Works* specifies that the work must remain intact. *Share Alike* ensures that any work

derived from the original is redistributed under the same license.

As well as bringing the Open Source ethos to other areas of our culture, the CC mark also tries to make choosing your licence intuitive and legally watertight using a series of website questions. You then download three iterations of the licence, the first (the *Commons Deed*) expresses the licence in layman's terms, the legal stuff is taken care of in the *Legal Code* and, finally, there is a machine-readable version (*Digital Code*) for embedding into work to allow search engines to identify it by the terms of use.

CC is also implementing a Founders' Copyright system to allow authors to donate their protected works to the public domain on a 'death + 14 years' basis, as was originally specified in the first US copyright acts, rather than the current 'death + 70 years' basis.

<http://creativecommons.org>

SUBSCRIPTION WORKS

Red Hat stay in the black

Red Hat has posted positive results for the third successive quarter – a small profit of US\$305,000 on revenues of US\$24 million; the latter figure is a jump of 14% on the previous quarter, slightly better than predicted. Sales of the flagship Red Hat Advanced Server (RHAS) product reached 12,000 units, up from 8,000 for the second quarter. When Red Hat announced the subscription costs for RHAS (US\$900 p.a. approx) it was thought they would scare off customers used to the low

costs, something that doesn't appear to have happened as a large percentage of sales are thought to be users upgrading. Chief Financial Officer Kevin Thompson said that he expects the initial costs to reduce by a few hundred dollars in the next year, but that the decrease in per-user revenue would be offset by a larger subscriber base. He also said that the company would be looking to a more stratified pricing structure in the future, aiming to cash in on the needs of high-end users.

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

Spare a paradigm?

“ Columnist John Dvorak recently chided Linux for looking too much like MS Windows while simultaneously pleading for a replacement for the desktop metaphor. Perhaps it is ageing, but the desktop works well and is familiar to millions. To convert the masses, you can either give them something revolutionary – risky – or something familiar – ‘better’.

With a recent strategy change, the Redmond Gang have decided that name-calling (Linux as “cancer” and “un-American”) isn't effective, so non-user-friendly and expensive are the new mantra. Just who are they kidding? What could be more unfriendly and expensive than their new licensing scheme? Sadly, ersatz studies commissioned by Microsoft will spew this disinformation for some time.

Meanwhile, José León and his friends at www.xpde.com are cloning the WinXP interface. I can hear Dvorak *et al* screaming, but the hoi polloi don't really care what OS is on their workstation. A familiar interface will go a long way to introducing these users to their new Linux desktops.

MS also claims that using Linux requires “costly” IT people, but a good foundation in Unix enables one to work with *any* Unix variant, including Linux. With Linux being a choice of younger entrants to the IT world, a cadre of talent awaits.

As always, power users have a choice in desktop environments for Linux and there may yet be a new desktop paradigm out there. The power and flexibility of Linux and the ingenuity of its developers provides us with not only a wide range of desktop choices, but the potential for a new desktop paradigm as well. ”

DMCA TOOTHLESS?

Elcomsoft not brought to book

A California jury has returned a not-guilty verdict on ElcomSoft, the Russian company accused of illegally selling software to crack Adobe's ebook format. The result is being seen as the first dent in the all-powerful Digital Millennium Copyright Act (DMCA) which outlaws developing tools that could be used to circumvent copy protection.

The case goes back to July 2001 when ElcomSoft employee Dmitri Sklyarov was arrested at a Las Vegas conference after a demonstration of

his company's *Advanced eBook Processor*, which removes the copy protection on files stored in Adobe's eBook format. When Adobe shied away from prosecuting Sklyarov personally, the case against him was dropped and he agreed to testify against his employer.

Because the software is obviously designed to break Adobe's copy protection, the trial hinged on whether or not ElcomSoft understood that by selling the software via the Internet, they were violating the DMCA. ElcomSoft President Alexander Katalov

said in his testimony that after Adobe had raised concerns about the software, it was removed from the online store. He also said the processor was intended to be used on legitimately purchased eBooks. Under cross-examination an Adobe programmer said that, despite hiring expert companies to scour the Internet for cracked eBooks, not a single item was found.

Jury foreman Dennis Strader said he and his fellow jurors could understand why a large company would risk its future by distributing



Protests were held in cities across the world following the arrest of Russian programmer Sklyarov.

software it knew was illegal. Strader also said in an interview following the trial that the panel found the DMCA itself confusing.

PUNY PROCESSORS

How small will your PC be?

Axis Communications has squeezed a complete Linux-based computer, including an Ethernet transceiver, 2MB Flash and 8MB SDRAM onto a single 27x27mm chip. The AXIS ETRAX 100LX MCM 2+8 is designed for integration into multimedia hubs, security systems and wireless access points and includes a Memory

Management Unit (MMU) that enables memory protection and full software compatibility with Linux apps. The company claims the chip is capable of running the full 2.4 Linux kernel rather than having to use the uLinux patches other embedded solutions rely on. The software development kit (SDK) is also based on a standard Linux 2.4 kernel.



MULTIMEDIA

Music to your ears

There have been a number of notable advances in the multimedia arena over the last year, with open source sequencing and editing applications spouting features that make Linux a serious contender when it comes to making music. The big problem is that there are so many elements to a really good sound system that you could spend days, or even weeks, configuring ALSA, Jack and countless other utilities and services.

Configure no more, as Kevin Ernste of the Eastman Computer Music Centre (ECMC) has release what he calls a *Turn-Key Linux Audio* solution. This consists of a single script which builds, installs and configures tools to do almost any sound-based task you'd care to try with your computer.

The package, which includes applications such as *Audacity*, *Ardour* and *Ecasound*, is a significant 116MB download and, at the time of writing is specific to Mandrake 9; however, it does exactly as specified and is well worth the small amount of effort.

<http://lulu.esm.rochester.edu/kevine/turnkey/home.html>

Embedded Linux News

● The Embedded Linux Consortium has finally published the **Embedded Linux Platform Specification**. Murry Shohat said they wanted to make sure the document was made as widely available as possible. "It's our plan to make this vendor-neutral document available through our website in a manner consistent with the distribution of version-maintained standards," Shohat said. Until the final version is ready, Shohat says he is willing to email copies in PDF or HTML format (requests should be sent to murry@embedded-linux.org and should contain a full name and company name if appropriate).

The consortium is now working on a Test Policy so that a full testing suite can be assembled later in the year. Shohat said that community involvement in the

process has been vital so far and will continue to be so as the specification and test policy is refined. "The ELC wants more participation from both Open Source developers and vendors at all points in the supply chain. Now is the time to convert potential market energy to kinetic effort!" <http://www.embedded-linux.org/>

● **Dreambox** is a very interesting looking home entertainment product from



Germany. Built around a 250MHz IBM processor, the device is fully DVB compliant and has a gaggle of standard PC ports (ethernet, USB, RS232 etc) meaning it could also be used as a 'net access device, video streamer, MP3 player and webcam broadcaster. There's also space in the box for a standard hard disk, making it ideal for PVR functions. <http://www.dream-multimedia-tv.de>

● Motorola's Metrowerks, which develops the Linux-based *CodeWarrior* IDE, have acquired the assets of **Embeddix inc** (the company formerly known as Lineo). The buyout includes not just bricks and mortar, but also intellectual property such as *Embeddix Plus* for Digital TVs and *Embeddix Plus* for Smart Handheld Devices, which is at the heart of Sharp's Zaurus product line.

NEWSBYTES

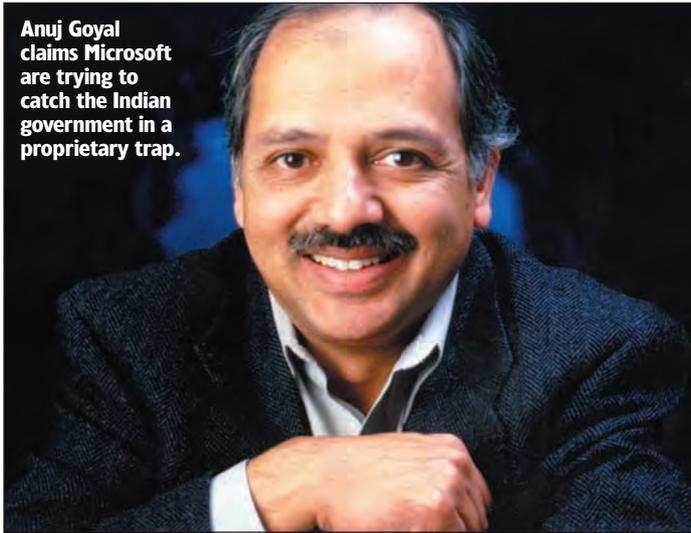
■ A number of vulnerabilities in the *Common Unix Printing System (CUPS)* can, in some cases, allow attackers root access to affected systems. The worst of the problems is a race condition which allows a local attacker to overwrite any file as root. **Easy Software**, the developer of CUPS, says the latest release (1.1.18) resolves these problems. Users who've yet to upgrade (including Mac OS 10.2 users) should ensure printer sharing is not enabled.

■ **MySQL** also has a worrying vulnerability allowing a potential cracker to use unpatched systems to launch a denial of service (DOS) attack. Upgrading to the latest version of the database server will fix this security problem.

Teratext have ported their XML-based text database and search engine to Linux on Itanium servers. Teratext claim the system can handle thousands of concurrent queries.

■ On the grounds of cost, the **Israeli army** are pondering replacing their Windows kit with Linux. Colonel Avi Kochba, who runs the army's IT infrastructure, told Ha'aretz that it was almost inevitable that Linux would find a place on their servers. "Today there is a trend to move over to Open Source code systems. We are examining Linux servers as a replacement for Microsoft servers, and I assume that some of them will be replaced." Whether the Israeli use of an 'un-American' OS will affect the preferential prices at which the Israeli military purchases equipment from the USA remains to be seen.

Anuj Goyal claims Microsoft are trying to catch the Indian government in a proprietary trap.



UPDATE

IBM's Passage to India

The open/proprietary conflict in India has pulled in Big Blue, as IBM executives have countered a recent Microsoft PR push by offering to put India at the heart of the global software industry if it goes down the Open Source route. Amuj Goyal, vice president, IBM software group, told the Indo-Asian News Service that IBM was willing to invest heavily in the enablement centres based on the decision taken by the government. "It should standardise on Open Source

and we will create centres for the software industry to flourish," he said.

Citing examples in Great Britain, Singapore and Germany, Goyal said that the only way customers can retain control of their systems is by adopting open standards and open source. Getting caught in the 'proprietary trap' will only damage the Indian software industry. "India does not have a software business. We want to build the future SAP or a BEA in India by focusing on open standards," he said.

David Cartwright

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



COMMENT

On the road again

“ As an IT consultant who often works hundreds of miles from home, being connected is essential. So, I upgraded my mobile phone and got Orange to turn on GPRS on my account.

The weird thing is, you forget just how handy it is to be able to stay connected to the Internet all the time without having to pay a per-minute online charge (GPRS is metered not on time but on how much traffic you shift up and down the link). So you can leave your laptop in the corner, stick the TV on, and be alerted as soon as email comes in. No more dialling up every hour!

Even nicer is that getting GPRS connected is a Linux-friendly concept. You don't need any posh gadgetry, your laptop should be able to talk to your cellphone via Bluetooth, infrared or even a boring serial cable. As far as the PC is concerned, the GPRS link is just another PPP modem link. There's only one tiny 'Gotcha' – you may need to add a bit to the modem initialisation string to make the phone connect correctly to the cellular network but this is no big deal to add to the dialup configuration script.

As with ADSL and Cable Modem, always-on Internet connections are a godsend. In the case of GPRS you have to be a little wary that, unlike fixed-link variants, you're being charged for everything you transfer. But the point is that having your PC check for email every five or ten minutes only costs you a few bytes-worth of transfer fee and it's just so cool in your hotel room and to have something beep at you when someone calls you on IM or sends you an email.

OFFICE SUITE

Lycoris gets productive with OpenOffice.org

Lycoris is augmenting its *Desktop/LX* distribution with a new office suite based on *OpenOffice.org*. Featuring the usual selection of

applications, each element has been moulded to follow the look of the Lycoris desktop, complete with XP-style cute icons and a default layout

designed to appeal to general users. Lycoris are hoping that a combination of tight integration into the OS, Office file compatibility and good looks will lure potential customers to investigate the US\$50 application suite.

Lycoris have also announced a Game Pak that integrates TransGaming's *WineX* into *Desktop/LX*, offering the ability to play over 150 *DirectX*-based Windows games. The GamePak, expected to cost US\$35, will include five full games. GamePak subscribers will also get discounts on major Linux gaming titles such as the *Kohan* series, as well as the long-awaited eventual ability to play cross-OS titles from Windows. www.lycoris.com

Lycoris's *OpenOffice.org*-based suite adds a user-friendly sheen to an already well-rounded package.



FREE THINKING...

PART 2 With an incredibly diverse range of speakers lined up for 2003's FOSDEM, **Alain Buret** caught up with a few of them in our last issue; here's another three he interviewed on behalf of *Linux Format*.

Interview

RAPHAEL HERTZOG

RAPHAEL HERTZOG – I'm an engineer in computer science. But in my free time I'm more of a 23-year-old Debian developer ...

LXF – What is your “Debian history” ?

RH – That's too much to explain in a single paragraph. :-) Apart from maintaining several packages (including Debian-cd where I'm the main author) I have done several projects which relates to the internal organisation of Debian. I invented the sponsor mechanism, the package tracking system, and I'm regularly involved in quality assurance tasks. This year I even participated to the election of the Debian leader.

On this occasion, I wrote a more detailed text describing my Debian history: www.debian.org/vote/2002/platforms/raphael

LXF – How did Debian-Edu start ?

RH – It started within Logid e. We were in contact with a group called Linux & P edagogie which listed good free software available for educational use. They needed a distribution that could be easily installable by any student or teacher. We decided to contribute that work by developing the initial (French-only) Debian-Education. Later we decided to change that in a Debian subproject to give it a wider audience and to ease to co-operation with other contributors.

LXF – Education is something with no clear limits : who is the target audience for DebianEdu ?

RH – Debian's motto is “The universal operating system”. DebianEdu is in the same path ... basically anyone who needs education-related software since our primary job is to package education-related software.

However some contributors are interested in solution to ease the administration of network of machines in schools ... and we'll happily accept their contribution of course.

LXF – What do you expect of your FOSDEM presentation ?

RH – I hope to make DebianEdu better known, to attract more contributors and to create relationship with other projects interested in free software and education.

LXF – In similar fields, the Debian-Edu project collaborates with the Ofset project. How do the two projects interact, and who decides what to do ?

RH – Hilaire Fernandez (who is behind Ofset) is a friend of mine and he naturally joined the DebianEdu project when I launched it. :-) We interact with the debian-edu@lists.debian.org mailing list which is open and publicly archived. I hope that any of your interested readers pay it a visit.

LXF – There are plenty of education oriented projects : yours, KDE... Do you think it's a good thing to have so many different projects ?

RH – Not all projects have the same objectives... However DebianEdu clearly wants to federate the effort of :

- selecting and packaging the best education related software
- improving them (translation...)
- documenting their use in teaching

LXF – Do you interact with other education projects ?

RH – Yes, I try to have representatives of each other educative project within DebianEdu so that we can effectively share our experiences and our work. DebianEdu is also part of Schoolforge, please visit www.schoolforge.net for our latest updates.

Over the sea

Travelling from the UK

Last year a dozen UK developers shared their transport and accommodation costs to FOSDEM. If you are interested in joining a similar

arrangement this year then go to <http://wiki.vmlinux.org/cgi-bin/moin.cgi/Fosdem2003InterestList> or email fosdem@lug.org.uk

Interview

JON “MADDOG” HALL

LXF – For those readers who haven't encountered your name before, the traditional “please present yourself” question!

JON “MADDOG” HALL – I am the Executive Director of Linux International. I have been in the computer industry since 1969, and have worked as a programmer, systems designer, systems administrator, product manager, technical marketing manager and educator at the college level.

I have authored *Linux for Dummies* as well as many magazine articles and have given talks all over the world.

I have a BS in Commerce and Engineering from Drexel University and a MSCS from RPI of Troy, NY.

LXF – This question is probably often asked, but why are you called “maddog” ?

MADDOG – My students gave me the name. Often my arguments with the Dean of Instruction were “too hot for maddogs and Englishmen” (from a Noel Coward play). The Dean was British, I was the maddog.

LXF – You advocate Linux and related software: what is Linux for you ? How would you explain this software model to a newcomer ?

MADDOG – People contribute their time and programming expertise to this project just like an amateur athlete runs a race or plays a sport. They are not paid for doing it, they just like doing it. After a certain amount of

software is created, this creates a pool of very useful, freely available code that allows people to solve their problems very quickly. Since the users of this code do not have to pay for software licenses for this code, they can use that money to tailor the code that is there EXACTLY to their needs.

LXF – You are a key member of Linux International. What do you consider to be the biggest achievement of this organisation.

MADDOG – In the early days of Linux we helped to publicise it and get people to understand the model. We protected (and continue to protect) the Linux trademark for use by people, and we help businesses understand the licensing behind open source.

Interview

ANN HARRISON

LXF – Please introduce yourself.

ANN HARRISON – The *Firebird* project works by consensus between the administrators so there is no one chief. I'm one of those administrators – my role is resident dinosaur. In 1985, I was the third person to join the company that produced the first versions of *InterBase*, the original source from which *Firebird* was forked. The first person in the company was my husband, Jim Starkey. He implemented version one and designed the next two. So, I know the original architecture of the product pretty well.

When Borland expressed interest in releasing the source to *InterBase* I was already involved through the support lists and worked with Borland to get the source ready to release. In the end, the release was less harmonious than might have been wished and as a result, Mark O'Donohue and a few others created the *Firebird* fork.

Mostly I answer questions like "Why was the limit on indexes set to 64?" and "What happens if we change this?" As more people learn the code, I get more free time.

LXF – *Firebird* came from a commercial company. Can you tell**us how this came about?**

AH – Yes. It started as an independent company, Interbase Software Corporation. That company was sold to Ashton-Tate, in hopes that they could lend some respectability to the product. They ran it as a wholly owned subsidiary, also called InterBase Software Corporation. After acquiring Ashton-Tate, Borland moved InterBase from Massachusetts to California and integrated it into their development organisation. Later they spun it off as a subsidiary also called InterBase Software Corporation, then brought it back into the main development group again, then released the source and created yet another InterBase Software Corporation, then abandoned the Open Source project and brought the developers back into the Borland corporate fold again.

Perhaps you can see why we use a phoenix and *Firebird* as symbols for the project...

LXF – What was the main goal when starting *Firebird*. It can be seen as "another database"...

AH – Different groups had different interests.

When Borland first suggested releasing the source to *InterBase*, they

wanted to take advantage of the popularity of open source companies. Many of the users of *InterBase* were interested in Open Source because they could be sure that the product would continue to exist and could be maintained – by them, if necessary. The code itself interested a group of developers, as did the chance to make their contribution to the database industry. For me, I like the fact that *Firebird* is small, transaction-based, and portable. Unlike the major commercial databases, it is designed to run without the help of a database administrator, which makes it a good choice for many applications that can't support much overhead.

LXF – Is there a commercial project behind *Firebird*, like consultancy or something else ?

AH – Yes and no. IBPhoenix, the company I work for, offers support and consulting for *InterBase* and *Firebird*. There are a number of other individuals and small companies that also offer those services, and more who specialise in consulting. In addition, there are a number of companies that offer add-on products.

However, there is no commercial organisation that controls the code.

LXF – You clearly say you want to have backward compatibility with Borland's Interbase. Why ?

AH – Because we want their customers. Borland has sold millions of copies of *Delphi*, most of which include at least a limited version of *InterBase*. We want to attract those developers towards the Open Source world and its attendant benefits.

LXF – What kind of arguments would you give someone to persuade him to use *Firebird* instead of another (free) database like MySQL or PostgreSQL ?

AH – We try not to argue about which open source database people should use, but to argue that open source databases have the same capabilities as the big-name commercial systems. *MySQL* is tuned to serve as a "read-mostly" web back-end. *PostgreSQL* is very extensible. *Firebird* is available on several platforms, free even for commercial products, low-maintenance, and has been used commercially for more than 15 years.

LXF – What are your expectations from you FOSDEM talk ?

AH – I'd like to raise the visibility of *Firebird* in Europe.

LXF – LI helps lots of developers by promoting them : how do you select these projects ?

MADDOG – We don't "select these projects". We help to promote Linux and Open Source in general.

LXF – What kind of relation does LI have with the Free Software Foundation; with the Open Source Initiative? How does it compare to these two organisations ?

MADDOG – The FSF believes that all software should be free. The OSI believes that all software should be open. LI believes the most important thing is choice and competition between vendors. Therefore we allow non-free and non-open software vendors to be members. We do

encourage them to generate free and open code whenever they can, but we understand their issues.

LXF – What's your opinion: How does Linux (the kernel) and GNU/Linux (the system) compare with proprietary Unix offers?

MADDOG – Proprietary Unix systems can still do some things that Linux can not. However, as even those with only a passing acquaintance with Linux are aware, the system is moving forward at a very fast pace, and I believe that it will catch up in a year or so to the proprietary systems.

By the way, I assume that you are talking about the Linux/GNU/XFree86/Sendmail/ BIND/SANE/BSD system when you say "GNU/Linux"? :-)

LXF – Two years ago, there was lots of money for Linux related projects in the expectation of IPOs and floatations. What do you make of this sudden interest?

MADDOG – Yes, it was a lot of stupid people with a lot of money to invest who did not look at the business as a true investment (ie start a company and grow it). Linux is a long-term investment, not a 'get rich quick' scheme, although some people did make out very well in the stock market madness.

LXF – What are your expectations of your talk at the FOSDEM ?

MADDOG – It will be brilliant, of course. ;-) Since it's a historical talk and a 'fun' talk, yet has real meaning and value, it should go over well.



Jon 'Maddog' Hall – buy him a nice foaming flagon of Chimay!

LXF – What's your favourite Belgian beer ?

MADDOG – I tend to like the Trappist Ales. I would really like to go to a monastery that makes these ales while I am in Belgium. [LXF](#)


Mailserv

Mailserv

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

Great distros

I have to admit I was wrong – only last month (after receiving the Demo-Linux coverdisc) I wrote to say that I thought that bootable CD distros were a waste of time and then you come up with an absolute winner!

The Freeduc distro on the cover of LXF36 is brilliant. As you pointed out, the hardware detection is really good (though my soundcard and modem – both of which work with Linux, but require considerable fiddling to get them to do so, are not configured by the *Knoppix* startup), and is useful on its own as well as making a good rescue disc. It is also good to see the *XFce* desktop getting some exposure, but it's really useful to see some of the programs that you don't normally get with other distros (*Celestia*, *Yacas*, *G'compris* etc) – it'd be nice to see some of them appearing on the coverdisc alone in future.



I have shown a few friends the disc and run it on several x86 machines (from an Athlon 1.2GHz to a Pentium 200MHz) – all with no problems. This is a disc I will be able to use to show off a bit of what GNU/Linux can do. The other two discs are pretty good too. Thank you.

While on the subject of coverdiscs, I fully appreciate your position with regards to putting all libraries required by packages – it just isn't practical. However one thing I would appreciate is that the version number of the package is printed in the listing on the coverdisc pages. It would be handy to be able to see at a glance if it is worth updating a program or

library that you already have installed without having to delve into the CD itself via *Lynx* or *Nautilus* or something.

Glen Tyler, *via email*
We're glad you found the discs useful. On the subject of version numbers – great

suggestion – this is something we have wanted to include for some time, but there is so little space for all the software names that we'd have trouble fitting them in! Redesign of the discs is on our 2003 ToDo list anyway.

No colours please

The anonymous writer of the Jan Letter of the Month comes over as aggressive, abrasive and arrogant.

If this was not their intent, I apologise, but it made me angry. I feel that this person has missed the point of "communication". To communicate effectively, you must use the medium that both yourself and the OTHER person are capable of using.

The writer is lucky to live in a colour world, there are more and more visually impaired and blind people who use computers. For many of these people colour is irrelevant, for others the bad use of colour is more than just a nuisance, it can make the text unreadable.

A Personal Computer is now making life more bearable and easier for many disabled people. For many, plain text is the ONLY format they can handle. I believe that screen readers have difficulty with web pages and HTML emails.

For the visually impaired, plain light backgrounds and plain fonts make life much easier, so usually plain text is best for them as well.

Many people in developing countries have basic computers, they may have basic screens and low hard disk space, they may only be able to use plain text.

So, it seems as if the writer does not want to communicate with anyone who does not have the same equipment as themselves, or

★ Letter of the month

This month's winner receives a copy of **Sun's StarOffice**

Scanners

I purchased SUSE 8.0 some months ago and have been slowly getting to grips with it. I buy your magazine on an irregular basis, primarily because I do not understand 75% of what is written in it. I think I am trying to gain knowledge by osmosis.

Both SUSE and Mandrake 9.0 installed easily from your cover discs, most of my equipment being recognised and configured automatically. But not my Visioneer scanner. So far as I can determine there are no drivers available for any of their scanners.

With Visioneer as an example, I come to the point of my missive. Do any manufacturers, not just of scanners, provide Linux drivers for their hardware or are all drivers obtained by reverse engineering? Is there intrinsically any reason why a manufacturer cannot provide a Linux driver along with the others on the CD supplied? Assuming that the absence of a Linux driver is down to manufacturer indifference, then is it not about time that magazines like your sister magazine *PCPlus* (to which I do subscribe) mentioned the fact in equipment reviews? I suspect that other

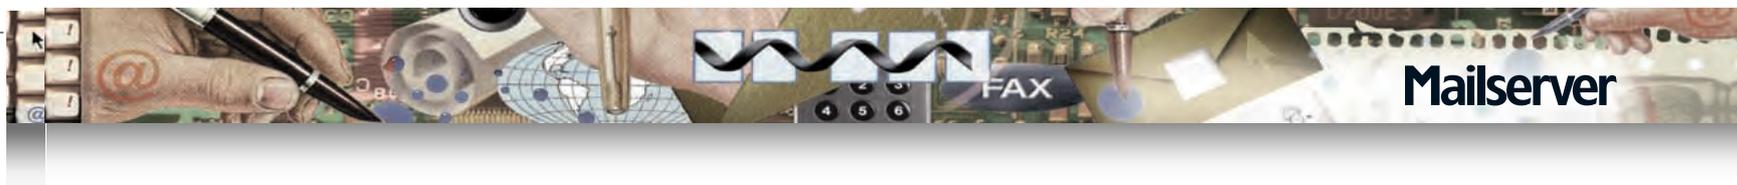
magazines would soon pick up on the initiative. The implied criticism might encourage manufacturers to better support of the Linux OS.

Kenneth Muir-Hunt, *via email*
I don't know of any scanner manufacturers that actively develop their own Linux drivers, although some are more forthcoming with the necessary info for Open Source developers than others. I think it could be a huge advantage for the manufacturers to actively develop a SANE backend for their scanners – if you were looking for any peripheral and knew it supported Linux, that would of course be one of your over-



riding reasons for choosing one brand of hardware over another.

I hope you find you are becoming more familiar with Linux now. To help you get more from Linux, we're giving you this month's star prize of a copy of *OpenOffice*. You didn't supply your address, so please contact us again to claim your prize! The point you make about hardware support is important, perhaps you should write another letter on this to our sister mag, who knows – you might win another prize!



has the same viewpoint as they do. That's considerably narrowing down any audience.

The use of plain text for emails is still advocated, despite ISDN, ADSL, fibreoptic and huge hard disks, because it is the lowest common denominator. ANYONE can use it, regardless of modem link, harddisk space, operating system or world view. It allows two-way communication without snobbery.

I take issue with the writers point that email is about personal expression. They (the writer) may prefer HTML, but what about the person receiving the emails? What do they think of the writer when they get a screenfull of gobbledeook that they just paid money to download. I don't think the recipient will have friendly feelings towards the sender.

So just because some people are able to send emails in colour and html with sound and visuals, doesn't mean they SHOULD send everything out this way.

If you really want to use HTML for your emails, make sure first the person you are sending to can use it and *wants* it. I don't.

To be effective communication has to be simple. Please, lets keep it simple, so everyone can understand and contribute.

Andrew Joseph, via (plain text) email

Black & white Paris

The world must be a difficult place for Paris (LXF36, Letter Of The Month). When he phones a friend,

he is unable to communicate all his thoughts. When he reads a book, he can only grasp half its meaning. (What a pity for us that our great authors were not published in colour. What expression has been lost forever to the world by the failings of technology?) Maybe he tries to conduct all communication face-to-face. That way he can flesh out his limited vocabulary with a system of colour semaphores.

Sadly, the world is restricted in many ways to reliance on the written or spoken word, without embellishment. Somehow we manage to get by. Perhaps, if Paris honed his verbal skills, he would find that he too could express himself adequately in black and white.

Bruno Prior, via email

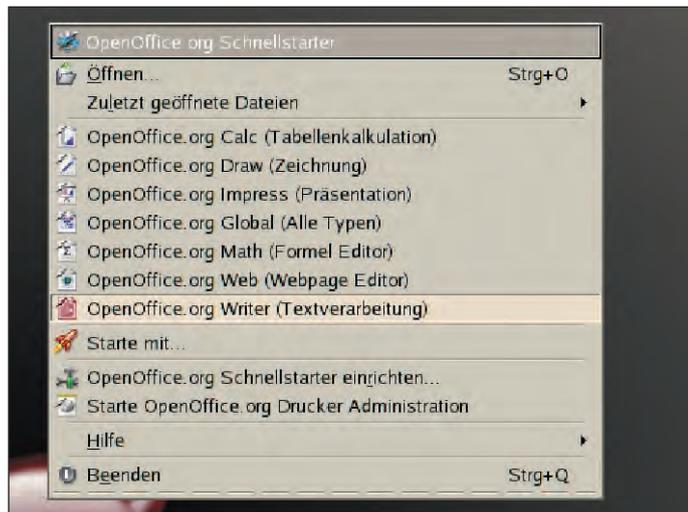
But just think what Chaucer might have created if he was able to imbed tables into The Canterbury Tales. On second thoughts...

Office best?

I have been reading with interest the articles on *OpenOffice.org*. I first used *StarOffice 5.2* on my Windows PC, then once I got *Ooo* on the LXF cover disc this went on to the Linux partition, so I bought a CD of *Ooo*. I have been using *Ooo* ever since.

Also I have taken this disk to work (I am IT manger); this has been installed on 4 PCs, and I am hoping to get the whole office changed over as soon as I can on fourteen PCs and two Suns.

I have found a few funnies as we save the data as Microsoft files,



OpenOffice.org – easier for novices to pick up than *Microsoft Office*.

fonts some times change size and change to bold. I think I know why – if you hit Return after a bold text, the blank lines below are also bold, so I think this gets taken to the next line of text when converted.

Also we had a few problems with pasted images and anchoring them so that they would not move when converted, but we got there in the end. Some of the DXF importing has not worked, but hey, at least it's got it! *MS Office 2000 Pro* has not got this import filter.

to be there to replace other pieces of software or to get the software companies to port their software. But until more people use Linux the software wont get ported so people won't change. I think the best place to start is in schools – as the ad men say, "Catch 'em young".

We also have a big database on Access. Ideally, I would like to get away from this as well. It would be nice to have an article on how to migrate from Access. This will in my case have to be a Windows-based

“For the most part, it is permissible to copy LXF CDs, but not in cases where the included software forbids it. You will have to check individual licences.”

In the end I have had people who have had no experience of *Office* picking up *Ooo* very quickly. I was hoping to get the server changed over to Linux mainly due to problems with *Windows 2000* not supporting the file permissions that are used in Unix (*Catia* high end CAD system). This has been killed due to the fact we now have networked accounts software (Windows-only at present).

Some of the other articles in the LXF news sections have been very interesting. I changed to Linux because I did not like Microsoft having 95% of the market. But as one of your columnists said this will only take you so far. Solutions have

version. Does *StarOffice* have this capability? I know data can be exported as a CSV file but what about reports etc?

Can we make copies of the cover CDs? I ask because I would like to try to push Linux and Open Source with other people.

Peter Cole, via email

There is no direct replacement to Access for Linux, mainly because the availability of free relational databases has reduced much of the demand. You could easily port your data to something like *MySQL* or *PostgreSQL*, and use some of the web-driven GUI front ends. These tend to just list entries rather than having properly customised reports though.



Highly contrasting text and backgrounds are a must for visually impaired users.

Mailserver

« The mag and the CD are complimentary, which is why we go to so much trouble putting things like software from the *HotPicks* section on the discs. You're selling your contacts short if you give them the CD on its own! The more paying readers we

“Let's get rid of the snobbery in the Linux community, and start getting Linux accepted for what it is – a damn good OS that can increase productivity.”

have, the better the discs will be for everyone. For the most part it is permissible to copy the LXF CDs, but not in cases where the included software forbids it. We have featured plenty of software not released under freely distributable licences, such as *Kylix*, *VisualRoute*, various Sun software. So you will have to check the individual licences yourself – if you're in doubt about a particular title, email us at the usual address for a “yay” or “nay” to copying it.

No Java?

Having recently decided to take a fresh look at Linux, I have installed

the Mandrake 9 distro from the LXF35 cover DVD. I then embarked upon a search for a Linux equivalent of my favourite windows editor (*UltraEdit*). I was interested in Neil Bothwick's mention of *JEdit* and decided to give it a go.

I managed to install the package you provided on the cover DVD, but *JEdit* then refused to run because the version of Java in my Mandrake installation is only 1.0.7, and *JEdit* requires 1.3 or later.

I then dove back into the DVD hoping to find the Java 1.3 or 1.4 *JDK* or *JRE*, but guess what? it wasn't there.

Since I don't have broadband yet, and I'm too impatient to wait for the next cover disc, I guess I have a rather lengthy download in prospect. Is there a reason why Mandrake don't supply the latest *JDK*? (Is it a Sun Licensing

issue?) Any chance the latest full *JDK* could be included on a future cover disc, or are there Licensing issues to contend with?

Glenn Knowles, via email

The reason why it isn't in the Mandrake download edition, or on our coverdiscs regularly, is as you guessed, a licensing matter. The Sun *JDK* is not freely distributable. We have included it before, but it took about three months of paperwork to get it sorted just for one issue! We do hope to include it again in the future though. In the meantime we have often included other Java implementations, such as *Kaffe*, which is certainly good enough for running *JEdit*.

MS WIN-tegration

I'm a software engineer and web developer for an internet solutions company, and buy Linux Format (amongst



other Linux / internet magazines) every month. I'm all for the use of Linux in and outside of the workplace, but I'm getting

increasingly annoyed with certain members of the Linux community. It seems to me that the strongest advocates of Linux are becoming more and more like the Windows community – “Our system is best, and don't you dare to use another one”. It seems as though if you want to use Linux, then you must use Linux alone and nothing else – the whole office must be converted to the Linux way of thinking.

In theory, of course, this is a sound idea. I'm sure that many businesses would significantly benefit from a move to Linux – decreased costs and increased productivity being two immediate benefits that spring to mind. In practice however, this is not always feasible. What we need to do is not shout about how everyone should be using Linux, but we should be more positive about integrating

Linux with existing operating systems. Yes, Linux is easy to use once set up, but many of us in small businesses have not got either the time or the resources to train our staff to use the new systems – let alone spend the time converting our accounts data (for example) from Sage into a Linux-friendly format.

Don't get me wrong – I love Linux, and use it 99% of the time

Helpdex

shane_collinge@yahoo.com





Evolution not revolution is needed to integrate with Microsoft Windows.

whilst at work. It lets me do my job easily and efficiently. I use *Apache*, *PHP*, *MySQL* and *Perl* amongst other tools to develop internet solutions and websites – I can test scripts locally as well as whole sites. It's perfect for me. However, the accounts department has been using *Sage* on Windows for a long time, and at the moment has no need to change. The admin department uses Windows too – and it works perfectly for them. The graphic design department all use Macs, and have done for years.

Why should they have to change? The beauty of Linux is how it can integrate all these systems without any disruption to the individual users. The graphic designers can share files with the windows and Linux users seamlessly – a machine running *Mandrake 8.2* with *Netatalk* and *Samba* handles that very effectively. Another Linux machine shares an internet connection with all the users in the office and provides firewall protection using *Smoothwall*. Both of these machines sit happily in the corner and also run *VNC*, so we don't even need monitors, mice or keyboards for them. We can access them through any PC (Linux, Mac or Windows – we have no preference!) if any settings need changing.

The linux PC sat on my desk (*Mandrake 9.0* – thank you very much for the discs!) communicates quite happily with the Windows PC also sat on my desk.

I'm writing software for many different companies, most of which use Windows, so I test software and

scripts on a Windows PC – that's not laziness, or inability to test under Linux, just good business sense. I want to be 100% sure that the software I write works on the platform that the customer uses.

Yes, I could run *Wine*, or something similar to check websites on *IE* or *Netscape* on *Windows 98* (and sometimes do), but I prefer to use the actual system that my paying customer is using.

So come on, Linux Community, let's get rid of this snobbery of "Our system's better than yours" and start getting Linux accepted for what it is – a damn good OS that can increase your productivity with minimal cost or disruption to your existing systems and staff.

Simon Moore, *via email*

Some fair points. I think most of the arguments for abandoning proprietary software are sound, but on a realistic, business level, what you say makes absolute sense as well. Of course, when it comes time to upgrade your accounts system, that might be a different matter!

Great for desktop

At last I think I'm there, it's taken me three distros, and 6 months, of dabbling and frustration, but finally I've migrated from newbie to nerd. I'm now running *SuSE 8.1*, on a dual boot, dual hard drive, with *Windows 98*, and it's great!!

SuSE lives up to expectations. A lot of your readers have been discussing whether Linux is really suitable for proper desktop use, I can now assure readers that this distro is the best ever! It's a truly multimedia operating system, the

manual, and installation instructions are very easy to follow.

I now have automated support for my digital camera, just plug into the USB port, and click on the desktop icon. For the video clips, *Mplayer* is superb, I even managed to view the *Lord Of The Rings* preview CDROM given away with a certain national newspaper.

Niggles – I had just two, my DVD drive wouldn't configure properly, so I disconnected the power lead from my CD writer drive, and reinstalled with just the DVD drive powered up. After installation I shut down, and reconnected the CD writer, reboot, and install the "new" hardware. The other niggles – the *OpenOffice.org* template wizards kept getting mixed up in the default folder with my newly created templates, and corrupted the program. So I installed *Ooo* locally in my Home directory – sorted.

Many thanks for a really great magazine, the articles on *OpenOffice.org* are very useful, I've now got lots of new fonts to use. One other *Ooo* tip; if you want a logo design, or a letter heading, use *Open Office draw* to design it then use *File> Export*, this creates a *JPEG* of the whole page, use your favourite graphics editor to crop, and resize as required.

Derek Smith, *via email*

The DVD problem sounds a little odd, but I'm glad you solved it. It may have been easier to save your own templates in a different folder, but installing *Ooo* in your home directory is fine if only you want to use it! Why not log on to our forum to share your tips with our other readers? [LXF](#)

SUBMISSION ADVICE

WHAT WE WANT:

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

WHAT WE DON'T WANT:

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

WRITE TO US AT:

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



Developers Solution Provider

Reseller Package

Ideal for Web Designers, Web Developers and Entrepreneurs

5GB Fast UK Web Space,
Each VS £20 p.a.

10 Virtual Servers included.
Expert telephone & email technical support, regular tape backups, 24/7 server monitoring, same day setup.

Assigned Technical Manager
99.7% uptime guarantee

Each Virtual Server includes:

Supreme Control Panel,
Unlimited Mail Boxes,
Announcement Lists & autoresponders, secure server,
3GB Bandwidth per month,
Graphical Statistics, Real Audio & Video, Password Protected Directories, Spam Filtering, 24/7 FTP & SSH Telnet access.

Powered by Zeus Web Server

Advanced features include:
CGI, ASP, PHP4, JSP, Java Servlets, ODBC, MySQL Databases.

£45 per month £450 p.a

Hosting your web site with CWCS lets you benefit from our state of the art data centre, speedy network, expert technical support and friendly customer service.

Our Services are your Business Solution

tel 0870 703 1000

info@cwcs.co.uk

www.cwcs.co.uk



prices exclude VAT.

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.



WHAT'S NEW...

Arkeia

For many users, backing up is a tiresome chore. *Arkeia's* redesign hopes to change the face of saving your data **p20**

Lindows

A whiff of marketecture in the Open Source world? Whatever you think, listen up as we test the first distro not aimed at 'insiders' and OEMs **p22**

Jool Kwartz >>

Its selling point is the highly capable *RADMIN*, but is this a proper small footprint server or merely a desktop box with ideas above its station? **p26**

Unreal Tournament 2003 >>

It installs straight from the box! Strangely though, the box it comes in doesn't boast a single mention of the Linux version. How very odd... **p30**

Books

Mastering PHP 4.1
Official Red Hat publications, including their first home-grown: Linux Admin Guide Users Guide **p34**



COMING UP SOON...

United Linux

Released at last, the enterprise Linux distribution that many of the big names are pinning their hopes on. We'll be taking a thorough look at this promising new development and assess its chances in a market already bursting with distros.

FilmGimp

With the development version already picking up screen credits, this fork of *The GIMP* looks better and better as time goes on.

Gnucash

Handling your finances may never have been easier.

Kaptial 1.1

The Kompany's personal finance software attempts to redraw the Linux application map.

Opera 7

The Windows version is in beta, the Linux variant of the third most popular browser is coming.

BACKUP

Arkeia 5.0.8

With a radical redesign, Arkeia hopes to change the face of backup applications. **Nick Veitch** saves everything...

Whatever your choice of backup device for larger networks, its revised GUI makes this Arkeia a winner.

■ **PRICE** From £350, free version available

■ **PUBLISHER** knox software

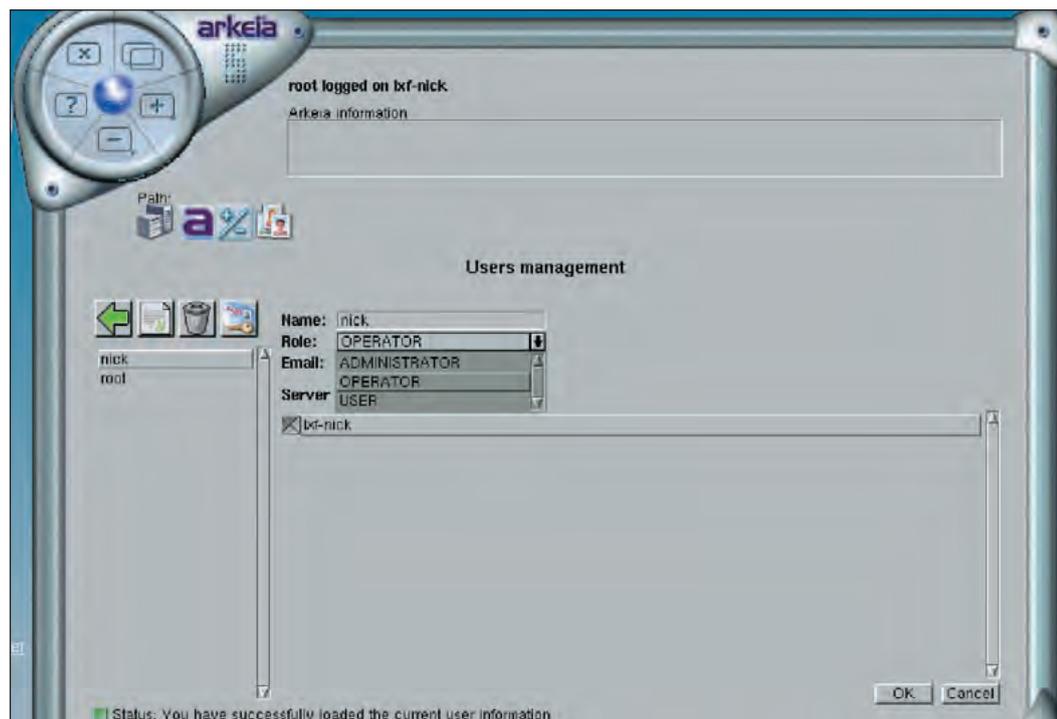
■ **WEB** www.arkeia.com

Arkeia has been one of the big names in Linux backup for some time now. Although the software has proved reliable in previous iterations, it also had a less than pleasant GUI and it wasn't terribly easy to use. The command line tools are still favoured by many, but the complexity of configuring and managing backup tasks only increases, and it is one area where a clear and easy-to-understand GUI can be a great benefit.

Arkeia 5 is the result of a complete redesign and rewrite of the software. While new features have been added, a lot of the emphasis has been on creating a system that is easy to use. The install process for Arkeia 5 is even more straightforward than previous versions. There are packages for Debian, Mandrake, Slackware, Red Hat, SuSe and now United Linux. These are mostly as rpms or binaries with install scripts. Installation takes a matter of minutes. As the full version of Arkeia is available on a free 30-day download,

Plugins

Another great new feature of Arkeia is its new plugin capability. These are specially designed to make it easier to backup specific types of data. At the moment, the plugins concentrate on databases, which makes sense. IBM's DB2, Oracle and MySQL are all supported, and this system allows live backups of the data while in use. The plugins have been developed with co-operation from the software vendors, and provide a useful and reliable way of making sure your databases are reliably archived.



The administration interface allows you to add users of different levels to each server

you can test that it works on your system before purchase too.

Drive Support

In order to use Arkeia, you will obviously need a backup device. For reasons of reliability and throughput, Arkeia only supports SCSI devices at present. There are IDE, USB and even Firewire tapestreamers available, but they aren't supported by Arkeia at all.

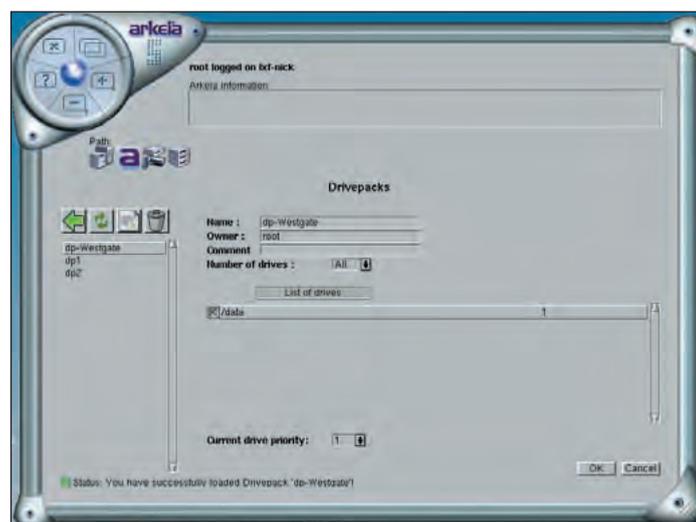
There is a full list of the 40 or so specific drives supported on the www.arkeia.com website. These are the specifically supported drives, and you may find that if you have a similar mechanism, or a rebadged drive, it is also supported. The usual suspects of tape backup are there, including the various shades of Quantum DLTs, Sony AIT drives, the Ecrix/Exabyte VXA mechanisms, various Tandberg and IBM drives and the Ultrium. Notably, the Onstream ADR range of drives are not supported, due to reliability issues, we were informed by Arkeia.

However, some help is at hand to those confused by all the options. The

Hardware detection system can now identify various types of drives and robots and set them up for you. This is a new and welcome feature addition to this version of the software.

As well as single drives of course, Arkeia also supports the use of 'robots' or automated tape libraries. There is a

huge range of these available on the market, but Arkeia does a creditable job of supporting them (again, you can check the website to see if a specific model/controller is supported, there's few that aren't). If needed, you can even share a library between multiple Arkeia servers if required.



It's necessary to create drivepacks, even if they only contain a single drive.

In some ways, *Arkeia* isn't really suited to smaller scale backups, at least in terms of the complexity of setup. Even if you only have a single backup drive and a couple of clients, you will be forced to create drivepacks, tapepools and other hierarchies, which are a little confusing. For the more extensive backup setup though, such features are de rigeur, so their inclusion is necessary.

Scheduling

In previous versions of *Arkeia*, in common with a lot of backup software, scheduled backups were handled by the standard Linux cron daemon. While this system worked, it was sometimes more difficult to manage. The new version of *Arkeia* comes with its own scheduling daemon which can initiate periodic backups independently of cron. The administration for scheduled backups is included in the main *Arkeia* interface, and is at least as easy to understand as a cron job!

As well as actually setting the repeat times (with the usual options of specifying a particular time, and whether to repeat daily, weekly, monthly or yearly) you can also set exceptions, to prevent backups being performed on particular dates. It is also possible to have the periodic backup expire on a certain date too, adding a little to the flexibility of the system. With a lot of periodic backups, it may be difficult to determine what is supposed to be happening. *Arkeia* solves this with a neat calendar interface. Click on a day and it will show you the list of all the jobs due to be executed. One shortfall of this system is that it isn't easily possible to

create multiple backups on the same day, without having to start a new process to do so.

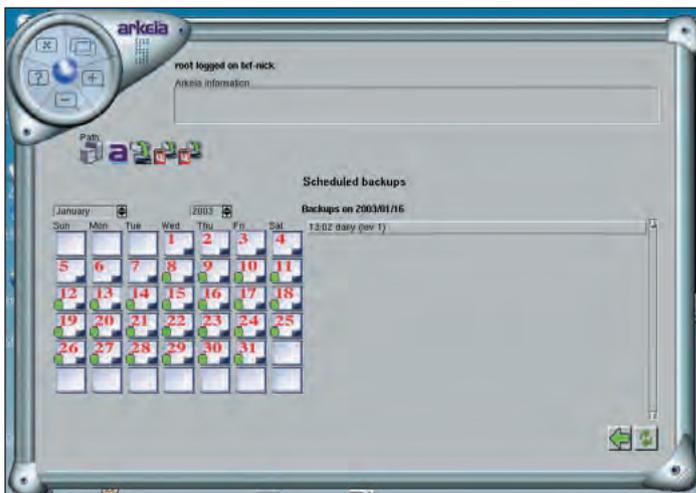
Logging

The X client for *Arkeia* also has an interface for viewing the log files. These are split up into logs for Backups, Drives, Tapes and Restore functions, as well as 'Log Consult' to allow you to filter out the messages.

The only problem is that some of the error messages and log entries are less than explanatory. A 'connexion error' sounds unpleasant, but is it an error connecting to the client/server, or the drive? Although there is the option to email messages to a designated account, the onscreen messages aren't that great and often disappear before you have had a chance to register what they were, necessitating more visits to the logging interface.

GUI

The GUI is radically different from previous versions, with an irregular shaped X window containing a control panel and various icon-driven screens. While it certainly looks 21st century, the lack of a standard menu bar and familiar paradigms may make it a little harder to get used to. Context sensitive pop up menus operate on right-mouse button clicks which makes the navigation easier, and a kind of hierarchical trail of icons lets you know whereabouts in the interface you are. The first screen lists the known servers – you will have to log on to one of these to proceed. By default a root account is set up, but you can add further users to each server with specific levels of access.



The calendar shows what days backups are due to be run in a simple, straightforward and easy to understand way.

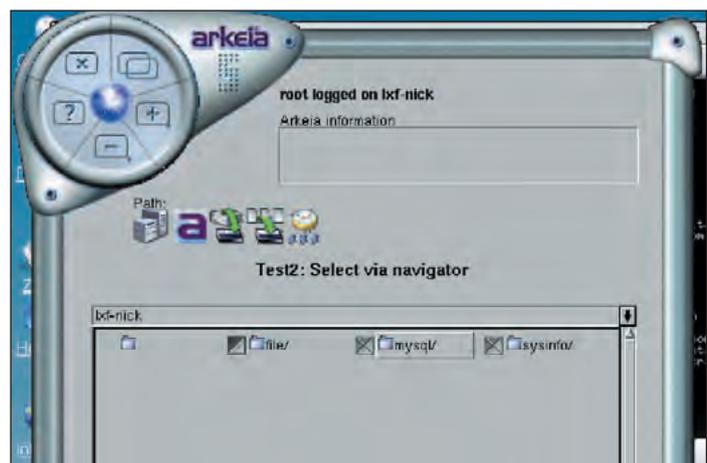
Disaster recovery

Quickly recreate your data on other hardware

A companion plug-in to *Arkeia 5* is the new Disaster Recovery software. This essentially comprises a bootable CD that can be used to recover a complete backup of a machine from the *Arkeia* server. The CD will completely rebuild the machine that has suffered the setback, including formatting and partitioning drives. The only information the user needs to provide is the IP address of the backup server, and the IP address and network details for the

client machine being restored. This provides an extra level of recovery.

Most backup solutions are fine, so long as the original client machine is usable, but in some cases, like complete drive failure, or the physical theft or destruction of the machine, restoring data could take a little longer. *Arkeia's* disaster recovery is able to recreate the data on the client machine completely, which could even be on different physical hardware if necessary.



The database plugins cause relevant folders to appear at the root level of the tree view, so you can navigate to and select individual databases.

Further sublevels give access to the creation of Tapes, drives, and savepacks as well as the configuration of periodic backups. Context-sensitive help is available by clicking on the '?' in the circular panel, but the manual is really necessary to understand the more detailed screens. You can spawn multiple views on the system, which can be extremely handy if you want to monitor a backup in progress whilst performing other administration tasks elsewhere in the system.

Overall, the GUI is a tremendous improvement on the previous versions, though there is probably more that could be done to make it really easy to understand. All of the functionality is also available through the various command line tools, which is also useful if you need to further automate or log specific tasks.

Licensing

As already stated, there is a free trial version of *Arkeia*, but there is also a completely free version which is restricted to one server and two clients, which may be suitable for

many users. This does exclude the possibility of using some of the extra features for which specific licences are required (eg tape libraries). This is one aspect of the licensing which is less than transparent – some functions such as data encryption, the use of specific hardware etc require separate licences – which is something all users should be aware of.

Arkeia 5 is certainly a big step forward from previous versions, and is bound to stay very popular for Linux users. It isn't perfect, but at least it is being very actively developed. **LXF**

VERDICT

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

A reliable performer which should handle the most complex backup demands.

LINUX FORMAT RATING
8/10

Everything about LindowsOS appears to be designed to make Windows users comfortable, from the menu layout to the icon placement on the desktop.

OS OF THE FUTURE?

LindowsOS 3

After two limited releases, aimed at 'Insiders' and OEMs, Michael Roberts' latest project finally gets a public airing. **Andy Chappelle** pushes through the crowd of muttering cynics to peer beneath the spin and see if its worth your time and money.

Whatever its faults, should Lindows be lauded simply for encouraging more users to switch to Linux? Also try Xandros.

- **DEVELOPER** Lindows
- **WEB** www.lindows.com
- **PRICE** US\$119 (\$129 for CD edition)

Maybe it's impossible to come to a review of a product that has generated so much publicity (from legal battles with Microsoft to an sadly non-existent tie up with AOL) without a number of preconceptions. In this case, Lindows is the GPL-baiting 'OS of the future' that had more than a slight whiff of dotcom marketecture from a company more adept at generating headlines than great software. It was also the product that would make us all pay – via "Click-N-Run" – for the applications we would normally expect to be included in a standard distro; and promised to

seamlessly integrate Windows applications with the Linux desktop.

Quietly dropping the "Windows compatible" tag wasn't a handicap – Lindows.com (the company behind the product) has had notable successes in convincing retailers of bargain price PCs to begin shipping LindowsOS (the operating system) as a rival to the established market leader. Wal-Mart in the US and Evesham in the UK have products based on the OS, but version 3 is the first available 'off the shelf'.

What is it?

LindowsOS is a Debian-based distro built around a 2.4.x kernel, Xfree86 and KDE. Stripped to its essentials, the download version (tested here) is a mere 350MB – a small download on a cable modem – and installed in just under 10 minutes on an entry-level PIII desktop PC. Such a swift installation seems, and is, remarkable, especially when the user intervention involves little more than a couple of clicks and the entry of one password. There is, however, a reason for this which we'll come to in a moment.

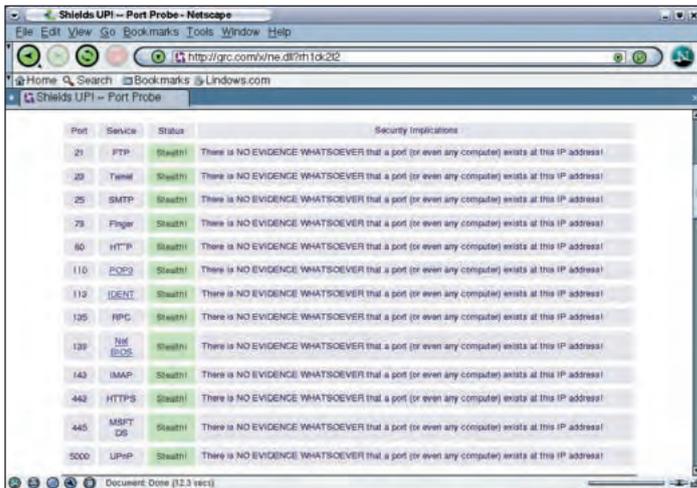
For your US\$119 (\$129 for the CD edition) you get the distro, web-based technical support and a year's membership to the Lindows Click-N-Run warehouse, a software repository containing over 1,700 items including a number of non-free applications such as *StarOffice*, *Photogenics* and the non-GPL *Tuxracer*. At the end of your year you'll need to cough up another (currently) \$99 to remain within the Click-N-Run program, though generously you have "a lifetime license to run the software already installed."

Right, back to the marvel that is the Lindows installation. It really is trivial, simply asking where you want to put the OS – there are no partitioning options, you can install on an existing partition or take over the whole disk – and then requesting a system password. Remove the CD, reboot and you're running Linux and a tightly branded version of KDE 3.0 with a look that (if you can believe it) actually improves on the already impressive *Keramik/Crystal* combination. It also kicks off with lovely alpha-blended menus which, though almost

completely pointless, do make me smile whenever I use them. The boot sequence provides only a few lines of feedback which, considering the target audience, is fine.

One of the big controversies about Lindows was that you would be forced to run as root and, as this is designed to be a 'connected' OS would lead to the world and his dog having access to your precious system. While this is still the case, the first dialog that appears on screen after installation gives you the option of adding users. This has the air of a compromise, and an explanation of the dangers of running as root would be a welcome addition to this dialog, but it's not the disaster I was expecting. You can also add users in the normal fashion through the K- or rather L-panel. Another option on the opening dialog box also allows you to change resolution, but as Lindows correctly identified and configured my graphics card to its realistic limits I didn't need to venture there.

The most notable feature of the installation is that, like a similar Windows install, you get the OS and



Shields Up reports that LindowsOS is running in Stealth Mode, and is 'invisible' to the rest of the Internet.

that's about it. Browsing and email is taken care of via Netscape, word processing (very basic) with *Kwrite*, *XMMS* for MP3 playback and there are a few other games and utilities, but no office suite, graphics package etc. After the options in Mandrake 9, this can be a bit of a disappointment, but again this is aimed at the mainstream user weaned on Windows, so a package selection section on the installation would probably be seen by some users as unnecessary clutter. And anyway, we've got access to all those Click-N-Run apps haven't we?

Well, er no. Not on the first system I decided to check anyway. The problem was that Lindows failed to recognise my Netgear network card (the first distro to have such a problem) due, I finally found out, to it being built on a NatSemi chipset. This is not just disappointing, it's a serious problem because there is no graphical 'Add New Hardware' section within the excellently constructed Lindows menu system and the only support option requires a working network connection. Of course this would have been noticed in the more verbose installation systems that Lindows eschews. To most users, who are likely to end up with LindowsOS pre-installed on a new machine, this won't be a problem, but for me it was a disaster. Time to break out a new machine – at least we know that the installation will be simple!

Round two

Next challenger for LindowsOS is a Toshiba Satellite 3000X4 laptop, untouched but for a little additional memory and pre-partitioned disk.

Again, the installation went smoothly and after rebooting I clicked the Netscape icon with not a little trepidation... The next moment I was directed to the *Lindows.com* website, a quick check to make sure it wasn't a locally cached page, and I'm away. Click-N-Run, here I come!

The menu structure of Lindows is great, even if it is sparsely populated. Applications are grouped under logical headings such as Audio and MP3, Email and web browsing, like a more refined version of Mandrake's Kpanel. After listing the installed application entries there's a 'Click-N-Run More' option which, when clicked, fires up the Click-N-Run application (a *Mozilla*-

Steamy Lindows

Interaction rather than compatibility

Lindows received a lot of press at its inception for claiming it would allow Linux users to run Windows apps straight from the box. In the first two versions of LindowsOS, this ambition was scaled back somewhat and there was a pretty stable *Wine* release which, indeed, allowed limited running of applications designed for Win32. The situation has changed in version 3.0 with the company, perhaps sensing an impending rush of support queries on potentially unsolvable issues, abandoning any pretence of Windows compatibility altogether and slashed *Wine* from the main release. There are instead promoting software that allows users to interact with their Windows-using colleagues – whether over a network (via *Samba*) or using interchangeable file formats, ie by using *StarOffice* instead of *MS Office*. *Wine* is,

however, still available via Click-N-Run and runs a few applications well while simply ignoring others, as with a standard *Wine* install.

The other big Lindows story of recent months was there 'tie up' with Netscape/AOL which led some to suggest the media giant was on the verge of embracing Linux for its main Internet access client. This was revealed to be an embellishment of high order. Lindows.com had, in fact, signed a deal to include Netscape's branded version of *Mozilla* (which does include AOL's messenger client for Linux) in their first general release.

Both these tall-ish tales demonstrate that Lindows' developers know the benefits of decent marketing which, if many industry pundits are to be believed, is where Linux has been going wrong all these years!

based affair) in a relevant section of the warehouse. From that screen you just click an application and it's installed. And that's it – no dependencies, no conflicts, no missing libraries. It's eerily similar to *apt-get*; in fact, it is *apt-get*, it just has a brilliant front end tacked on.

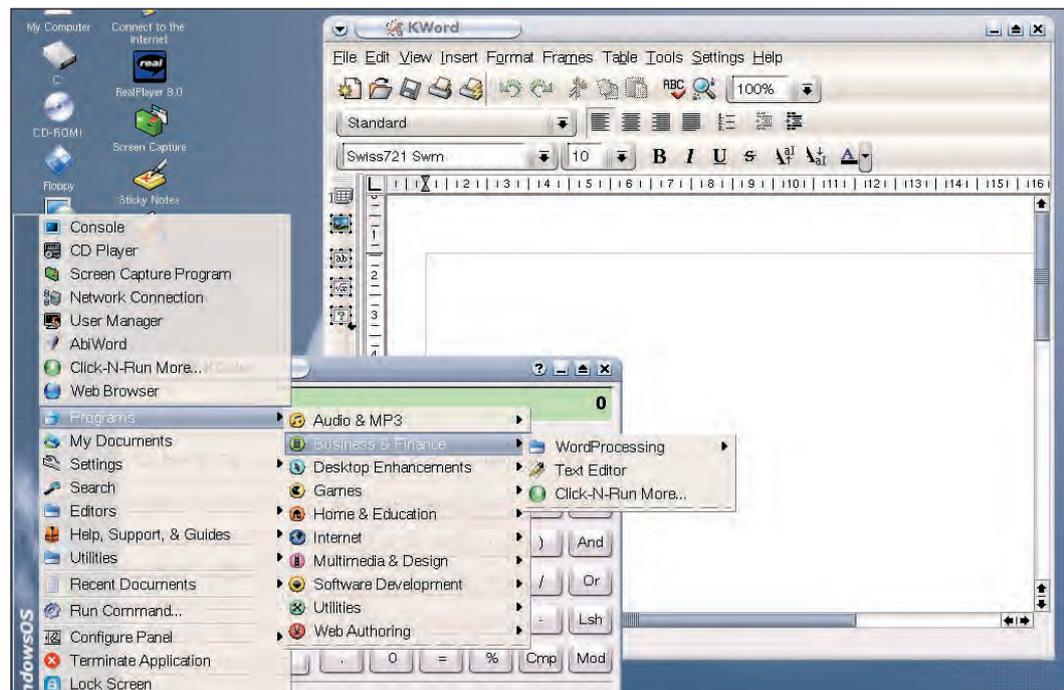
On a whim, I opened up a console and typed **apt-get update**. Nothing. Now Click-N-Run is definitely based around *apt* so it had to be installed. I poked around a bit and discovered

that the apt hosts list was incomplete (or broken), so after adding a few lines to the list I was able to do:

```
apt-get update
```

```
apt-get install synaptic synaptic
```

and I had a rival to Click-N-Run. Sure it's not quite as elegant or integrated and it doesn't have the likes of *StarOffice* available, but it is free and the range of applications available dwarfs Click-N-Run enormously. You can see why the Lindows team might



Desktop defaults to an elegant Keramik/Crystal combination and transparent menus. God bless you KDE coders!

MINI-CUBE SERVER

NetServe Kwartz

Need a compact server to run a small office without the hassle of installing distros etc? **Richard Drummond** test drives a compact little number.

Small footprint server box that is easy to manage. See also Cubit; or RaQ, Qube, manufactured by Cobalt.

- **SUPPLIER** Jool
- **WEB** <http://www.jool.biz/>
- **TEL** 020 8457 9860
- **PRICE** £799

Cobalt pioneered the server appliance with their RaQ, a device which you could plug into your network and get up and serving your website, mail or whatever with a minimum of knowledge or fuss. They took this idea to new levels with the Qube, a tiny cube-shaped box which did the same but was small enough to fit into any office. While it looks cool and is easy to use, the problem with Cobalt hardware is that you don't get much bang for your buck. The solution: the Jool Netserv Kwartz.

Like the Qube, the Kwartz is a Linux-based server housed in a small form factor case and managed with a web-based interface. But, whereas a top-of-the-line Qube will set you back £1400 and offers a 400 MHz processor, with the Kwartz you get a comparable system with 1 GHz of processing power for only a ton more than half the price of the Qube.

Clocking the Kwartz

Afficionados of the Qube will also appreciate the aesthetics of the Kwartz. It's housed in a sleek, aluminium chassis with a stylish plexiglass front. In many ways it bears more than a passing resemblance to the Cubit device that we reviewed last issue, and it too offers an effective antidote for those beige box blues.

Unlike the Cubit's case, which is custom machined, the Kwartz's is an off-the-shelf part. It loses the bombproof robustness of the Cubit, but is a whole lot lighter. Also, consequently, it provides a much more

conventional layout, with the motherboard being sited on the floor of the chassis; hence all the other components are much more accessible. It boasts more storage potential, too, with three standard drive bays: one external 5.25inch bay, populated with a CDROM drive; one external 3.5inch bay, populated by a floppy drive; and one internal 3.5inch bay which houses a 40 GB IBM drive.

Usefully, this bay is located at the rear of the unit and is mounted on sled fastened by thumbscrews. Removing the drive is a simple case of removing the lid (also fastened with thumbscrews) disconnecting the drive cables, undoing the screws and sliding out the sled. Also, the case offers two (of the four) USB ports and audio ports on the front of the unit – great for accessibility.

The Kwartz's motherboard is also very similar to the Cubit's mini-ITX board. The Kwartz however has a

slightly larger footprint Flex-ATX board, the Shuttle SpaceWalker FV25. The extra real estate is put to good use by offering a Socket-370 ZIF processor socket rather than the surface-mounted C3 processor in the Cubit.

This is populated by a 1GHz Celeron unit in the Kwartz, which has a slight edge in performance (putting out 1930bogomips compared to the 800MHz C3's 1600) and has the advantage of being upgradeable. The downside is that it requires active cooling, and the review unit was fitted with a particularly loud fan which, coupled with the noisy PSU fun, was quite a contrast to the Cubit's virtually silent running. Other extras on the

Kwartz include dual Firewire ports and a floppy drive header.

The Cubit and the Kwartz share similar VIA chipsets. The Cubit has the PLE133 chipset, which boasts on-board Trident Cyberblade/i1 graphics, while the Kwartz has the PN133T chipset (Twister T) with on-board Savage4 graphics core. This offers much better performance, although, like the Cyberblade, 3D acceleration is unsupported under Linux. TV-Out is provided on the motherboard, but with the Savage4, TV-Out support is better under Linux, since the XFree driver can use the BIOS for choosing TV-compatible screen modes. Like the Cyberblade/i1 the Twister uses a unified memory architecture and steals main memory for the graphics



framebuffer, but the Savage device supports up to 32MB.

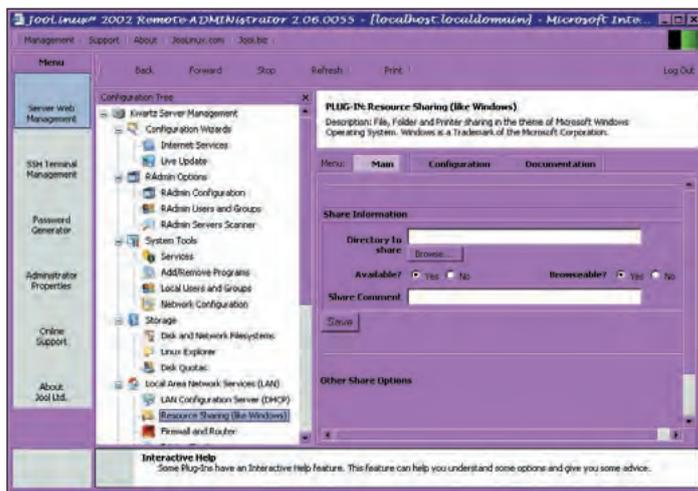
Of course, for server tasks, graphics performance is irrelevant. Much more important is memory and disk bandwidth. Like the Cubit, the Kwartz uses SDRAM DIMMs, but the Kwartz ships with a (for a server) stingy 256MB included as standard – but up to 2GB is supported. Disk bandwidth is impressive with the same UDMA100 VIA IDE core, and, according to hdparm, it manages 45MB/s for raw disk reads.

The Kwartz motherboard has one PCI slot, which, thanks to the PCI riser, supports a full-size PCI card mounted horizontally. In our review machine, this had a Realtek 8139-based NIC to supplement the on-motherboard Ethernet adaptor.

Remote admin

What makes the Kwartz a server isn't its hardware, it's the software that's pre-installed. It uses a customised version of Red Hat 8.0, called *Joolinux*. This comes preloaded with a stack of server applications, including *Apache 2*, *MySQL*, *Postgres*, *Squid*, *Samba*, and more. Although the Kwartz in this review can be administered like any other Unix box – either locally or remotely through the shell – its selling point is its *RADMIN* web-based remote administration system.

RADMIN is a lot like the well-known *Webmin* system, although *RADMIN* interface is perhaps better organised and so easier to navigate. A problem with *RADMIN*, however, is it only works with *Internet Explorer 6*, due to the component technology used to build the interface. It seems like a bit of a restriction in any case.

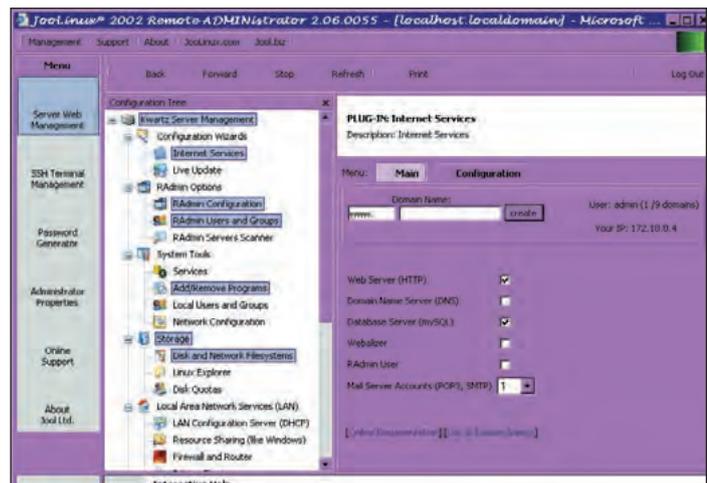


Adding a Samba share is straightforward with the *RADMIN* system.

Another problem I found was that many icon images failed to be displayed properly in *RADMIN*, but I don't know whether this was *IE*'s problem or *RADMIN*'s since I was unable to test it on another browser.

Once you have logged into *RADMIN*, you're presented with a main menu down the left hand-side of the screen. The default option is Server Management, which displays a tree of configuration modules next to the menu, and the currently selected module's dialog is placed in the main portion of the window on the right. Other options from the main menu launch the Java-based *MindTerm* SSH applet and a handy password generator is also included.

The tree of configuration modules is split into six classes: Configuration Wizards, *RADMIN* Options, System Tools, Storage, LAN services and WAN services. The first of these offers an Internet Services wizard that makes light work of adding a virtual domain to your Kwartz box and setting up services for that domain. *RADMIN* Options lets you configure *RADMIN* itself and administer *RAdmin* servers. A cool feature here is the Server Scanner, which will probe your network for other boxes with *RADMIN* installed and so let you configure them all from within a single interface. The System Tools group includes dialogs to control which services are run, to perform user account administration and to configure your network settings. Storage includes tools to manage your *fstab* and set up quotas. It also provides a Java-based, graphical file manager. The LAN services class of modules lets you configure services, which are available on the LAN side of



This wizard lets you set up virtual domains, and so get up and web-serving quickly.

your Kwartz server. This includes Mini-DHCP, file sharing with *Samba*, firewall settings, printer sharing, and the web proxy *Squid*. The WAN section, then, is obviously for configuring services presented to the Internet, such as HTTP (with *Apache*), DNS, SMTP, PPP (for dial-up access), and so on.

On the whole the *RADMIN* system works well, but is not as easy to use as it might be. It's undoubtedly a lot simpler than editing a host of config files manually, but I don't think it really lives up to the server appliance role. Like *Webmin*, *RADMIN* still requires the user to know the jargon and the services they are configuring. A useful 70-page printed manual is included with the Kwartz and most configuration pages in the interface have online documentation, but usually this consists simply of links to the relevant man pages, etc. Much more could be made of the system for new users by adding task-based wizards. The Internet Services wizard for setting up domains is a good start, but more should be added to lead you through tasks such as setting up *Samba* or *sendmail*.

At your service?

Besides the *RADMIN* system, another neat tool included with the Kwartz is *jorec*, an automatic back-up system. What this does is to make a complete backup of your *filesystem* and write it to a series of CDRs along with a bootable image of Linux (requires a CD writer). If your system goes down, you can boot up from this backup and completely restore your system.

Hardware-wise, the Kwartz is an impressive little device. It offers all-

round better performance than Mini-ITX units such as the Cubit, with the added plus of an upgradeable processor, and it trounces the Cobalt Qube. But, while the Kwartz would make an interesting and capable desktop machine, the hardware just doesn't excite in the server realm. It will quite happily take on the server requirements of a small office, though.

The Kwartz comes preloaded with all the server applications most of us will ever need and the *RADMIN* system makes the job of administration more manageable, but the Kwartz falls short of its aim of being a server appliance. It is great for those that need to install a server, but don't want the hassle of installing a distro and then all the services they need.

The price of the Kwartz, at a whisper under £800, looks expensive when you tot up the prices of the components used in the system, but consider you're getting the remote administration system and free support, too. When measured up against the Cobalt Qube though, it looks like a bargain. [LXF](#)

VERDICT

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

Despite hardware more at home on the desktop, the Kwartz is a capable little server. The remote admin tools and its cool looks further sweeten the deal.

LINUX FORMAT RATING
8/10

Reviews **Unreal Tournament**

Is that *The Predator's* lawyers on the phone? Tell them I'm not in!

SHOOT-'EM-UP

Unreal Tournament 2003

System requirements

x86 processor (750MHZ)
128MB RAM (minimum)
3GB HD
Accelerated 3D card (Gforce 2+/ ATI 8500+)

The sport of the future is here, and it involves quite a lot of guns, apparently. **Nick Veitch** is press-ganged into trying it out.



3D first person shooter, with a few extra 'sporty' modes. Also consider Return to Castle Wolfenstein.

- **DEVELOPERS** Epic/Atari
- **PRICE** £24.99
- **WEB** www.unrealtournament2003.com

One of the best features of *Unreal Tournament 2003* is the background story – it's so totally dispensable you don't need to give it a second thought. There is no great plan, underlying subtext or goal to work toward, you just have to shoot people, and occasionally pick up things.

There are five playing modes to the game. Deathmatch is a straightforward frag-fest. Pick up weapons, shields, health, and then stalk the arena gunning down your opponents when they least expect it. The familiar range of powerups and actions are available, including Crouch and Prone to improve your accuracy. Acrobatics are the order of the day, with unfeasible sideways leaps around corners being a sure way of confounding your enemies.

Team Deathmatch, naturally, involves working with a team against a team of opponents, but is otherwise much the same. Friendly fire is an option that can be turned on, and team strategy can be communicated with the handy voice menu, which also includes taunts to the opposition.

Capture The Flag requires a little more teamwork – you have to defend your own base while attempting to capture the opponent's flag. Most of these levels are symmetrical, so there is no particular advantage to either team (though it is a lot easier to get lost), but some are subtly different. Most of these levels include at least two main routes to the opposing base, one of which often requires a bit more skill and agility than the other.

In single-player mode, you can specify actions for your AI teammates (Offence, Defence, Support or Auto) which can significantly alter the outcome of the match, depending on your individual tactics.

Domination is a capture and hold game. This has been updated since

earlier versions of *Unreal* and now requires the capture of two control points. The points must be held for about 10 seconds to score a point.

Finally, Bombing Run actually does have some element of sport in it. Teams compete to deliver the ball to the opposing goal. The twist here is that while you have the ball, you can't use your weapon (the fire button is used to pass or shoot the ball instead), and thus become a sitting duck. Well, perhaps more like a running around and jumping about duck. This last mode is the most different to the standard Deathmatch, in that it does really require co-ordinated teamwork and strategy. Strangely enough, you need to score goals to win, which means protecting the ball carrier and, er, distracting the opposition.

The weapons themselves haven't changed much from earlier versions, though the different rate of fire, blast



Firing many of the weapons will give away your position if you're hidden.

area and range combinations seem to have been adjusted.

More Big Guns

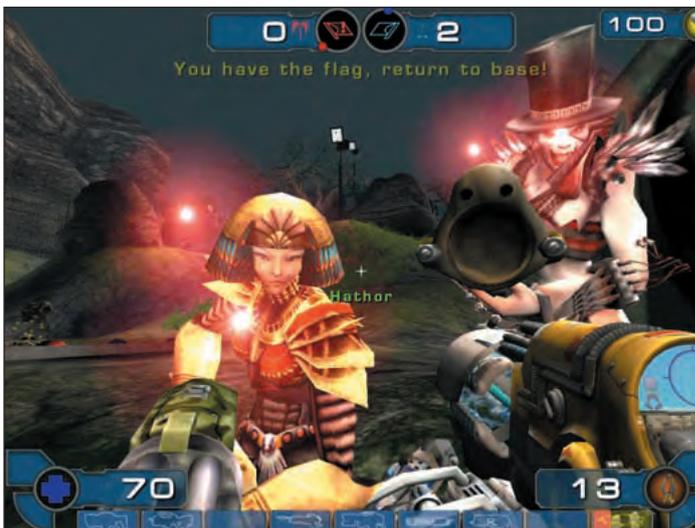
Typically, devastating weapons have a shorter range. The Rocket Launcher is fairly powerful, but rockets are often slow enough to dodge at a reasonable range. The Flak Cannon sprays out shrapnel (which even ricochets off the scenery) but is very limited range-wise.

The sniper's favourite weapon will be the Lightning Gun. Firing a bolt of electricity, it gives an instant hit with no possibility of dodging. The secondary mode activates a scope with infrared capabilities for seeking out figures lurking in the shadows. The most innovative weapon is the Link Gun. The standard weapon fires a pulse shot or a beam as a secondary mode; however, you can also use the beam on team-mates who also have a Link Gun. Instead of harming them, it

boosts their own weapon's power, for more devastating effect. Although it might require a little more organisation, this is a pretty effective strategy for defending static points, and any humming of the *Ghostbusters* theme while you attempt it just adds to the atmosphere.

Design

The graphics themselves and the design of the supplied levels is superb. Sure, there are rather undemanding "inside of a spaceship" maps which are basically metal walls with a few ducts and pipes thrown in for ornamentation. But the 'outside' maps are superb. It would be easy to suffer a few frags because you stopped to admire the view from the brow of a hill. The grass effect is particularly great, and not only looks good when you are dashing through it with an enemy flag, but also provides good cover. The audio is good



This is closer than you usually want to get to your opponents.

Stealth Release?

It does exactly what it says on the box... and a whole lot more!

Unreal Tournament 2003 is a bit of a landmark for Linux gaming. It installs directly 'out of the box'. You don't need any additional downloads as with *RTCW* for example. It is the first mainstream game to be released simultaneously for Windows and Linux, in the same box.

The interesting thing is that there is no mention of Linux on the packaging at all. All the system requirements are for the Windows version. The name 'Linux' doesn't appear in the user manual either. In fact, if you didn't know it was there, you might never find the Linux installer.

There is no mention at all of the Linux version on the official *UT2k3* website either, which is even stranger. It may be that the Linux version was included on the discs at the last minute, but surely the website could have included additional information. There is plenty of info on the demo version for Linux after all.

It seems likely that Epic perhaps didn't want to get into support issues by officially recognising the Linux version, but as they haven't replied to any of our attempts to contact them, so I guess we'll never know.

Graphics cards

It isn't true to say that only Nvidia cards work with *UT2k3*, although these will probably provide the best results. The game makes extensive use of the S3TC (Savage Texture Compression) routines which are only supported under the proprietary Nvidia drivers, not the one included with Xfree86 (they're on the

DVD this issue!). You'll also need to have a GeForce2 or better card to run the game on your box.

Alternatively, the new drivers for the ATI 8500/9000 cards also support S3TC with the new driver available from the ATI site. We have not been able to test how well these work at the moment.

too, but mostly you'll be hearing explosions and the screams of pain from (hopefully) your opponents.

There are maps to suit all types here. Some are more open, others provide more vantage points for snipers. Some are complicated and confusing, favouring those with good memories and the patience to find the quickest route between different points. A few maps are basically a multi-layered collection of open platforms that the player can jump between using the jump points. Aerial

acrobatics may be required to prevent drifting off into space on occasion.

Playability

Unlike many games of this type, there is a deal of playability in the single-player game too – there are plenty of challenging levels and the opponent AI is pretty good – perhaps a little too good on occasions.

The style of the game is a little more fast-paced, even in the Deathmatch modes, which normally lend themselves to camping. Good



Indoor arenas are generic, but you'll want to send postcards from outside.

Reviews Unreal Tournament



Zooming in with the Lightning Gun has an almost X-ray vision effect.



Deathball! looks more like a Sunday-league nightmare than a shooter.



Domination now requires the capture and control of two locations.



"D'oh! Why me?" The vain search for the hot-dog vendor continues...

level design ensures that a range of disciplines is required to truly master the game. You may be able to shoot straight and know where to hide, but that doesn't really help you on the more open levels where speed and

agility are at more of a premium, especially in the Bombing Run games.

As well as the usual integrated multiplayer options (and there are plenty of servers on blueyonder and jolt for *UT2k3*), there is an integrated

stats capability. Authorised servers transmit completed game stats to the main *UT2k3* stats server, which gives a breakdown of your personal scores, as well as the players of the week/month.

Mods

The original *Unreal* game spawned many add-ons and modifications, and *UT2k3* looks set to follow suit. There are already any number of new player models and skins to add to the game. Since these are simply generic texture files, they will work just as well on Linux as they will on the Windows version of the game. Since the majority of servers for multiplayer games inevitably seem to run on Linux these days, it's in the interests of everybody that engine modifications are made available for Linux too, so expect to see simultaneous updates..

At the moment a number of mods are already available, including *DeathBall!* – an amusing take on the 'Bombing Run' style of the game.

Conclusion

The fact that it installs pretty easily, it's identical to the Windows version and that it can be updated and modded too means that, to all intents and purposes, it's just as good as the Windows version. At last a great game for which there's no reason to dual boot. *Unreal Tournament 2003* is a great game, definitely the best in the genre for Linux, and set to spawn a whole host of modded fun. [LXF](#)

Installation

Run the script from the command line

There has been some discussion on websites and newsgroups about the Linux installer for *UT2k3*. The Linux installer script lives on the disc labelled 'Disc 3' and is a script file that launches the graphical installer (the Loki one),

Since there is no documentation at all for this in the manual, on the disc or on the *UT2k3* website, it has caused some people a few problems. Firstly, you shouldn't 'cd' to the disc to run the script. Change to the root user and run the script from the command line using `sh /mnt/cdrom/linuxinstaller.sh` This is important, because the installer will need to swap discs during the install,

and on some distros you may not be able to eject the disk easily.

The second point of note is that the installer names the discs incorrectly. The disc the installer calls 'the Play disc' is the one labelled Disc 1. The disc the installer calls 'Disc 1' is the one labelled Disc 2. We'll leave it to your imagination what happens with the last disc.

The install takes quite a hefty 2.7GB of hard disk space, as it uncompresses and stores all the data locally on your drive. This is because the uncompressed files include the server version of the game, should you want to set up your own LAN or internet server.

VERDICT

Graphics	10/10
Playability	10/10
Features	9/10
Value for money	10/10

Disappointing lack of documentation for Linux install, could've been handled by a big disclaimer. Otherwise – Wow!

LINUX FORMAT RATING
10/10

Official Red Hat Linux User's Guide

Biagio Lucini investigates the benefits of buying a printed compilation of official documentation.

- **AUTHOR** Red Hat doc team
- **PUBLISHER** Wiley Publishing
- **PRICE** £29.95
- **ISBN** 0-7645-4967-7

This compilation of freely available official documentation, with two CDROMs of Red Hat 8.0 is intended as a comprehensive guide for novice to intermediate users.

The book covers aspects from basic install/upgrade to common tasks. It explains how to use the default graphical environment (based on the GNOME 2.0 desktop) as well as introducing basic shell commands. Elementary introductions to more advanced topics like updating packages and structure of system directories also feature, if maybe in a too shallow a fashion to be useful to intermediate

users. Video, audio and image processing are also given proper amounts of space. In addition, information less commonly needed by the target audience like rescuing a broken box and partitioning the disk are discussed in Appendices.

The book is easy to read: critical info is highlighted in the form of notes, tips, caution or warnings that readers can't miss – important for installation. It doesn't make any assumption about the user's knowledge, and all the basic apps chosen as default for a particular task in Red Hat Linux are often well explained, but sometimes (as is the case for printer/network configuration or user management) more emphasis is put on graphical tools rather than on the command line. Moreover, possible alternatives to some tasks and processes also included in the distro, are unforgivably forgotten or relegated



to an appendix (eg KDE). Some advanced tools like *The GIMP* are briefly over-viewed, without getting detailed.

Complete beginners or people with little knowledge of Linux in general will benefit most from this book. Some may question the usefulness of charging for reprinted information that is freely available, but essentially you're paying for the convenience. If you are looking for a book describing the basics of Linux in general, this book is probably not the most useful available, as it is completely focused on Red Hat. Perhaps the intermediate user can glean the odd bit of useful information

from these pages, but a lack of depth in such sysadmin bread-and-butter areas as command line and configuration files mark this out as strictly a beginner's reference to a particular distribution..

VERDICT

Misses almost half of its target audience, but is a must-have for those installing Red Hat without any knowledge of Linux who want to conserve their printer cartridge use.

LINUX FORMAT RATING
 7/10

Official Red Hat Linux Administration Guide

Jonathan Kent thumbs through Red Hat's first home-grown publication and first admin guide

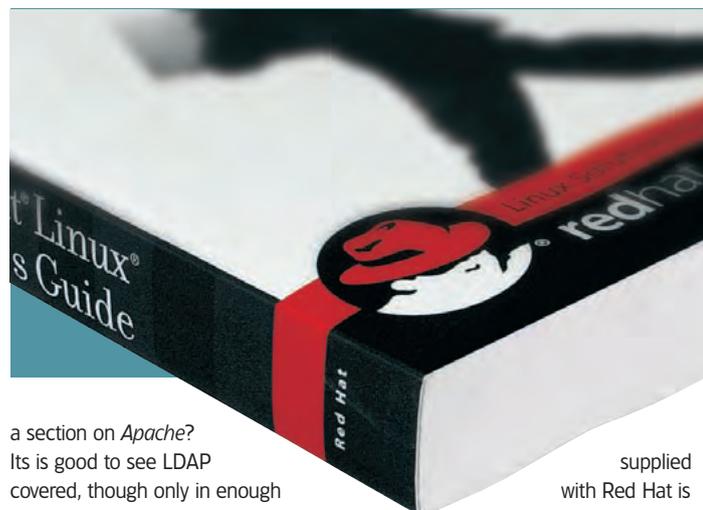
- **PUBLISHER** Redhat Press
- **AUTHOR** Red Hat
- **ISBN** 0-7645-1695-7
- **PRICE** £37.50

Red Hat's first Admin book is aimed at intermediate to advanced users, and split into four major sections – Essential Information, Basic System Admin, Network Admin and Security. The Essential Information section starts off well with a chapter on the configuration of the excellent *Kickstart*, then goes on to discuss the boot process that Red Hat uses and the configuration files and scripts. This detail provides all the info you would need to understand and modify the services that are started at boot time.

The short chapter on GRUB feels a bit out-of-place, considering the target audience, but is well written. Ignoring the title of the oddly named Red Hat Linux File System chapter, it's a good explanation of how the file system is laid out and includes a well explained outline of the */proc* file system. Basic process and system monitoring like GRUB, feels out of place and the section on *ps*, *top* and *df* commands have no real place in a book like this.

Basic System Admin is sadly an apt section title, covering topics any reasonably competent administrator could perform. RAID configuration and Kernel Management are featured, but not in enough detail to be really useful.

Network Administration, is not as exhaustive as it could be; and why does every Linux book always include



a section on *Apache*? Its is good to see LDAP covered, though only in enough depth to get you started. There is no mention of NIS or how to integrate your server into an existing NIS environment; email, NFS and *Samba* are only just touched upon.

Finally, the security section does provide a good intro to basic security concepts. The sections of OpenSSH and PAM provide a good installation and basic configuration guide but not much more. The 70-page Installation Appendix is redundant, as the installation documentation that is

supplied with Red Hat is very good. On the whole, this is a disappointing book, apart from the *Kickstart* chapter.

VERDICT

Those just getting started may find it useful, but most administrators will have to look elsewhere due to the lack of technical depth. Well laid out though...

LINUX FORMAT RATING
 5/10

Mastering PHP 4.1

Paul Hudson takes a look at another contender in the PHP book market

- **AUTHORS** Jeremy Allen and Charles Hornberger
- **PUBLISHER** Sybex
- **ISBN** 0-7821-2924-2
- **PRICE** £37.99

Owing to the fact that PHP 4.3 will be released by the time you read this, you might be forgiven for thinking that a book entitled *Mastering PHP 4.1* is probably limiting its market somewhat. It's geared towards newcomers to programming as a whole – veterans may be put off by the topics discussing what variables are, and quite rightly so – it doesn't stray all that far beyond the minimum requirement for one to gain a good grasp of PHP. However, what it does cover is covered in style – each chapter begins with a thorough explanation of theory, jargon, and appropriate PHP configuration.

Readers will also undoubtedly be happy to see a great deal of info regarding topics beyond standard PHP development, including debugging scripts, using CVS, deploying web applications, and managing security on your server. Each of these topics are presented in a structured, easy to read manner, so the book flows very smoothly from one area to another.

The book is sadly pulled down by a few omissions and oversights that perhaps should have been picked up before release. Specifically, output of code samples in the book is rarely shown, which is something that frustrates many people. Furthermore, there is no conclusion or bibliography, so users wanting to learn more about most of the topics covered therein will have to rely on their search engine of choice. A particular niggle is the use of **\$HTTP_POST_VARS** and associated variables, which have been deprecated in place of the new super-global



arrays. The peculiar thing is that the super-globals are mentioned in the book, but not used in the examples.

If you don't mind that some of the content is a little out-of-date, and the subject matter isn't quite as extensive as it could have been (generating Flash content, for example), then this book does stand out amongst the crowd for ease of reading and depth of the information that it does include. A nice touch is that all the examples use XHTML-compliant output.

If the publishers were to release a second edition which resolves as much

of these highlighted oversights as possible, I wouldn't hesitate to recommend it highly to one and all. As it is, though, it's a darn good read let down only by a few niggling holes in crucial areas. **LXF**

VERDICT

An excellent introduction to the world of PHP, benefitting from being aimed squarely at newbies. Not bad for the price, but could be so much better.

LINUX FORMAT RATING

8/10

RHCT and RHCE

Two great ways to become trained and certified on Red Hat Linux



redhat.
CERTIFIED
TECHNICIAN

Core sysadmin skills

- Install and configure new Red Hat Linux Systems
- Attach new systems to an existing network
- Perform core system administration

Fewer classes required for certification. RHCT is the first step towards RHCE.



redhat.
CERTIFIED
ENGINEER

Senior sysadmin skills

- RHCT-level skills plus
- Set up networking services
- Configure security
- Diagnostics and troubleshooting

RHCE is the #1 certification for overall programme quality (*Certification Magazine, January 2002*).

To find out more about Red Hat Training or to enroll in one of our courses, please email training-uk@redhat.com, telephone +44 1483 734909 or visit www.redhat.uk/training/, citing code TRUKIC

Roundup

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

OUR SELECTION AT A GLANCE

- XFce
- GNOME
- KDE
- Window Maker
- FVWM
- AfterStep
- Fluxbox
- Enlightenment
- IceWM
- QVWM
- Sawfish

Window Managers

Fancy a change of desktop? **Mike Saunders** pushes his X server to the limit with eleven of the best Linux window managers.

Veteran *LXF* readers will recall our window manager roundup way back in issue 2, and since then we've ran various features and tutorials on getting the most out of your desktop. Much progress has been made in the world of WMs too – fully fledged desktop suites like GNOME and KDE continue to improve in ease-of-use and stability, and lighter alternatives offer more unique ways of working.

Linux newcomers may balk at the vast amount of software available in this area. Essentially, the GUI in Linux (and most Unix) systems is separated into clearly defined layers; the X Window System (usually XFree86 in Linux) is the raw engine which communicates with the graphics card,

draws images and text on the screen, and creates and destroys windows etc.

Something needs to organise these windows, though – there needs to be a way to move them around, resize and close them, launch programs and switch between them, and so forth. This is where the window manager program comes into play. It provides titlebars, buttons, launchers and other goodies. You can use other WMs and the contents of the windows

will remain the same, but the decoration and way of organising them will change. Although a simple WM is supplied with X (TWM) – the designers didn't tie it to the system, giving the user maximum flexibility.

We examine eleven of the most popular and powerful WMs available for Linux, considering the feature set, stability and speed of each. We planned to cover fourteen WMs, but space constraints prevented us. You

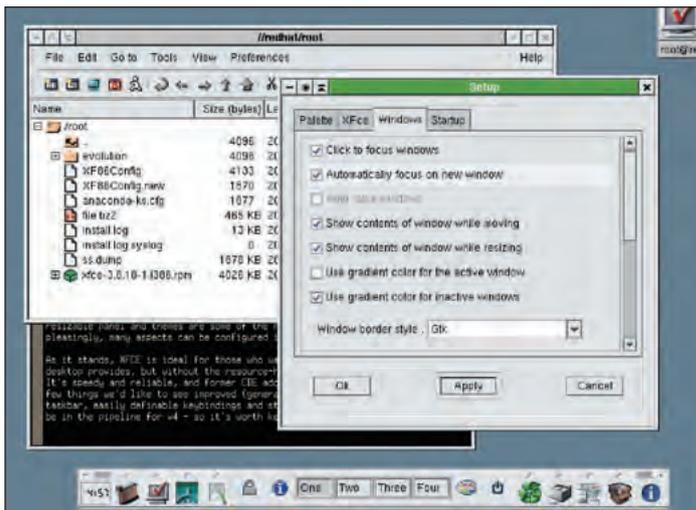
can still find the prodigal progs – *Ion*, *Oroborus* and *AmiWM* – on the cover disc along with the eleven from the magazine! For home systems, the ability to theme and configure a WM endlessly can be a positive point, whereas those deploying Linux in the workplace will prefer something more consistent and approachable. Whatever your taste and needs, you may find something to make your computing life that much better.

'Much progress has been made in the world of WMs – fully fledged desktop suites like GNOME and KDE continue to improve, and lighter alternatives offer unique ways of working'

XFce

CDE meets FVWM in this lighter desktop alternative – pronounced ‘ex-face’

■ **VERSION** 3.8.18 ■ **WEB** www.xfce.org



Although X established itself as the standard windowing system for Unix machines many years ago, the commercial vendors were shipping their own preferred desktops. The threat of Windows NT on the workstation market led larger players to adopt a single consistent WM suite, CDE (Common Desktop Environment), to provide a sane way of moving between Sun, HP, IBM etc boxes.

XFce borrows heavily from CDE's appearance, and along with the FVWM-derived WM it provides a small file manager, SMB browser and various other tools. It's reasonably compact (requiring only *GTK* as a major dependency) and the default setup is easy to get started with. Drag and drop, virtual desktops, resizable panel and themes are some of the more notable features – and pleasingly, many aspects can actually be configured in the WM itself.

XFce's default styling, with the file manager and setup window active.

As it stands, XFce is ideal for those who want some of the goodies a larger desktop provides, but without the resource-hogging that goes along with it. It's speedy and reliable, sports some nifty window decors, and former CDE addicts will feel right at home. The few areas we'd like improvement (general polish, modular panel and taskbar, easily definable keybindings and stronger configurability) appear to be in the pipeline for v4 – so it's worth keeping an eye on.

VERDICT

Features	6/10
Documentation	8/10
Performance	8/10
Stability	7/10

A decent middle-ground with a good feature/size ratio. Excellent to run on lower-spec systems.

LINUX FORMAT RATING

7/10

GNOME

Widespread industry-backed desktop keeps KDE coders on their toes.

■ **VERSION** 2.0 ■ **WEB** www.gnome.org

GNOME's developers received a major boost not long ago when Sun Microsystems, vendor of meaty Unix servers and workstations, announced

that GNOME would be the default desktop on their Solaris OS. With Sun funding documentation and interface guideline work, and Red Hat also

putting money into the pot, the future of this *GTK*-based desktop certainly looks very bright indeed.

The arrival of GNOME 2 a few months back was a huge event – earlier releases had been accomplished desktop suites with some terrific applications, but lacked a certain polish (and the horrendously slow *Nautilus* didn't help matters either). However, the development of GNOME 2 saw a great deal of work on both the front-end and back-end, *Nautilus* was given heaps of performance improvements, and the final result is a powerful desktop that corporate types are taking a shine to.

As with KDE, we're focusing on the WM components here; the *Panel* and *Metacity* provide the main functions, with the latter replacing *Sawfish* as the default WM in recent releases. The stock setup delivers a familiar Windows-like arrangement, with opaque window moving, and a panel containing a taskbar, virtual desktop pager, clock and program launchers. This can be resized and hidden, and an auto-hide feature is available too.

Unfortunately, a menu editor isn't supplied with the standard GNOME 2 setup, so altering items will require manual editing of config files. Similarly, few tweakable settings are available for *Metacity* – the respectable idea is that

consistency needs to be maintained across workstations, without users changing too much. As often happens with perceived shortcomings in more minor operational aspects of applications, there have been a number of vitriolic flames from coders over this, so a careful balance will need to be found to keep more demanding users happy.

Still, keybindings can be modified easily, performance is sound and we came across no major crashes in our tests. *Metacity* is supplied with a good array of themes and the whole desktop is supremely well documented. Ultimately, GNOME 2 is a great step forward in presentation and stability, and gives KDE a very good run for its money. Here's looking forward to 2.2...

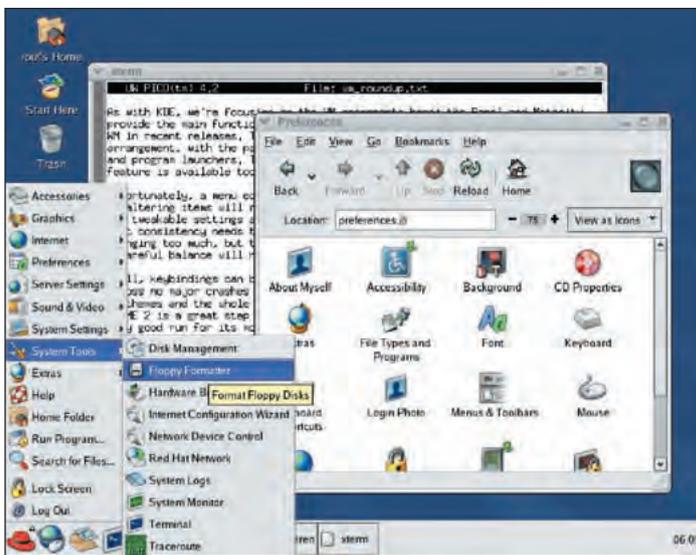
VERDICT

Features	7/10
Documentation	10/10
Performance	6/10
Stability	8/10

Versatile and polished, well documented and stable – GNOME keeps getting stronger.

LINUX FORMAT RATING

8/10



Red Hat's GNOME 2 'Bluecurve' desktop in all its anti-aliased glory.

RoundupWindowManagers

KDE

For those folks who think one click is enough...

■ **VERSION** 3.1rc5 ■ **WEB** www.kde.org

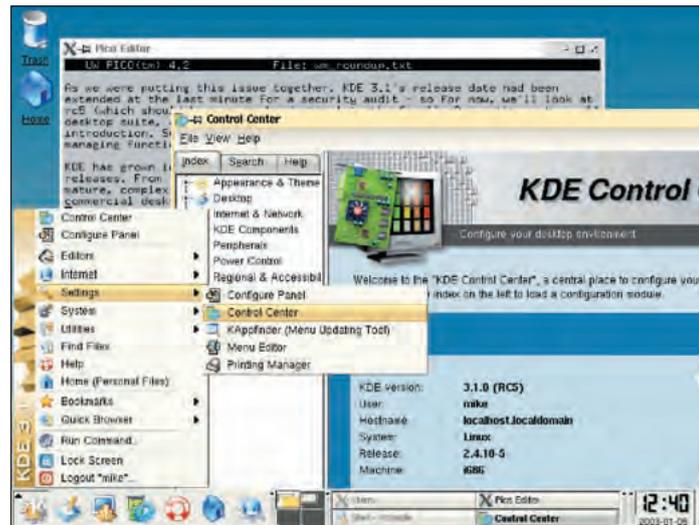
As we were putting this issue together, KDE 3.1's release date had been extended at the last minute for a security audit – so for now, we'll look at rc5 (which should be virtually identical to the final). Being the most popular desktop suite, and supplied as the default on many distros, KDE needs no introduction. Suffice to say, though, we're only looking at the window managing functions here, to make a fair comparison with the smaller progs.

KDE has grown in leaps and bounds since the (admittedly impressive) 1.x releases. From the chunky but functional early days it has evolved into a mature, complex and attractive programming framework which easily rivals the commercial desktops. We built rc5 from source – binaries for your favourite distro should be available at release time.

The most important part of KDE's WM functions are *Kwin*, the simple window manager, and *Kicker*, the

applet-swallowing taskbar-containing panel. Both interact decently; windows snap to one another, maximised ones resize automatically when the panel is hidden, and in general it's extremely consistent. If you use a small resolution, you'll find the ability to resize the panel to an exact height (rather than in large pre-set steps) is a truly wonderful touch.

KDE's pager uses miniwindows to offer a preview of the virtual desktops, while the taskbar can be set to group similar running apps – a useful clutter-eliminating feature. Everything can be configured through the superb *Control Center* and dialogs, and the supplied online help is almost invariably first-class. Power users who prefer to do most actions via the keyboard aren't left out either – in addition to the supplied schemes (emulating Windows, Mac etc), all bindings can be tuned precisely. The focus model can be altered for those who make heavy use of the mouse.



Keramik, the default theme, with all its yellow loveliness.

In all, the Qt-driven suite continues to improve in all directions – 3.1 appears to be one of the fastest releases we've seen (although naturally quite significantly slower than most of the more lightweight WMs that are available). It's robust and flexible, easy to configure, has some gorgeous themes and is hugely satisfying to use.

VERDICT

Features	9/10
Documentation	9/10
Performance	6/10
Stability	8/10

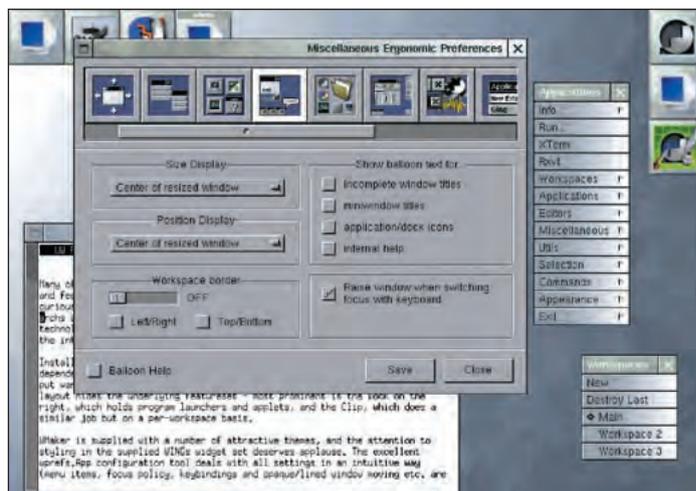
Heavier on resources than its smaller siblings, but packed with goodies galore.

LINUX FORMAT RATING
9/10

Window Maker

Strictly speaking, it doesn't actually make windows, but you're not that pedantic, surely?

■ **VERSION** 0.80.2 ■ **WEB** www.windowmaker.org



One of *WMaker*'s sweeter themes, with the Prefs box open.

Many of the programs we're examining here are attempting to emulate the look and feel of other environments. One such desktop is *NeXTSTEP*, a curious platform originally built around 68k machines, then released for other archs as *OpenSTEP*, and finally being bought up by Apple (and some ideas and technologies remain in Mac OS X). Bowman, a *FVWM* derivative, was first to copy the interface, and now we have *Window Maker* – see www.gnustep.org.

Installation from source or binary packages is a breeze with no esoteric dependencies; once compiled, it's advisable to run the `wmaker.inst` script to put various home directory files in place. Initially, the smooth and clean layout hides the highly capable underlying featureset – the most prominent that you will notice is the Dock on the right, which holds program launchers and applets, and the Clip, which actually does a similar job to the Dock, but on a per-workspace basis.

WMaker is supplied with a number of attractive themes, and the attention to styling in the supplied *WINGS* widget set deserves applause. The excellent `wprefs.App` configuration tool deals with all settings in an intuitive way (menu items, focus policy, keybindings and opaque/lined window moving etc. are all easily modified). Supplied documentation is minimal, aside from some workable guides on the Net, but speed and stability is up there with the best. A good choice if you want something different – pop the cover disc in your drive and give it a whirl!

VERDICT

Features	8/10
Documentation	6/10
Performance	7/10
Stability	8/10

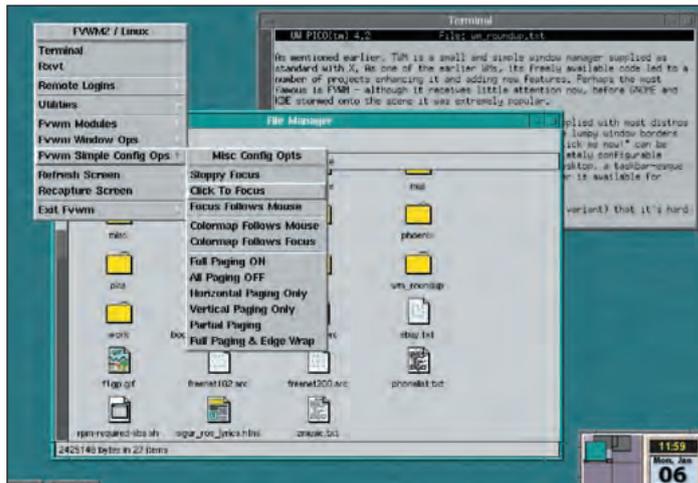
The best of NeXT – robust, quick and pleasing to use.

LINUX FORMAT RATING
8/10

FVWM

What the F stands for is lost in the mists of time...

■ **VERSION** 2.4.14 ■ **WEB** www.fvwm.org



As one of the earlier WMs, *FVWM's* freely available code led to a number of projects enhancing it. Perhaps the most famous is *FVWM* – before GNOME and KDE stormed in, it was extremely popular.

Being a favourite of Free Software hackers, *FVWM* is supplied with most distros and is a reasonably sized download. At first glance the lumpy window borders and lack of buttons screaming out “click me now!” can be daunting, but the purpose of *FVWM* is to provide a completely configurable environment. Menus are brought up by clicking on the desktop, a taskbar-esque button box aids switching apps, and a graphical pager is available for moving around the virtual desktop.

FVWM installations vary so widely (consider the *FVWM95* variant) that it's hard to rate the functionality of its

The straightforward *MWM*-influenced style, with *FvwmButtons* bottom-right.

default settings. There's no supplied graphical config tool, but the text files behind provide a wealth of options – in particular, whole functions containing other commands can be assigned to single operations, making it enormously flexible. Keybindings, window decorations, focus policy and almost everything else can be tweaked and tuned to perfection, and each option and module is heavily documented.

VERDICT

Features	9/10
Documentation	9/10
Performance	8/10
Stability	10/10

Not easy to get started with, but rock solid, mature and outrageously configurable. An oldie, but a goodie.

LINUX FORMAT RATING
9/10

AfterStep

More NeXT tomfoolery with this old-timer...

■ **VERSION** 1.8.11 ■ **WEB** www.afterstep.org

AfterStep, like *Window Maker*, is another WM attempting to emulate the look and feel of the NeXT GUI. Originally appearing as *Bowman*, a

FVWM fork, it's now one of the oldest WMs here and still has a very loyal following thanks to its depth of configuration and wide range of themes.

Tweakable options and settings take longer to get used to with than with *WMaker*, but there's a decent copy of the Dock, and some sweet touches not in *WMaker* like a built-in graphical pager, expanding drawers with more launchers, and a window list along the top.

AfterStep has a hugely configurable desktop with enough bits to tune to satisfy anyone. Those wanting NeXT style should try *WMaker* first, but expert users may feel more at home here.

VERDICT

Features	9/10
Documentation	7/10
Performance	7/10
Stability	8/10

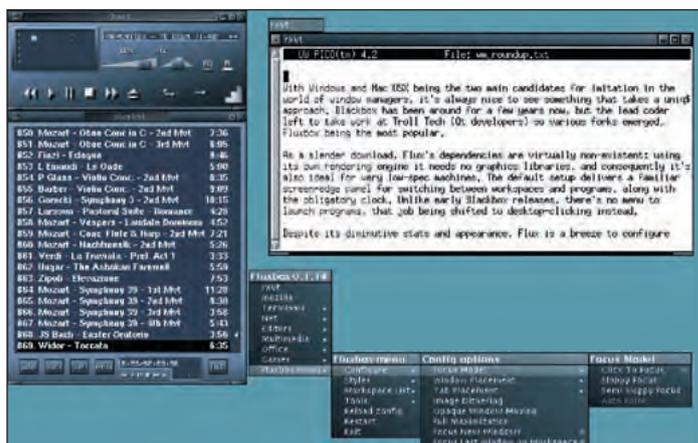
Daunting at first, but a comprehensive and mature alternative to *WMaker*.

LINUX FORMAT RATING
8/10

Fluxbox

Is your Linux box in a state of flux? You need this!

■ **VERSION** 0.1.14 ■ **WEB** fluxbox.sourceforge.net



With Windows and Mac OS X being oft-imitated by other window managers, it's nice to see something that takes a unique approach. *Blackbox* has been around for a few years now, but the lead coder left to take work at Troll Tech (Qt developers) so various forks emerged, *Fluxbox* being the most popular.

As a slender download, *Fluxbox's* dependencies are virtually non-existent; using its own rendering engine it needs no graphics libraries, and consequently it's ideal for low-spec machines. The default setup delivers a familiar screen-edge panel for switching between workspaces and programs, along with the obligatory clock. Unlike early *Blackbox* releases, there's no toolbar menu to launch programs, that job being shifted to desktop-clicking instead.

Most Flux themes try to use as little screen space as possible.

Despite its diminutive stats and appearance, *Fluxbox* is a breeze to configure through the menus. Window placement, focus, workspaces and all the usual settings can be modified, and the single flat man page explains everything well enough. Overall, in terms of sheer speed, *Fluxbox* flies with its tight codebase, speedy drawing engine and simple but attractive themes.

VERDICT

Features	7/10
Documentation	7/10
Performance	10/10
Stability	9/10

Lightning fast and simple to work with – a no-nonsense WM.

LINUX FORMAT RATING
8/10

Sawfish

LISP-driven WM seeks new desktop to shine in.

■ **VERSION 1.2** ■ **WEB** sawmill.sourceforge.net



Sawfish's decoration of individual windows can help organise your work.

Originally known as *Sawmill*, but renamed due to a clash with another program, *Sawfish* superseded *Enlightenment* as the default GNOME WM for a few releases. Its aim, according to the developers, is "simply to manage windows in the most flexible and attractive manner possible", and this is achieved with thorough use of the LISP scripting language at its core.

Built around the *GTK* toolkit (with versions for both *GTK1* and 2), *Sawfish* is supplied with most distros and requires a few extra libs to get started. It doesn't provide any taskbars, icon docks or similar tools – those are left to other external programs. Pleasingly though, almost all of the settings can be altered through the simple config dialogs – virtual desktops, focus policy and the like can be changed, and one particularly smart feature is the ability to assign decorations on a per-window basis rather than globally.

Sawfish's continuing popularity no doubt stems from the LISP engine underneath; having a full programming language to modify the WM gives extreme flexibility for those who like to get their hands dirty in code. As a stand-alone WM it does a fairly decent job when simplicity is all that's required, so check it out if you want to see what goes on beneath a WM. If attractive appearance of your windows is a priority, *Sawfish* will raise a smile.

VERDICT

Features	8/10
Documentation	7/10
Performance	7/10
Stability	8/10

Solid and simple, with plenty to keep hackers happy.

LINUX FORMAT RATING
 7/10

WINDOW MANAGERS THE VERDICT

The sheer amount of WMs available for Linux – good or bad thing? Some claim that the number of choices can overwhelm new users; others rightly point out that GNOME and KDE are becoming the standard Linux GUIs now and only the inquisitive will try to run something else. Whatever the case, there's no denying that the flexibility of X caters for everyone – those who want something unique have plenty to choose from, and the lightweight WMs can turn that old 486

box into a useful little workstation. As with most app genres in the world of Free Software, WMs have made stunning progress in the last few years. Gone are the days of hand-hacking config files to get something remotely functional; now we have two slick, mature and professional desktop suites to make Windows refugees comfortable, and distro vendors like Red Hat, SuSE and Mandrake funding the coders can only help more.

These eleven all fit nicely into three categories: the fully-fledged

environments of GNOME and KDE, the middleweight desktops of *XFCE*, *Enlightenment* and *WMaker*, and the small and zippy WMs which provide great alternatives for older hardware or long-time users.

Of the larger duo, we find that KDE just pips GNOME to the post, but the differences are becoming negligible for general users. *WMaker* is our pick of the mid-range WMs, offering a smooth and pretty desktop which truly feels unique, while *IceWM* comes up trumps at the low-end for

being fast, configurable and easy to use. Equally, *Fluxbox* is a good choice on low resolution setups.

Finally, the selection of one's WM is a very sensitive area and flame wars abound on the Net (even rivalling the old *Vi* vs *Emacs*). We've tried to avoid getting bogged down in the purely personal matter of aesthetics, instead concentrating on the technical side, so please don't send us hour-long rants if you disagree with us on a minor cosmetic point! Now, where did we put our asbestos suits...? [LXF](#)

Table of features

We tested the window managers on a low-end Pentium and a newer PIII box, both running Red Hat. Note: package sizes and memory stats below are approximate. Memory usage varies from system to system, and is affected by the program's configuration and other factors.

Name	License	Toolkit	Inspiration	Package size	Avg. mem usage	Configurability	Alternatives
XFce	GPL	GTK	CDE	4MB	5MB	Medium	GNOME
GNOME	GPL + LGPL	GTK	Various	N/A	10-15MB	Medium	KDE, OEone
KDE	GPL	Qt	Various	N/A	15-20MB	High	GNOME, OEone
Window Maker	GPL	WINGs (supplied)	NeXTSTEP	2MB	3MB	Medium	AfterStep
FVWM	GPL	Inbuilt	MWM, others	1.7MB	2-3MB	V. high	FVWM95, E
Fluxbox	Distributable	Inbuilt	Blackbox	370KB	2MB	Low/Medium	Openbox
Enlightenment	GPL	Inbuilt	Amiga, others	10MB	5MB	V. high	FVWM, AfterStep
IceWM	LGPL	Inbuilt	Win9x, OS/2	800KB	2.5MB	Medium	QVWM, PAWM
Sawfish	GPL	GTK	Unique	1.6MB	4MB	High	SCWM, FVWM
AfterStep	GPL	Inbuilt	NeXTSTEP	1.2MB	2-3MB	V. high	WMaker, Bowman
AmiWM	Distributable	Inbuilt	Amiga WB	140KB	1MB	Low	IceWM, WB theme

HotPicks

The best new Open Source software on the planet!



Richard Drummond

As well as the Java series, Rich finds time to review new Linux apps.

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

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

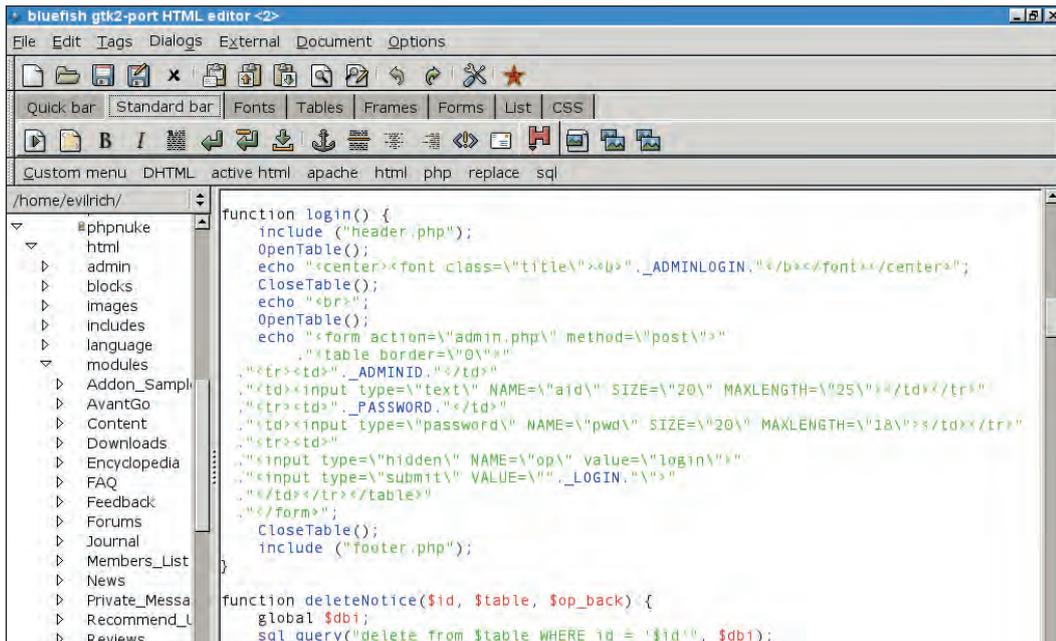
If you have any suggestions for things that we should cover, email us at linuxformat@futurenet.co.uk

HotPicks at a glance

Bluefish 0.8	42
Bochs 2.0.1	43
Gambas 0.4.1	44
KBudget 0.41	45
Phoenix 0.5	45
Epiar 0.2.0	46

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!



The new **GTK+2-based Bluefish** feels a whole lot crisper and cleaner than its predecessors.

HTML EDITOR

Bluefish



■ **VERSION** 0.8 ■ **WEB** <http://bluefish.openoffice.nl/>

While some people still swear by a plain text editor such as *Vi* for hacking HTML, for larger projects a dedicated HTML editor can make web development quicker and less error-prone. *Bluefish* is just such a beast. It's not a frilly, WYSIWYG homepage builder, beloved of web designers on Windows and Mac, but a text editor equipped with tools to aid web development. It boasts syntax highlighting for HTML, PHP, XML and Java (and is easily extensible to cover additional languages), and has dialogs and wizards to quickly enter the most common HTML constructs.

Bluefish 0.8 has been ported to the new *GTK+2* toolkit, and substantially benefits from the upgrade. This means that *Bluefish* can now use anti-aliased fonts, and the application now feels a whole lot crisper, cleaner and more precise. The port isn't all good news,

though 0.8 has features not seen in 0.7, several haven't made it to the *GTK+2* branch yet. Most notably, there's no preview function (you must test code in an external browser) and no project management features.

Bluefish still offers some powerful features. Its Search And Replace function is particularly strong and optionally supports POSIX and Perl regular expressions. Likewise its support for multiple character encoding is good and easily allows conversion between encodings. Another neat touch is its ability to auto-generate image thumbnails.

Without a doubt, *Bluefish's* star feature is its extensibility. You can create your own dialog-based macros, which are able to prompt the user for parameters and generate and insert code based on those parameters. This is a great feature for building a library

of frequently-used code fragments to speed up development, and the default set-up includes parametric macros for doing mouse roll-overs in Javascript, various database functions in PHP and a lot more. Also useful is *Bluefish's* ability to call and integrate with external programs. This obviously includes things like browsers, but this feature has much wider possibilities than that. For example, in the default setup, you can launch *tidy* (the W3C's HTML beautifier) to check, clean up and reformat your code – or use *weblint* to validate your code.

The most obvious competition for *Bluefish* is the Qt-based editor, *Quanta*. *Quanta* has the edge in terms of ease-of-use and features, but *Bluefish* does offer functionality which *Quanta* doesn't have, such as the parametric macros and thumbnail generation. Both packages have their faults – for example, with both, the HTML entry dialogs fall over when the HTML is embedded in a scripting language – and which you prefer probably comes down to a matter of personal taste. *Bluefish* is capable tool and will prove a real boon for time-pressed web developers who can live without frills.

PC EMULATOR

Bochs

■ VERSION 2.0.1 ■ WEB <http://bochs.sourceforge.net/>

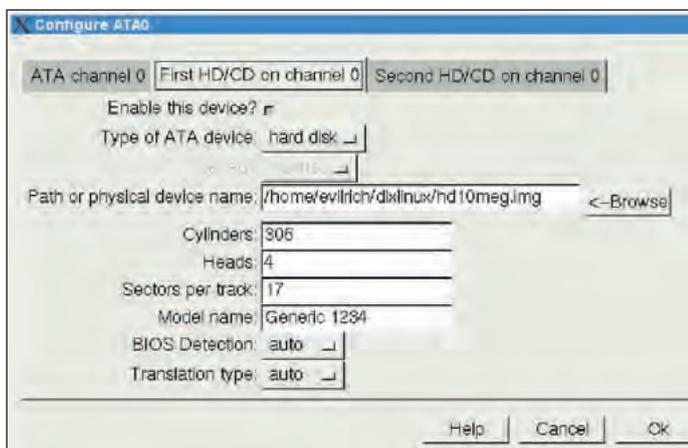
PC emulation is an application that has come of age in recent years – partly due to the ready availability of sufficient processing power to make emulation practical, but also due to increased sophistication of the emulators themselves. The market leader on Linux is arguably *VMware*. Actually, *VMware* isn't strictly a PC emulation; it's a virtualization. In *VMware*, the PC's CPU isn't emulated; software runs on the host processor in a virtual environment. If you don't want to shell out for a commercial application, then the Open Source projects *Plex86* and *Bochs* are gradually becoming a cost-effective replacement for *VMware*. While *Plex86* is a virtualization in the manner of *VMware*, *Bochs* is a full-blown PC emulation. While it runs considerably slower than either *VMware* or *Plex86*, it has the advantage of portability to other processor architectures and can even simulate an SMP architecture (with up to 15 processors) on a uniprocessor box.

With the release of version 2.0, *Bochs* has made a quantum leap in ease-of-use and functionality. One reason for this is its new plug-in architecture. You can build *Bochs* with the support for all the devices that it supports, and pick what you need to use at run-time by modifying your configuration file. This makes it simple to create a series of different emulated environments should you need to. This applies to emulated devices – such as Ethernet, parallel and serial interfaces – and to display libraries and configuration interfaces. *Bochs* 2.0 supports a variety of display libraries, including plain *X11*, *libSDL*, *wxWindows* and *curses*. Similarly for configuration, you can stick to the old console-based configuration menu or use the new *wxWindows* configuration GUI.

Bochs used to be a pain to configure, but the new *wxWindows* interface makes life much easier. You can tailor your emulated environment through a series of dialogs which are accessed via a drop-down menu. This lets you pick your CPU and memory settings, add virtual storage devices,

configure I/O devices and so on – all through point-and-click. It's still not as easy to use as *VMware*, but it's a big improvement over the old *Bochs*. Problems include a lack of online help (so far) and an unintelligent default set-up (which means you need to do a lot of work to get, for example, your virtual network interface to fly). It would be handy if some task-based wizards were added to lead you through the setting up of drives, networking and so on.

The new *Bochs* 2.0 is faster. It's still no match for the virtualizations, but double the performance of *Bochs 1.x* is claimed. It's hard to gauge the real world performance of the virtual environment – it depends on too many variables – but *Bochs* 2.0 feels a lot more responsive. The sluggishness of emulated devices tend to be more noticeable than the lack of processing power of *Bochs's* virtual PC. It's not recommended to use real hard drive partitions with *Bochs*, so you must use slower image files; also, the emulated display is just a dumb framebuffer – there's no graphics drivers to provide accelerated graphics in a guest operating system. Still, *Bochs's* CPU emulation is flexible, offering integer



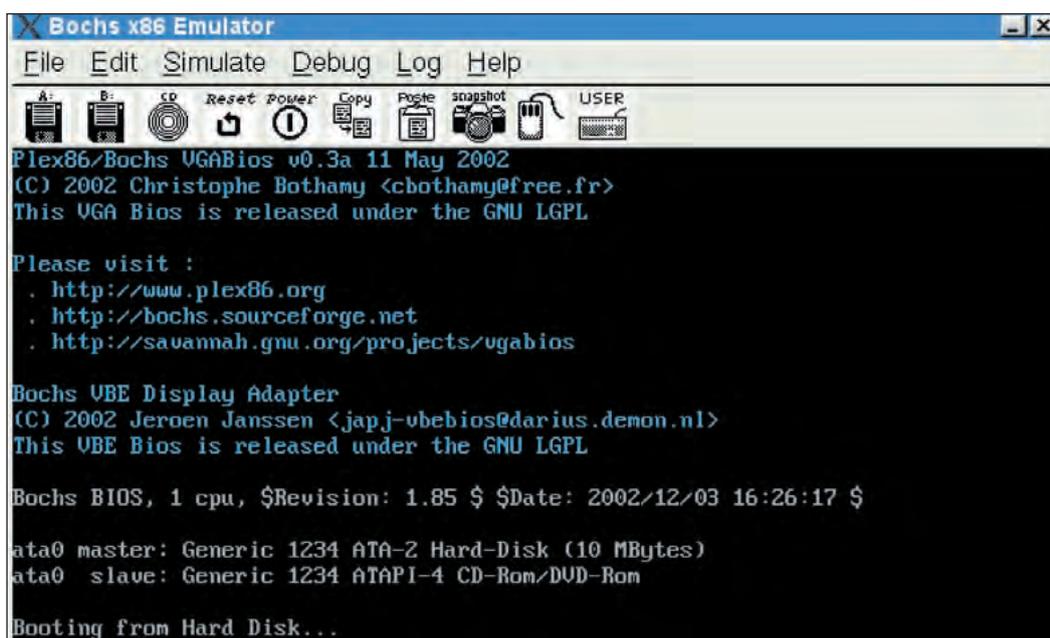
The new *wxWindows* GUI makes configuration a lot less painful.

and floating point emulation and it supports MMX, SSE and SSE2 extensions. Interestingly, partial support for AMD's x86-64 architecture is also provided.

Device support in *Bochs* 2.0 has also improved. You can now have up to four ATA channels for a maximum of eight ATA devices, and large file support boosts the size of virtual drives to a maximum of 32GB rather than 2GB. As we said above, *Bochs* can only safely work with hard files rather than physical drives, but access to real CDROM drives works reliably. In conjunction with the El Torito support in the *Bochs* BIOS, this makes installing a guest operating system straight forward. Networking is more flexible too, but configuration isn't easy. *Bochs* can now use the kernel TUN/TAP network device for virtual

networking, and, if you have masquerading enabled in your host kernel, your virtual machine can gain full access to your network without you having to assign it its own network address. Networking with *Bochs* 2.0 is nowhere as near as transparent as in *VMware*, though.

Better performance, easier configuration and an all-round boost in flexibility mean that *Bochs* 2.0 is a vast improvement over earlier releases. It's not a replacement for *VMware*, but can be used on other CPU architectures to furnish a useful environment for running Windows and DOS binaries when required. However, *Bochs* is much more useful as a developers' tool, to create a range of virtual environments to deploy and test code. As such, its advantages of flexibility and portability begin to shine.

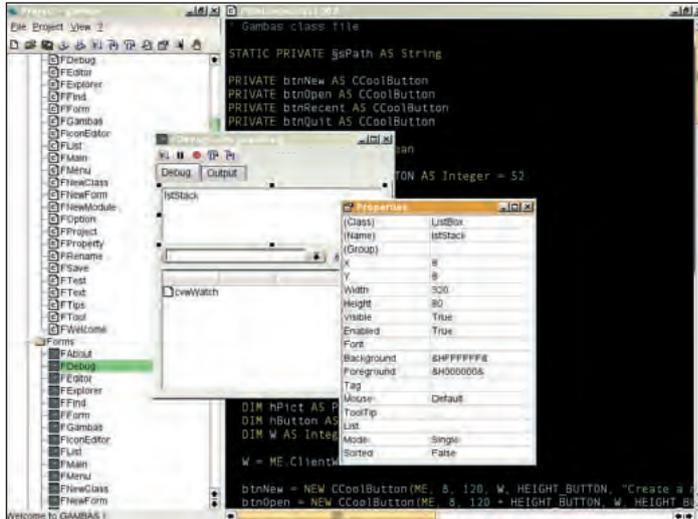


Bochs provides a flexible and portable environment for running x86 operating systems.

VISUAL BASIC CLONE

Gambas

■ VERSION 0.4.1 ■ WEB <http://gambas.sourceforge.net/>



The *Gambas* IDE makes developing simple GUI-based applications a doddle.

Despite whatever ideological differences you may have about the programming language BASIC, it would be hard to dispute the success and popularity of Microsoft's Visual Basic platform. The lack of a similar environment on Linux could well be a factor deterring many developers from adopting Linux. With *Gambas* – standing recursively for Gambas Almost Means BASIC – Linux at last has the beginnings of an effective and usable alternative to Visual Basic. It is not meant to be source compatible with Visual Basic, and, currently, *Gambas's* scope is a lot narrower than VB – its focus is the development of client-side, console and GUI-based applications – but *Gambas's* ease-of-use and the professionalism of its implementation means that the barrier of entry to Linux developers is now lower than ever before.

Gambas feels superficially a lot like Borland's *Kylix*. For one, it is also based on the Qt toolkit; for another, its interface is broadly similar, although a good deal simpler. *Gambas* isn't architecturally tied to Qt, though. The *Gambas* class library abstracts away the details of Qt, and it will be possible (and, in fact, intended) to replace the existing Qt implementation with other GUI toolkits that you want to use. It just needs to be implemented.

Gambas is based on a command-line BASIC interpreter which understands an object-based flavour of BASIC. This implements all the usual constructs of a structured BASIC plus object-based extensions. The *Gambas* object model is simple, but this simplicity is not a drawback by any means – it's one of *Gambas's* key weapons. It is designed for component-based development rather than true object-orientation. Thus you cannot define or inherit custom classes with *Gambas*; rather you subclass the supplied components and override methods to implement desired behaviour. This ties in closely with *Gambas's* form editor, which provides a visual environment to piece together user interfaces or forms.

The power of the *Gambas* environment is demonstrated by the *Gambas* IDE, which itself is implemented as a *Gambas* project (and browsing the source code to this is an excellent way to learn the capabilities of the system). The IDE provides the usual elements: a project browser, which gives a tree-shaped view of the components in your project; the source editor, in which you write code; the form editor, with which you build GUIs; and the property editor, which allows you to manipulate the properties of the components used in a GUI.

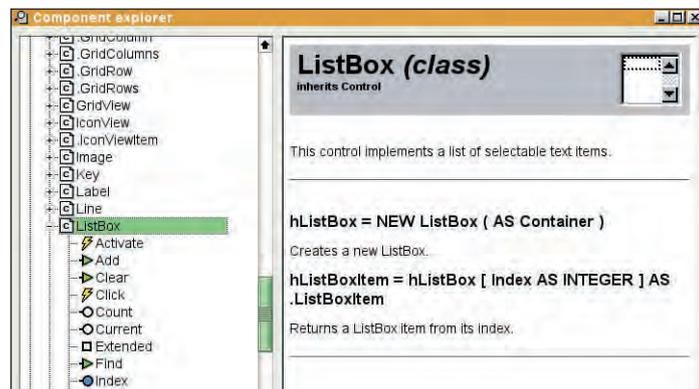
Gambas provides a familiar set of widgets with which you can build forms, including buttons, lists, trees, tab panes and so on. A simple pixmap editor is also included for creating icons and button images. *Gambas* provides a limited control over the grouping and positioning of widgets within a form and its event model is very basic. Each interface event has a corresponding method in a component's class; you simply implement the method to react to that event. For example, if you wish to create a GUI which adapts to window resizing events, you must handle these events yourself and manually reposition widgets within the form.

The *Gambas* source editor offers few bells and whistles but is quite adequate for the job. Naturally, it provides the usual syntax highlighting, undo, and search and replace functionality. More advanced features such as code-completion or dictionary look-up are lacking. *Gambas* has an integrated symbolic debugger, which lets you trace execution through a

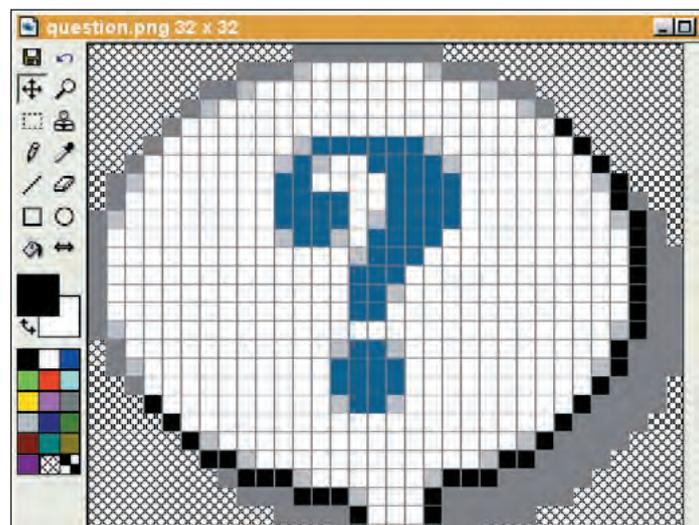
project, watch variables and expressions, and so on. Surprisingly, you cannot set breakpoints, though.

Deploying *Gambas* projects is laughably simple, since the system includes a compiler which supports the creation of executable binaries. The executables it produces are incredibly small in size and require only the *Gambas* run-time environment to work.

Gambas provides an impressive and slick development environment. The attention to detail clearly shows that it has been a labour of love and catapult it into a league above many more advanced development tools. Its ease of use, completeness and stability make it a useful environment, despite its simplicity. Touches – such as the built-in component explorer, which lets you browse the documentation for the supplied component library, and the *Gambas Language Encyclopaedia*, which documents the language syntax and the component library – mean that this is an easy platform for beginners to adopt.



The integrated component explorer provides a handy reference to browse through the whole of the *Gambas* component library.



There's even a built-in pixmap editor for knocking up icons and images.

PERSONAL FINANCE MANAGER

KBudget

■ **WEB** <http://www.garandnet.net/kbudget/> ■ **VERSION** 0.4.1

Back when home computers first appeared, the sceptic would often ask the early-adopter, "Great. But what can you do with it?" A frequent reply was, "Erm. It can help manage my bank accounts." I don't suppose many people ever did use their first home micros for this, or, if they did, they quickly realised it was more effort than it was worth. With today's software, luckily, such applications are commonplace. While you can use a spreadsheet and some macros to the job, a dedicated tool such as *KBudget* can help speed things up dramatically.

KBudget is a KDE-based application which lets you keep track of money flowing through back accounts and helps you set and (hopefully) keep to a budget. The main display shows the totals and budgets for your accounts for

the current month. There are three types of account you can set up. Income and expense accounts are not real accounts, they're categories to help you organise money that goes into and out your real accounts. An income is just as it sounds and covers things like your salary; expenses are your outgoings such as bills you have to pay; cash accounts are your real accounts, such as a current account. You enter transactions to record the flow of money from one account to another.

KBudget is a simple application. It doesn't let you set up recurring transactions; it won't let you generate reports; and it can't produce appealing pie charts and graphs. What matters is the pure and simple functionality; such as the ability that *KBudget* gives you to set budgets for each of your various accounts. Say that you expect to spend £500 on bills that month. You

Income			
Account	Budgeted	Received	Remaining
Salary	£ 2,000.00	£ 2,000.00	£ 0.00
Total	£ 2,000.00	£ 2,000.00	£ 0.00

Expenses			
Account	Budgeted	Spent	Remaining
Bills	£ 1,950.00	£ 600.00	£ 1,350.00
Total	£ 1,950.00	£ 600.00	£ 1,350.00

Accounts			
Account	Start	Change	End
Current account	£ 500.00	£ 1,400.00	£ 1,900.00
Total	£ 500.00	£ 1,400.00	£ 1,900.00

Remaining from Dec: £ 0.00
 Budgeted for Jan: 1,950.00
 Total: £ 1,950.00
 Ready, Expected change: £ 50.00 Expected balance: £ 550.00

KBudget will help you become your own iron chancellor and reign in your purse strings.

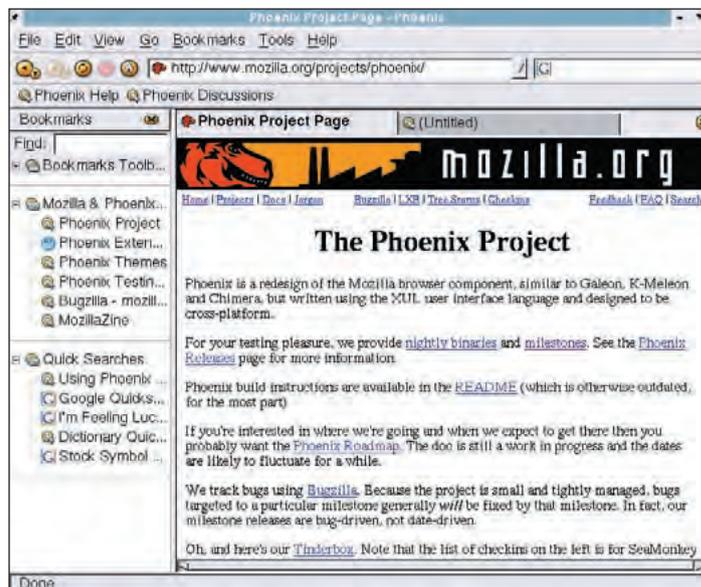
would set the budget for your Bills expense account to £500. As each expense transaction is entered, *KBudget* will remind you of how much you have left in your budget and give you an estimate of how much will be

left in your cash accounts at the end of the month. Simple, but effective monitoring of your financial comings and going. Sticking to your budget will no doubt prove more difficult than using the program!

WEB BROWSER

Phoenix

■ **VERSION** 0.5 ■ **WEB** <http://www.mozilla.org/projects/phoenix/>

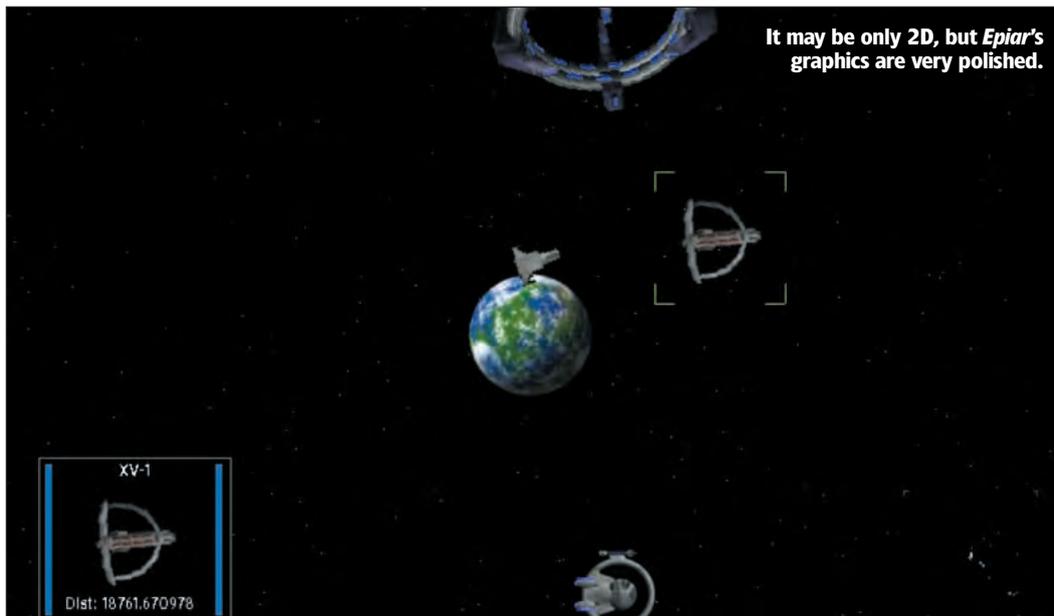


Mozilla is a great browser, but it's just too large. If you agree with that statement, then you'll love *Phoenix*. *Phoenix* is a web browser which is based on *Mozilla*, but has the goal of being lighter and faster than its older sibling. Currently the Linux version weighs in at a download size of 8.7 MB (by comparison, *Mozilla* 1.2.1 is 13.5 MB). The difference between *Phoenix* and other *Mozilla*-based browsers such as Galeon is that *Phoenix* doesn't just embed the *Mozilla* HTML engine *Gecko*; it uses the core of *Mozilla*, *Gecko*, the XUL interface language, and all. *Phoenix* is noticeably quicker than *Mozilla* and starts up in less than half the time. It is just a browser, however, and doesn't offer any of the mail, news, or composer features that attract users to *Mozilla*'s all-in-one capability, which might make some users wonder what's actually in those 8.7MB.

The *Phoenix* web browser is a slimmed down *Mozilla* which is light and easy to use.

Phoenix reaps many of the benefits of its *Mozilla* foundations. It has a world-class rendering engine and, since it can use *Mozilla* and *Netscape* 7.0 plug-ins, is supported by a wide range of plug-ins. It offers all the features you would expect from a browser, such as bookmark and cookie management, and embellishes on this with tabbed and full-screen browsing. Thanks to XUL, *Phoenix* is easily extensible. Currently available extensions offer such additions as mouse gestures and a spoofing function to modify the user agent string that *Phoenix* sends to web servers. Naturally, theming is also supported, and plenty of themes are available for download, including ports of many *Mozilla* themes.

It may feel a little spartan compared to the full-blown *Mozilla* or to *Konqueror*, but *Phoenix* is a very capable browser. Its speed, smaller footprint and uncluttered interface will appeal to those who want the advantages afforded by *Mozilla*, but don't want the bloat.



SPACE GAME

Epiar

■ VERSION 0.2.0 ■ WEB <http://epiar.net/>

Epiar is a game of space exploration and combat in which you play the role of a shuttle pilot bent on amassing fame and fortune. The game and universe of *Epiar* are obviously inspired by the all-time classic, *Elite*. (and its descendant, the net-based *JumpGate*) The main difference in *Epiar* is that you get a 2D, top-down view of your shuttle craft, rather than a first-person view from within the cockpit. So much for progress, right? However, in a game like this, depth of gameplay and the level of immersion are what count.

Depth of gameplay is something *Epiar* currently lacks, simply because it's very much a work in progress. Potentially, *Epiar* is a game of epic proportions; alas, in the latest CVS version at the time of writing the universe consists of a mere 10 planets.

You may ask why are we talking about *Epiar* in *HotPicks* if it's not finished yet? Well, most of the progs we discuss is not finished to some degree or another. *Epiar* as it stands is fun shoot-'em-up, and with more work could become an engrossing game. What makes it stand out at the

moment is the work that has been expended on presentation. Unlike many Open Source games, *Epiar* oozes quality – from its use of the Loki installer for easy installation, to its background music, to its well-written built-in piloting tutorial.

The game begins with you in space orbiting the stable planet Junonia. You have some credits and an unarmed craft. You can land on the planet to buy weapons or, when you have earned more money, upgrade your ship. Weapons available include the usual range of lasers and missiles, but also useful is the status sensors. This will give a read out of a targetted ship's shield strength; when the shield has been sufficiently depleted by your attack, you can board the ship and loot it. Currently, piracy is the only way to make money in *Epiar*, trading has not yet been implemented. This is one

area we would like to see developed, because one aspect of *Elite*'s enduring longevity was the variety of careers you could embark on.

Combat is simple – you have a targeting system, with which you can lock on to your prey, and then any shots will automatically be aimed. If the target goes out of range, the system will lock on to the next target in range. This takes a bit of getting used to, and, because of the flat representation of space, actually tracking down what you want to target can be tricky. Your radar can help here, but there's no way to identify those blips on your radar screen.

Navigation needs some work, too. While you can pop up and zoom a star chart, finding your way between solar systems isn't easy. It would nice to have some way to locate planets by name. Currently, the only way to travel through hyperspace is to pay to go via jump gates, but the network of jump gates isn't marked on your chart. No doubt, some kind of device for creating your own hyperspace jump point will be included as a ship upgrade in a future game; perhaps a mapping device will be available, too.

There are lots of great ideas on the drawing board for *Epiar*, including missions and network play. It would be great to have some historical details about the various political alliances in the *Epiar* universe, specifications of the ships you will encounter, more detailed statistics on the planets you will visit and so on – all important to create immersiveness. This is something that everybody can contribute to – without any programming knowledge. If you want to name a planet, why don't you lend a hand? [LXF](#)



Pay careful attention to the political status of a planet before engaging in piracy in its solar system – you don't want to start robbing your allies.



Upgrade your ship with a multitude of weapons of mass destruction – but you'll need plenty of currency for all the really nasty stuff.

Ready for the desktop?

Is Linux ready for the DESKTOP?

With help from Linux luminaries and industry insiders, **Andy Channelle** and **Nick Veitch** assess the state of Linux on the desktop.

cover feature



With the server sector smitten – the most recent Netcraft figures available gave *Apache* some 66% of the web server market – thoughts within the Linux community are turning to the desktop, where accepted wisdom says Unix-like operating systems are too difficult, the dominant player is far too powerful, and users are resistant to change. But accepted wisdom is often wrong – just look at the success of *S Club 7* – and things are changing. A recent International Data Corporation (IDC) study suggested that while Windows still dominates the desktop, Linux is poised to unsettle Apple's *MacOS* in second place, perhaps as soon as the end of 2003. Corporate muscle in the form of Sun, IBM, Oracle and others is giving Linux a significant gloss of acceptability in enterprise; and Microsoft's recent licensing changes are pushing cash-conscious buyers

into the arms of the opposition. In short, the vicious-circle barrier to entry – there aren't enough users to justify building applications and there aren't enough applications to justify becoming a user – in the OS market looks breakable for the first time since the explosive rise (and disheartening fall) of *Netscape* in the 1990s.

So what has changed?

At the forefront of this challenge to the status quo is the concept of Open Source distribution. The source code of everything from applications and desktop environments to the kernel itself is given a life beyond a single individual, group or corporation, and something that is not 'owned' in the traditional sense is impossible to 'buy'. To pick a business phrase at random: you can't 'cut off the air supply' to an Open Source competitor. If, for example, Sun decided to back off its commitment to *OpenOffice.org*, the

app suite would suffer a setback, but certainly wouldn't disappear. Sun could stop selling *StarOffice*, but that would not endanger its Open Source cousin. AOL could decide to can its development effort on *Mozilla*, but the nightly builds would keep on coming. These projects have a life of their own, and only indifference can stop them.

Is Linux ready for the desktop? We think the answer is, generally, yes; but it is not that simple – is it ever? The first issue is that there is more than one type of desktop user, and Linux may not be appropriate for everyone; the second problem is what exactly do we mean by 'ready?'

The first difficulty can be solved by simply dividing users into either home, education or business desktops; there are obviously other distinctions, but broadly speaking people use their PCs for these jobs. Some applications and tasks (and potential gaps in application availability) straddle all of

Mired in TCO

Just how do you define 'Total Cost'?

Total Cost of Ownership (TCO) is an area ripe for argument, with every study seemingly contradicting the one that preceded it. The problem is that quantifying the cost of running an OS/application combination over a period is difficult, taking in not just initial purchase price, but also hardware upgrades, support charges and the cost of keeping everything running smoothly. When you're looking at supporting tens, or even hundreds of thousands of users, a small change could make a big

difference. Factor in the cost of migration and training and the issue is not as cut and dried as one (or indeed the unimaginative bean-counters attached to your project) might imagine.

One recent study, commissioned by Microsoft from IDC, claimed that Linux is more expensive to support in all areas but web serving (where the MS solution has a six per cent higher TCO). In areas such as network infrastructure, file and print serving and security, the difference was judged to be between 11-22% in

favour of MS Windows. The study said the extra costs incurred by Linux were "related to the relative immaturity of the management tools available today for Linux systems;" but that "over time, the gap in support costs between Linux and Windows will contract."

One of the biggest (and most argued about) components in IDC figures is the cost of manually supporting users, with Microsoft arguing that as MCSE qualified administrators are more numerous and so cost less to hire. In response, Linux

consultant Brian Schenkenfelder claimed that a competent Linux sysadmin can handle three times more users than the average Windows admin. Researcher Chad Roberts from the Robert Frances Group went even further claiming that while the Windows admins that he engaged with could only comfortably deal with 10 – 15 systems, a Solaris, NetBSD or Linux admin can manage around 1000 servers at a time.

"Sure, the salary's more expensive, but I get more life out of them," he said.

Ready for the desktop?

“Will Linux ever be ready for every desktop, or should advocates be aiming for a niche market like Apple and the publishing industry?”

these areas, but some are more appropriate to one than another. For instance, Jason Spisak of Lycoris, developer of *Desktop/LX*, suggests there are six big needs that home users have: Web, email, word processing, games, digital photos and music. The first three tasks are pretty universal, but business users are more likely to need to use spreadsheets, databases and presentation software, plus cross-platform groupware/calendaring apps. In addition adventurous home users will probably require personal finance management software or easy DTP tools, and students or academics may need esoteric applications covering outlining or bibliography management, as well as graphics, scientific analysis and research. Those on the cutting edge may be intending to edit video or audio, create dynamic web pages or capture episodes of *Buffy The Vampire Slayer*. As you can see, the ‘desktop’ is not that easy to define!

To cover the second issue, we’d need to find out how able Linux is in each of these disciplines. “Ready” would mean that appropriate apps (proprietary or Open Source) are not just available, but stable and useable in the intended context, and that mainstream hardware works as expected. If elements are missing, what needs to be done to resolve the issues, and are they being tackled? Will it ever be ready for every desktop or should Linux advocates be aiming for a niche market, in the same way that Apple have successfully reigned (so far) in the publishing industry?

Ready: Examples

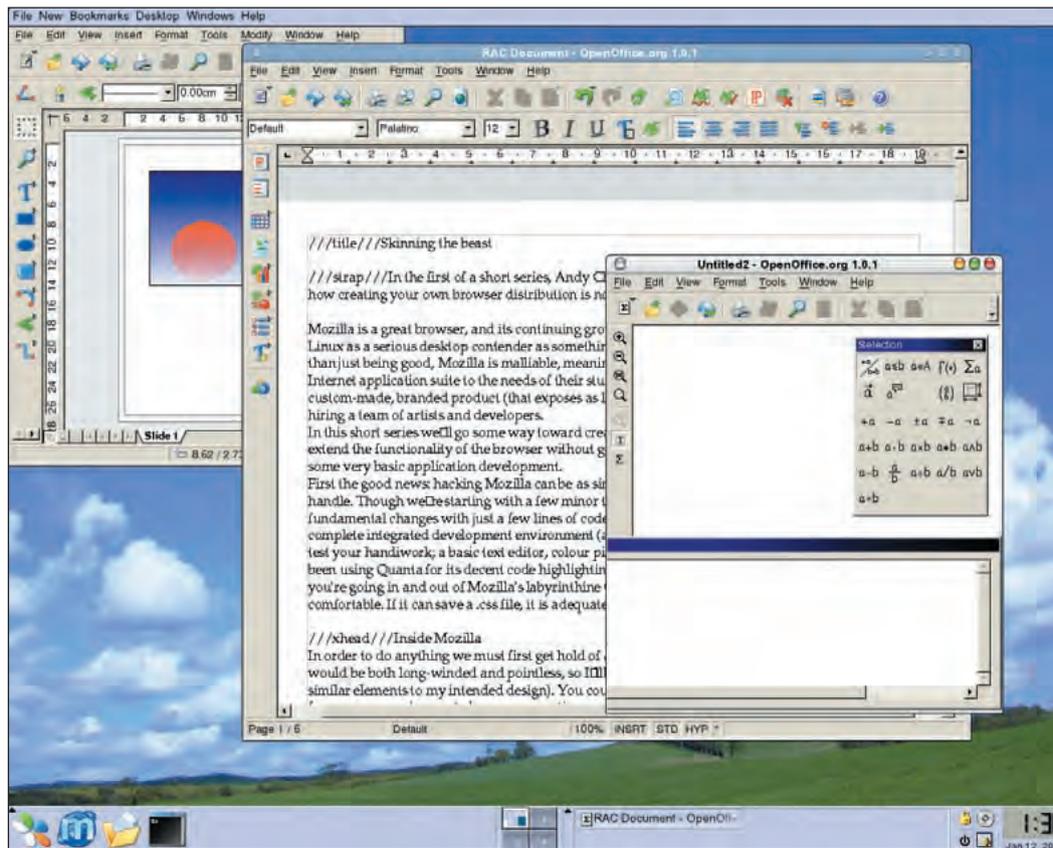
Imagine that as CTO of a fictitious insurance company, *LXF* has to sequentially update 5,000 desktop users. Mainly, their needs are basic: email, standard productivity apps (word processor, spreadsheet), controlled Internet access and the ability to query and amend our *SQL*-based customer database. The

company is keen to keep hardware and support costs to a minimum. The situation is complicated by the fact that the upgrade will take place in stages (so new desktops must be cross-compatible with legacy ones) and that company managers want to retain the existing server infrastructure which includes *Exchange 2000*.

It is the latter which draws us to Ximian’s *Desktop Professional* product which includes both *StarOffice* from Sun – to handle the millions of documents that have been created with MS applications over the company’s history – and *Connector* for *Microsoft Exchange* which allows *Evolution* to act as a standard client for *Microsoft Exchange 2000*, letting



Ready for the desktop?



OpenOffice.org provides all the standard productivity software in a free, open package.



staff share calendars, task-lists and contacts. *Evolution* will also be our default mail client. As well as attractive support options, much of the updating and patching process can be handled by Ximian's *Red Carpet Express* which automates conflict and dependency resolution, potentially reducing the need for extra admin staff. Database access can be handled via any number of apps with support options

ranging from barely there to full service. Linux is ready for this desktop.

However, if two-thirds of these users were Customer Service Advisors (CSA) and our company had invested heavily in Call Centre Management software, the whole enterprise falls apart, as solutions providers such as Rockwell have yet to port their range to Linux. In this case, Linux is not ready for the desktop, but then

neither is MacOS; Windows is your only option for now, though with both Linux and the call centre industry flourishing in places like India, it's only a matter of time (and customer pressure) before the required tools are available to make the change.

IDC's Dan Kusnetzky says when it comes to business critical applications, buyers tend to choose the software first, "Organisations have preferences in applications. If their preferred application is only available on Windows, they won't take the time to investigate the alternatives – they'll just acquire Windows to run the application," he says. One solution is coaxing an application to work under *Wine*, but, Kusnetzky says, confusion arises when it comes to support, one of the major areas of concern highlighted in a number of recent IDC polls. "It is not clear that the supplier will support the application running anywhere but Windows."

And the situation looks similar on the home front: Joan Public uses *MS Word* at work, so when it comes to choosing a home PC she notices that the application is often included in the package. As Ms Public expects to occasionally bring .DOC files home, the OS is a secondary consideration to the availability of *MS Word*.

OpenOffice.org (OOo) developer Nick Richards says that this is the place where high-profile and, more importantly, cross-platform applications such as OOo or *StarOffice* can have the biggest impact, stressing the ability to open and save common file formats. Stable and robust support for the most widespread formats is essential for the success of any current office suite but, Richards says, in the future there needs to be a shift in emphasis: "it's nice to see proprietary companies like Corel get involved in Open standards. I'm thinking particularly about the OASIS Office File Format Technical Committee." The solution, he says, is widespread adoption of Open standards across the whole computing industry. "File format compatibility has to be made irrelevant in the long term, then people can actually choose their computing platform based on quality and choice, not vendor lock-in"

Of course, not everyone is likely to be happy at the prospect of this ideal.

Mind the gaps please!

We ask some industry insiders a potentially controversial question...

LXF – Where do the most significant holes in the Linux application set lie?

"On the client side, a decent integrated DVD player would be great, *DeCSS* isn't for everybody. An easy to use groupware server with client support on the desktop would make a huge difference in business. As Apple have noticed there's no better way to get innovation on your platform than by giving out free developer tools. Linux has some damn good developer tools!"
Nick Richards, OpenOffice.org

"It seems that there are certain key applications that people need to use to do their work. This includes the

Macromedia three (*Flash*, *Director*, *Dreamweaver*) and *Adobe Photoshop*. Other than these applications, Linux seems to provide a competent alternative that is freely available... The main thing to remember though is that Linux improves every day and becomes more and more competent for its users."
Jono Bacon, KDE

"Applications in the MS Windows environment have had a few years head start. It's unlikely that the major players like Adobe and Intuit will port their software anytime soon, so independent software companies can get a head start in this market with some great tools and stake a claim. A good personal money

manager that's had some usability testing and a good UI with a strong feature set built specifically for *Desktop/LX* would sell extremely well to our userbase."

Jason Spisak, Lycoris

"Once you've covered things like web browsing, Instant Messaging, PIM, office software and personal finance (all of which we address other than browsing), you've covered a huge amount of what people use on a daily basis, but then you get into other things like an *ACDSee* kind of app, and then all those special-purpose apps and this is where easy to use development tools come in to play!"
Shawn Gordon, TheKompany.com

Dan Kusnetzky: "Controlling file formats has been one of the tools Microsoft has used to lock in its customers. While it is possible that XML will help in some cases, it won't make it any easier for all of the documents that exist today."

What you see

For many desktop users, Linux is (and will remain forever) hidden: what they see is the user interface, the applications. For the next generation of users – those who think the command line is a military term – KDE or GNOME is Linux, just as the Explorer shell is Windows and Aqua/Dock is OS-X: these are the people who will use what they're given. Is this where the real battle is to be won or lost?

Both of the main desktop environments have made great strides in the last few years, adapting and improving the Windows paradigm. To some diehards, the focus on being "more Windows than Windows" is tantamount to heresy – a race to the lowest common denominator, but it seems that the endeavours of both KDE and GNOME coders are definitely making an impact.

Jono Bacon, KDE hacker (and *LXF* contributor) says we shouldn't underestimate the contribution that KDE, GNOME and other interface builders have made to Linux on the desktop: "[They] have been fundamental in putting an interface on what is essentially a command line OS," he says. Highlighting the "distinct shift in culture" that has been quietly bringing Linux to the mainstream, Bacon says, the latest breed of users have grown up in the GUI generation. "Linux is still quite a different system to use, and much of this is below the desktop interface and involves things such as mounting disks, getting hardware to work and other things. New users from the Windows world would probably find things such as compiling a kernel and installing modules quite challenging, and in this respect, distribution vendors need to

X11 on OS-X

Install with just one click

Steve Jobs is known for using his keynote speech at AppleExpo shows to introduce new product. The iMac, Cube and OS-X were all unveiled to a rapturous reception at one show or another. This year, his presentation had a few points of interest for Linux users too, although in quite a roundabout way.

First up was the introduction of Safari, a new default web browser for OS-X that had been rumoured for a while. The unusual thing about *Safari* is that it isn't built on *Mozilla*, as was expected, but on the *KHTML* rendering engine from KDE. The second notable news from Apple

was the release of a single click installation of X11 for OS-X. The MacOS has been able to use Linux applications for a while – via *Fink* – but this is the first time an officially supported X distribution has been available.

So what does this mean for the Linux desktop? Well, it's hard to say. Mac users will most obviously gain access to the cream of Linux software, but it doesn't work the other way round. What we can hope for is that application developers take note of the numbers and, perhaps opt to create their software for X on OS-X rather than *Aqua*. A recompile for x86

Linux will almost double their potential audience. The likes of Adobe, Macromedia and Quark could easily dip their toe into a new market with little risk to their current revenues or development effort, and professional users would get another option on both hardware and software. This is more long-term than short-term thinking though, and whether these companies take that path or not remains to be seen. If they do, it will be because the market has become too big to ignore. Hey, perhaps Apple will even release a version of *Quicktime Player* for Linux...

work particularly hard at hiding this technology from those who don't want to tamper with it."

In the office, where expertise is just a call or email away(!), the needs tend to be more routine, and Linux's famed security, configurability and stability are extremely desirable. "I think that what differs with the office is that the environment is one that can make the best use of the technology, and there is less reliance on some of the multimedia aspects of Linux that are still a little behind those in the Windows world."

Massive effort has ensured that a huge range of hardware is supported by most distros, but Linux is still seen as a niche product. This won't change, says Bacon, until you can go to your local electrical store and buy a digital camera, graphics card or MP3 player with Linux drivers and instructions included with the products.

"The free software community has made great progress in developing an incredibly efficient and powerful operating system and constantly refining and improving what it can do, but we need the vendor support to push it further. The good news is that vendors are indeed listening to consumers, and the fact that Linux is moving closer to the number two slot for desktop popularity is helping promote the fact that people are using Linux and liking it."



Not all websites are as accessible as they should be.



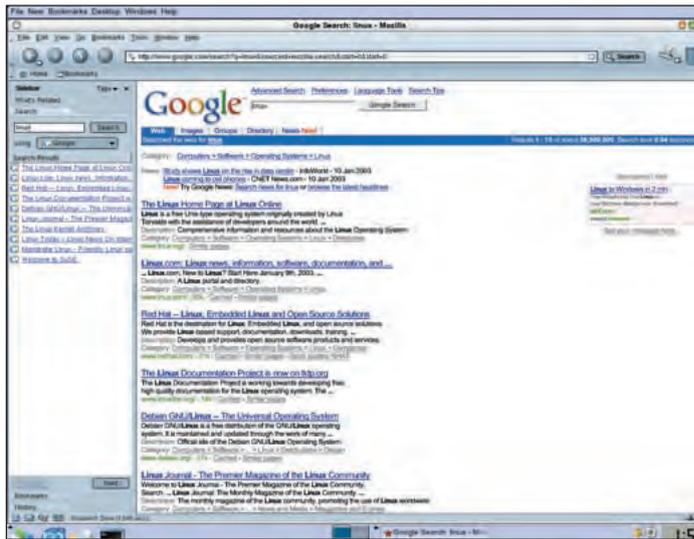
The GIMP is capable of basic image editing, but its interface is not the most intuitive design when compared to some proprietary image editors.

Nick Richards concurs, saying the support of hardware vendors is absolutely critical in making desktop Linux a reality. "Unless we can be preloaded there isn't a chance of breaking out of a small potential community into people who are less familiar with computers. I was excited to see Evesham sell their first Linux



"Vendors need to work particularly hard at hiding Linux technology from those who don't want to tamper with it."

Ready for the desktop?



« box a month or so ago and eagerly await to see who's next. The box-shifters have got to do something to improve their margins and the low-cost licensing of Linux may be the best solution for them."

However, it would be wrong to suggest that the barriers are just with the hardware, especially when it comes to home computer users where ease of use in software is more important than cramming in every feature under the sun. Video editing, PVR, educational software, DVD authoring and simple image

Mozilla's web browser is fast, innovative and accessible.

manipulation are all growing areas among consumers, and while Linux software is available to cover these areas, it tends to be sparse, difficult to acquire and install, and less than intuitive to use. Where is the equivalent of Apple's iLife suite, Microsoft's *MovieMaker* or *Photoshop Elements*? The tools to create these applications, and others, are freely available for Linux, so why is no one creating them?

The commercial developers of these type of tools tell us the Linux home user market is almost non-existent and is unlikely to grow significantly unless something can be done to raise its profile among the non-technical public. So we're back to the by now traditional Catch 22: no home users equals no applications equals no home users. Non-corporate Open Source developers, who have done much to improve the OS, are usually attracted by cutting edge projects that are going to stretch their abilities. So while *The GIMP's* next release promises a plethora of innovative new features, no-one has attempted to make a *GIMP*-based, script-driven application capable of removing red eye with a few clicks on big pastel-coloured buttons. It is regarded as a less sexy job.

The "no demand" argument is also one of the points raised by many of the ISPs we spoke to about their unwillingness to support anything but Windows and Mac. The other problem, they say, is a lack of standardisation. A Blueyonder spokesman told us that at present it would be impractical to train support personnel in the intricacies of every distribution.

"While we're happy for our users to connect their Linux boxes to our network, we don't offer any tech-support for the same reason we won't offer advice on home network problems. There are just too many permutations of things that can be influencing a problem."

Conclusion

So... is it? The answer (unsatisfactory though it may be) is: *it depends*. For many users, Linux is indeed adequate or, in some cases, ideal. There are some clouds in the otherwise sunny sky though: while Linux has a world-class selection of web browsers, it

Comment

Chris Sontag, SCO

Do you think Linux is ready for mainstream desktop use?

For certain types of desktop use, Linux has proven to be a successful dedicated purpose deployment. For example, in a call center, where several replicated desktops are all serving the same purpose, Linux is especially well suited. However, work still needs to be done before Linux is ready to be a complete replacement for the masses.

If not, why?

There are a few key obstacles that Linux must overcome before it is ready for mainstream desktop use. Clearly, more sophisticated technology users have a strong desire to use Linux. However, for the less sophisticated class of users, Linux needs to be made simpler – it needs to be easier to use and configure.

In order for Linux to become ready for mainstream desktop use, end-users must also change their attitude towards the technology. End-users must be more willing to embrace a new environment and new alternatives.

The demand for alternative desktop solutions is on the rise, however, some users are delaying upgrades because of a fear of change. After these fears are overcome, Linux will thrive.

What is the main impediment to Linux use in the Office?

The main impediment to Linux use in the office is that it is not yet 'user-friendly' enough for the average end-user. Configuration and ease of use still need to be made simpler, and the Web browsers need more work. Plus, a fully robust office suite needs to be readily available. In

addition, Linux lacks an application comparable to Outlook with full email and calendar capabilities.

Are there any major applications missing from the Linux lineup?

While there are many beneficial applications available today for Linux, it still lacks many vertical applications customised for specific business needs such as those used in dental and medical offices and in the manufacturing industry. In contrast, Unix and Windows already have many of these applications.

Do you foresee this situation changing?

Most certainly – progress on Linux is inevitable. It is continuing to improve and develop and is becoming a more viable solution, fully capable of competing with existing mainstream desktop platforms.

Comment

Frederick Bastok, MandrakeSoft

Do you think Linux is ready for mainstream desktop use?

For the home users, it truly depends on what the user does with their computer. If they type some text documents, manage personal finance, listen to MP3 or surf the Internet, Linux is ready. But if they want to play games or use education programs, they will still have to use Windows, and a dual-boot is seen by some as undesirable.

For the corporate world, Linux can be used as far as users don't need any vertical applications. We still lack some specific applications (like *AutoCAD* for CAD users). As most users only need office suite software and Internet access, Linux is suitable.

If not, why?

We lack applications. The problem is no longer the perceived difficulty of Linux, thanks to KDE and GNOME. But we need games for the home users and vertical applications for the corporate users.

What is the main impediment to Linux use in the Office?

Vertical applications.

What would it take for Linux to be 'Ready'?

We rely on the publishers that make those applications. I'm confident that this will change in 2003. Several large accounts are evaluating Linux on the

desktop and will migrate soon.

Are there any major applications missing from the Linux lineup?

AutoCAD or *CATIA* for CAD users, specific applications for logistic and transport, accounting programs (like *SAGE*) for SMB.

Do you foresee this situation changing?

Yes, I do. Everyday, we see that companies are evaluating Linux on the desktop. Consequently, they put pressure on the software publisher to port the applications they need to Linux (many of them already have a Unix version so that's not a big deal). I do believe that things are going to change in 2003.

Key applications

Standard applications suitable for any Linux desktop

Last year saw the launch of three distributions aimed squarely at the desktop market, as well as desk-centric updates from the likes of SuSE, Mandrake and RedHat. Xandros, Lycoris and Lindows.com have taken a different approach to the '2000 applications plus every conceivable server' distributions we're used to seeing. Instead, concentrating on making the user experience as pleasant and simple as possible and paring down the number of applications to whatever is needed to do the basics.

All three, to an extent, have attempted to woo the Windows user with distinctly MS-styled user interfaces and installation routines that even a chimp could manage. In the case of *Xandros*, you can even install a copy of *Microsoft Office* 'out of the box'. The others have stressed ease-of-use and broad file compatibility. In addition to tailored distros, the advent or maturity of a handful of applications is also helping to convince users and enterprise buyers that there is another way. Next-generation users raised on P2P and free (in the illegal sense) software will also increasingly influence things.

MOZILLA

The *Mozilla* project rose from the ashes of Netscape after it was eviscerated in the browser war by Microsoft's decision to integrate *Internet Explorer* into Windows. The *Mozilla* suite is intended as a development tool for creating browser distributions and other applications. It is built around a number of web standards including XML, CSS and DOM, and features an extremely swift rendering engine called *Gecko*. There's a *Mozilla* browser available for just about every operating system you could think of.

www.mozilla.org

STAR OFFICE

Sun purchased an also-ran office suite and first started giving the application away before finally donating most of the codebase to the open source community. The result was *OpenOffice.org*, which then formed the basis for Sun's own *StarOffice*. Comprising word processing, spreadsheets, presentation and illustration tools, as well as a pretty good HTML editor, *Star/OpenOffice* triumphed in our recent office suite roundup. Like *Mozilla*, this productivity suite is available

on all major architectures, making it an excellent balwark for companies investigating a wholesale migration from Windows to Linux.

www.openoffice.org

KDE 3

With a little tweaking, KDE3 is the prettiest GUI on earth, but it also boasts a fantastic file/web browser, a good selection of integrated applications and the ability to change almost every aspect of the interface with just a few clicks. Mandrake, Lycoris, Lindows and Xandros all favour KDE.

GNOME

GNOME seems to be playing second fiddle to KDE at the moment (protesting letters to the usual address), but is still no slouch when it comes to day-to-day work. It is the default environment for Ximian and Red Hat, though both have made major changes to both the look and feel to fit in with their own branding.

CROSSOVER

Codeweavers hacked their way into the *Wine* project, and released packages

enabling users to install standard Windows browser plugins, including *Quicktime*, into *Mozilla* or *Konqueror* and run *Microsoft Office* without having to dual boot into Windows. Rumours abound that the Codeweavers's next project will let Linux users run *Photoshop* and/or *Dreamweaver*.

www.codeweavers.com

EVOLUTION

Evolution is Ximian's flagship product, aiming to bring groupware capabilities to the corporate Linux desktop. Many users have expressed their dismay at the sheer 'Outlookness' of the application, but that could be regarded as a good thing if your target user has only ever used *Outlook*.

www.ximian.com

THE GIMP

The Gnu Image Manipulation Project brings pro-level image editing to the Linux desktop. Extendible, intuitive and powerful, everybody loves *The GIMP*. Except people who've spent their working life using *Adobe Photoshop*. And Adobe aren't all smiles either.

www.gimp.org

means nothing if the likes of Natwest or Tesco refuse to recognise the existence of anything but IE on Windows; *The GIMP* is an incredible piece of software, but it's winning in a field of one; we want to be able to watch *Spinal Tap* on DVD without the risk of copyright police hammering down our front doors.

These small niggles, and others, are the 'but' in our answer. For

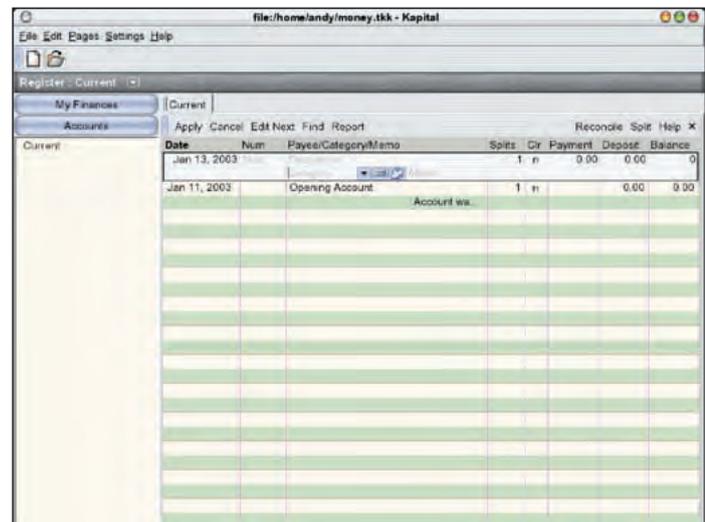
“For corporate users, there is a problem of support which, whether perceived or real, needs to be addressed.”

corporate users there is a problem of support which, whether perceived or real, needs to be addressed. Users can, for instance, get bespoke support from the likes of SuSE or Red Hat – that is, after all, their primary revenue

maker – and patches or security problems are usually resolved quickly, but the perception of an 'ad hoc' support structure remains. A spokeswoman from Rockwell FirstPoint Contact, for example, told us



Blender brings 3D design and rendering to the masses – once the masses can get past the user interface.



Personal finance is one of the areas where Linux is lacking. *Kapital* should change that, and make handling your overdraft a little more transparent.

Ready for the desktop?

Xandros offers users a very Microsoft Windows-like desktop experience... as does LindowsOS.



that though the company hadn't ruled out Linux (and the move towards web services and VoIP was making 'the OS' less relevant), customers weren't convinced of its suitability.

"One of the concerns that vendors have is the level, quality and availability of technical support for Linux. Technical resilience and zero downtime is more important for the

call centre than for other areas of business. Generally when you are relying on open source developers for fixes, the wait can be longer than chasing a company that develops its own operating system," she said. This is a question of perception, and the entry into the market of respected names such as Sun will have a major impact on buyers' opinions.

Now the good news: things are changing. After a concerted campaign by both Linux and Mac users it is possible to fill, file and even pay your tax return (and access other government services) online; pre-installed Linux systems are a growth area for one of the world's largest retailers; a generation of computer science students – the SysAdmins of the future – are being given a choice; governments and corporations driven by cost efficiency are seriously investigating alternatives to Windows – which will in turn expose millions of new users to something different; Even Microsoft are making their WMP codecs and DRM technologies

available for developers interested in integrating them into Linux software.

Jill Ratkevic, of advocacy/news site www.desktoplinux.org, says the level of interest in Linux has been growing steadily over the last 12 months, and not just for governments hoping to get the best value from their tax revenues. "Recently I was speaking with the IT director of a Fortune 500 company that does business annually in the US\$25 billion dollar range and they are looking at Linux for desktops." Like many others we spoke to, Ratkevic says that the most compelling reasons for this willingness to look at all the options are improvements in the user interface, the seemingly belligerent licensing changes on the Microsoft product line and their decision to timetable the breaking of support for previous operating system and office software versions. "We have experienced a tremendous amount of businesses, consumers, and educators interested in Desktop Linux in 2002. I expect 2003 to be a real breakout year" [LXF](#)

Application map – The view from here

How popular Linux software measures up to common tasks performed with Closed Source competitors

CATEGORY	APPLICATIONS	SUITABILITY	COMMENT
PRODUCTIVITY	<i>StarOffice, OpenOffice, Koffice</i>	Good	Excellent, inexpensive productivity suites capable of most day-to-day tasks. The only problem may be opening very complex or graphics-heavy <i>MS Office</i> files.
GRAPHICS	<i>The GIMP</i>	Good	<i>The GIMP</i> is a good choice for pro-level image editing, with an extensive range of tools and plugins available. Its scripting abilities makes it the ideal choice for batch processing. Potential problems include the non-standard interface and lack of CMYK colour space.
AUDIO	<i>XMMS, noatun, Kaboodle, RealPlayer</i>	Excellent	A wide range of players take care of your MP3s, Ogg or Streaming media. No <i>WMP</i> (as yet) or <i>Quicktime</i> .
VIDEO	<i>Xine, Ogle</i>	OK	Video players are capable and attractive, but a lack of a official CSS licence mean you must resort to potentially criminal behaviour to watch <i>Spider-man</i> . Absurd. Video editing is available though it's a complex job with something like <i>Broadcast 2000</i> .
CD BURNING	<i>XCDRoast, Gcombust, CDRRecord</i>	Average	No problems with the functionality of the software, though there are some ease of use issues. DVD burning needs to be radically simplified.
DTP	<i>Scribus, Sketch, TeX</i>	OK	Scribus looks promising, but still needs a lot of work to compete with the like of <i>QuarkXPress</i> or Adobe InDesign. Lower-end users may be better off sticking with OOo.
DIGITAL CAMERA	<i>gphoto, GIMP</i>	Average	Software for managing, altering images okay, but plug and play support for USB mass storage devices still hit and miss in many distributions.
3D	<i>Blender, SoftImage, Renderman, Maya</i>	Excellent	Linux is the number one desktop operating system in the film/3D industry and many pro-level tools are available. <i>Blender</i> was recently relicenced under the GPL.
GAMES	<i>Quake, Doom III, Creatures, Return To Castle Wolfenstein, UT2003</i>	Poor	Despite having an excellent implementation of <i>OpenGL</i> , Linux suffers from lack of profile in the games arena. <i>WINEX</i> from TransGaming brings many current releases to Linux.
BROWSERS	<i>Mozilla, Opera, Konqueror, etc.</i>	Excellent	An adherence to W3C standards makes for a robust trio of browsers, but support from web developers weaned on the IE breast needs developing.
MAIL CLIENTS	<i>Evolution, KMail, Mozilla</i>	Excellent	They send, they receive, they all do HTML mail.
MUSIC	<i>Audacity, Rosegarden, Broadcast 2000</i>	Good	Music software has come on in leaps and bounds recently, and Linux now makes sense for small project studios.
WEBSITE DESIGN	<i>Bluefish, Quanta, Mozilla, Amaya</i>	Good	Though not in the same league as <i>Dreamweaver</i> , these applications offer a useful set of tools across the spectrum of abilities from <i>Composer's</i> simplicity to <i>Quanta's</i> CSS skills.

The changing face of

FONTS

Marco Fioretti introduces the new font management system that will soon be common to all GNU/Linux platforms

Some of the GNU/Linux font components have been rewritten from scratch, and the newest distributions are finally deploying them all over the desktop. This may seem just another spoonful of eye candy, but there are many reasons why this is a Good Thing (TM) for both end users and developers. Let's look at what is actually changed, and how.

The old font system

The font management system which is being replaced is located on the server side. The graphic display server, XFree86, can use the fonts available in the directories specified with the `FontPath` keyword: whenever an application needs to draw a character, the server fetches the corresponding glyph from that directory, and does the actual drawing.

If there are several fonts directories, they can be declared with the corresponding number of `FontPath` entries: they will be searched in the order in which they were declared. This is a completely static system, which can be faster, and easily maintainable, on older computers with little RAM, slower processors and a small number of fonts (see Linux Format 34 Dec 2002 for a roundup of mini-distros ideal for running on 'heritage' hardware).

The next step in this direction was to add support for centralised font management. There are programs that can serve fonts both to the computer actually running them, and to the other hosts on the local network. The

first font server was `xfs`, followed by `xfsft` and `xfstt`, which were developed to support TrueType fonts. This didn't necessarily mean that you needed more than one server to deal with all fonts. Red Hat, for example, ships a version of `xfs` which is patched to also support TrueType.

The advantage here is that the system administrator must only install new fonts, and/or upgrade the server, on one single machine. In practice, probably to simplify packaging and installation scripts, the font server is set up by most distros even on stand alone machines. In any case, such servers can peacefully coexist with statically declared font paths.

As a matter of fact, it is a standard technique to use both, declaring the server first, and, in case that fails, one or more local directories. An old (pre-2003) XFree86config file could then have a font section like this:

```
Section "Files"
FontPath "unix:/7100"
FontPath "/usr/X11R6/lib/X11/fonts/"
EndSection
```

The first line says to ask for fonts to the server listening on port 7100, the second to search for them, if still needed, in the specified directory.

The drawbacks

This architecture, which is still the default on almost all distributions released through 2002, started to fall behind the expectations of users and developers some years ago.

One first problem is security. An ill-configured font server may answer to `/any/` remote request coming to port

7100, practically begging for a denial of service attack. The counter-measures are both global, ie the properly configured firewall that everybody should have, and specific to `xfs`: it can be configured to not listen to remote TCP connections and/or to only answer to trusted clients. Apart from security however, there are also, and much more frequently, problems in the flexibility, performance, and internationalisation areas.

For example, the fonts must be named *exactly*, or the `xfs` server will not find them. What is much worst, however (as far as font managers are concerned, of course!) is the fact that the GNU/Linux desktop has made giant steps in recent times: more exactly, applications and desktop environments and even operating systems have evolved much faster than the graphic core system. Delegating all font management to an external server has the result that users and developers of specific applications must wait until the **server** has been upgraded to satisfy their specific font needs.

For this reason many developers just gave up and reinvented the wheel: being on their own, they also came, unavoidably, to slightly different wheels. Even not considering KDE and GNOME, there are now several applications who want to install and configure their own font stuff, from *Mozilla* to *TeX* to *Wine*, and then make any sort of things with them. We all know what the end result is. You managed to display correctly a non-ASCII web page? Good: now try on a

different browser, then save it as text, and open it with `cat`, `Emacs`, or `OpenOffice.org`. Oh, and since you already have all these windows open on the same file, try to **print** it from each program.

Of course, all this could already be fixed in the past, but only with a series of unrelated voodoo hacks: it really should not be that hard.

Enter `xft2` and `fontconfig`

How does the new system solve these problems? Synthesising (extremely) it just moves the detection of the font, and the generation of the corresponding images, from the server to each graphical client. In this way, the latter is finally free to do whatever it likes with fonts. Of course, if this were all, it would just tangle the current mess further, and beyond repair. The real news is that now there are standard tools and libraries to do it, so that once *they* have been polished, and widely deployed, every application will be able to use them with the same results, and configuration will happen once, regardless of the method used.

Figure 1 on the right shows a simplified view of the new architecture, and how its parts relate to each other. The new components are two new libraries written by X guru Keith Packard, `fontconfig` and `xft2`, placed *inside* each generic application. I will explain in detail their exact purpose and internal behavior later on in this article, but let us look at the bigger picture first.

Figure 2: **Stair-step effect in text rendered normally**

As it will be explained later, *fontconfig* tells which fonts are available, their characteristics and how to use them to the application core logic, which actually decides what should be drawn on the screen. Once this decision is taken it is communicated to the other new library, *xft2*.

This second piece acts as a layer of glue logic between freetype and the X server. The result is that what to draw is found out locally, and the X server, which is the one still in full control of the actual screen, is then told what to do. This can happen quite quickly with XFree86 version 4.0 or higher; with older servers the end result on the display is the same, just slower. Let us now look at these new objects in more detail.

The X Render extension, designed quite recently, adds several capabilities to the X protocol, like alpha blended

image drawing and anti-aliasing support. Its main feature, as far as we are concerned here, is that it also makes possible accelerated communication of font drawing instructions from *xft2*-enabled clients.

Server extensions

This happens, as already mentioned, by driving the actual font rasterizing engine, and feeding its results to the X server, in the faster form that the latter can support. To easily find out if your server supports this extension, just type the following:

```
#xdpyinfo|more
```

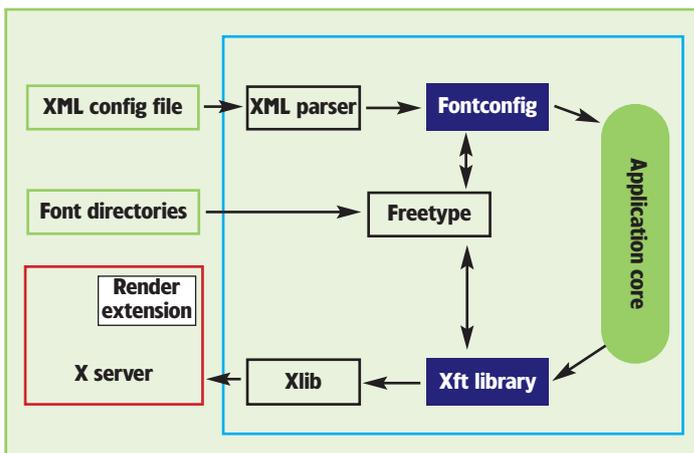
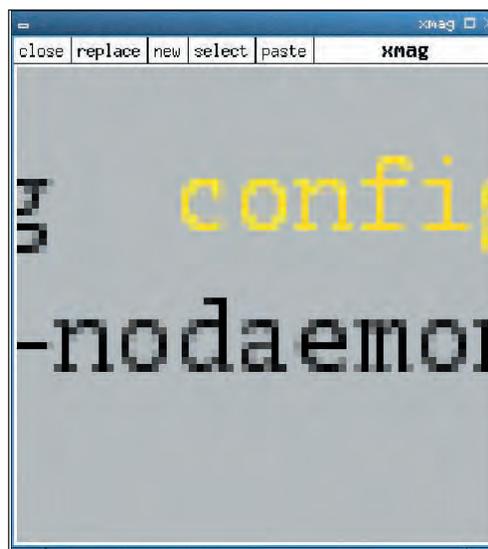
at a shell prompt, and check to see whether RENDER is listed in the extensions section.

For every character to be displayed, the freetype library can fetch the corresponding glyph from the font files. This happens inside the application, remember: previously,

ie using a font server, each program had to ask to the X server which fonts were available, and X answered either directly, or passing the request to the font server. It was still the font server which eventually said to X what was to be displayed.

Keep in mind that in this way an application is enabled to do much more than just find out *which* glyph should be drawn: the modern font files, now directly accessible, also contain a lot of other information, from the exact size of the glyph to detailed positioning information for advanced layout. Two features widely used by some languages, for example, are the so-called Glyph Positioning Table (GPOS), and Glyph Substitution Table (GSUB). The first allows to control exactly, and in accordance with the language system currently used, the placement of each glyph. The second knows how to substitute one glyph to another.

Once the application, thanks to these libraries, has finally decided what should be drawn, where and how, *xft2* has two choices: if the X server supports the Render extension, this one will be used to communicate quickly what to show on the display. If the extension is not available, *xft2* won't quit, just slow down. The obsolete server will receive much longer sequences of low-level graphics instructions at the xlib level, and execute them, with the same final result (xlib is the lowest level X API, providing calls like XDrawArc or XDrawLines).

Figure 1: **Client-side font architecture**Figure 3: **The same text, antialiased**

Font buzzwords

A micro glossary

GLYPH: the actual picture to be drawn on screen or paper to represent a character.

HINTING: description of the appropriate ways to modify a glyph for proper rasterizing at small sizes.

OPENTYPE: a font specification supporting advanced internationalization and typographic features. For more explanation, see: <http://freetype.sourceforge.net/opentype/>

POSTSCRIPT: an object-oriented language, developed by Adobe Systems, used primarily to produce vector graphics of images and fonts.

RASTERIZE: to represent an image with a matrix of properly coloured dots. Raster graphics are also called bit-mapped.

SCRIPT: a set of characters covering a particular writing system, such as Cyrillic or Latin.

SERIF/SANS SERIF: serif typefaces use several decorative marks to look nicer and easier to read. Sans serif typefaces are composed of much simpler lines.

TRUETYPE: scalable font technology introduced in 1991 (also a registered trademark of Apple Computer, Inc.)

TYPEFACE: A design style for a set of characters, like Times or Helvetica. The two broader categories of typefaces are serif and sans serif.

VECTOR GRAPHICS: the other way to represent images: it uses mathematical formulas that define all the lines in the image. Unlike raster graphics, vector graphics look the same, not worst, even when scaled to different sizes. Fonts represented in this way are called scalable fonts.

The advantages of this compatibility with older display servers may not be evident at first sight, but it would be really unfair to dismiss them. Having an application able to do its own nice fonts regardless of the server is not just eye candy: it can be a life-saver for all users with eyesight problems, or dealing frequently with formulas or other unusual things. It also saves money: sometimes



FontManagement

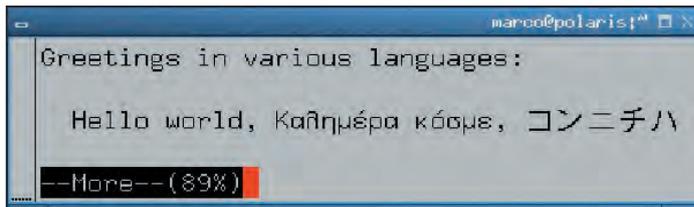


Figure 4: A classic: “Hello world”, UTF-8 style

◀ upgrading the X server to obtain better fonts means a new graphic card which means a new motherboard which means... you get the picture!

This application is best illustrated considering school or university labs – students with obsolete terminals can run any algebra or other font-intensive application on a server and still see what they are doing.

Fontconfig

The most interesting piece for the end user, and also for several developers, is probably the other half of the

picture, *fontconfig*. This independent library has nothing to do with rendering, but can be embedded in *any* application which has the need to locate, select and configure fonts. While *xft2* cannot work without the information made available by *fontconfig*, the reverse is not true. *fontconfig* is usable even by non-graphic applications, and doesn't require the X server at all. This means that print drivers, for example, can finally use the same font setup used by desktop clients: for the end user, this is a big step toward a really WYSIWIP desktop (What You See Is What Is Printed).

Example fontconfig Commands

```

<!-- Font directory list
<dir>/usr/X11R6/lib/X11/fonts/Type1
</dir>
<dir>/usr/share/fonts</dir>
<dir>~/.fonts</dir>

<!-- Rescan configuration every 30
seconds -->
<rescan>
<int>30</int>
</rescan>

<!-- Accept deprecated 'mono' alias,
replacing it with 'monospace'
-->
<match target="pattern">
<test qual="any" name="family">
<string>mono</string>
</test>
<edit name="family"
mode="assign">
<string>monospace</string>
</edit>
</match>

<!-- Alias well known font names to
available TrueType fonts -->
<alias>
<family>Times</family>
<accept><family>Times New
Roman</family></accept>
</alias>

```

```

<!-- If the font still has no generic
name, add sans-serif -->
<match target="pattern">
<test qual="all" name="family"
compare="not_eq">
<string>sans-serif</string>
</test>
<edit name="family"
mode="append_last">
<string>sans-serif</string>
</edit>
</match>

<!-- Load per-user customization
file -->
<include ignore_missing="yes">
~/.fonts.conf</include>

```

Through a raster library (freetype in *figure 1*) *fontconfig* loads the characteristics of the available fonts. It also loads, via a parser like the *expat* library, its own settings, stored in one or more XML files: all the font selection is performed according to these instructions. The Example *fontconfig* Commands code listing on the left shows some actual setup commands: a real configuration file would be much longer, but the example still helps to see the potential of this tool.

The first thing to notice is that both system and user level font directories are supported: furthermore, they are checked periodically, and the results are maintained in a cache: adding new fonts to the system becomes as simple as copying them in one of those directories, and let the system find them by itself! The font files can even be symbolic links, or (in future versions) be grouped in subdirectories.

The several fonts in the example are named, referred to, and matched to the application requests using their actual attributes, like monospace, serif, and so on. This is done in such a way to guarantee that every time, the application gets the font most closely matching its current expectations. For example, names and attributes can be internally changed to standard ones, or aliased to substitute similar fonts: in the example, “mono” is changed to “monospace” and “Times New Roman” is replaced by “Times”. *fontconfig* can even be told to make specific assumptions on some not specified attributes (e.g. the “sans-serif” test).

There are a lot of other checks and instructions like this, all in a

format which can be easily edited by hand or by any third-party configuration utility. This is what makes it possible to say once and for all, to every program using *fontconfig*, “Use this, and if it isn't available use that”, or “these are all and the only fonts which completely support your language”.

Last but not least, additional configuration files can be included, as shown in the final statement. This enables each user, even without the root password, to add and setup fonts at will, safely trying different settings: just remember to back up the original `~/.fonts.conf` file!

This comes at a price, of course. The algorithm is so smart and flexible because it creates a cache of all the information parsed from the configuration file, and this costs memory. However, the advantages already listed are well worth it for most desktop users, and it gets even better for people who routinely deal with non ASCII characters or non alphabetic languages, often mixing them in the same document.

Such cases may bring the traditional system to its knees, and are quite frequent too, even in Western countries. Imagine a generic program which has already managed to display or print half a page of English text, and all of a sudden is faced with a complex equation, or some Asian name or trademark; what about all the spreadsheets, price lists, and the many other documents where Euro, Pound, Dollar, and many other currency symbols are expected to live side by side?

When this or similar things happen, *fontconfig* can list to the applications all the fonts which contain the desired glyph and are similar, in size and other attributes, to the one used until that point: there is no need to know it in advance.

For these reasons, the new system is also a natural match for UTF-8 (see *boxout*), and what finally makes it widely usable. UTF-8 is a binary character encoding which covers all existing languages and then some, with the drawback that one never knows if a new font may be needed to display the next character in a file. At the same time, there is practically no font covering the whole range of

Unicode/UTF-8

One world, one character set

Unicode is an universal character set, introduced to support all alphabets and new symbols, like the Euro, without reconfiguring each application, each time. The Unicode format more popular right now, and used by default in Red Hat 8.0, is UTF-8, which uses from 1 to 4 bytes for each character, depending from the language: its main advantages on other Unicode formats are that it doesn't require changes to many text processing programs, or extra disk space to store US ASCII or European (ISO-8859-1) text files. UTF-8 codes all non-ASCII characters in more than one byte, and each byte in such sequences has the highest bit set to 1, to avoid confusion. The main Unicode resources are www.unicode.org and the corresponding Linux HOWTO at www.tldp.org/HOWTO/Unicode-HOWTO.html

Antialiasing

Antialiasing (also called oversampling) is a technique to smooth lines which, because of resolution limits, would otherwise look like stairsteps. The effect is achieved surrounding the stairsteps with pixels of intermediate colours. Figures 2 and 3 show how this happens in practice

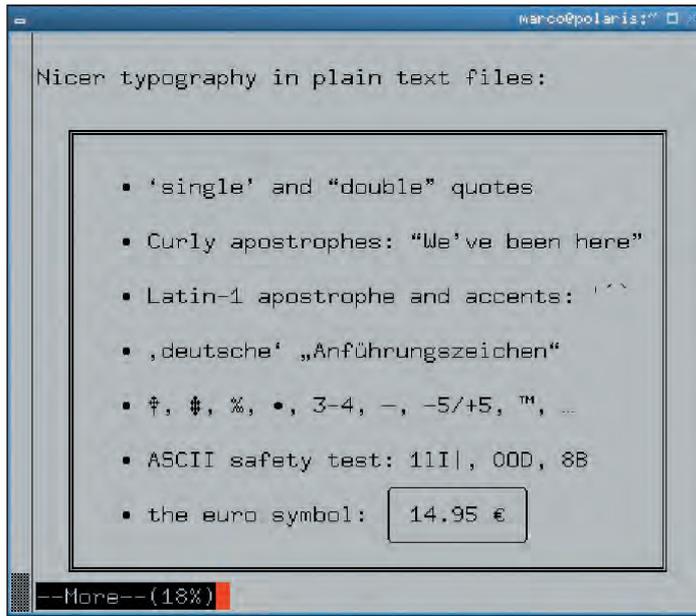


Figure 5: UTF-8 makes even plain text nicer

Unicode characters today, and this is not likely to change. However, since *fontconfig* already has information about Unicode coverage for each font, the inverse approach becomes possible, at least in principle: the application declares all the characters it needs to display, and *fontconfig* can return the font or fonts which have all the relevant characters.

To check if and how your GNU/Linux system supports Unicode and advanced font rendering, download the file <http://www.cl.cam.ac.uk/~mgk25/ucs/examples/UTF-8-demo.txt>, and display its output in a shell with **more**. Figures 4, 5 and 6 show how several parts of this file are correctly rendered in Red Hat 8.0 in a console started with the following command:

```
#uxterm -fn -misc-fixed-medium-r-normal--18-120-100-100-c-90-iso10646-1
```

uxterm is a simple shell script provided by Red Hat: it just checks the current values of the `LC_ALL`, `LC_CTYPE`, and `LANG` variables, and then starts an *xterm* with the proper encoding.

What is the status?

At the time of this writing (December 2002) the deployment of applications using these libraries in mainstream, prepackaged distributions is just started with Red Hat 8.0. Of course, developers and adventurous users

were already experimenting for several months. By the time you read this, there will have been certainly many bug fixes, as well as packaging and performance improvements, both in single applications and newer distributions. Keep up with them by regularly updating your software.

Several existing programs will benefit from being ported to client side fonts also at the maintenance level: if the program was already rendering fonts by itself in some “proprietary” way, replacing that code with a standard library will eventually make life easier for its developers.

Inside *fontconfig*, maybe the most promising development is the announced modification to the font matching algorithm to make it work as much as possible like the CSS2 (Cascading Style Sheets 2) standard used in the latest web browsers: this should be another important step toward really consistent document rendering, whatever the application that is used to do it.

In parallel to the mainstream development, there are also people trying alternative approaches, like the *Xft* (and *FreeType*) Hack at http://www.cs.mcgill.ca/~dchest/xft_hack/.

The toolkits around which KDE and GNOME are built, *GTK* and *QT*, are still working to integrate the libraries and debug the result (*Qt* seems to be behind right now, but this

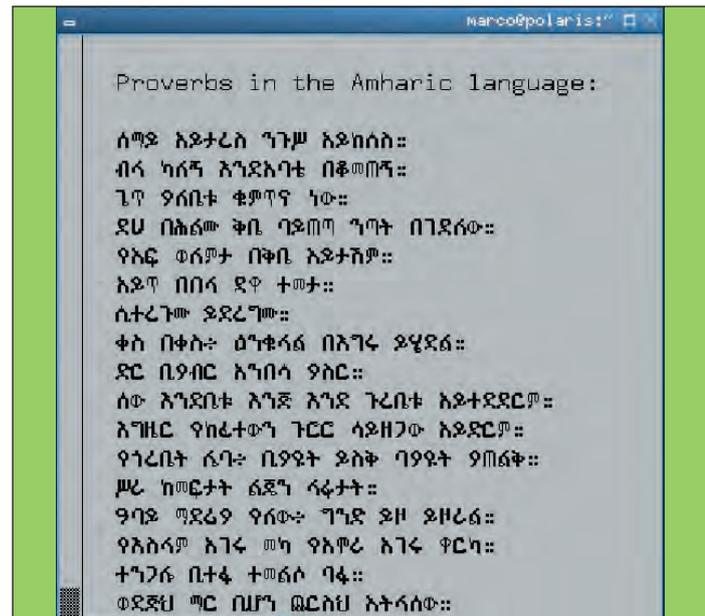


Figure 6: One xterm for every language

will certainly change soon): the same applies to multi-language layout engines such as *Pango*

(www.pango.org). In the printing arena, the *libgnomeprint22* library, which uses *fontconfig*, is also under active development.

On the desktop there are still major applications like *Mozilla* and *OpenOffice.org* which have not yet converted to the new system. More exactly, *OpenOffice.org* plans to move to *fontconfig* in a major release after 1.1, while there are some semi-experimental versions of the browser patched to use *xft2*: check the *Mozilla* home page for details. Of course, since both these programs are cross-platform, their developers face a much more complex situation, so some patience is due.

In short, Linux desktops mostly based on *xft2* and friends are already available and quite usable, but don't expect them to be 100% clean, reliable and as fast as possible for a little while yet.

Credits

The credit for making these tools available goes to the *freetype* team, and, for *xft2* and *fontconfig*, to Keith Packard. As far as this article is concerned, any error it may contain is my responsibility only. I am grateful to Keith, but also to Havoc Pennington and to all the other folks on the Red Hat 8 (Psyche) list, for explaining to

References

Xft2/*fontconfig* Home page
<http://www.fontconfig.org>

Font Configuration and Customization for Open Source Systems
<http://www.xfree86.org/~keithp/talks/guadec2002/html/>

The Xft Font Library: Architecture and Users Guide (partly obsolete, but interesting)
<http://www.xfree86.org/~keithp/talks/xtc2001/paper/>

Freetype
<http://www.freetype.org>

XFree86 Font De-uglification HOWTO
<http://www.tldp.org/HOWTO/mini/FDU/index.html>

Adding Fonts to *XFree86*
<http://www.halley.cc/ed/linux/howto/fonts.html>

me how *xft2* and *fontconfig* work, and why: I encourage you all to help Keith, the other developers and the packagers of your preferred GNU/Linux distribution by testing the libraries in as many applications as possible, and reporting any bug or requests to the proper forums, i.e. the users groups of each specific distribution or program (there is also a *fonts on XFree86* mailing list at <http://www.xfree86.org/mailman/listinfo/fonts>). [LXF](#)

Build your own Linux PC BACKUP SERVER

To show that even those that live in ivory towers need to fix the plumbing occasionally, Professor **Les Hatton** of the UKC Computing Laboratory dirties his hands and building a cheap Network backup PC from scratch.

Old enough to recall the cost of proprietary mainframes and their national-debt sized maintenance contracts? The PC built here will comfortably out-perform a Cray super-computer I deployed fifteen years ago on behalf of a previous employer for approximately 0.001 % of the cost.

The first reaction of most people curious enough to take the side off a PC is to put it back on again quickly. However, in reality, there are very few pieces, they fit together in mostly logical ways and it's no more complex than mending a push-bike (though a lot cleaner!) In the first half of this article, I describe exactly what to do and what to watch out for those like myself who are dexterously challenged. Once you have built one, you also realise that they aren't that hard to fix either when they go wrong!

Building & backing up

When you've built your PC and got it up and running, I'd like to address the subject of backup, a vital subject for home and small business users. Although modern hard discs are exceptionally reliable (about 10,000 times more reliable than 20 years ago), they still fail occasionally, so you can't rely on one hard disk as the sole repository of your valuable data.

One of the standard answers to this over the years since we outgrew floppies has been to use some kind of attached tape device, for example, DAT, Exabyte, Onstream or whatever. The idea was to keep multiple copies

of your data on simple cartridges in these formats and various schemes can be used. For a long time, I used three rolling backup copies and an archive copy kept at another site. The principle is simple: always have at least two recent copies of your data, since the odds of losing both of them is really very small. Keeping another copy offsite protects you against disasters like fire and floods. I have to say though that the quality of this sort of media can be quite disappointing. For example: of seven Onstream cartridges used in a backup cycle in the last year, two gave write errors after less than 10 uses – inexcusable at nearly £30 a cartridge. Exabyte cartridges have also given problems when used on drives other than the one on which they were written and of course most tape drives will set you back £500 upwards. In contrast, hard drives are stunningly reliable and generally of comparable cost to the cartridges, let alone the tape drives required to use them. They are also very much quicker than any tape drive, therefore raising the possibility of backing up your data on spare hard drives in a purpose-built PC.

Building the PC itself

To actually build a PC, all I have used are a small cross-head screwdriver, preferably slightly magnetic to pick up rogue screws (do we need to tell you not to wave it around near floppies?), a standard flat-head screwdriver with a fairly small head and some insulating tape to tidy up the wiring.

Also, you need good light to work in; and a magnifying glass to read the motherboard hieroglyphics.

You can order parts for PCs by mail-order from the large numbers of suppliers found in any PC mag. They are generally rather cheaper to buy at a computer fair though – that's what I do. Be a little bit careful at computer fairs as some dealers are rather more opaque than others. Wander round the fair checking prices first (they are usually pretty close), and only deal with suppliers who give you a decent-looking invoice. I would recommend going to a fair a couple of times before you buy anything expensive to get a feel for which suppliers are regulars and choose one that is geographically pretty close to you should something you buy not work. Having said this, I have had no problems with the suppliers I now use and they all tend to sell more-or-less

the same stuff anyway, and are even generally clued up about Linux compatibility these days. Showing up early gets you the widest choice – so let's go shopping!

The bits

First of all, you will need a case. These vary from around £20 for your entry-level ATX to £200 or more for those looking for something with deep blue fluorescent lights to bathe their anorak in a sultry glow. It's a false economy to opt for the very cheapest, as the cases may flex a bit and this can lead to problems with PCI card seating whenever you pick it up. My last case was like this, I tried to shave a few pounds off my budget but ended up throwing it away. For about £30 you can get a good case with a couple of fans already installed and most importantly a power supply. These are usually 300W or 350W which is more than adequate. Also, make sure you choose a case big enough for a standard ATX format motherboard. Avoid small cases like the plague unless your day job is a watch maker. >>

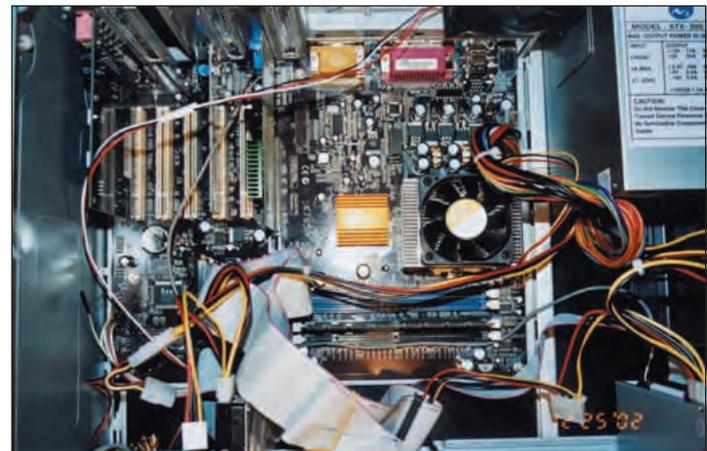


Figure 1. The fan sitting on top of the processor.

BackupBox

« You will thank me for this advice. If your case hasn't got any fans installed, you will have to buy one or two. They are only about £3 (don't forget to buy the fitting screws at the same time) and a cool computer is a happy computer. Note that the case should contain all the tiny fitting screws you will need, (except for the fans).

Next, you will need an ATX format motherboard. There are many kinds of motherboard. For a mirror server such as we are building, you don't need anything fancy. I've used Gigabyte, ECS and Jetway motherboards satisfactorily. Note that some motherboards are designed for Intel chips and some for AMD. I much prefer AMD products as being cheaper and faster than their Intel equivalents, but you may have a different preference. For this project, I used an ECS K7S5A motherboard suitable for Socket A (Athlon and Duron) processors. This supports either SDRAM or DDR memory, (DDR is faster but a bit more expensive – I'll come back to this) and has a AGP graphics card slot and 5 PCI slots. Look out for deals. Most suppliers will do a deal consisting of a motherboard, processor and fan of your choice and for a few pounds off marked price. (Note that the processor fan is an absolute MUST and is not to be confused with the case fan I discussed above). You can see it sitting on top of the processor in **figure 1**.

I went a little over the top and got a K7S5A + Athlon XP1900+ + fan + 256MB SDRAM for £150, but I notice that this price has already gone down nearly £10 since. Dammit! An Athlon 1.2 GHZ would shave quite a few additional pounds off to allow you to get DDR memory, which is likely to have a much bigger effect on your system performance (most modern machines are memory access time-limited not processor-limited). Now a word of advice. Fitting the processor fan is the single most nerve-wracking moment in building a PC. Most suppliers at computer fairs will fit both the processor and the fan if you ask – it's a false economy not to. I built four PCs before I dared to do it myself – it's your choice and you have been warned! If you decide to do it yourself, you will need a flathead screwdriver and nerves of steel.

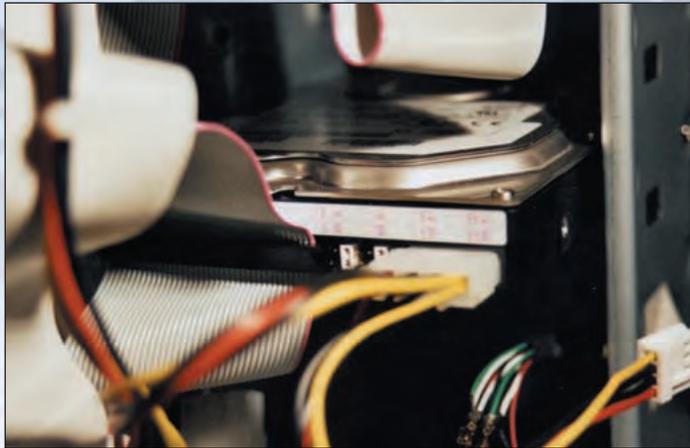


Figure 5: A close up of the hard disk and its jumper and cabling.

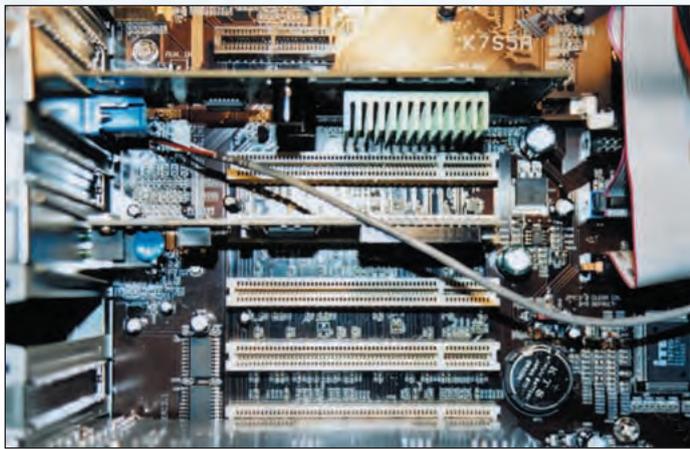


Figure 3: The graphics card seated in its AGP slot and five PCI slots, the second of which is occupied by the 10/100 network card.

The rest is easy. Shop round and buy yourself the following: a hard disk (you can get IDE discs between 40 and 120GB readily, you will have to decide how much storage you might need). I've used IBM and Seagate usually and a 7200 rpm 40GB disc is around £50; an AGP graphics card for £10-£20, (I usually use a Pine Nvidia TNT2 of various memory sizes for cheap machines); a 10/100 ethernet PCI card for around £6-8, (I used Realtek NC-100TX-R), a PS/2 keyboard and mouse as nearly all motherboards support them (£7 the pair); a floppy drive (about £7 new, of usually unrecognisable origin) and last but not least, some kind of CDROM. I usually use Liteon DVD players (£30). As a small note, although we are only building a server, it only costs a few more pounds to add a reasonable graphics card, the DVD drive and a reasonable sound card (I use the Xwave-7000 which is Linux and SoundBlaster compatible) to make the PC far more general-purpose. Its your

choice. If you are just using it as a server, you can do it for up to £40 cheaper than I paid. Note that the motherboard will include the cables you need to connect the IDE devices (hard disk and DVD or CD) and the IDC floppy but in my experience the supplied IDE cable may not be very generous, a separate IDE cable of around 75cm is only about £4, and you can always tie it up at the end if it's too long. That's it, shopping done, I spent a grand total of £267!

The assembly

Find yourself a flat workspace with a really good overhead lamp. Remove the screws on the back of the case which hold the side on, (usually 2 or 3 down the right-hand edge as you look at the back). Remove the side, (usually by sliding it about 2cm towards the back of the case). Lay the case down flat, removed side upwards and unravel the cables from the power supply (top left as you look inside the box) and the front panel

cables which come from the right. Hang them over the side of the box out of the way.

Next, plug the supplied power lead into the back of the computer and plug it in a wall socket that is not switched on in order to earth the case. Touching the case before you handle each component discharges any stray static, which can damage electronic components (I have never damaged any, but this is probably because I am a dull computer scientist; if I was a TV quiz show host however...) Take the motherboard carefully out of its sleeve, lay it flat on its insulating mat on the desk with the ports, (serial, PS/2, USB and so on) on the left. Inspect the ports and then look at the back of the case, where you will find matching templates. Extract the appropriate ones by pushing them gently with a screwdriver and then waggling them until they come off. You should also now choose which PCI slots you are going to use (see **figure 3**) and remove the blanking plates on the back of the case the same way. Note that the CDROM will have a sound cable included with it which usually fits on the motherboard immediately underneath the AGP slot, so fit the CD-ROM sound cable at the motherboard end now, then hang the other end over the side of the case.

Motherboard first

I am assuming that your motherboard has the processor and processor fan already mounted. If not, then you will have to fit them. First make sure the motherboard is well supported underneath. To fit the processor, lift the ZIF (Zero Insertion Force) lever on the side of the processor socket, touch the case then rotate the processor until it matches the holes on the socket, drop it in gently and lower the lever. The motherboard manual will show you how to do this. *Do not* force it. After fitting the processor, fit the processor fan, a big chunky thing. You have to line up the metal tags on the side of the fan with the plastic tabs on the side of the socket holding your processor. Place on top of the processor, fit one side and then use a screwdriver to fit the other side. At this point, you might be wishing you had asked the supplier to

do it! If you manage to do this without plunging the screwdriver through the motherboard, then read on. At this moment, take the opportunity to attach the small power lead from the processor fan onto its socket on the motherboard (nearby and called something like CPU FAN – you can see it just to the right of the processor in **figure 1**). That's the tricky bit, from now on, it's plain sailing.

Looking into the case, you will see numerous holes in the bottom as you view it. If you now look at the motherboard you will see six mounting holes in the motherboard. In the packet which came with the case, you will find lots of little odd-looking screws and 'stand-offs'. You will have to use your imagination to choose the right ones. If you can't figure out which to use, take all the parts back now and ask for asylum! Put the stand-offs in the case to line up with the motherboard to suspend it a couple of millimetres above the case bottom. Use all six to keep the motherboard stable when you have to push the cards in later. Now carefully pick up the motherboard and rest it on top of these and screw it to the six stand-offs. Don't fit the memory yet. Lastly, fit the rather funny-shaped plug on the power supply into the motherboard, (you can see this on the left of the processor in **figure 1**).

CDROM, discs etc

Fitting these is usually easy. The only complex bit is making sure they are jumpered correctly. If you look on the back of the CDROM and hard disk, you will find a series of little connections with a jumper fitted over one. You can see the hard disk jumper in **figure 5** near the centre next to the power lead.

They will be probably be configured correctly, but the CDROM should be a slave device and the hard disk should be a master device. It will tell you on the back of the device in almost invisible writing – a magnifying glass helps at this point. Next push out the plastic cover on the front of the case from the inside where you want to fit the CDROM. Sometimes there is a metal blanking plate in the way, which you should waggle loose and remove. Then slide the CD-ROM from the front into the case and secure it

with two attaching screws from the stash that came with the case. Most cases have guiding rails so this should be easy. Next slide the floppy drive into its cage further down, this time from the inside, making sure it is pushed in as far as it will go or you will have problems ejecting the floppies. Fix the floppy drive in position with its attachment screws. Finally slide the hard disk in above it, also from the inside and secure it. At this point you will see why it is a good idea to fit the memory last. Note that the CDROM drive will have a sound cable included with it which usually fits on the motherboard immediately underneath the AGP slot, so fit the CDROM sound cable at the motherboard end now and hang the other end over the side of the case.

Boards & PCI slots

Fit the graphics board first. This goes into the AGP slot on the motherboard above the white PCI slots as shown in **figure 3**. Fitting boards into slots on the motherboard is always a bit unnerving. You have to push surprisingly hard but make very sure the board is lined up properly in the slot first. You will hear them click home. You should also know that incorrectly fitting boards is the primary suspect if your creation won't boot properly later. When you have fitted the graphics card, screw it to the back of the case with a single screw on top, (although this depends on the case).

Fitting of the PCI boards is done in exactly the same way. Push them in carefully but firmly into the slots you chose at the beginning and fix them with the fixing screw to the back of the case. Finally fix the memory to the appropriate socket on the right-hand side of the board. SDRAM or DDR are fitted the same way and the slots are now designed in such a way that you can't fit them into the wrong slots or the wrong way round. Again push in firmly and then attach the catches at each end to engage with the slots in the ends of the memory module. How are you doing so far?

Cabling up

We're now ready to cable it. First fit the IDC cable (one of the two ribbon cables that came with the motherboard) to the floppy. This is

about the only thing you can accidentally fit the wrong way round, (at the floppy end). As you look at the back of the floppy, the cross-over bit in the cable should be on the left. Then fit the IDE cable into the socket marked IDE1 on the motherboard and then fit the other two plugs to the CDROM and the hard disk. This is where your longer IDE cable can help. Either order of plugs is OK, as the jumpers on the devices tell the motherboard which is the master and which the slave. You can tweak a bit of extra performance by using a second IDE cable in socket IDE2 and using IDE1 for the hard disk and IDE2 for the CDROM but I have no exact figure as to how much real extra performance this buys you. **figure 4** shows the IDE and IDC cabling.

I'm assuming that you can also fit the CDROM sound cable (it comes with the CDROM) now or you have already fitted it as noted earlier. A finally fiddly bit is fitting the wires which come from the front panel onto the motherboard. These are labelled things like HDD LED (disc light), POWER LED (power light), RESET SW (software reset) and so on. Fit them into sockets on the motherboard usually down the bottom according to the manual which accompanies the motherboard. This is where a magnifying glass (and ideally a tame leprechaun) comes in handy. The three wires I've mentioned above are always present. Some of the others may not have anywhere obvious on the board to plug them. Just tie them up in this case.

Nearly there now. The last thing to do is to attach the power leads from

the power supply to the various devices. These will only go in one way round and there are more than you need, so just plug away. Make sure you power the hard disc, the CDROM and the floppy. The floppy is a bit fiddly. If you have two case fans, one can be powered off the motherboard (look for the SYS FAN connector) and the other directly from the power supply. This is a good time to tidy everything up with insulating tape. The end product should look something like the example in **figure 2**.

Booting

Now for the moment of truth. You are going to boot it. You might as well be optimistic and have the boot floppy and CD of your favourite distro handy. Most times, the machine will boot straight away, (although leaving the side off to start with just in case is a good idea, so long as you don't go groping around inside with the power switched on). You will normally need to press the F1 key or somesuch on bootup to get into the BIOS to set things like system time. On a few occasions I have had to tell the BIOS to detect the IDE devices because it didn't find them automatically, but I have never had to do anything else. If you get this far, consult past issues of Linux Format to see how to install your distro. I am currently using SuSE 7.3 and 8.0. Installing SuSE 8.0 is a complete no-brainer – wonderful. The Linux community has made huge strides in recent distro installation. In comparison, I recently had to install Windows 98 three times reformatting each time before I got it to take. I've no idea why, I did the same things



Figure 4: The IDE and IDC cabling.

BackupBox

« each time. With SuSE 8.0 you'll be up and flying in 30 minutes. I'm assuming that you're capable of setting up the network configuration so that it can talk to the other machine(s) that are on your network.

However, the finished article might not boot :(In every case in my experience, this has been due to the graphics card not being seated properly. The symptom is that the machine flashes its lights as it starts but nothing appears on the screen. Most home constructors will tell you that ill-fitting AGP or PCI boards cause most boot failures. Just switch your machine off and reseal them carefully. So far, this has never failed me. If it won't boot, you might try books like *How To Build Your Own PC* by R.A. Penfold, ISBN 0-85934-479-7 for much more detail than I can cover in this article.

Network backup

At this point, you should now be basking in the satisfaction of having built your very own computer. So let's set it up to do something really useful like backing up selected directories on your network. The code below is a paired down listing of a Bash script

which achieves this goal. The script runs on each machine which needs to be backed up, although of course the directories to be backed up will differ in general.

The strategy is explained in the script comment header. It's very simple. The key is the power of the find command in section a) of the script, which checks for changes against the last time the directories dir1 and dir2 were backed up. It then zips them up, (the -@ option tells zip to use the standard input), ships them over to our new creation and then unzips them there, mirroring the original file structure. Finally, it touches the file last_update_xxxx; where xxxx is the name of our new backup machine to update the 'time of last update'. I have left out error recovery as an exercise to the reader. As it uses the secure shell copy scp to move files between machines, you will either have to enter the password to allow this (twice), or set up the machines as trusted. You can also add other frills like starting the backup script automatically at regular periods using a cron entry.

The process is extraordinarily fast. I back up a 2GB archive several times a

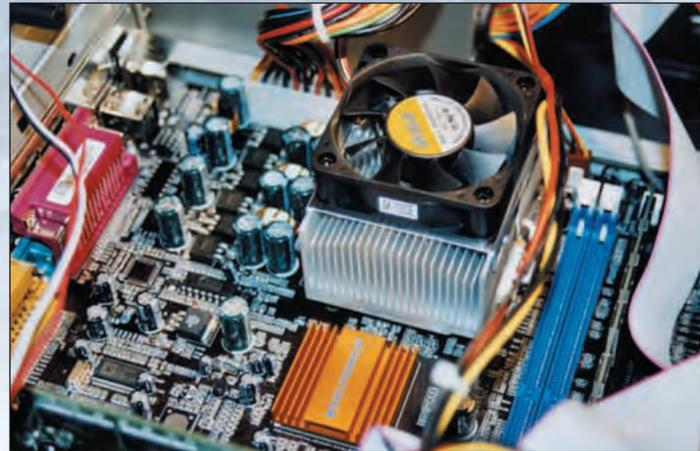


Figure 2: The finished article, rotated through 90 degrees to fit.

day across several machines in less than five minutes. I still have a 10baseT network, so it would be much faster if I had a 10/100 as the zip / unzip time is relatively short. What a wonderful program Zip is !

For security, I have three backup machines that I use to protect all my data for work (and a few personal files too!) Two are onsite and the third is a laptop which is kept offsite.

Supposing I have two machines, one called original and the one we have just built called 'eureka'. Simply install this script on original in a file

called update_home say and use it by typing:

```
% update_home eureka
```

One final wrinkle - the above script will fail if the initial zip file size exceeds some OS dependent value like 4Gb. In this case just copy the entire directory system first using **scp -r** (recursive copy). Thereafter the above script will be fine so long as incremental changes do not exceed this which should satisfy most users. That's all there is to it. Hope you have fun building your machine, and that your backup saves your data. [LXF](#)

Script for network backups

```
#!/bin/sh
#Back up any changes in selected sub-
directories of the /home directory to our new back
up machine. This is accomplished as follows:
#a) Find everything which is newer than the
touch_file.
# b) Create a zip file of these files.
# c) Copy the zip file to the PC to be updated.
# d) Unzip the zip file on the PC to be updated.
# e) Update the local log and touch files.
# Revision: $Revision: 1.6 $
# Date: $Date: 2002/12/28 18:42:49 $
#-----
if test $# -ne 1
then
echo "Usage: update_home
[machine_name]"; exit
fi
# Set up various values and temporary files
bdir=/home
touch_file=$bdir/PROJECTS/last_update_$1
log=$bdir/PROJECTS/update.log
upd=do_update
zipa=update_$1.zip
rm -f $upd $zipa
```

```
# If this is the first time, create the last update
file in pre-history so that everything gets copied.
if test ! -f $touch_file
then
touch -t 197001310000 $touch_file
fi
# Define the directories to back up from this PC.
We exclude the home directory because the
environment stuff may be different on the backup
and we don't want to overwrite that.
dirs=" \
dir1 \
dir2"
# Section a) The all important find command.
for dir in $dirs
do
# Find files which have changed since the
last time the touch_file was updated.
echo " .. $dir"
find $bdir/$dir -newer $touch_file -type f -
print >> $upd
done
echo ""
echo
"++++++"
```

```
++++++" >> $log
echo "Update of $bdir on machine $1 at
`date`" >> $log
chmod +x $upd
#
# Section b) Building the zip archive.
#
zip $zipa -@ < $upd
#
# Section c) Transfer the zip archive to the
backup machine
scp $zipa $1:/home/lesh
# Section d) Unzip with the silent options.
ssh $1 "cd /; unzip -o -/$zipa"
echo
"++++++"
echo "$bdir updated on machine $1"
echo
"++++++"
# Section e) Update the touch file and clean up.
#
touch $touch_file
rm -f $upd $zipa
```

What on Earth is... SERIAL ATA?

Trying to escape from a parallel dimension? **Richard Drummond** casts light on the new, backward compatible, successor to ATA.

»» So what on earth is Serial ATA?

Serial ATA is a much-needed update to standard parallel ATA technology that takes advantage of high-speed, serial connections to boost transfer rates up to 150 MB/s.

»» Hang on a minute. What's parallel ATA, then?

Parallel ATA is the name that has been retrospectively coined for good-old fashioned ATA, the technology that is used in most PCs today for interfacing with mass storage devices such as hard drives and CDROMs.

»» I see. Don't you mean IDE?

No. I mean ATA – that's its proper name. ATA stands for AT (Advanced Technology) Attachment, the system first introduced by Compaq in 1986 in their PC clones (it was actually developed by Compaq, Western Digital and Impris). It's called AT Attachment, because ATA is based on the 16-bit ISA bus first used in IBM's PC-AT. IDE stands for Integrated Drive Electronics and actually refers to the type of drive that is attached to an ATA interface. In an IDE drive, the drive controller logic is integrated onto the drive itself and thus is more reliable than earlier drives where the controller was on a separate card. Calling the ATA interface or even IDE drives 'IDE', however, is a bit redundant, because all drives these days – including SCSI devices – are really integrated drives in this sense.

ATA became an ANSI standard (after several years of deliberation) in 1994, and has been enhanced and revised many times since then. The latest Parallel ATA standard, ATA/ATA-7 specifies a maximum transfer rate of 133 MB/s, using the Ultra DMA 133 mode.

»» So ATA is an ANSI standard. What about Serial ATA? Who has developed that?

Serial ATA has been developed by the Serial ATA Working Group, a body whose members include Intel, IBM, Dell, Maxtor, Quantum and Seagate (see www.serialata.org). The Serial ATA specification is free for anybody to download and to implement.

»» You say serial ATA can do 150MB/s. That isn't much faster than my Ultra/ATA133 system.

True. But Serial ATA has more room to grow. Serial ATA 2 has already been proposed with transfer rates of 300MB/s and SerialATA3 is expected to manage 600MB/s. Parallel ATA, on the other hand, has been extended as far as it will go. Nobody even expected that Ultra/ATA133 would actually be possible.

»» Why do we need Serial ATA? Current drives can't saturate the data bandwidth provided by even Ultra/ATA-66

You're right. The main problem with parallel ATA isn't currently speed, it's the high-voltage, parallel conductors that the technology requires. Those horrible, old-fashioned, 40-wire (80-wire in Ultra/ATA33 and above) ribbon cables are inflexible, unreliable and – with a maximum length of 40cm – too short. They make positioning drives in a case a real pain and they restrict airflow through the case. Serial ATA uses flexible 7-conductor cables that can be up to 1m in length, and so has none of these problems.

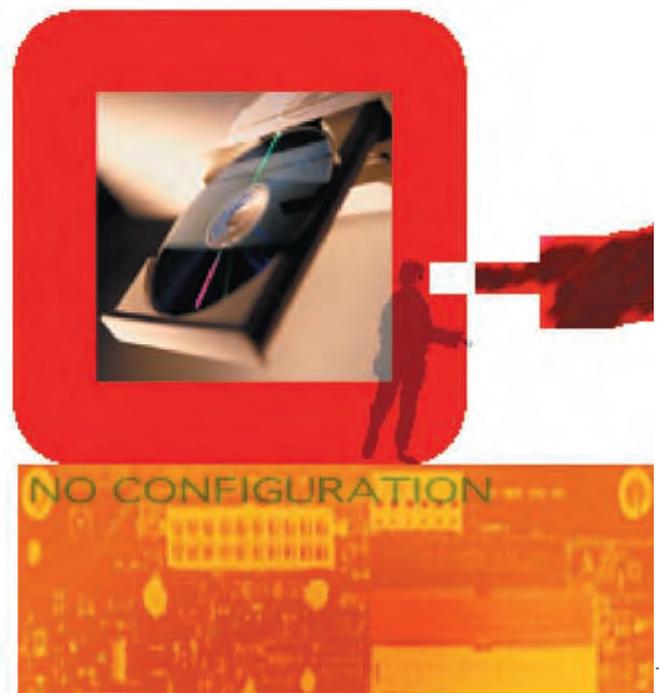
The advantages aren't just limited to cabling, however. The parallel 5V signals that parallel ATA uses are difficult and expensive to route on a

motherboard, require more expensive componentry and those IDE headers take up too much real estate. Serial ATA uses low-voltage signals and the Serial ATA data connector is only 8mm wide. This opens the door to much cheaper, smaller form-factor, lower power motherboards.

At the drive end of the cable, the same connector is used no matter the physical size of the drive. With Parallel ATA, 2.5" drives require a different, more compact connector from standard 3.5" drives. This is no longer the case with Serial ATA, and, in fact, the smaller connectors may enable a market for smaller form factor drives.

»» It still doesn't make sense. How can a serial connection be faster than a parallel one?

It does seem counter-intuitive to start with, but, yes, serial connections can be made to go faster than parallel ones. If you think about it, all of today's



fastest bus technologies are serial, such as Ethernet, Fibre channel, USB 2.0, and Firewire.

With a parallel connection there is always the problem of electrical interference between the individual conductors of the connection. The 80-way cables introduced with ATA/33 – which each signal conductor separated from each other with a ground – was an attempt to limit this, but with the existing high-voltage parallel ATA technology, it's always going to be a problem. Parallel ATA also makes the interface more complex, as the data has to be managed so that it's correctly synchronised.

Serial ATA solves this problem by using low-voltage (250mV) signals in differential pairs. A serial ATA connection is a point-to-point connection with two channels – one for transferring, one for receiving – and each channel is made up of two conductors. Each conductor in a channel carries the same signal with opposite phases. Thus any interference affects each conductor equally and because they have opposite signs, the interference cancels out.

» **You mentioned USB2.0 and Firewire. These are already high-speed serial connection which support storage devices. Why bother with Serial ATA?**

Glad you asked. Serial ATA is a more specialised technology than USB 2.0 or Firewire. It is designed for internal use only and for use only with storage devices. USB and Firewire are primarily for external storage, can use much longer cables and support more diverse types of device, including printers, scanners, cameras, etc. However, Serial ATA allows much higher transfer rates than either either USB 2.0 (which has a maximum of 60MB/s) or Firewire (which currently allows up to 50MB/s). But like ATA, this is an internal device protocol only, so USB and Firewire are likely to remain as the standards for external connectivity to scanners, cameras, printers and the like.



WhatOnEarthSerialATA

» You say Serial ATA is a point-to-point connection. What does that actually mean?

It means that it's the end of master and slave devices. In Parallel ATA, drives are connected in pairs

on a single channel (daisy-chained on a single ribbon cable) with the first drive designated as the master, the second as the slave. Not only does this make it difficult to site drives in a case and require the user to configure drives to be either the master or slave by setting jumpers on each device, it also has implications for performance because the ATA host may only issue commands to one drive on the same channel at a time. There was also a problem that badly behaving drives often wouldn't work effectively with their partners on a channel, leading to reduced performance.

With Serial ATA, each drive is connected to the ATA adapter via its own cable. One of the design goals of Serial ATA is that drives should be much

easier to install. Thus there's no more daisy-chaining of master and slave devices, no more messing about with various jumper settings, and the adapter can talk to all the connected drives simultaneously and independently.

» So Serial ATA uses different data cables from Parallel ATA? What about power connectors?

They are different, too. Serial ATA uses a new 15-pin power connector.

» Why so many pins? Current drives use only four pins for power.

One reason for this is that Serial ATA introduces a new 3.3V power for drives – in addition to the current 12V and 5V lines. Lower voltage means lower power, more efficient drives are possible.

» Lower power drives? That sounds like great news for laptops.

Yes indeed. Also crucial to mobile computing Serial ATA includes enhanced power management. Serial ATA devices may eventually support new reduced-power 'partial' and 'slumber' states. This will not only save battery life on laptops, but may also help combat noise-pollution.



RELIABILITY

LOW VOLTAGE

NO JUMPERS

NO MASTER/SLAVE
NO DAISY CHAINING

HIGHER
TRANSFER RATES

»» What about the other pins on the power connector?

Each voltage requires three lines – positive, negative and ground – and with three voltages, that totals nine pins. The remaining six are there to support hot-plugging of devices.

»» Higher speeds, more flexible cabling, no configuration and hot-plugging – sounds too good to be true! What's the catch? What about driver support? Will it break existing software?

Not at all. Serial ATA was designed from the outset to be fully software-compatible with Parallel ATA.

»» How is that possible?

When a Serial ATA adapter is powered-on it works in a Parallel ATA emulation mode, with full support for Parallel ATA transfer modes, interrupts and so on. This will work in one of two modes: master-only mode or master/slave mode. In the former, each Serial ATA device is presented as if it is the master device on its own channel on a Parallel ATA interface; with the latter, devices are mapped onto virtual master/slave pairs.

»» So my operating system will require no changes to be able to use Serial ATA technology?

No. Serial ATA is literally a drop-in replacement for Parallel ATA and requires no software changes. Of course, the additional features not present in Parallel ATA – such as hot-plugging and power-management – may not be supported without updated drivers.

»» I'm sold. When can I buy a Serial ATA adapter?

Companies are beginning to launch Serial ATA products. For example, Promise (see www.promise.com) do a range of PCI-based Serial ATA adapters, one with two-serial ports and one Parallel ATA port (supporting two drives), and one with four serial ports. HighPoint have also extended the range of PCI ATA RAID controllers with Serial ATA devices. (See www.highpoint-tech.com).

Both the Promise and Highpoint devices are actually based on Parallel ATA chipsets, with serial-to-parallel bridges to enable support for Serial ATA. Both ranges of cards are 32-bit 33/66 MHz devices.

3ware have produced a range of Serial ATA RAID controllers, the Escalade 8500 series (see www.3ware.com), with models supporting 4, 8 or 12 serial or parallel drives.

It's worth pointing with PCI Serial ATA adapters that the speed of the PCI bus may be the limiting factor for performance. Since Serial ATA performance scales linearly with the number of drives, if you have four fast drives attached to a Serial ATA interface, that could easily saturate the 133 MB/s bandwidth provided by a standard PCI slot. This will be less of a problem with cards that support 66 Mhz slots or 64-bit cards such as 3ware's devices.

»» What about motherboards with Serial ATA on board?

Several motherboards with on-board Serial ATA have been announced. No motherboard chipsets with integrated Serial ATA support are yet available, however, so again the current PCI-based solutions may not realize the full performance potential of Serial ATA. Intel and VIA both are working on updating their chipsets with Serial ATA support and you should see products for release in early 2003.

»» And all these products will work with Linux?

Yes. As we said, Serial ATA will work with existing Parallel ATA driver software. Seagate already demonstrated a working Serial ATA drive last year which ran fine under Linux.

»» So, what about Serial ATA drives?

They have been slower to appear, but expect Serial ATA drive to appear from all the usual vendors such as Maxtor, Seagate, IBM, etc. Initially these are likely to be geared towards the high performance end of the market (think servers and AV work). As there is no performance benefit for DVD/CD drives or tape streamers, these will probably only be available in serial ATA form once it is a more widespread standard.

»» Will I be able to use my existing Parallel ATA devices with a Serial ATA adaptor?

Yes. Some Serial ATA adapters are actually hybrid devices, providing a mixture of serial and parallel ports (such as the devices supplied by Promise mentioned earlier on this page).

The other solution is a parallel-to-serial converter which will let you attach a Parallel ATA device to a Serial ATA port. Both HighPoint and 3ware produce such converters. These should work for all ATA and ATAPI devices, including DVD/CD drives and tapestreamers.

»» I'm nervous about adopting a new technology? Do you think Serial ATA will catch on?

It is most certain to be a successful technology. Everybody wins with Serial ATA. A lot of effort has been made to make the upgrade to Serial ATA as transparent to the user as possible, but still offer all those great benefits such as improved performance, better cabling, hot-plugging and so on. Plus the cost of Serial ATA devices will not be any more expensive initially than Parallel ATA. (Serial ATA will probably eventually reduce costs).

The vendor wins because Serial ATA is easy to implement and the lower voltages and pin-counts required means its cheaper to design and produce PCBs for Serial ATA. The design cost may be minimised by using an off-the-shelf parallel-to-serial bridge with existing Parallel ATA chipsets.

»» Is this the end of the road for Parallel ATA?

Not quite yet. Parallel ATA and Serial ATA are expected to co-exist for the next couple of years. But don't expect any further improvements to Parallel ATA technology.

»» What's the roadmap for future updates to Serial ATA?

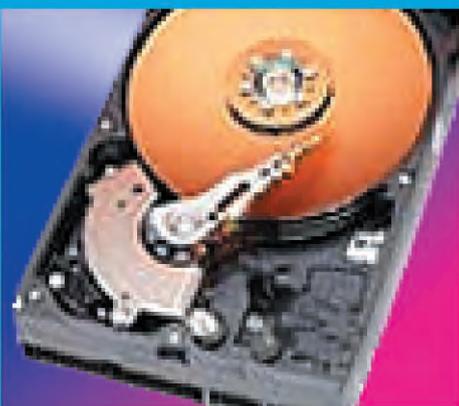
Serial ATA 2.0 is currently on the drawing board and will offer bandwidths of up to 300 MB/s per second with the same cabling and connectors used by Serial ATA 1.0. Serial ATA 2.0 is due to hit the market place in 2005. Serial ATA 3.0 is slated to scale up to 600 MB/s, but is not expected to be introduced until 2007.

»» Will Serial ATA finally kill off SCSI?

The performance and low-cost of Serial ATA will certainly make SCSI unattractive except for the most demanding applications. The latest Ultra320 standard still has edge in terms of raw performance, but it is an expensive technology and has all the problems of generally associated with SCSI, and an Ultra640 standard is expected. SCSI won't be left languishing in the archaic realm of ribbon cables, however, because a consortium, the Serial Attached SCSI Working Group, has been set up to develop a serial SCSI standard, based on a similar signaling technology as Serial ATA. This new Serial Attached SCSI standard or SAS is expected to hit the marketplace in 2004. (There's more information at www.serialattachedscsi.com.)

»» Where can I get more information on Serial ATA?

The best place to start is the Serial ATA Working Group site at <http://www.serialata.org/>. Here you can find FAQs and the Serial ATA 1.0 specification itself. Storage Search also have an excellent Serial ATA page (see www.storagesearch.com/serialata.html) with lots of information about Serial ATA announcements and products. Also checkout your favored storage vendor, because most provide information and resources on Serial ATA. [LXF](#)



Tutorials >>>

Our experts offer help and opinions on a whole host of Linux applications

Your guide to getting things done!

Whether you are just starting out in Linux, or an experienced veteran, there's always more to learn. Every issue of *Linux Format* is packed full of practical advice, and nowhere is it more concentrated than in our tutorials pages.

Here you'll find expert guides to all sorts of things, from Basic Linux usage to understanding and deploying network solutions, from simple script coding to the complexities of Perl regular expressions, Java server apps and more. We aim to bring a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact us, by post, by email (linuxformat@futurenet.co.uk) or log on to our website and post your suggestions in our special forums? (www.linuxformat.co.uk). Hope to hear from you soon!

Nick Veitch EDITOR

How code is represented

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap
between lines:
begin
mniWordWrap.Checked := not
end;
```

Usually, you'll find the code on our CD/DVD too.

THIS MONTH TEACH YOURSELF...

Perl

End-of-series collection of learning resources **p72**

Squid

Ever thought of using a proxy web cache? This is the Open Source solution for individuals or small networks **p74**

Mozilla >>

New Series – how to build your own *Moz* skins like our lovely Rhapsody In Orange **p78**

Java

Compiling Java to native code sounds tricky, but it isn't **p82**

OpenOffice.org

You don't need a dedicated package to make 3D designs – *OOoDraw* is here! **p84**

PHP

Use Ming to make generating Flash animations a breeze **p88**



TIP OF THE MONTH!

Perhaps you would like to create PDF files of documents, images or other pages from your applications. The snag is that Adobe don't do a version of Acrobat for Linux (they do support Linux with Acrobat Distiller, but at a hefty, server-oriented price that only corporate users can afford).

Well, the good news is that you don't need to. You may already have come across the Ghostscript tools for viewing Postscript files (and PDFs), but you may not be aware that the software also comes with many command line tools for manipulating Postscript and PDF files, including

Generating PDF files

ps2pdf, which, as you might guess, converts Postscript files into PDFs.

This is good news indeed, as virtually every application supports postscript output, and printing to files. Simply print to a PS file from your application, and use the ghostscript tool to convert it to a PDF.

```
ps2pdf input.ps output.pdf
```

It's as simple as that. The output PDF will have the same document size as your original PostScript output.

You can also use all the common PDF output parameters you might expect. For instance:

```
ps2pdf
-dColorImageResolution=72
-dDownsampleColorImages=true
input.ps output.pdf
```

The above example activates the downsampling of images contained in the document, and sets the resolution to 72dpi, a typical value for viewing the document on screen.

Loads of other options are available! You can find more documentation in the HTML docs that come with Ghostscript (probably to be found in `/usr/share/docs/ghostscript`, depending on your particular distribution). [LXF](#)

LEARNING MORE

Going further with Perl

In his final tutorial on the language, **Charlie Stross** looks at where you can find further sources of Perl information.

Perl is a big, gnarly language; if you include the core modules it's probably comparable in size and complexity to Java or C++, and has a similar learning curve to become proficient. To get to the stage of being able to effectively develop complex applications in Perl requires 6-12 months of full-time work.

Luckily, Perl starts off easy; a one-line "Hello, World" program in Perl really *is* one line (as opposed to the huge lumps of class declarations and wrappers you need to cough up in Java or C++). Because much of Perl's syntax is shared with other UNIX-family languages, it is possible to slip into Perl sideways. When I first started trying to use Perl seriously in 1993, I had mainly programmed in the Bourne shell and *Awk*, with some C on the side. From that background, Perl 4 (basically Perl 5 without lexical scoping and objects) came naturally – it was possible to become effective, certainly to the level of re-writing basic shell/*Awk* scripts in Perl, within a day or so.

Size is important

In fact, if you're intimidated by the size of Perl, you may find this the best way to get a handle on it: learn the UNIX shells (in particular Bash), read *The Awk Programming Language* by Aho, Weinberger and Kernighan, write a couple of small *Awk* scripts ... and Perl will seem eerily familiar. (Perl is such a comprehensive superset of *awk* that there's an *Awk*-to-Perl translator in the core Perl distribution that translates *awk* programs into idiomatic Perl 4 almost perfectly.)

Perl.com is just one instance of O'Reilly & Associates broad support for the language.

The screenshot shows the perl.com website interface. At the top, it says "O'REILLY perl.com THE SOURCE FOR PERL" and "The 2002 Map of the Internet is here!". Below this is a navigation bar with links like O'REILLY.COM, O'REILLY NETWORK, PERL.COM, etc. The main content area features several articles and sections:

- Downloads**: Documentation, CPAN Mirrors, FAQs, Training, Resources, Article Archive, O'Reilly Books, Search.
- Articles**: "Improving mod_perl Sites: Performance: Part 5" by Steve Berman, "Class::DBI" by Tony Bowen, "Sinky SAMP3: An Affordable MP3 Stereo Component" by Jon Local, "Scripting Groove Web Services" by Jon Local, "This week on Perl 6 (11/24-12/01, 2002)", "Managing Bulk DNS Zones with Perl", and "Controlling iTunes with Perl".
- Columns**: "Perl 6 News" and "Perl Recipes of the Day: Friday".
- Right sidebar**: "Earn University of Illinois Certification in System Admin. Linux/Unix from UserActive" and "Courses include: The Unix File System, Networking & DNS, Unix Services, Scripting for Administrators".



The other approach to learning Perl is to attack it as a first language, without reference to its roots. This is probably faster if you don't have the Unix background, but won't give you a feel for Perl idiom, much of which is Unix-specific and more or less intuitive if you've worked with other Unix shells and languages. There are a lot of books about learning Perl as a first language, and there's no substitute for browsing a bookshelf in your local computer bookshop to find one that you can get along with. A couple of general words of warning apply though; there's a tendency among some publishers to try to sell books by the kilogram, and weight/page count does not equate with quality. If you see a tutorial or introductory book with more than three authors listed, or more than five hundred pages, then it's probably an omnibus assembled in a hurry, and the editorial or quality control oversight may be a bit suspect.

For many years the canonical book on learning Perl was *Learning Perl* by Randall Schwartz (O'Reilly and Associates, 3rd edition, ISBN 0-596-00132-0). It's not particularly huge, it reduces Perl to an approachable scale, and the single author is one of the Perl ancients and knows his subject very well.

When learning a language, there comes a point at which a tutorial isn't enough. In addition to an introduction to a language construct or concept, you need an exhaustive reference. For Perl, the only canonical reference is 'the Camel book' – *Programming Perl* (O'Reilly and Associates, by Larry Wall, Tom Christiansen, and Jon Orwant, 3rd edition, ISBN 0-596-00027-8). This meaty doorstep is probably too terse to learn from unless you come from a Unix programming background, but it provides an exhaustive reference to the core language and the standard modules, along with chapters on important topics such as object oriented programming and inter-process communications.

Perl online resources

One point to note before you go out and spend hundreds of pounds on huge tomes is that Perl itself comes with a huge pile of online documentation. Look in your Perl installation directory for a subdirectory called 'pods' and you'll find close to two-thirds of a million words of documentation in pod format. This is accessible via the **perldoc** command; in a terminal window type **perldoc perldoc** for the manual page of the perldoc tool. There's an intro text ("perldoc perlintro") that gives a brief tutorial and flavour of the language, but most of the documentation is more technical; **perldoc perl** will give a list of the various pod files, their subjects, and the command to read them.

The pod documentation set evolved out of the Perl man page. Back in Perl 3, the man page ran to over a hundred pages; in Perl 4 it split into different files (**perlrun** covers command-line options to the perl binary, for example, while **perldoc** covers the debugger), and in recent Perl 5's (such as 5.8) a lot of

supplementary stuff has been added. The usenet newsgroups for Perl – **comp.lang.perl.misc**, **comp.lang.perl.moderated**, and so on – sprouted FAQ lists, and these are included (and run to about 50,000 words in their own right).

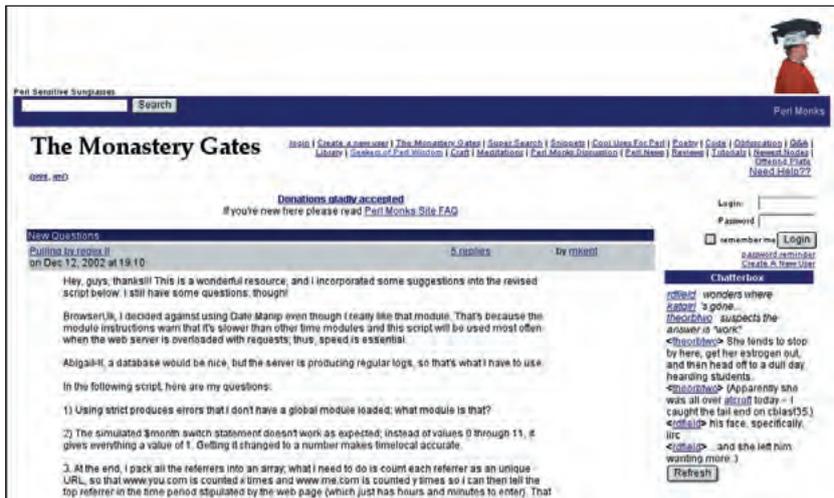
In fact, if you're happy reading on a screen, you can do without a printed copy of *Programming Perl*. The information in *Programming Perl* is all somewhere in the pod documentation (though Christiansen's writing style is more accessible than some of the pods, which suffer to some extent from Unix man page disease). There is also a hypertextifying utility that allows you to convert the documentation to HTML or *LaTeX* formats (for online hypertext or hardcopy output); when installing, use **installhtml** to build a copy of the Perl documentation for your web server. (It's described in the INSTALL file in the Perl source distribution.)

In addition to the core documentation, there's a lot of more specialised information on the public web. If your primary interest is using Perl as a CGI programming environment, then visit **http://perl.apache.org** for all mod_perl related issues. For generic Perl modules, CPAN can be browsed on the web via **www.cpan.org**. There are some useful community sites: Perl Mongers (the Perl advocacy group) at **www.perl.org**, the Use Perl portal website at **http://use.perl.org** (news updates and discussion areas), The Perl Journal at **www.tpj.com** (Perl's community mag, available online for a subscription), and Perl Monks **http://perlmonks.org/**, a slashdot-like arena for Perl.

Finally, **www.perl.com** is maintained and run by O'Reilly and Associates – publishers of many of the books recommended in this column, employer of Larry Wall (as a research fellow, to work on Perl), and owners of the most impressive commercial Perl web portal on the planet.

Polishing your Perl

If you're already programming in Perl, and use it for day-to-day data mangling tasks, you may want to take a look at *Data Munging With Perl* by David Cross (pub. Manning, ISBN 1-930110-00-6). Data munging focusses on the one task for which perl is unequivocally the tool of choice – taking data stored in one format and filtering, sorting, reducing, and translating it into other formats. Essentially a book about data structures, parsing, and data analysis, this is a really useful text because munching on



data is about the one job all Perl programs do (be it reading and writing databases via DBI, generating HTML via CGI, parsing HTML received via LWP, or just reading text files and doing creative things to them).

Perlmonks is a vibrant group of users featuring new news, opinions and discussion every day.

Exploit the power

On a more abstract level, *Effective Perl Programming* by Joseph Hall (with Randall Schwartz) (Addison-Wesley, ISBN 0-201-41975-0) is a terse collection of fifty lessons that will take you from being a competent programmer to having a deep understanding of some of the highly elegant tricks implicit in Perl's repertoire of data structures. It's an advanced tutorial, covering lessons that you won't find in *Learning Perl* and which are implicit in the coverage of *Programming Perl*, and this book (or something like it – Sriram Arinivasan's *Advanced Perl Programming* and Orwant, Hietaniemi and Macdonald's *Mastering Algorithms With Perl* cover the same territory and a bit more, albeit in ten times as many pages and less elegantly) is a great way to learn how to fully exploit the power of idiomatic Perl. It's not as rigid as a Design Patterns-based methodology, but using standard Perl constructs intelligently can bring order of magnitude performance improvements over a brute-force attack such as might be written by transliterating raw C into Perl. **LXF**

An indispensable bookshelf

Trees are proud to die for these top titles!



If there's one Perl book that you really need to keep handy – if you use Perl for a living – it's the *Perl CD Bookshelf* (version 3, O'Reilly and Associates, ISBN 0-596-00389-7). O'Reilly have made a speciality out of Perl documentation (Larry Wall is a research fellow there), and this is a compendium of seven of the key texts on the language. In addition to a paperback copy of *Perl In A Nutshell* (the second, improved edition – itself an

indispensable desk reference), the CDROM that comes in the binder includes *Perl In A Nutshell* and the third edition versions of *Programming Perl* and *Learning Perl*, as you'd expect. There's also

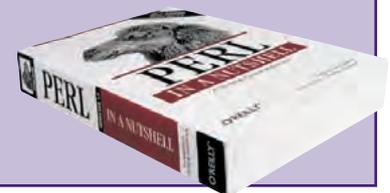
a copy of Tom Christiansen's *Perl Cookbook*, a packed collection full of useful procedures and algorithms for accomplishing day-to-day tasks in Perl, including an impressive range of file maintenance operations.

In addition to these core books, this version of the CD bookshelf drops the Windows content to focus on some more general Perl development books. *Perl and LWP* covers programming the web using LibWWW-Perl or LWP, a vital

toolkit that lets you download information from the web, parse HTML to extract information, and even build small web servers into your own applications. Then there's *Perl and XML*, extending the utility of Perl as an Internet programming language to the next level with coverage of processing XML data in Perl, including PerlSAX, XSLT, and the Document Object Model. Finally, *Mastering Perl/TK* shows up; the only respectable book so far about the only decent cross-platform GUI programming kit for Perl is a welcome addition to the collection.

About the only criticism that can be made is that O'Reilly are only putting seven books on the CD – *Programming*

the Perl DBI and *Advanced Perl Programming*, and *Perl For System Administration* would all be useful additions, even with the resulting penalty of a higher cover price. But it's hard to exaggerate the value of this book too highly – *Perl In A Nutshell* is itself vital to a jobbing Perl programmer, and the combination with six other core books provides a level of coverage that can't easily be equalled.



WEB PROXY CACHE

Squid 2.5 Proxyserver



Giving everyone on your network direct access to the Internet is not necessarily a great idea. **Chris Denton** looks at the classic Open Source alternative.

Squid is a stable, widely used and highly flexible web proxy cache. Its main purpose is to provide Internet browsing for machines that are not directly connected to the web. This can be usefully subdivided into two main elements. Firstly, the *Squid* server accepts requests for web pages from clients and retrieves them on their behalf. Therefore, it acts as a proxy. Secondly, the *Squid* server stores retrieved pages locally so that after the first request for a web page has been fulfilled, subsequent requests can be met without having to go onto the Internet at all. This is the caching function.

The National Laboratory for Applied Network Research, based in the US, looks after *Squid*, with a lot of help from unpaid volunteers across the globe. The lead programmer is a guy called Duane Wessels and he has done an excellent job of ensuring the efforts of a large team are suitably coordinated. *Squid* is released under the GNU GPL and is one of the great Open Source apps, although not as high profile as *Apache*, for example, or *Samba*. Nevertheless, it deserves just as much respect and is an essential component of many a network, large and small.

Proxy vs NAT

Assuming that you do not have an Internet IP address for each of your desktops and servers, getting everybody onto the web will involve some trickery. The most popular solutions are either some kind of proxy or alternatively a system of Network Address Translation (NAT). In the Linux world, NAT is often referred to as masquerading. Which you will go for depends on your policies on certain configuration and security issues. It also depends, of course, on the type of network you have.

The twin headaches of security and performance tend to scale up so that organisations with lots of users, perhaps on a variety of platforms, are more in need of a web proxy cache like *Squid* than smaller outfits who may find it easier to get away without one. Of course, at the top of the market there are solutions available which can be integrated into an existing NAT-based network to address security concerns. These tend to be commercial products, and deep pockets are usually required. For most scenarios, *Squid* is a perfectly good method of allowing Internet access on your network in an inexpensive, reasonably safe and generally manageable way.

A common complaint with the use of proxies is that you need to add them into the browser settings as well as configure any other web apps (*RealPlayer*, *GAIM*, etc.) that might be required. It

can be frustrating when some web-based programs either don't support HTTP proxies, or, when they do, only support them without authentication. However, the situation is improving and the drawbacks of having to use a proxy, as opposed to getting the appearance of a direct connection, are not really that severe.

Installing Squid

There's a minimal preparation needed on your system. You will need GCC installed and also Perl, or you won't get very far. However, in all likelihood both these pieces of software will already be in place. It is also a good idea (though not absolutely essential) to have *Apache* running as well. Again, this program is probably included with, or at least easily available for your platform. Better still, it can be readily compiled from source. You can verify that you have everything with the following commands:

```
type gcc
gcc is /usr/bin/gcc
type perl
perl is /usr/bin/perl
type httpd
httpd is /usr/local/apache/bin/httpd
```

As always, the first thing to do when installing new software is to get a copy of the latest stable release. *Squid* lives at www.squid-cache.org where the source code and a lot of useful information can be obtained. The present version 2 strain has been live since 1998 and had reached 2.5 at the time of writing. In January 2003 the next edition, 2.6, is scheduled to be released. Development tends to be slow and steady, so even though 3.0 is being worked on, it may not arrive for some time.

RPMs are widely available with many distros including it in their server builds by default. Whilst these are generally fine, as with any RPM they can lag behind the latest source code release, put files where you don't want them and not always have the features you want already compiled in. For these reasons we are going to install from source, however, if you prefer to use a binary it doesn't matter as the rest of the article will still be relevant, even if file locations may be slightly different.

So, once we have obtained the source code from the *Squid* website or one of its mirrors, and ensured it is in a suitable location, it can be uncompressed and extracted like this:

```
tar xvzf squid-2.5.STABLE1.gz
```

The next step is to run the configure script. This prepares the installation and controls what features of *Squid* are included or excluded. There are many possible options to this script, but for the moment we shall ignore the majority of them and go for a default setup. The only thing explicitly specified here is the location of our installation. All *Squid* files and directories will be placed in a structure beneath this point.

```
cd squid-2.5.STABLE1
```

```
./configure --prefix=/usr/local/squid
```

Some of the other possible configure options will be covered below. Finish the installation by typing:

```
make all
```

```
make install
```

If everything has worked successfully, the main *Squid* binary, helpfully called “squid”, will be placed in `/usr/local/squid/sbin`. To make things nice and simple, this directory should be added to your PATH environment variable in the file `~/.bash_profile`.

Configuration

Now its time to get *Squid* working the way we want it. Edit the `squid.conf` file, which is to be found, in our case, by going to the `/usr/local/squid/etc` directory. It is rather large, but very well laid out and follows a similar syntax to *Apache*. You shouldn't have many problems doing what you need to do by pointing your text editor of choice at `squid.conf`. However, for those that prefer something a bit more graphical, the *Squid* module that comes as part of the standard *Webmin* installation is really very good.

For the most part, the system defaults will be fine to getting on with. Even at this stage, though, we can make a few well-chosen modifications at the correct point within `squid.conf`:

```
http_port 8080
```

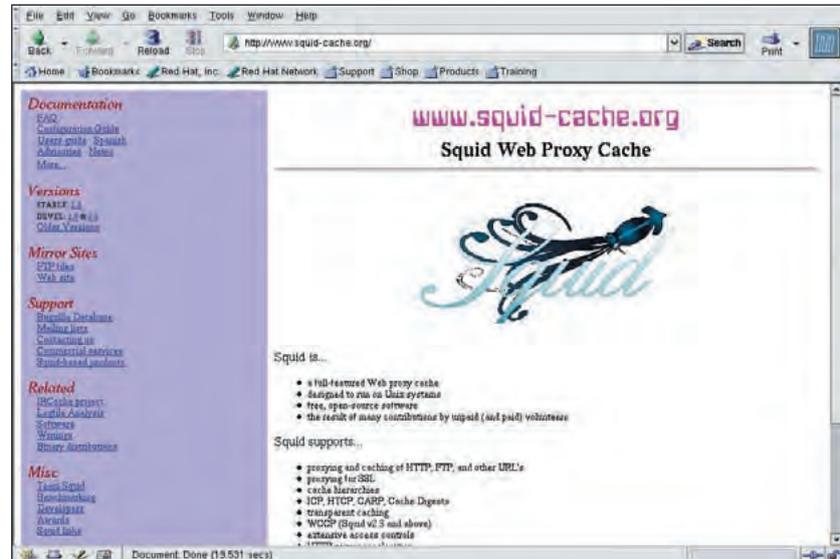
This is the port that your users' browsers will connect to. The default is 3128, but 8080 is also frequently used. We could also specify a hostname or IP address here, which would be useful on servers with more than one network card.

```
visible_hostname brasseye.mydomain.com
```

I've found it necessary to add the fully qualified domain name of my *Squid* server here. Though the program should be able to work this out for itself, it can fail to do so, which will cause *Squid* to abort at startup, so it's sensible to add the entry in.

```
cache_dir ufs /usr/local/squid/var/cache 500 16 256
```

The cache directory is the home of the web pages that *Squid* stores and as such it is very important that it is configured as appropriate for your system. The first of the three numbers is the cache size in megabytes, and should ideally be as big as possible. The default is 100MB but it's not unusual to see this changed to



one Gigabyte or more. Multiple entries of “cache_dir” and *Squid* will use all of them, spreading the load across multiple disks.

The third field names the cache directory itself, and for a server that's likely to be heavily-used, this should really be a file-system in its own right, if not actually on a separate disk with its own controller. The final two numbers relate to first and second level subdirectories of the cache. With a very large cache it might be a good idea to increase these too.

Well, with those changes saved away, we should have a workable implementation of *Squid* more or less ready to go. Before we kick the program off, though, there are a couple of final steps. Permissions can cause a problem, because by default *Squid*'s effective user when executed by root is actually nobody (as in the system account of that name). Nobody will need write permissions to the log and cache directories before *Squid* will start properly. This can quickly be achieved by doing the following:

```
cd /usr/local/squid/var
```

```
mkdir -m 777 cache logs
```

Almost there now, but before we can start using *Squid* the cache needs to be initialised:

Not the prettiest of websites, but packed with useful info.



Keeping Tabs on Users

Sarg - The Squid Analysis Report Generator

Each and every HTTP request made by your users is logged, but without a decent method of interpreting that information it's not going to do you much good. *Squid* does provide the option to perform logging in an *Apache*-style format, so that reporting tools intended for webservers, such as *Analog*, can be used with *Squid*. If you want to do this simply amend the relevant entry in `squid.conf` to:

```
emulate_httpd_log on
```

Much better, though, to use a log analyser that is specifically designed for *Squid*. You could go for *Calamaris*, a popular tool obtainable from <http://calamaris.cord.de>. It's a capable enough program and is supported by *Webmin*, so *Calamaris* is especially handy

if you have gone down that route. However, for a detailed and well-organised look at what your users are up to, *Sarg* is very hard to beat.

You can get the source code from <http://web.onda.com.br/orso>, and links to RPMs are provided as well if you prefer your software pre-compiled. *Sarg* is easy enough to install in any case, and with just a couple of configure options can be nicely integrated with your current *Squid* setup:

```
tar xvzf sarg-1.2.2.tar.gz
```

```
cd sarg.1.2.2
```

```
./configure --enable-bindir=/usr/local/squid/bin --enable-sysconfdir=/usr/local/squid/etc
```

```
make
```

```
make install
```

The binary is called *sarg* and that's what you execute to produce the reports. The configuration file, `sarg.conf`, is as logically named as it is well set out, and closely controls *sarg*'s behaviour. The output is in HTML, so once again we find ourselves in need of *Apache*. There are a good number of settings in `sarg.conf` that control the generation of the HTML code, allowing for a nice amount of customisation. There are also some more fundamental settings, such as log and password file locations, that should be addressed. You will be rewarded for reading right through `sarg.conf`, but to be going on with here are some worthwhile changes:

```
access_log
```

```
/usr/local/squid/var/logs/access.log.0
```

```
title "My Access Reports"
```

```
password /usr/local/squid/etc/passwd
```

```
date_format e
```

```
lastlog 12
```

We've amended the log file to `access.log.0` because we don't really want to run *sarg* on a log that's still being written to. Ideally, *Sarg* should be run under cron, not long after the *Squid* logs are rotated and with the same regularity. The date format defaults to `u` for USA and should be changed if you prefer to see the European layout. By **lastlog**, *Sarg* means the number of reports to generate before overwriting. The default is never overwrite, but unless you have a bottomless supply of very large disks, you will probably not want to leave it at that.

Escape Tunnel

Bouncer - Any Port Through Squid

Although *Squid* is primarily intended for web traffic only, it can be used as a proxy for almost anything else you care to think of. This is done by a process known as SSL Tunneling, and one program specially designed for this task is called *Bouncer*. While not Open Source, *Bouncer* is freeware and is available to download as a binary from www.r00t3d.org.uk. As well as Intel and Sparc versions of Linux, it is also

provided for a handful of other platforms, including Windows.

To give you an specific example of how *Bouncer* works, let us assume that we have a server on the Internet that is running *SSH* server. Let's call it **jam.myotherdomain.com**. From behind a proxy we would not normally not be able to start a *SSH* session with this box. However, it becomes possible if we run *Bouncer* with the following options:

```
bouncer --port 2222 --destination
jam.myotherdomain.com:22 --daemon --
tunnel brasseye.mydomain.com:8080 --
t_user chris --t_password mypassword
```

This would open a tunnel between port 22 on the remote server and port 2222 on the client's local machine, going through the *Squid* server. Unfortunately, the password field is entered as plain text, rather than keypresses being hidden by the usual stars or hashes, so is

a bit of a security worry. So, to get the desired *SSH* connection the only thing that we need to do is:

```
ssh -p 2222 localhost
```

There's a lot more you can do with *Bouncer* than this and further documentation is available from the website. You may very well want to keep *Bouncer* away from your users. However, you should be aware that it is out there, and they could well find it.

◀◀ squid -z

That may take a few minutes, but when it does we can finally run *Squid*, and this is done by simply executing the program without any options:

squid

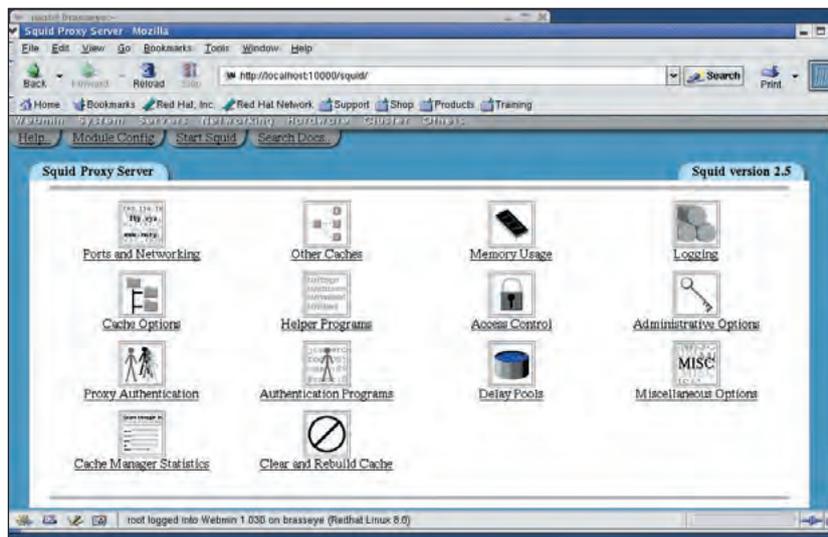
By doing this, *Squid* automatically starts in daemon mode, so you should get your prompt right back. Then it's time to point a browser at your new web proxy cache and start surfing. If you have any problems, the logfile `cache.log` should be the first place you look, and will usually give you human-readable error messages that are quite straightforward to use for debugging. The other main logfiles are `access.log`, which keeps a record of each HTTP page request received; and `store.log`, which tracks each item retained in the cache.

Authentication

Before we get too pleased with ourselves, it's worth remembering that the *Squid* server we've just set up has absolutely no security features whatsoever. Anyone whose connected to our LAN/WAN can use it as long as they know the hostname and port number of our server. There is logging taking place, but it is only trackable by the client PC, not by the user doing the surfing. No URL filtering is taking place either, so these users could be going to any number of disreputable sites and we can't stop them doing it nor can we reliably track them down after they've done it.

In fact, since what we've done is little better than give every machine on our network a potential unmonitored connection, it

Webmin conveniently comes as standard with just about everything you need for Squid.



LXF Squid User Access Reports									
Period: 13 Dec 2002-13 Dec 2002									
User: chris									
Sort: BYTES, reverse									
User Report									
ACCESSSED SITE	CONNECT	BYTES	%BYTES	IN-CACHE	OUT	USED	TIME	MILISEC	%TIME
data/zone www.linuxformat.co.uk	53	325,521	79.13%	0.00%	100.00%	00:00:15	15,032	57.90%	
data/zone sk.bliznet.net	2	27,939	6.80%	0.00%	100.00%	00:00:00	310	1.19%	
data/zone sk.bliznet.net	5	21,272	5.17%	0.00%	100.00%	00:00:00	330	1.24%	
data/zone www.linuxformat.co.uk	4	17,107	4.16%	75.04%	24.96%	00:00:00	125	0.46%	
data/zone www.linuxformat.co.uk	3	10,709	2.60%	0.00%	100.00%	00:00:03	3,684	14.19%	
data/zone sk.bliznet.net	7	3,366	0.82%	0.00%	100.00%	00:00:01	1,929	7.43%	
data/zone www.linuxformat.co.uk	4	2,478	0.60%	0.00%	100.00%	00:00:03	3,841	14.79%	
data/zone www.linuxformat.co.uk	1	1,527	0.37%	0.00%	100.00%	00:00:00	22	0.08%	
data/zone smch.questionnet.co.uk	2	1,364	0.33%	0.00%	100.00%	00:00:00	491	1.89%	
TOTAL		411,185	100.00%	3.12%	96.88%	00:00:25	25,364	100.25%	
AVERAGE		411,185				00:00:25	25,364	100.25%	

A typical Sarg report. Much easier to read than a 10,000 line logfile, we think you'll agree!

would be a really good idea to turn it off again ASAP:

squid -k shutdown

For peace of mind, we should make sure there's some sort of authentication process in place. This basically means tying the users of the proxy down to certain properly validated individuals. There are quite a few different methods for doing this, and in fact *Squid* version 2.5 has completely overhauled the way it handles authentication, offering much improved functionality. You could go for LDAP, for instance, or maybe use Linux's own `/etc/passwd` file. *Squid* has now been integrated with *Samba*'s excellent *Winbind* daemon, so if you're running a Windows network users can be recognised from their domain accounts without even the need to enter a username or password.

The method we're going to use here, though, is the traditional NCSA form of authentication. Basically, this involves creating and maintaining a simple encrypted password file for *Squid*'s exclusive use. To get this working we must of course first to install it from the source code:

```
cd squid-2.5.STABLE1/helpers/basic-auth/NCSA
```

```
make
```

```
make install
```

In fact, we could have avoided these step if when we were running the *Squid* configure script we'd specified that we wanted NCSA like so:

```
./configure --prefix=/usr/local/squid --enable-basic-auth-
helpers=NCSA
```

After this is done, the `squid.conf` needs to be amended to reflect our new aims. We have to tell it about the authentication parameters we want to set:

```
auth_param basic program /usr/local/squid/libexec/nscs_auth
/usr/local/squid/etc/passwd
```

```
auth_param basic realm Our brand new web proxy cache
```

The first line is the important one, as it tells *Squid* which authentication program to use and where the corresponding password file is kept. The second line is more or less a comment field, and it will appear on the clients' login box whenever they need to start using the proxy. The only other thing that we need to do in *squid.conf* is add the relevant access control list (**acl**) and set up a corresponding **http** access rule:

```
acl password proxy_auth REQUIRED
http_access allow password
http_access deny all
```

The purpose of those entries to ensure that anyone who correctly provides a username and password is allowed to use the HTTP proxy, but if this is not done then access is denied. To create the password file and put some users into it we need to use the **htpasswd** command. This comes as part of *Apache*, but will work on its own if you really don't want to permanently install a webserver alongside *Squid*. The password file is created when you add the first user, which should be done like so:

```
htpasswd -c /usr/local/squid/etc/passwd chris
```

You will be prompted to enter and then confirm the desired password. Subsequently, users can be added and have their passwords modified by using the same command but without the **-c** option. To delete a user you will have to manually remove the line relating to them from the password file. If that sounds like a unfriendly process, you could just use *Webmin* as *Squid* users can be perfectly well managed from that. To make things even easier, there's a CGI script available from www.squid-cache.org which enables users to change their own passwords from a web page. It's called *chpasswd.cgi* and is well worth a look.

URL Filtering

Something else you might want to do, especially if you're operating in a business environment, is to disallow access to certain sites. This is fairly easy to configure using an **acl** and then an **http** **access** statement:

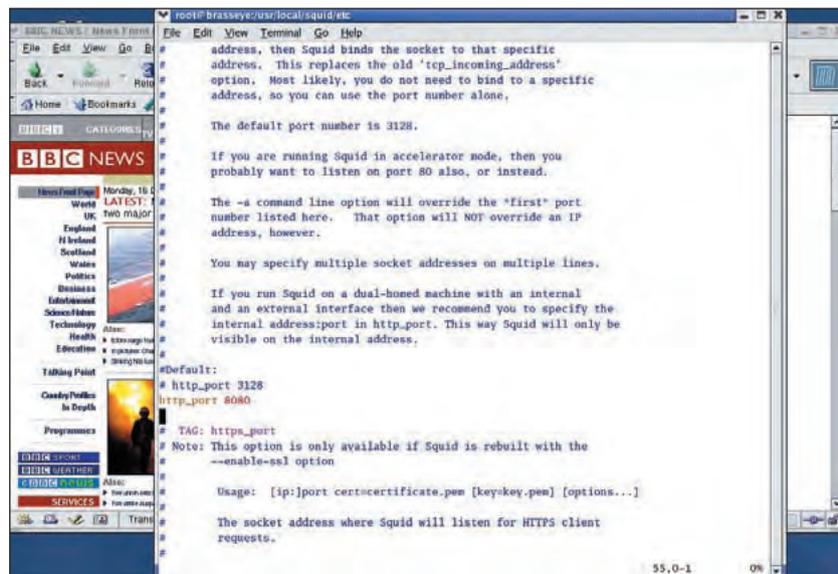
```
acl bansites url_regex "/usr/local/squid/etc/bannedsites.txt"
http_access deny bansites
http_access allow password
http_access deny all
```

The file specified in the **acl** is a simple text file listing all the URLs that are out of bounds. It should be noted that banning is performed by regular expression pattern matching, so you don't have to have complete URLs on the list, but it is a good idea to not be too general. For instance, banning the word "sex" may seem sensible, but that would also take out www.sussex.com and www.davidsex.co.uk to name but two. The order of the **http** **access** statements are important, as they are applied in that order until a match is found. So something cannot be rejected if a previous rule has already approved it.

We can make our filtering a bit more sophisticated – it is possible to only ban sites at certain times, say during normal business hours, and allow traffic to go through on evenings and weekends. Here's a basic example of how this could be achieved:

```
acl restrictedsites url_regex
"/usr/local/squid/etc/restrictedsites.txt"
acl restrictedtimes time MTWHF 09:00-17:00
http_access deny restrictedsites restrictedtimes
http_access allow password
http_access deny all
```

It is possible to combine the entries for restricted and banned sites and implement both policies simultaneously, and indeed any



The squid.conf file.
Where all the real work is done.

others along similar lines, if so desired. You may not wish to trawl the Internet building up your own list of dodgy webpages. Prepared files are obtainable from various sources, such as www.squidblock.com or www.squidguard.org/blacklist. Even these, though, are by no means exhaustive and with the speed at which domains come and go on the web it's unlikely that they ever will be.

More Things to do with Squid

One possible further modification that might be required is the activation of SNMP monitoring for your proxy. This could be very useful if used in conjunction with a decent monitoring tool such as *Cricket* or *MRTG*. It would also be a neat way of circumventing *Squid's* own native analysis program, the distinctly uninspiring *Cache Manager*. SNMP support needs to be compiled in, so to get it you should return to the source directory and:

```
./configure --enable-snmp
make clean
make all
make install
```

Access to the SNMP port is denied by default, so to get it working you'll need to add this in *squid.conf*:

```
snmp_access allow all
```

If "allow all" seems a bit on the open side, access rights can be refined by the use of **acls**. Another situation that could well come up, indeed almost certainly will, is that the log files grow to a large size. These can be rotated using an option of the *squid* command, which is best scheduled under *cron*:

```
0 0 * * 6 /usr/local/squid/sbin/squid -k rotate
```

Instead of acting as an HTTP web proxy cache, *Squid* can also be setup as an HTTP accelerator, acting on behalf of a webserver rather than clients' browsing. It would be most useful if a webserver sits behind some kind of network bottleneck. The *Squid* server running as an accelerator is then placed beyond the weak point, and it satisfies incoming HTTP requests itself. All the real webserver has to do is be accessible to the accelerator.

This tutorial has covered the basics, but *Squid* is an application with many different possible uses and a mind-boggling amount of potential configurations. In most situations, it can almost certainly provide some real benefits to a network of any size that is connected to the Internet. [LXF](http://www.linuxformat.co.uk)



BESPOKE BROWSERS

Mozilla: Skinning the beast

In the first of a short series, **Andy Channelle** strips away the flesh of Mozilla to show how creating your own browser distribution is no longer a crushingly difficult job.

Mozilla is a great browser, and its continuing growth will probably do as much for the acceptance of Linux as a serious desktop contender as something like *OpenOffice.org*, KDE or *The GIMP*. But more than just being good functionally, *Mozilla* is malleable, meaning schools, ISPs and offices can tailor an entire Internet application suite to the needs of their students, customers or workers, providing a truly custom-made, branded product (that exposes as little or as much functionality as is required) without hiring an expensive team of artists and developers.

In this short series we'll go some way toward creating a bespoke browser suite, discover how to extend the functionality of the browser without getting bogged down in XML coding and look at some very basic application development.

First the good news: hacking *Mozilla* can be as simple or as complex as your ambition and skills can handle. Though we're starting with a few minor tweaks, you may soon be making quite

fundamental changes with just a few lines of code. The other great thing is that you won't need a complete integrated development environment (although you can use one if you like) to make and test your handiwork; a basic text editor, colour picker and graphics package should suffice. We've been using *Quanta* for its decent code highlighting and file navigation (which is a godsend when you're going in and out of *Mozilla's* labyrinthine CSS files) but feel free to use whatever is most comfortable. If it can save a .css file, it is adequate.

Inside Mozilla

In order to do anything, we need to start with a base skin to build on. Creating one from scratch would be both long-winded and rather pointless, so I'll be starting using the standard Modern theme (which has similar elements to my intended design). You could of course use Classic or download something from www.mozdev.org to base your creations on, or if you have the time and inclination, start from nothing. If you're at all concerned about ruining *Mozilla*, it's useful to have a second installation in your home directory. This is good for testing purposes and saves having to mess around with permissions.

The first problem is that, since a fairly early version, *Mozilla* has been configured to use skins archived into .jar files. This is

basically a zip file cleverly disguised as a different format, so we can open it with any zip enabled tool – though you might have to rename it modern.zip if your utility doesn't recognise .jar files. Extract it to /chrome/modern/ but keep the rest of the directory structure. Take a while to look at some of the contents: you'll find a collection of GIFs, CSS, DTD files, a few XML documents and a smattering of other files arranged quite logically: /global is used for configuring widgets, buttons, scrollbars and the like; /communicator is the repository for sidebar, taskbar and profile information and the others are application specific, covering the browser, messenger and composer.

Now you could immediately go in and start changing things, but you'll find it doesn't make any difference to the browser because Mozilla is still looking at **modern.jar** for its configuration information. To change this we fire up a text editor and load the file /mozilla/chrome/installed-chrome.txt. The relevant section that we have to change is this:

```
skin.install.url.jar:resource:/chrome/modern.jar!/skin/modern/communicator/
skin.install.url.jar:resource:/chrome/modern.jar!/skin/modern/editor/
skin.install.url.jar:resource:/chrome/modern.jar!/skin/modern/global/
skin.install.url.jar:resource:/chrome/modern.jar!/skin/modern/messenger/
skin.install.url.jar:resource:/chrome/modern.jar!/skin/modern/navigator/
```

Simply take out the 'jar' references to end up with:

```
skin.install.url.resource:/chrome/modern/skin/modern/communicator/
skin.install.url.resource:/chrome/modern/skin/modern/editor/
skin.install.url.resource:/chrome/modern/skin/modern/global/
skin.install.url.resource:/chrome/modern/skin/modern/messenger/
skin.install.url.resource:/chrome/modern/skin/modern/navigator/
```

Bizarrely, you might do this and still end up making no difference. In that case, scroll down the document and you'll find another section still pointing to the .jar archives. Change that as above and you're in business.

To demonstrate some of the basics of skinning, this month we'll be making a few rudimentary changes to the *Mozilla* web browser. Specifically we'll alter the menu bar, drop-down lists, toolbar and some of the toolbar buttons. The results aren't hugely spectacular but they do demonstrate how just a few changes can make a big difference.

Rhapsody In Orange

If you've ever visited the *Linux Format* website, you'll know some of the team have an inexplicable fondness for bilious shades of orange, so we'll use that as a basis for our theme. First, navigate to /global/toolbar.css and scroll down to the section marked 'menubar' where you can configure the background and borders of the menubar:

```
menubar {
  cursor: default;
```

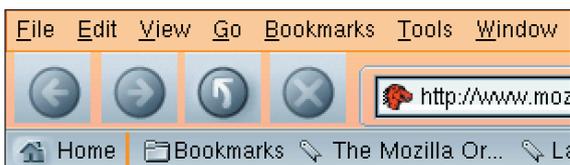
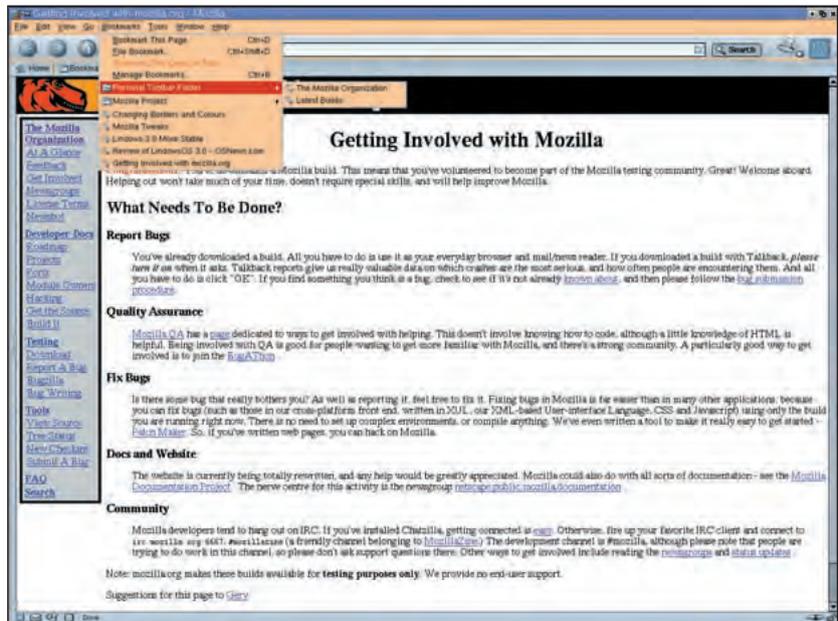


Fig 2: The 'File Bookmarks' dialog box proudly displaying its newly minted White buttons.



```
min-width: 1px; /* DON'T DELETE!
Prevents hiding of scrollbars in browser when window is
made smaller.*/
background-color: #FFCC99;
color: #000000;
}
```

The relevant section here covers the **background-color** and **color**, configuring the background and the menu label text. In the above example we've changed the traditional blue/grey to a more pleasing biohazard orange (using the websafe colours from *Kcolour*) while keeping the text black. Remember that as this is in the /global directory, anything you change here will also be reflected in the other applications of the browser suite.

```
menubar > .toolbar-holder {
  border-top: 1px solid #FF9933;
  border-right: 1px solid #FFCC66;
  border-bottom: 1px solid #FF9900;
  border-left: 1px solid #FF6600;
}
```

These entries define the areas surrounding the menubar; in this case all four borders are a single pixel in width and are rendered in a range of terrifically tasteful tangerine hues. If you launch *Mozilla* now, you'll notice we have the menu bar and its surroundings coloured, but the toolbar buttons look awful and clicking on one of the menu entries brings down the regular grey options box. We'll tackle the latter problem next. >>

Image formats

The mantra of Mozilla.org is that everything within the application window is rendered by *Gecko*, therefore when designing skins, we're not just restricted to using GIF images, even though the standard skins are both replete with them. GIF files are useful – they're small and capable of transparency – but it seems a shame to waste all that rendering power, especially as all the images are stored locally and don't need to be downloaded over a slow Net connection. The best alternative is to use PNG files for your

graphics. While keeping file sizes down, PNGs don't look as fuzzy as GIFs, meaning your buttons and backgrounds can have better definition and you can also use alpha-blended transparencies to add realistic drop shadows and glows to your artwork. The future should bring even better options as *Mozilla* will gain a native SVG rendering subsystem, meaning we could have a fully scalable, cross-platform UI along the lines of *Ximian's Gorilla* theme.

TutorialMozilla

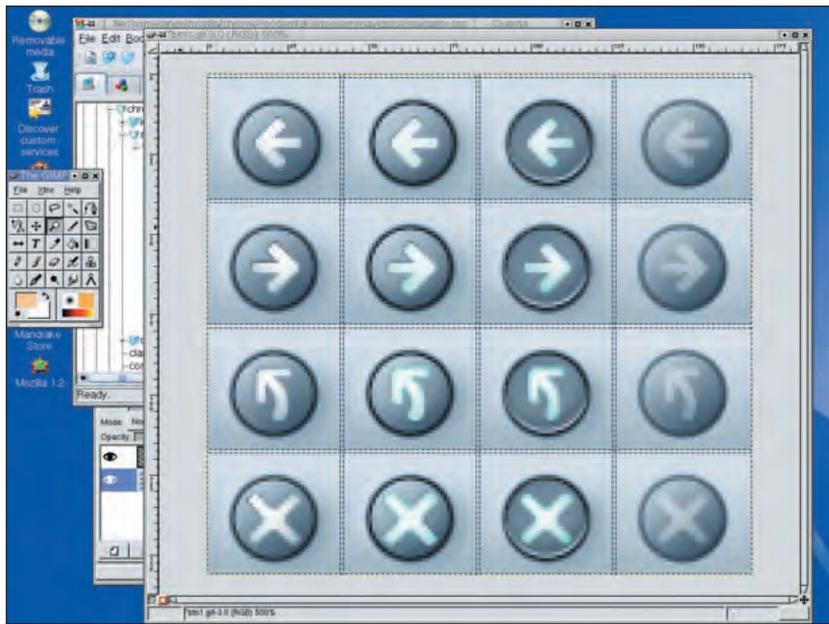


Fig 3: Here you can see the image regions used by Mozilla for the toolbar buttons.

Open `/global/menu.css` and scroll down to the section marked `/* :::: menu/menuitem :::: */`. The first colour reference you'll see here defines the text of the drop-down list, while the second sets the colour for unavailable menu items. However, there is no background colour set, so within this first **menuitem** section, add the line:

```
background-color: #FFCC99;
```

To change the colour of the menu items as you mouse over them, scroll down and find `/* :::: menu/menuitems in popups :::: */`. Here you'll find the two familiar tags: **background-color** and **color**, which change the roll over and text colour respectively. Your changes can be checked in an instant, so it really encourages experimentation. Save your file, close and reopen *Mozilla* and you should have something similar to **fig 1**.

It is also worth noting that all of the elements we have configured with hex-reference colours can also be rendered with a background image, which is why in **Figure 1**, we have a greyish toolbar with a messy orange border. Using the Cascading Style Sheets model, the browser's style (a toolbar with an image) overrides the global style (toolbar with hex reference), but a discrepancy in the image size leads to an unpleasant border at the bottom of the section.

Open up `/communicator/toolbar.css` and scroll past the license to the **toolbar-primary** section. The graduated tint on the toolbar is called by the line:

```
background-image: url("chrome://communicator/skin/toolbar/prtb-bg-noline.gif") repeat-x top;
```

We have two options here. We can either point this URL to a different file, in which case just create a new graphic file one pixel wide by 40 pixels deep in *GIMP*, colour it however you fancy and change the reference in the CSS file. Or we can simply change the existing image. Either way, the result is a repeating image running the length (**repeat-x**) of the toolbar.

Push all the right buttons

There are a couple of ways we can handle buttons in *Mozilla*, either as standard *CSS/Mozilla* widgets or as graphical files. The first method is simply a variation on the work we've already done; navigate your way to **button.css** in the `/global` directory to find the configurable elements here. The advantage of designing buttons this way is that with minimal effort you can change so much. For a swift demonstration, open **button.css** and in the **button** section, change the hex value in the **background-color** to White, save the file and start *Mozilla*. Firstly you'll see the Search button has become white. Now open Bookmarks/File Bookmarks, or launch the mail client to see how you've changed the buttons across the suite (**figure 2**, previous page).

The disadvantage of defining buttons this way is that restricts your modifications somewhat. If you are aiming for a specific look, you may have more luck creating buttons in the form of image files, as demonstrated by the standard toolbar set in **figure 3**.



Fig 4: Testing out our first attempt at the toolbar buttons. Using PNGs on a transparent background also means we don't have to fiddle around aligning the buttons with the toolbar background.

Phoenix

A helping hand from UXL

Trying to skin the whole of *Mozilla* is a bit daunting – there are literally hundreds of image files and references that need to be built/edited – even if the result can be very satisfying. If you only really need to brand the browser, perhaps for navigating a corporate/university intranet and accessing webmail, consider hacking away at *Mozilla's* little sister project, *Phoenix* (or whatever it's called this week).

This is a stripped down web browser, in a similar vein to *Galeon*, but has the advantage of being built entirely on XUL, *Mozilla's* eXtensible User interface Language. Using the same techniques outlined in this tutorial, a branded *Phoenix* distribution is easily within the range of anyone capable of changing a few lines of code and designing (or acquiring) some images. Download the browser from <http://www.mozilla.org/projects/phoenix> and uncompress it into a directory – installation is as simple as that. Again you'll need to unzip your base skin (*Phoenix* only ships with Classic) into the skin directory and edit `installed-chrome.txt` to point to the uncompressing a whole lot of pointless graphics and CSS files onto your hard disk, most of the work done in *Phoenix* takes place in the `/browser` directory. In

here you'll find a selection of graphical files and a single CSS file, a little confusingly also called *browser*. *Phoenix* also makes it simpler to package and distribute your skins, hence the growing number of skins on third party websites such as <http://mozdev.org> (an online development collective which eclipsed the number of *Mozilla's* own officially available chomes within just a few weeks of the 0.3 release).

Moz-tech 101

A brief look at some of the underlying technologies that make Mozilla so flexible

CASCADING STYLE SHEETS

In the context of a website, Cascading Style Sheets (CSS) enable the webmaster to define many elements of their site in a single document; for instance backgrounds, colours and font formatting. Simply placing a link to that document in the `<head>` of an HTML page would then force that page to inherit the styles previously defined. The big advantage of this is that you can refine your design and make site-wide changes by altering a single document. It separates the design from the content.

If, though, you wanted a specific page to have a distinct look, you can add style and formatting information to the `<head>` of individual documents and this would over-ride the link to the external stylesheet. Furthermore, if you wanted to style a single element of

a page you could format it in the normal fashion and that formatting would override both the external link and the document `<head>` formatting: hence the term 'cascading'. The main thing to remember is that the formatting nearest the element to be styled takes precedent; it's like being in a family where the people you're closest to have the greatest influence.

This, unsurprisingly, is also how it works for Mozilla's UI – as we have seen with `toolbar.css`. In the `/global` directory we edited `toolbar.css` to select a solid colour for the background of the toolbar, and then subsequently eclipsed this by defining a background image for the bar in `toolbar.css` in the `/communicator` directory, and then layering button images on top of this in the `/navigator` folder. The `/communicator` directory is judged to be 'nearer' the UI element we're

designing, while `/navigator` is seen to be even closer still, so it dominates.

Very thick books have been written about CSS, and explaining it fully is beyond the scope of this series, but a comprehensive grounding in the technology will certainly help you make the most of Mozilla. A useful primer is available at www.webdesignhelper.co.uk/design_elements/css/css3/css3.shtml, while you will find a dissection of the official CSS definition at www.w3.org/TR/REC-CSS1

NEXT MONTH

While CSS defines the look of your browser project, the actual workings of your buttons, toolbars and menus are based around another standard web technology – Javascript – which we will probe next time.

To do any work on these buttons, you'll need to open up `/navigator/icons/btn1.gif` in *GIMP*. You may be surprised to see that the graphic file (see [figure 4](#)) contains each of the four buttons in all four states – active, highlighted, depressed and inactive. Replacing these 16 icons with your own creations and saving (with transparent background) will complete the transformation of the main navigation section of the browser toolbar. For *Rhapsody In Orange* I've begun with a simple orange disc, which is bevelled in *GIMP* to make it look like a button. Then I've added a back arrow in white. To create the various states, I simply duplicated the first layer three times, moved them into position and then changed the colour balance (highlighted), inverted the colours (pressed) and altered the transparency to 65% (inactive). As I'm going to save out as a PNG, that last one will keep its transparency value, and let the toolbar background image show through.

This is all well and good, but still rigidly follows the scheme already set out by the Modern theme. What if you want ultra high-resolution buttons for running on a big monitor, where 32x32 pixel icons all but disappear; or very thin buttons for a minimalist-looking toolbar?

To meet these demands, you'll need to open up your text processor, find `/navigator/navigator.css` and scroll down to the section headed `* ::::: primary toolbar buttons ::::: *`. We'll just pick apart the first section.

```
.toolbarbutton-1 {
list-style-image: url("chrome://navigator/skin/icons/btn1.gif");
min-width: 0px;
}
```

The first lines name the buttonbar and tell *Mozilla* where to find the source image – you'll need to change the extension to `.png` if you've opted to ditch the GIFs. The following four entries define the various states of the Back button and point the browser at the relevant section of our 16-icon grid using the `-moz-image-region` command. The last part of each entry tells Mozilla that the button is rectangular and sets the boundaries in a series of pixel values. The sequence corresponds to top, right, bottom and left edges. With *GIMP*'s rulers set to pixels, it's light work measuring your buttons and altering the CSS to fit.

```
#back-button {
-moz-image-region: rect(0 41px 38px 0);
}
```



Fig 5: And here are the results of this month's endeavours! There are still some rough edges and elements that need to be themed, but it's a pretty good start.

```
#back-button[buttonover="true"] {
-moz-image-region: rect(0 83px 38px 42px);
}

#back-button[buttondown="true"] {
-moz-image-region: rect(0 125px 38px 84px);
}

#back-button[disabled="true"] {
-moz-image-region: rect(0 167px 38px 126px) !important;
}
```

On the last entry, the `!important` element tells us that this state (disabled) is the natural position for the image – after all when you first fire up the browser, there is nowhere to go back to, so it makes perfect sense that this will be the primary image.

Of course, you're not restricted to using this method, the only limit is your imagination (and maybe the time you've got for your project!) You could use individual graphic files for each button set – find the graphic file for the Print icon for a demo – or even for every image. The advantage of this method is that it is very memory-efficient, as only one graphic file needs to be loaded up for the whole toolbar.

The biggest problem about the graphical approach is that even in the simplest of browser configurations there is so much to skin that most sane individuals (or those not being salaried by AOL, Lindows etc) will give up long before reaching such fripperies as the mail client. See the *Phoenix* boxout on the opposite page for a somewhat simpler solution. [LXF](#)

NEXT MONTH

Using the methods set out here, every single element of *Mozilla*'s user interface is subject to the whims of your imagination. Scrollbars, dialog boxes, the URL box and everything in between can be tweaked, overhauled or abandoned depending on your needs – in fact we've only just scratched the surface. Next issue we'll venture into the content space of our browser distribution to see how you can integrate the browser with your Intranet or Internet site. We'll also take *Rhapsody In Orange* one step further than this tutorial and get animated with our own throbber.

GOING NATIVE

Speaking Java

Learn how to compile Java to native code with this introduction from **Richard Drummond**.



Java's most touted feature is its write once, run anywhere capability. Since Java is compiled to architecturally neutral bytecode, rather than instructions for any particular processor architecture, it can run on any platform with a Java run-time environment. Many vendors however, produce compilers that translate Java code to native code. Why is this? Well, you lose the advantage of portable code, but you win on startup time, execution speeds, and the memory requirements of your compiled code. This is incredibly attractive for embedded market, if you can't afford to use the behemoth that is Sun's hot spot *JIT* engine.

The GNU choice

You don't have to fork out on an expensive commercial compiler to experiment with native Java compilation. It can be done with GNU's *gcj* - or GNU Compiler for Java - a front-end for *gcc* that has been part of the GCC distribution since release 3.0. *gcj* knows how to compile Java code to bytecode or native code - and pre-compiled bytecode to native code.

Since *gcc* already supported a number of languages, adding Java translation was not that difficult (Java in many ways is a simpler version of C++, after all). The difficult part was implementing the standard Java class library, and, indeed, in *gcj* 3.2 this is still not complete. *libgcj* provides a fairly complete implementation of the JDK 1.1 class library (and includes parts of *JDK1.2* and *JDK1.4* such as new I/O) compiled as a native shared library. *gcj* can also make use of bytecode libraries, since it includes a Java interpreter, accessible on the command line through the command *gij*.

Gcj does have a few other limitations. At the moment it cannot handle inner classes. On the plus side, since it effectively

treats compiled Java code the same as compiled C++ code - it shares the same ABI (Application Binary Interface) - interfacing Java with native code is simpler and more efficient with *gcj*. It does support the standard JNI (Java Native Interface) API for calling native code, but also defines its own API called CNI (or Cygnus Native Interface).

Using GCJ

The *gcj* command is designed to be, as much as possible, a drop in replacement for the standard Java compiler, *javac*. Take, as an example, our familiar Hello World class.

```
class Hello
{
    public static void main( String args[] ) {
        System.out.println( "Hello World" );
    }
}
```

This can be compiled to bytecode with *gcj* using the command below:

```
gcj -C Hello.java
```

The **-C** switch here tells *gcj* to generate bytecode rather than the default, which is to generate native code. The result is bytecode file, *Hello.class*, which can be executed as normal. You can also use *gij* to run this simply with

```
gij Hello
```

Since it is designed to be compatible with *javac*, *gcj* takes similar parameters. For example, use the **--classpath** switch to specify your classpath, the **-d** switch to specify the directory to store generated code, and the **-D** switch to define Java properties.

To compile to native code, you must specify the class which is the entry point of your program. The GNU linker needs a **main()** routine, and, with Java, there's no way to work this out automatically. You tell *gcj* your main class with the **--main** switch. Thus, to compile the above class to native code, use

```
gcj --main=Hello Hello.java -oHello
```

Just like *gcc*, if you don't give a target filename, *gcj* will generate an executable called *a.out*. Use the **-o** switch to provide a name for your program. The executable file *Hello* is a plain, garden variety ELF (at least on Linux), which you can invoke as you would any other binary. They can even be debugged with GNU debugger, *gdb*. When it comes to deploying your code, the only significant run-time dependency (unless you brought in external Java libraries via your classpath) is the shared library *libgcj*.

A benchmark

Compiling to native code should make your Java programs run more quickly. Code generated with *gcj* will be on a par with compiled C++ code in performance.

This is much better than any Java interpreter, and comparable with the speed of good JIT compiler such as Sun's HotSpot

GCJ tools

There's more in the toolbox

The GCJ suite includes the following tools:

gcj The compiler front-end itself, which can compile Java code to bytecode or native code, and is a replacement for *javac*

gij This is the front-end to the *libgcj* bytecode interpreter as is used similarly to the */java/* command to execute bytecode.

gcjh A tool to generate C++ headers from bytecode and is used when employing CNI to integrate Java code with native code.

jav-scan Examines a Java source file and dump informations regarding it.

jcf-dump Examines a Java bytecode file and dumps information. This is similar to the standard */javap/* command.

engine. (There are those that argue that JIT compilers can potentially offer better performance, because they can dynamically compile code and adapt to run-time requirements. Compilers such as *gcj* can only perform static optimisations.)

How about some numbers? The following class implements the class Eratosthenes Sieve, an algorithm used for locating prime numbers and often used in benchmarking suites.

```
public class Sieve {
    static String results1, results2;

    public static void main(String[] args) {
        System.out.println("Running Sieve benchmark.");
        System.out.println("This will take about 10 seconds.");
        runSieve();
        System.out.println(results1);
        System.out.println(results2);
    }

    static void runSieve() {
        int SIZE = 8190;
        boolean flags[] = new boolean[SIZE+1];
        int i, prime, k, iter, count;
        int iterations = 0;
        double seconds = 0.0;
        int score = 0;
        long startTime, elapsedTime;

        startTime = System.currentTimeMillis();
        while (true) {
            count=0;
            for(i=0; i<=SIZE; i++) flags[i]=true;
            for (i=0; i<=SIZE; i++) {
                if(flags[i]) {
                    prime=i+i+3;
                    for(k=i+prime; k<=SIZE; k+=prime)
                        flags[k]=false;
                    count++;
                }
            }
            iterations++;
            elapsedTime = System.currentTimeMillis() - startTime;
            if (elapsedTime >= 10000) break;
        }
        seconds = elapsedTime / 1000.0;
        score = (int) Math.round(iterations / seconds);
        results1 = iterations + " iterations in " + seconds +
            " seconds";
        if (count != 1899)
            results2 = "Error: count <> 1899";
        else
            results2 = "Sieve score = " + score;
    }
}
```

Enter this and save it as `Sieve.java`. Compile it with `javac -target 1.1 Sieve.java`

Here are the results for a variety of different Java run-time engines, running on my Debian 3.0 (and a bit) box with a 1200MHz Athlon processor. The score is the number of iterations per second.

```
Shell - Konsole
Session Edit View Bookmarks Settings Help
evilrich@confucius: /root$
evilrich@confucius: /root$ gdb Hello
GNU gdb 5.3-debian
Copyright 2002 Free Software Foundation, Inc.
GDB is free software, covered by the GNU General Public License, and you are
welcome to change it and/or distribute copies of it under certain conditions.
Type "show copying" to see the conditions.
There is absolutely no warranty for GDB. Type "show warranty" for details.
This GDB was configured as "i386-linux"...
(gdb)
(gdb) run
Starting program: /root/Hello
[New Thread 16384 (LWP 993)]
[New Thread 32769 (LWP 994)]
[New Thread 16386 (LWP 995)]
Hello World

Program exited normally.
(gdb)
(gdb)
(gdb)
(gdb)
(gdb)
(gdb)
(gdb)
(gdb)
```

On the web

<http://gcc.gnu.org/java/> - The GNU Compiler for Java home page.

<http://www.gnu.org/software/classpath/classpath.html> - Alternative free implementation of the Java class library which is gradually being merged with *libgcj*.

<http://www.kaffe.org/> - The Kaffe JVM, an lightweight, open source Java run-time with a JIT engine.

Code compiled with *gcj* can be debugged with a native debugger such as *gdb*.

JVM	Score
gcj-3.2 (native)	2625
gij-3.2 (interpreted)	132
Sun JDK 1.1.8	394
Sun JDK 1.4.0	3835
Kaffe 1.0.7	1544

As you can see from the results of this rather simple test, code compiled with *gcj* acquits itself rather well. It is slower than the *JDK 1.4 JIT* engine in this test (it may be faster in other benchmarks; a different Linux distro or processor would certainly make a difference) and is faster than Kaffe's *JIT* engine. Both the interpreters, *gij* and *JDK 1.1.8*, gives scores which are an order slower than the *JITs* and *gcj*.

What does all this mean? Very little really – due to the variables mentioned above, Benchmarks should always be taken with a pinch of salt. But it gives some justification to the assertion that natively compiled code is fast. Sun's latest *JIT* engine is fast, but at times when you can't afford a 60–70MB overhead just for the Java engine itself and when the start up times of a big *JIT* are unacceptable, technologies such as *GCJ* begin to look interesting.

The other great advantage of *GCJ* is that it is Free Software (Free as in freedom, not in the Sun sense of a zero-cost download – the capital 'F' is the giveaway). If you are ideologically opposed to Sun's licencing policy, or you are working on an unfunded project, then *gcj* can offer a remedy. It's more suited to server work than client side at the moment – due to the incomplete AWT implementation – but when *libgcj* matures it will a valuable tool there too. [LXF](#)

DRAWING

OpenOffice.org for power users

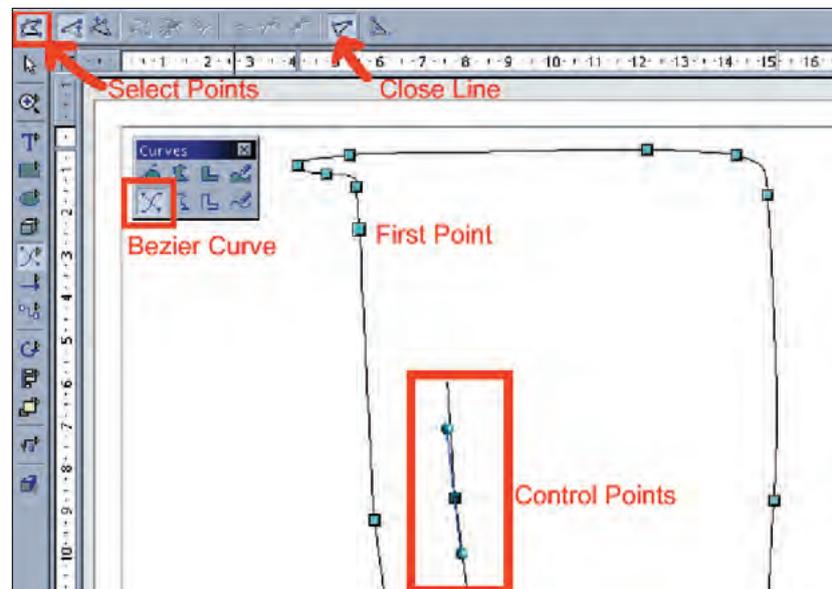
In the last of our OpenOffice.org (OOo) tutorials, Neil Lucock tries to get the wax crayon off his monitor after showing what graphical goodies Linux's killer application can deliver.

There are two types of drawing programs in common use. The *GIMP* and *Photoshop* are bitmap editors – they manipulate coloured pixels. *Draw*, *Freehand* and *Corel Draw* create lines and filled shapes mathematically. The advantage of these vector programs is that you can scale a drawing without losing any details. Zoom in on a bitmap (or a JPEG) and eventually you'll see coloured squares. Zoom in on vector artwork and it scales smoothly.

This tutorial was put together in OOo version 1 running in KDE 2.2.2. on Mandrake 8.2. I assume that you know how to open and close files, make directories and hold a mouse. However, vector drawing packages are not always very intuitive, you click on something and apparently nothing happens. Although not intended to be an introduction to vector drawing techniques, we'll have to cover a few basics to make sure what follows makes sense, so please excuse anything that seems remedial.

In this series over the last five issues of *Linux Format*, we've made comparisons with MS *Office*, as most people who use computers at work tend to be familiar with the programs in that suite. There is a drawing applet available in MS *Office*, but it in no way compares to what you can do in OOo *Draw*. We'll go through the key tools you need to produce your own illustrations, do a couple of drawings and show you what it can do.

The Bezier tools.
Understand these and you can draw anything.



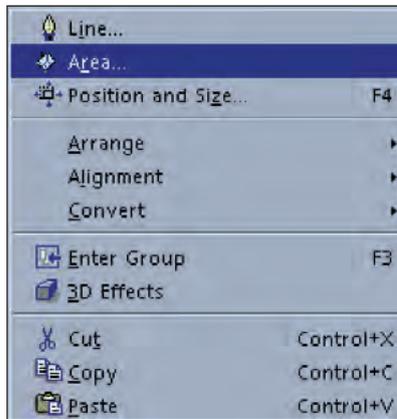
Start OOo and click File>New>Drawing. OOo gives you a blank drawing area. I'm going to draw the view in front of me (a couple of computers and a monitor). It helps if you think of an object as a series of shapes. First, I want to draw the coloured panel on the front of my Linux PC. It's roughly rectangular and the temptation is to find the rectangle tool. Not a good idea – once a rectangle, always a rectangle. The best way is to use the Bezier tool (found under Curves. See the boxout). Click the tool to select it, put your mouse cursor over where you want the shape to start (it doesn't have to be exact, you can change it later) and click and hold the left button as you move the mouse. If you just click and release, you'll get nothing.

Select Points

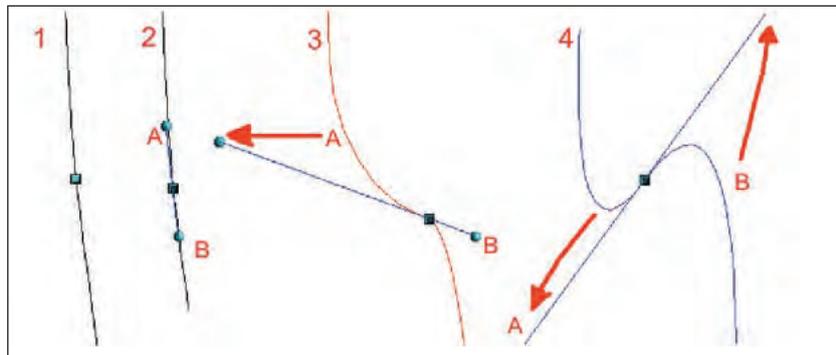
It should start to draw a line and it will continue drawing the line until you double-click to end it. Move the mouse and click to put more of your line in, particularly where there is a change of direction. What you are doing when you click is putting Control Points in, and you can use these to modify the final shape. Draw a shape that's fairly close to the real thing and end the line. Click the line to select it and small green boxes appear around the shape. You then click on the Select Points icon. There are two – one at the top of the left side tool bar and one at the bottom; use either. These change the line into a series of Control Points. There's one at each place you clicked when you made the shape. When you click the Select Points icon, the top tool bar changes to give you some useful point editing tools. The one I need here is Close Line, which joins up the two ends. This will eventually allow you to fill the inside with colour. Before you do that though, you will need to get the shape right.

Each Control Point can be selected and moved. The line between it and its neighbours stretches to the new position. You can alter the shape of the lines on each side of the Control Point by clicking on it a second time. Two balls appear on screen, attached to the Control Point by short lines. The First Control Point and the last one you made only have one ball to move. Moving these balls further away from the Control Point affects the adjacent line over a greater distance. Moving them at right angles to the drawn line (ie away from it) alters the curvature. Sometimes the two balls are directly over the Control Point. Grab one and move it, then grab the other. Understand this behaviour and you have the key to making anything with a vector program. Move your Control Points as you wish until you have the shape you want, then click a second time on one of the Select Points icons to finish editing.

We made this a closed shape so we could fill it. Select your shape and right-click to bring up a menu and choose Area. You can also use Format>Area from the menu. This dialogue box is



The right-click menu. Different entries appear, depending on what is selected before you right-click.



1, a Control Point. 2, Click it a second time to show the control balls. 3 and 4, The results of moving the balls around – the lines change shape.

another major tool for making your art look good. There's a tabbed dialogue box. Choose Area and you get the a choice of various fill types.

Shadow works like the dropped shadow you can assign to text. It puts another slightly offset copy of the shape behind the original. You can set the offset to fall in one of eight directions, you can make the shadow any colour you choose and you can adjust the transparency. Higher values make it more transparent. Transparency can be set for the whole shape or as a gradient. The gradient has options to set the kind of gradient that you require and is very configurable.

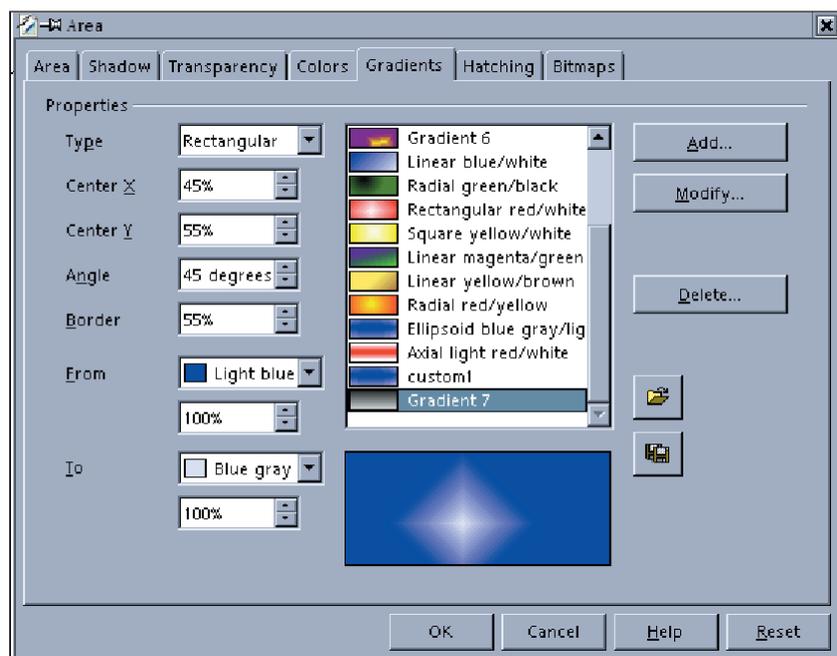
Colors allows you to choose from the 90 shades already in the program or make new ones. Gradients and Hatching both allow you to make new designs and use them in other parts of the suite. You can assign a bitmap to your shape, there are tools to adjust its position and size. I've used 3D programs that did not have such easy-to-use tools.

I coloured the shape blue, then used the same technique to make the case, the details and the smaller BSD machine that lives on top of the Linux machine. Once you have positioned parts, it's a good idea to group them. Select a part, hold the Shift key and select another. Then use the menu Modify>Group or right-click for a menu. Once grouped, the parts of your fledgling masterpiece can be moved and resized. You can always use Ungroup if you need to alter something.

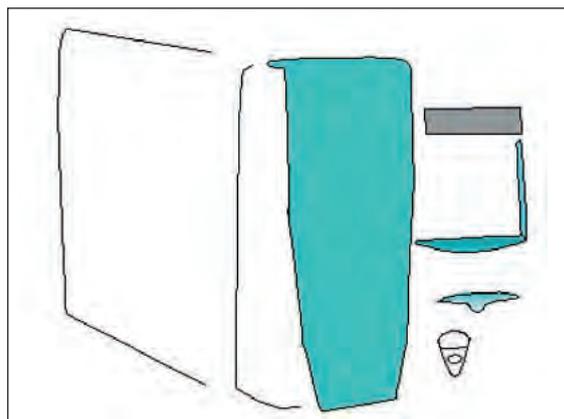
Into the third dimension!

The monitor started as a series of grey rectangles. I made the screen by taking a screenshot with *KSnapshot*, then pasting it in with Insert>Graphics. The lower part of the monitor has a pattern. I used Format>Area and made some new Hatching. This was applied to two boxes and the lines around the outside of them were made invisible. The tool to do this appears directly above the working area whenever you select a line or pre-made shape. I then drew another line around the outside of both boxes. I made a small box and made the contents the same hatching type I had used for the bottom part of the monitor, but with a transparent background instead of grey. Place it in front of the other part and it looks like the speakers. The manufacturer's names on the monitor and smaller PC were done with the "T in a box" text tool. This scales the text in the box when you scale the box. It's a good tool to use when you want to include text in your drawing but the font size won't go that small.

The illuminated button at the front was made with *OOo Draw's* 3D tools. The primitives provided can be modified (so, for example, a sphere can have the number of segments increased



The gradients dialogue, accessed by right-clicking and selecting Area.



Making parts to create the larger PC.

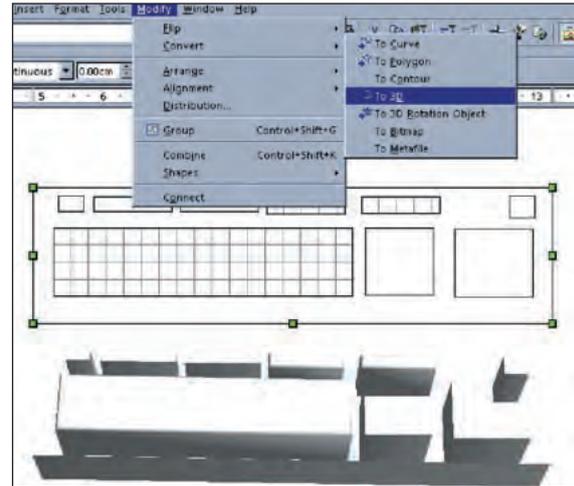
or reduced). You can change the lighting, the shadows and the colour of your model. Okay, it's not *3DS Max*, you are not going to do the graphics for a game in it, but it's nice to have. The switch is an illuminated sphere with a grey sphere in front of it. Select the shape you want from the 3D Objects tool bar, click and drag it on the work area. Once you have made it, you can use the Rotate tool to move it in three dimensions. Perhaps circles and squares are not much use to you, but *OOo Draw* can make 2D drawings >>>



The finished picture. The background is made simply by sending a coloured square to the back of the image.

◀ into 3D objects! I made the keyboard by drawing a series of squares. Modify>Convert>To 3D makes your shape into something that you can rotate. Once you have the object in the right position, use Modify>Convert>To Bitmap to make it 2D again. This saves some memory. I made one keyboard, then made a copy which was rotated and put on top of the second machine. At the bottom of the left-hand side tool bar is a 3D Effects icon. This gives control of the 3D lighting and geometry.

I also made a small space scene out of a pyramid, a sphere and a few squares. The texture map for the planet came from the Net, I drew the ship hull texture myself. The stars background was



The series of boxes that were used to make a 3D keyboard at the top of the page and the result below it.

mapped onto a square and I placed it behind the other objects. Remember that *Oo Draw* isn't a true 3D program – It's a drawing program with some nice extras. Even with half a gigabyte of memory, the screen display became very slow as the viewpoint was scrolled when looking at the drawing of the PCs. The same applied when moving the 3D objects.

I used the Text tool to write "Space is Dark" across my picture, changed the font to a pseudo-japanese style and used the Format>Fontwork tools to bend the letters into a semi-circle. The Fontwork tools allow you to put your text along a curved line, add shadows and other attractive effects.

Guide to the Oo Draw Tool Bars and interface

Expandable submenus that stay open to help you work

Across the top the default is the Line tools. The Edit Points tool is associated with Bezier tools, mentioned opposite.

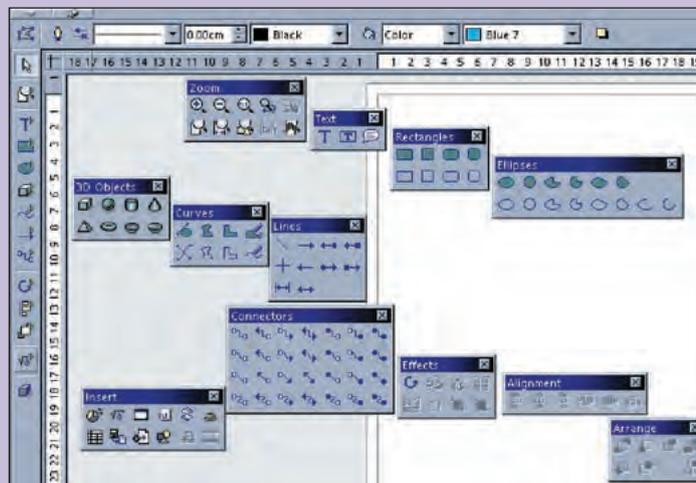
The 'pen nib' icon brings up a menu where you can adjust the properties of any line or shape in your drawing. Note that shapes have a line around the outside. Set the properties to Invisible to get rid of it.

The 'arrow head and tail' icon lets you put various shapes on the beginnings and ends of lines.

The drop-down box to the right sets the type of line. You can have dots, dashes and various combinations. The thickness is set in the next box. You can also type values directly in this.

The colour of your line follows. You have a broad range of ready colours in *Oo*, you have to add extra colours to *Oo* before you can use them. There's a Color dialogue under Area (if you right-click). Add a new colour and it will also be available to colour lines.

The Paint Bucket is the tool to fill the inside of shapes. The drop-down box has your fill options; Color, Gradient, Hatching, Bitmap. Change it to Hatching, then select the hatching type in the next box to the right and it will change. The



The sub-menus. The line of tools that are shown onscreen directly above the working area changes depending on what is selected.

last icon on the line turns on the dropped shadows. Again, you can set the shadow qualities simply by right-clicking and selecting Area.

All *Oo* tool bars that have a tiny green triangle can be expanded, there's a submenu that appears. If you click on the dark top bar of the menu, it will stay open, which is very handy.

THE ZOOM TOOL BAR

Useful to leave open while working. The "+" magnifying glass sign zooms, you have to click somewhere before it works. The "-" version works instantly. The other icons are specialised Zoom tools, giving you a 1:1 view, previous zoom level and an undo, a tool to display the entire width of a drawing (called a Slide by *Oo*),

show objects at their maximum and minimum sizes and a hand tool to move the display to one side or the other.

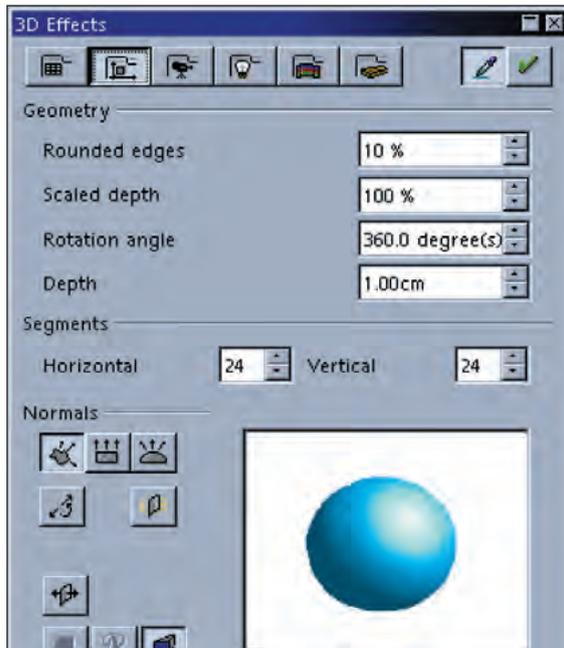
TEXT TOOLS

The "T" gives you a text insertion point. Either drag out a box to write in or just type and the box will expand to fit. The 'T' in a box' scales your font to the box size. Drag out the box, type in your text, then click on the arrow (or on any blank part of the screen to de-select the Text tool) and click and drag the box to see the text grow or shrink to fit the box. You can have the box either visible or invisible. The speech bubble (called a Call-out) is nice for labels. Click the item you want to label, drag your line to where you want the label to appear and click again. Next click the 'T' in the text submenu to type your description into the box. The label properties (thickness, colour and style) can be adjusted with the Line tools.

SHAPE AND CURVE TOOLS

The Rectangles tool draws either filled or empty squares and rectangles. You can change one to the other easily with the paint bucket tools. There are also rounded shapes. A right-click (or F4) will

OOo Draw is an amazingly capable program. We're aware that there are a lot of features that have not been covered, but there's enough here to get you started. The Help system is very well put together throughout the whole of *OOo* – remember to turn on the Extended Tips under Help and briefly hold the cursor over anything you don't understand.



The 3D effects dialogue. This is where you modify the 3D primitives and lighting. Experiment with different effects.

OOo Draw could very easily exist as a separate program, such as its capabilities. If you've even a passing familiarity with *Corel Draw* and *Freehand*, you'll agree that *Ooo Draw* isn't in the same class, but it's easy to use and has plenty of features for the average user. It likes plenty of memory too, but that isn't a problem – if a machine is running X, it should be fine.

It's hard to believe that Sun could Open Source such a brilliant tool as *OOo Draw*, but they are obviously convinced that software developed in this way has a myriad of benefits. *OOo* software runs on Linux, Solaris, Windows and a Mac version is currently under trial. Wait a year or so for the word to spread and with luck we'll be looking hard to find any Windows users who still use *MS Office*. *OOo Draw* is a great art resource, part of an excellent suite of office tools. Killer Application! [LXF](#)

Space Is Dark. The starship and planet are 3D, the text is OOo's Fontwork.



allow you to access the Position and Size menu. Here you can set the corner angle and make the shape slant. There's no preview, click OK to see the effect.

The Ellipses Tools draws circles and arcs, either filled or empty. The arcs (or pie-slices) are drawn by dragging out a circle, then placing one point, moving along the edge of the circle to the other and clicking. Again, you can alter the line and fill styles.

3D Objects places a primitive (that's a basic shape, not a monkey) on your drawing. Click and drag your square (or whatever shape you want) and alter it with the shaded 3D square at the bottom of the left side tool bar.

The Curves submenu is one that you really need to master. The "circle with a line across its edge" icon makes a closed Bezier curve. Click and drag to your first Control Point, then continue to click to define your shape. Double click to automatically join where you are to the starting point. The crossed line icon underneath the top icon in the submenu makes an open line. Double-clicking ends the line where it is. The filled polygon icon makes a closed shape with straight edges, the "L" shaped icon will make a

closed shape with 90 degree angles. I used it to draw around the base part of the monitor I made. Finally, there's a filled curve tool that would allow you to draw a shape, it would fill in the last line when you double-clicked.

LINE AND EFFECTS TOOLS

The Line tools are just click-and-drag. The cross is for drawing 45-degree lines. The dimensioning tool is here too. *OOo Draw* comes with 28 types of ready-made connectors to put into your drawings. Click and drag to place one.

The Effects submenu is worth keeping open. The Rotate icon places an target point in the centre of a selected shape or line – the centre of rotation. You can move it so that your shape rotates around a point that you define. The cursor changes to either a part circle with an arrow or two fine parallel arrows to show that a 3D rotation is possible, depending on what is selected and where the cursor is. The Flip icon mirrors the original along one axis, which you can move. Place your cursor over one of the green boxes that surround the selected object. It then changes shape, you click and drag the object to the new position.

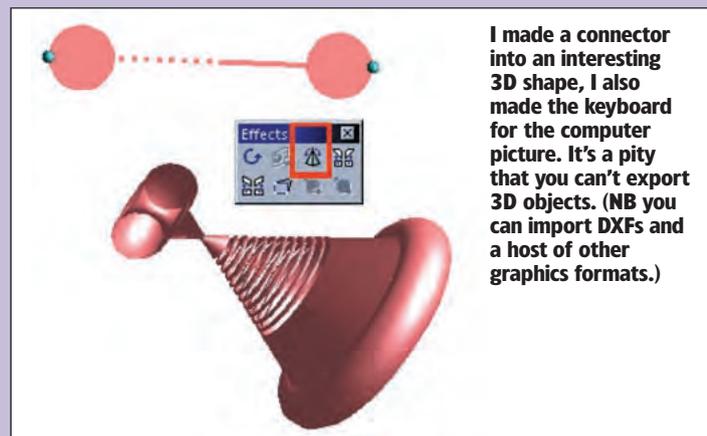
The icon to make a 2D object 3D can be seen in action in the screen shot shown below, highlighted in red. The other effects tools wrap a shape around a circle, distort a shape and can interactively make and alter gradients.

ALIGNMENT, INSERT AND ARRANGE

The Alignment tools place the selected object either along the top, side or bottom of your page. Keep the Arrange submenu open all the time if you are making a complicated drawing. Vector

drawing programs make each new shape or line appear in front of a shape or line drawn before it. If you want your latest shape to be partially covered by a shape drawn earlier, send it backwards until it is behind the earlier shape, or send it to the back so that it is behind every shape.

When I drew the space scene, I mapped the stars onto a flat square. It hid everything in the scene, so I sent it to the back. The Insert submenu allows you to put photos, graphs, calculations etc. into your drawing.



I made a connector into an interesting 3D shape, I also made the keyboard for the computer picture. It's a pity that you can't export 3D objects. (NB you can import DXFs and a host of other graphics formats.)

FLASH & MING

Practical PHP Programming



In the second installment of his three-part mini-series on multimedia, **Paul Hudson** shows how to add some whizz-bang to your site – using **PHP**, of course!

Flash generation is a hotly contested topic – while the creativity offered by Flash movies is all very well, many users are totally put off sites that insist on Flash content – sometimes due to its seemingly spurious animation possibilities, but often also due to the need for a browser plug-in which many do not have. If you're a Flash-hater, then re-join the series next month as we conclude with PDFs. Still here? Great!

Introduction to Flash

Flash is a proprietary format owned by Macromedia, which was opened a while ago, with Macromedia also offering large amounts of source code to developers who wanted to develop Flash-compatible programs – although *getting* the code required accepting the Macromedia licence agreement.

Since then, a lot of development work has been done by the *nix community to make a free-for-all version of a Flash generation library, and PHP has support for two Flash libraries: *libswf*, and *Ming*. *libswf* is fairly old, closed-source, and isn't developed any more, so generally people tend to skip by it. *Ming*, however, is LGPL, object-oriented, actively maintained, and heaps of fun to boot. If you're too young to remember *Flash Gordon*, you may not appreciate the *Ming* naming humour.

Flash movies themselves are vector-based – they remember the shapes you draw, the fills and gradients, etc, by way of mathematics; and can be resized and scaled with almost no loss of quality. Note that you *can* import bitmap pictures into Flash, but these don't resize well. For the purpose of this article, we'll be using the *Ming* extension to PHP. If you want to try out code examples as you read, read *Installing Ming* on the opposite page.

Before you begin, note that very often you need to tweak objects in your movie to get them to work just as you want them to. I've found in the past that a movie which looks perfect in *Mozilla* on Linux looks slightly off in *IE* on a Mac, causing me to have to make minor changes to bring them into line. Also, it's important to note that all values specifying some form of distance, length, height, or size are in 'twips', which is twenty units per pixel. Flash movies scale to fit their container, though, so these measurements are entirely arbitrary figures.

A simple movie

One of the biggest advantages to *Ming*, in my opinion, is that it is object-oriented – you create a shape object, tell it what colour it should be, then add it to the movie. In order to try to cram as much in as possible this month – Flash-generation code can be

very long – we're going to dive right into some code using some simple shapes, then pick it apart straight after. If your knowledge of Object Orientation is rusty or non-existent, refer to the *Object Orientation In A Nutshell* box overpage.

```
<?php
$mov = new SWFMovie();
$mov->setDimension(200,20);
$shape = new SWFShape();
$shape->setLeftFill($shape->addFill(0xff, 0, 0));
$shape->movePenTo(0,0);
$shape->drawLineTo(199,0);
$shape->drawLineTo(199,19);
$shape->drawLineTo(0,19);
$shape->drawLineTo(0,0);
$mov->add($shape);
header('Content-type: application/x-shockwave-flash');
$mov->output();
?>
```

Save that code as *ming1.php*. Generally speaking, you will want to embed your Flash movies inside web pages, and that requires inserting the following line somewhere in a HTML page:

```
<EMBED src="ming1.php" menu="false" quality="best"
bgcolor="#FFFFFF" swLiveConnect="FALSE" WIDTH="200"
HEIGHT="200"
TYPE="application/x-shockwave-flash"
PLUGINSPAGE="http://www.macromedia.com/shockwave/
download/index.cgi?P1_Prod_Version=ShockwaveFlash">
```

First we create a new instance of the **SWFMovie** class and assign it to our **\$mov** variable. An **SWFMovie** object allows you to manipulate attributes of the movie as a whole – size, colour, animation frame rate, etc. It's also used to add other Flash objects to your movie, so it's essential you hold onto the **SWFMovie** object created.

SetDimension() is an **SWFMovie** function that allows you to set the height and width of a movie by specifying values in the first and second parameters. Remember that Flash movies generally have their dimensions set in their host application – usually a web browser. The values you specify here are for the movie as you're creating it, however if the Flash movie is forced to display at a different size, your items will automatically be proportionally scaled to fit the assigned space.

Moving on to the core of the code, we have a new class – **SWFShape**. Unsurprisingly, we use objects of this class to manipulate shapes in Flash movies – the process is simply create, manipulate, then add to the parent movie object. If you forget to add your shapes to your movie object, the end result is that they'll be missing from the final output, so be careful. In the example above, the parameter that **SetLeftFill()** takes is the

PS

Too young to remember *Flash Gordon*? You probably will be – the original movie matinee series dates from 1936! If you haven't seen the 1980 movie or the long-running *King Features* comic strips, the humour behind the naming of Flash and Ming will be lost on you. A quote for the fans: "Flash, Flash, I love you – but we only have fourteen hours to save the Earth!"

return value of an **AddFill()** call. **AddFill()** is a function of the **SWFShape** class, and is overloaded (there's more than one version of it). The version used in the example above takes four parameters – the amount of red to use, the amount of blue, then green, and finally an optional alpha parameter. The fill returned by the **AddFill** function is used to supply the first parameter to **SetLeftFill()**, which is also overloaded. The end result is that the value passed to **SetLeftFill()** sets the fill on the left-hand side of the edge – in our example above, this is red.

Next we call **MovePenTo()** and **DrawLineTo()** several times. **MovePenTo()** lifts the drawing “pen” from the canvas and places it down at the X and Y points specified by the first two parameters respectively. **DrawLineTo()** moves the pen in the same sort of way, except that it doesn't “lift” the pen from the canvas first, meaning that a line is drawn from the last pen location to the X and Y parameters passed into **DrawLineTo()** respectively. **DrawLineTo()** is called a total of four times, giving us a box, and finally we call the **Add()** function of our **SWFMovie** object, **\$mov**, passing in our new box as the parameter – this adds the new shape to the final output.

The last two lines are crucial to the whole process, and must be used precisely as seen above. The first of the two calls the **header()** function, passing in the correct content type to instruct browsers that the information following is a Shockwave Flash movie. The very last line calls the **Output()** function of our **SWFMovie** object, which sends all the information you have prepared about your Flash movie out to your client. Once you've called this line, your script is complete.

To view your animation in action, load the HTML page you created earlier into your browser – all being well you should see something like the picture on the right. If your Flash movie doesn't load at all, it's generally the result of an error in the PHP script. However, when viewing the HTML page you *won't* see any PHP warnings, because the Flash movie is being sent direct to your browser's Flash player as part of a larger page. You can work around this by loading the Flash movie directly into your browser by visiting <http://yourserver/path/to/ming1.php> – you should see the errors printed as normal.

Flashy text

Through the use of a custom file format, Flash makes great use of vector text inside its animations. To get started, you need to download one or more fonts from www.neuralust.com/~mingdocs/fonts/getfonts.htm – these are a variety of popular fonts pre-packaged as Flash FDB font files. For this article, I'll be using Impact, so you might want that one at the very least, or you can choose your own.

Following the rest of the library, text inside your Flash movie is manipulated using objects. The two key classes here are **SWFFont** and **SWFText** – the former holds the actual font shape data, whereas the latter holds information about the text as a whole, including colour, position, string data, and the instance of **SWFFont** used to draw the letters. Place your chosen FDB file in the same directory as your Flash scripts, then save the following code into the file *ming2.php*:

```
<?php
$font = new SWFFont("Impact.fdb");
$text = new SWFText();
$text->setFont($font);
$text->moveTo(200, 400);
$text->setColor(0, 0xff, 0);
```

Installing Ming

It's not so merciless after all

Those of you with a Debian box can at this point simply type **apt-get install php4-ming**. You're done, and can get onto the trying out the code in this tutorial! Everyone else, follow these relatively simple steps, and you're all set to go...

1) Download Ming

Ming is available from its homepage at <http://ming.sourceforge.net/>. The download includes its own build instructions, but it's usually just a matter of a **make static** command, followed by **make install**.

2) Run **ldconfig** (might not be in your path, so run **locate** or **whereis** to find it)

3) Reconfigure PHP

```
cd [your_php_dir]
./configure --with-ming [your-other-normal-build-options-here]
```

4) Build and install php as usual, restarting your web server if you need to.

5) Add **extension=php_ming.so** to your *php.ini* file

6) Start using PHP to generate Flash!

Be sure to read the *Ming* installation instructions accompanying your download – it'll include more up-to-date instructions if they exist.

download:

- [version 0.2a](#)
- [+ \(.tgz\)](#)
- [+ \(.zip\)](#)
- [older versions](#)

other stuff:

- [function reference](#)
- [examples](#)
- [installation howto](#)
- [older news](#)
- [Ming wiki](#)

The *Ming* homepage on SourceForge is a mine of helpful information.

There is an installation HOWTO online at <http://ming.sourceforge.net/install.html> that explains various ways of installing it, and even offers a pre-built *php_ming.so* file for you to use if you'd rather avoid compilation.



```
$text->setHeight(200);
$text->addString("Text is surprisingly easy");

$movie = new SWFMovie();
$movie->setDimension(6400, 4800);
$movie->add($text);

header('Content-type: application/x-shockwave-flash');
$movie->output();
?>
```

Twelve lines of code and all we get is this lousy box? 'Fraid so – for now!

You'll need to alter your viewflash HTML file to display *ming2.php* rather than *ming1.php*, and I suggest you also change the WIDTH and HEIGHT attributes of the **<embed>** object so that the Flash movie is larger – otherwise you'll find the text is probably too small to notice.

Now that you've seen how it's done and what it looks like, onto how it works – we start with our two new classes, **SWFFont** and **SWFText**. The **SWFFont** class is remarkably easy to use – merely pass the name of the FDB file you want to use as a font, and save the return value for later use. You should already have downloaded an FDB font (or created your own using *Ming's* **makefdb**), so you should replace **Impact.fdb** in the example with your own font. In line two of our script, we create a new **SWFText** object and store it in a **\$text** variable. This object



works in pretty much the same way as our previous `SWFShape` object – we set various properties of it, then add it to the parent movie once we're done.

The first thing we do with our `$text` object is call its `setFont()` function, which makes this `SWFText` object render in the font used to create the `SWFFont` object specified as the only parameter. In our case, we created our `SWFFont` object using `Impact.fdb`, so calling the `setFont` using the new `SWFText` object will draw the text in this object using the Impact font. Next, we call the `moveTo` function to place the text inside the movie, then call the `setColor` function (the values are hexadecimal) to set the text to lime green. The `setHeight` function sets the height of the text in twips, but again remember the final size of the text is entirely dependent on the size the movie is played back, and also the dimensions of the parent movie object itself – the value you set here is just relative to the rest of the movie.

The final, and most important function we call for our `SWFText` object is `addString()` – this allows us to draw the string passed as parameter one to the position we set with our `moveTo()` call. It's important to note that the pen with which the text is drawn is set to the baseline – if you use `moveTo()` to set the position to `0,0`, the text drawn will be drawn outside of your movie. You should recognise the last five lines of code from our previous script – we create a movie, set its dimensions to be nice and big, add our text object, then output the movie. As you can see, the basic process is the same as when working with shapes, so this should be fairly easy for you by now.

Actions

Through its powerful ActionScript language, Flash provides a very flexible scripting environment to allow developers to take more direct control over the operation and flow of their script. For

Flash text is slightly more impressive than boxes, but only slightly.

Text is surprisingly easy

Object Orientation in a nutshell

In a class of its own...

Object Oriented Programming (OOP) is one of my favourite programming concepts – simply put, you define your program in a set of autonomous sections that interact with each other by calling functions. There are a few pieces of jargon you'll need to know...

A class is a definition, or a blueprint, of one part of your program. Classes are used to define "things" in your program – for example, in a program modelling a house, you might have a class for a room, a class for a TV, a class for a telephone, etc.

An object is an instance of a class. Your `TV` class covers a general `TV` definition – an object of that `TV` class might be the `TV` in your lounge. From the blueprint class of the `TV`, you can create multiple `TVs` and each are then individual – they can have different colours, positions, etc.

Each class has various properties and functions inside it, and each instance of that class has a copy of those properties and functions. You might say that your `TV` class has a

`ChangeChannel()` function and a `$Channel` variable, which would mean that each `TV` object you create from your `TV` class has a `ChangeChannel()` function which changes the `$Channel` variable. Note that calling `ChangeChannel()` doesn't change all the `$Channel` variables of your `TVs` – just for the `TV` you called `ChangeChannel()` on.

PHP's Flash implementation uses the *Ming* library, which is object-oriented. To create a Flash movie, you create an object of the type `SWFMovie`; to create a shape in your movie, you create an object of the type `SWFShape`, etc. As long as you have at least a fairly simple grasp of OOP, you should be able to follow the code in this article – if you encounter problems with classes and objects, you'll find a large collection of OOP tutorials on the web.

Lastly, you'll need to know that a constructor is the function that is automatically called by each object when it is created. It sets initial values, allocates memory, etc.

example, you can call `stop()` to stop playing the movie, then `play()` to continue; `gotoFrame()` allows you to jump to a particular part of your movie, and `getURL()` allows you to browse to a new web page. There's a large collection of actions available to you, and the PHP documentation has some very good (if very long) examples on how to make use of various functions.

In order to give you a quick start, I'm going to take a look at a quick example of actions. Save this next piece of code as `ming3.php` and modify your HTML to point to the new file:

```
<?php
function MakeActionBox($red, $green, $blue){
    $shape = new SWFShape();
    $shape->setLeftFill($shape->addFill($red, $green, $blue));
    $shape->movePenTo(-100,-20);
    $shape->drawLineTo(100,-20);
    $shape->drawLineTo(100,20);
    $shape->drawLineTo(-100,20);
    $shape->drawLineTo(-100,-20);
    return $shape;
}

$button = new SWFButton();
$button->setUp(MakeActionBox(0xff, 0, 0));
$button->setOver(MakeActionBox(0xff, 0xff, 0));
$button->setDown(MakeActionBox(0, 0, 0xff));
$button->setHit(MakeActionBox(0, 0, 0));
$button->addAction(new SWFAction("getURL('http://www.
linuxformat.co.uk;
'lxr');"), SWFBUTTON_MOUSEUP);

$movie = new SWFMovie();
$movie->setDimension(200,200);

$displayitem = $movie->add($button);
$displayitem->moveTo(100,100);

header("Content-type: application/x-shockwave-flash");
$movie->output();
?>
```

The first thing you will notice is that I define a custom function, `MakeActionBox`, to handle some of the grunt work you'll experience when working with the `SWFButton` class. The `SWFButton` class, which I make an instance of for our `$button` variable, has several 'states' that each require a shape – how the button looks when it's up, when the mouse is over it, when the mouse is clicked on it, and where the mouse can be clicked on it.

Each of these states require a complete shape of their own, so I automate the process of setting up a shape by using the function `MakeActionBox`. Going through the main chunk of code line by line, you can see the first thing I do is create an instance of `SWFButton` and file it away in the `$button` variable. I then call four functions – `setUp()`, `setOver()`, `setDown()`, and `setHit()` – to define how this button should look when the user interacts with it. My implementation is short to save space, but you'll generally find it's more visually appealing to have more than just the colour change between states.

Next we come to the important function of this particular script – `addAction`. `addAction` takes two parameters – the `SWFAction` object to add, and a flag – when the action should execute. Options include `SWFBUTTON_MOUSEUP` as you see above, or alternatively `SWFBUTTON_MOUSEDOWN`,



While it might look fairly similar to our last box, you'll have to trust us that this one is clickable!

SWFBUTTON_MOUSEOVER, and more – you'll need to see the documentation for a full list.

As the first parameter to **addAction**, we pass in **new SWFAction(...)** – the constructor of the **SWFAction** class takes a string that contains the ActionScript code you wish the action to execute. For this action, which will execute when the user clicks the mouse button on the object, we want to execute the **GetUrl ActionScript** function. In the example, **GetUrl** is passed two parameters – the URL to load, and the name of the window to load it in. If the named window doesn't exist, it will be created for you. To go much further into ActionScripting, I suggest you refer to the online **ActionScript** dictionary at www.macromedia.com/support/flash/action_scripts/

So, the **addAction** line translates to "Create a new ActionScript action that will load the *Linux Format* website into a new window, then attach that action to our button so that it executes whenever the user clicks the button". After the action code, there's a slight change to the normal procedure – we use **\$movie->add** as before, but you'll notice that this time we grab the return value and store it in the **\$displayitem** variable. The reason that this is done because, when adding shapes, text, buttons, and sprites to a movie, the **add()** function returns a special type of object, **SWFDisplayItem()**, which is a handle to the object inside the movie. This means you can add the same button (or shape, text, etc) to the movie several times over and manipulate them individually without much fuss.

This functionality is important here, because you can't manipulate the position of an **SWFButton** object directly – you need to add it to the movie first, then manipulate the position of the returned **SWFDisplayItem** object, and you can see the line after the **add()** call we do just that. Finally, the movie is sent to output as per usual. If you'd like to give your button more visual impact (we only have so much space here!), you might want to try combining the previous code regarding text with this button to make something more interesting.

Animation

And now, the bit you've all been waiting for! Adding animation to your Flash movies is at once a marvellously fun and incredibly tricky affair. If you're feeling a little unsure of your Flash skills so far, I would recommend you go back a little and read through the parts you need to brush up on before you continue.

The key to animation is the **SWFDisplayItem** object returned by the **add()** function of your movie object. **SWFDisplayItem** objects have a variety of functions that allow you to move, rotate, scale, and skew your objects easily. This next example (save it as *ming4.php* and amend the HTML again) demonstrates how animation works, by modifying our previous font script:

```
<?php
$font = new SWFFont("Impact.fdb");
$text = new SWFText();
$text->setFont($font);
$text->moveTo(300, 500);
```

```
$text->setColor(0, 0xff, 0);
$text->setHeight(200);
$text->addString("Text is surprisingly easy");

$movie = new SWFMovie();
$movie->setDimension(6400, 4800);

$displayitem = $movie->add($text);

for($i = 0; $i < 100; ++$i) {
    $displayitem->rotate(-1);
    $displayitem->scale(1.01, 1.01);
    $movie->nextFrame();
}

header('Content-type: application/x-shockwave-flash');
$movie->output();
?>
```

You should recognise nearly all of that from the previous coverage of fonts. Notice that the **\$movie->add(\$text)** line has now changed so that the return value is captured and stored in **\$displayitem**. The entirely new code block follows immediately afterwards – we run through a loop 100 times, each time calling **rotate()**, **scale()**, and **nextFrame()**. I'll cover **rotate()** and **scale()** in a moment, but first I want to explain **nextFrame()**. You use **nextFrame()** each time you want to move forward to the next frame of your Flash animation. Animation works by defining the initial state of the movie, advancing the frame, then specifying changes from the previous frame.

The **rotate()** function takes a single parameter – the floating-point value of the amount to rotate your **SWFDisplayItem** object from its current rotation. In our example, I've used **-1**, which means it adds **-1** of a degree of rotation with each frame, which, because of the way Flash rotation works, means that the text rotates in a clockwise manner. The **scale()** function takes two parameters – the amount to scale the object's width, and the amount to scale its height. Again, this is based upon its last state, which means that the scaling is compounded and therefore by **0.01** to the size of our text over 100 frames we are almost tripling the size of the object.

So, the entire contents of the **for** loop translates to "rotate slightly, scale slightly, next frame" one hundred times. See – animation is actually quite simple! Granted, you'll find things get harder when you have various different timings going on, but you have a good grounding now at least. There is a lot more to creating Flash movies with PHP than we've had time to look at here. I suggest you look into the **SetRate()**, **Save()**, and **StreamMP3()** functions of **SWFMovie**, as well as the **SWFFill** object as a whole – adding gradient fills is a doddle.

Conclusion

Creating a Flash movie in PHP involves first getting over the fairly steep learning curve. When you're getting started, figuring out all the classes, objects, animation, and jargon will take a good amount of time and effort on your behalf, but we believe you'll find it pays off – you can really get some cool results with Flash, and even more so when combined with PHP.

I strongly suggest you write your own functions to handle as much of the repetitive code as is possible – once you can get the boring stuff down to one function call, it becomes much easier (and a lot more fun!) to write Flash code. [LXF](#)



Animating text is only a fairly short jump from where we were before.

About Paul Hudson

Paul Hudson is a London-based web developer specialising in PHP and Perl programming. He can be emailed at hudzilla@php.net

NEXT MONTH

Next month we'll be continuing the multimedia creation mini-series with a look at how PDF files can be created dynamically using PHP. PDF files are in many ways a key component of the Internet in that they allow users to download and print off an exact copy of official documents and magazine articles, and you'll find that PHP has some really cool functionality to support PDF manipulation.

Any comments or suggestions about this PHP series? Please be sure to drop us a line!

Answers

If you are really stuck and the HOWTOs yield no good result, why not write in? Our resident experts will answer even your most complicated problems!

Our experts

Whatever your question is, we can find an expert to answer it – from installation and modem woes to network administrations, we can find the answer for you – just fire off a letter or email and it'll all be taken care of by the *Linux Format* boffins.

LXF answers guy
David Coulson is a networking and security guru with plenty of sysadmin experience to boot.



Nick Veitch is the editor of the magazine, and answers your easy questions! Or indeed anything to do with *Grub*, *LILO*, *netatalk*, vi...



MySQL server

Q I have installed *MySQL* together with Red Hat 8.0 or Mandrake 9.0, however, when I run *MySQL*, the system prompts this same error code:

Error 2002: Can't connect to local MySQL server through socket '/var/lib/mysql/mysql.sock' (111).

I already did **chmod 755 /var/lib/mysql**, it would not work! Your help and advice? Thanks for the attention.

Conrad Fong

A /var/lib/mysql/mysql.sock is a Unix domain socket via which your *MySQL* client will talk with a local server. In order to be able to use this, you will need to have your *MySQL* server running, which will likely require you to install another package containing *mysqld*, the server part of *MySQL*. If you have another *MySQL* server already running on a remote host, you will need to use the **-h** flag to set the hostname or IP. More information on the *MySQL* client can be found with **man mysql**.

00o dict dirs

Q Just to let you know that there was some problems outlined with installing the UK dictionary... unfortunately I tried your method and it did not work. I did a search via nifty *find* utility for *dictionary.lst*, firstly, this was actually a typo and should have read *dictionary.lst*. Secondly, the enclosed dictionaries that came with the CD does not overwrite this file from what I've followed in your method as the enclosed zip files does not contain *dictionary.lst*! Thirdly the location of where the dictionaries are stored is not in share.... to make it easier, issue this command inside the directory where *OpenOffice.org* is installed – off the top of my head it was (for

OpenOffice.org1.01):

- 1 **cd** *OpenOffice.org1.01*
- 2 **find ./ -name dictionary.lst -print**
- 3 the directory where it finds it, **Bob's your uncle**, so change to that directory.
- 4 **Unzip** the files into that directory.
- 5 **Make** a backup copy of *dictionary.lst* and edit it to include **DICTIONARY en_GB** or similar – check the original copy first beforehand.
- 6 **Start up** *OpenOffice.org* and go to tools, options, select languages and a check mark appears beside **English UK** and you're done.

Excuse the rough layout of what needs to be done, as this was the top of my head... and apologies if my criticisms seem harsh but hey, your magazine rocks! I Love it ... keep it up folks!

Tom

A It sounds like you've managed to work around some problematic documentation and achieved the goal – Well done! It's always a good idea to share your experiences with other Linux users as you have done, as nine times of ten someone will have beaten you to the problem and to the solution too, more often than not.

```
david@tailtiu:~ (pts/0)
Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> status

mysql Ver 11.18 Distrib 3.23.52, for pc-linux-gnu (i686)

Connection id:          73178
Current database:      mail
Current user:           root@localhost
Current pager:         stdout
Using outfile:          ''
Server version:        3.23.52-log
Protocol version:      10
Connection:            Localhost via UNIX socket
Client characterset:   latin1
Server characterset:   latin1
UNIX socket:           /tmp/mysql.sock
Uptime:                11 hours 34 min 40 sec

Threads: 11  Questions: 376557  Slow queries: 0  Opens: 1601  Flush tables: 1  O
pen tables: 64  Queries per second avg: 4.834

mysql>
```

MySQL can connect to a server via either TCP or Unix sockets, the latter being more appropriate for local connections.

We didn't hear from anyone else who had this problem, so either they were nice and quiet, or they didn't have the same issue as you with the dictionary. Hopefully, they may appear from the woodwork and thank you for taking your time to figure it all out.

Linux for Macs

Q I have an old Mac PowerPC 9500/120 for which I could try out Linux. However, I am not very sure about Linux for the Mac, even though I have bought the latest *Linux Format* mag.

A few questions to help me on my way:

- 1 Is there a version of Linux suitable for this Mac? If so how do I get it?
- 2 I've read somewhere about having to partition HDs to run Linux. Would I have to do this on the Mac?
- 3 I'm not a programmer so if the answer to 1 is Yes, does it come with simple installation instructions? Or, can I get these elsewhere on the Internet?
- 4 On the DVD for the December magazine there is *OpenOffice.org*. If I manage to install Linux on

my Mac, will I be able to install this software?

I'm sorry if all this is a bit basic but I am keen to try out Linux on both my Mac and PC!

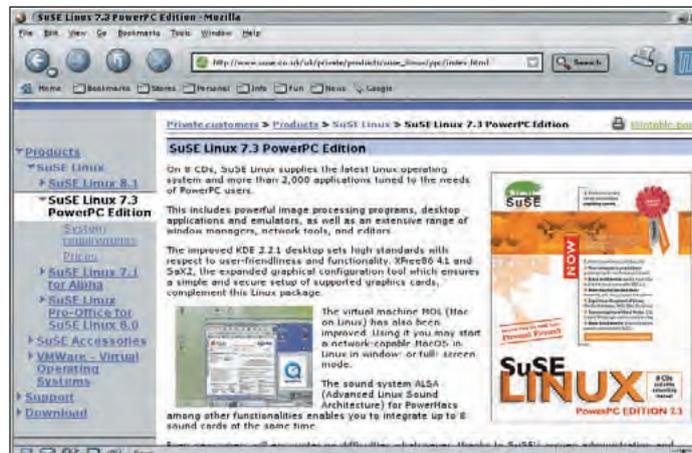
Keith A Rogers

Many Linux distributions are available for the PowerPC (PPC) architecture, including Red Hat, Debian and SuSE. If there is a particular distribution you prefer to use, then see if there is a PPC build of it that you can use. You do need to partition your hard disk, so you can continue to use MacOS, but this is all part of the installer.

As for OpenOffice, you will need to either obtain PPC binaries, or build it by hand using the source code. You can easily check if a binary release is for PPC, as it will contain 'ppc' in the filename. If it's for an Intel/AMD system, it will be called 'i386' or 'i686', or indeed any 'ix86' architecture available.

Lilo boot

Thanks for an inspirational magazine. I have used the DVD from May 2002 to install Mandrake 8.2 successfully



SuSE Linux is one of many Linux distributions which provides a PowerPC build of their distribution.

several times. Recently, after having changed the graphic card, the PC that I use is not able to download all the packages necessary from the DVD for a proper Mandrake 8.2 re-install. This problem recurs every time, in a different way each time. The graphic card is listed as compatible: nVIDIA RIVA TNT2 M64 Pro. There are a few micro scratches and some discolouration on the DVD, nothing that would affect a normal CDROM.

If I send you the DVD, could you check it out and send me a new one; even if I need to purchase it?

Additionally I recently discovered to my horror that this DVD wrote *Lilo* or *GRUB* to, not the boot sector on IDE-0 or IDE-1, but to the BIOS boot sector! Despite the BIOS boot sector being write-protected AND password protected, Trend Micro did not alarm! No virus has ever succeeded in silently hacking the BIOS on my PC before!

[Amibios, MSI motherboard MS-5169, AMD K-6 II 450MHz processor , 384MB RAM, IDE-0 has Windows 98 x 2 (2 partitions), IDE-1 has Windows 98 and Mandrake 8.2]

If I boot with the DVD, press **F-1**, type **rescue** <enter>, select: **Restore Windows Boot Loader [OK]**, and follow instructions... nothing happens – error message: **unmount / failed**. Unless I stumble upon some secret routine, I will never again be able to use this computer for “Windows only” configurations.

Is there anyway to force install *Lilo* only on the boot sector of the second hard drive IDE-1 so that Linux will only be available if the BIOS is set to boot to IDE-1?

I am not alone. Several of us have experimented for years with Linux and seldom gotten past the trial stage. We check each other's efforts, even video filming a sequence to re-check afterwards if the user really did tap the keys we thought he did. It's the weird stuff that stops us. The stuff that can't happen (except to everyone we know). The weird stuff that isn't possible (except in the real world). Time is at a premium, endless experimentation and daily reading of all relevant newsgroups isn't possible.

Paid support gives up when results are not readily forthcoming, often they imply we are making all this up, or suggest buying a new CD. All network problems are always the cable's fault: “buy the proper cables”. Still they still expect to be paid.

There seems to be a definite market for a magazine or database offering real world solution in English for impossible install and operation problems. The “weird stuff” that can't happen... except in the real world.

Mark

You've not told us why you are unable to access the DVD from within Mandrake. If you get different errors each time, it would suggest that you are using a different method to achieve access to the DVD. As always, an error message is very useful, since there are many, many reasons why you may not be able to access the DVD. If you can outline the exact method you are using to access the DVD and the error message you

A QUICK REFERENCE TO: PGP

PRETTY GOOD PRIVACY

Privacy on the Internet is a major issue, particularly now that Western Governments are looking at implementing various methods to monitor email traffic produced by both businesses and private individuals. While this was quite expected and it's difficult not to blame governments for wanting to keep track of electronic communications, as they currently do with letters and telephone calls, it does present the problem that the Internet can no longer be considered private for emails.

Pretty Good Privacy, otherwise known as PGP, allows us to generate a unique key for signing or encrypting our emails. We can send an email to an individual so that only someone knowing their private key is able to produce the original plain text version of the email. We can also sign an email with our key so that the recipient can ensure that the email has not

been modified by authenticating it with our public key.

As PGP is a public key encryption system, we distribute our PGP key on a system known as a keyserver, such as www.keyserver.net or www.keys.uk.pgp.net We can download someone's public key in ASCII format and add it to our keyring for use by PGP. PGP can download keys from a server on demand, so most mail systems will download a unknown key for us automatically.

PGP requires keys to be trusted, so it's important to have your key signed by other individuals to create the 'web of trust'. You can't just go around and sign people's keys, and usually it requires confirmation of identity with a photo ID, to ensure that people are who they say they are. If keys were signed by anyone without confirmation, then the web of trust would be useless and no one could trust any keys on the

PGP key server as they could not know if they were genuine keys with real trusted signatures or keys which they could not trust.

Many mail clients, including *mutt*, support PGP natively and only require the *pgp* binary to be installed on the system. Others, such as *pine* and *Mozilla*, require additional plugins, *pgp4pine* in the case of *pine* and *Enigmail* for *Mozilla*. Both of these can be found on freshmeat.net. Many other mail clients have PGP capabilities, so you will need to check the site of the client of your choice in order to m

PGP is available at www.pggi.org, but GNU have produced their own system which is compatible with PGP, known as *GPG*, or *GNU privacy guard*. Many distributions have packages of *GPG* available and nearly everything which can use PGP will work with *GPG* too without any modification.

FREQUENTLY ASKED QUESTIONS MOZILLA

FAQ What is Mozilla? It's sort of like Netscape, isn't it?

Mozilla was developed using the source code originally used to build *Netscape4*. However, it has practically all been rewritten, so apart from the general look and feel, it's a completely new browser.

Netscape, which is now owned by AOL, has taken the *Mozilla* code and modified it slightly to produce *Netscape 7*. However, *Mozilla* is still in development, so *Mozilla* should be considered over *Netscape 7* when looking for a browser.

FAQ Where can I get Mozilla?

Mozilla can be downloaded from www.mozilla.org and is also packaged by many distros for easy installation. Debian users need only do **apt-get install mozilla** and it installs everything for you. *Mozilla* can be downloaded from mozilla.org in both binary and source releases, and as the *Mozilla* team generate optimised binaries, it's often a good idea to just grab the binary tarballs, as it can take a matter of hours to build *Mozilla* on a reasonably powerful system, so it's not something you want to rebuild a few times every day on a PII-500.

FAQ Is it stable, or is it liable to fall over on me and burst into flames?

Mozilla released their 1.0 build quite a while ago and the current stable release is 1.21. There are also development releases and nightly builds from the code residing in the cvs tree, so you can select the level of 'bleeding edge' you prefer. The current development release is 1.3a, which may or may not work depending what you do with it.

FAQ What can I get done with Mozilla?

Mozilla supports everything *Netscape 4* does, including a web browser, mail client, address book and HTML composer to create web sites. There are also a number of other additions, including an IRC client. Due to its development, *Mozilla* also includes various debugging tools for Javascript, HTML and CSS.

FAQ I heard there was lots of cool stuff in Mozilla? What does it do that Netscape doesn't?

Generally, *Mozilla* is significantly more stable than *Netscape 4*, and the rendering engine known as *Gecko*



***Mozilla* (www.mozilla.org) is a very powerful web browser, which supports all of the current standards for HTML and CSS, making it a great choice for developers**

is completely compliant with the HTML, XHTML and CSS standards, as well as being much faster than the HTML rendering system which was used with *Netscape 4*. There are a number of HTML and CSS tests which are available with *Mozilla* to show off its swanky rendering capabilities and all the fancy stuff it can do with CSS to make your sites even more exciting.

FAQ Are there any third party addons for Mozilla?

There is a multitude of third-party projects for *Mozilla*, many of which can be found over at www.mozdev.org. Many of these add extra capabilities to *Mozilla*, or extend upon those which are currently available. Indeed, a number of mozdev.org projects have been merged into the main *Mozilla* tree, exposing these projects to a much wider audience.

Hop on over to mozdev.org and take a look at the top 50 projects and see if there are any which interest you. Some are quite silly and don't do anything particularly useful,

receive, then we can provide a more useful answer to your query.

On the *Lilo* front, there is no such thing as the 'BIOS boot sector'. *Lilo* will write to either the boot sector of a particular hard drive, usually /dev/hda, or to the boot sector of a partition, such as /dev/hda5. *Lilo* is not specific to a machine or an IDE channel, it is specific to a hard drive or partition. If you want to put *Lilo* on another hard drive, you simply need to change /etc/lilo.conf and modify the **boot=** option to point to another hard drive or partition.

Lilo certainly didn't write anything to your BIOS, and we would be interested to find out how you came to such a drastic and wild conclusion. There are numerous Linux users who have installed it on their system

without any problems and we've heard from many individuals who have written to us with their success stories. Searching newsgroups may sound like time-consuming work, but if you paste your error message into <http://groups.google.com/>, it'll usually spit out something useful. There is a wealth of information on Linux available on the Internet, with newsgroups, IRC channels, mailing lists and so forth, and most problems can easily be overcome by finding out what someone else has done to get around it.

make (no) install

Q On pages 102/3 of the latest issue, *Christmas 2002*, you mention the program *Stellarium*. In the text

describing this program you say "... if you want to try it out first, skip the make install part and run Stellarium with ./src/stellarium"

My question is; is this an option that is applicable for every program that is compiled?

Many times I have compiled and installed a program using the full make install process, only to find that it really wasn't the program I was after. I then have this laying around on my hard drive. Or is there a way to delete a program installed this way that I haven't discovered yet?

Thanks for a great magazine and also thanks for Mandrake9.0. I've installed it and everything so far is working just hunky-dory. :-)
Dave Le Huray

A Some programs can be executed from the build folder, as they have not built any libraries along the way. Usually, however, things build support libraries that need to be installed before the program can be used, so that the binary can link with the libraries correctly. It's really a matter of simply trying to run it from the command line and seeing if it works.

With source tree builds, many support **make uninstall**, which just deletes everything it installs. It does not do any checking, so if something else has overwritten something with a different version, then it will simply delete it. Packages are a much better way to handle things, or you can pass a **--prefix=/usr/local/<pkgname>** option to **./configure**, so that

but others are extras which make you wonder how you did without in the past.

FAQ **Oops! Mozilla died on me. How do I report this?**

The *Mozilla* guys developed a bug tracking system known as *Bugzilla* which holds all of the *Mozilla* bugs. Anyone can submit bugs to *Bugzilla*, at http://bugzilla.mozilla.org/enter_bug.cgi?format=guided or search for existing bugs and find out the status of them with a particular *Mozilla* build.

In order to submit problems to *Bugzilla*, you will need to register with the *bugzilla.mozilla.org* site, and then you can post bugs to the system. As always, it's a good idea to include as much detail as you can in order to ensure that your bug can be reproduced and studied by those who can fix it.

FAQ **Is Mozilla comparable to Evolution?**

Not at all. *Evolution* is a PIM, where as *Bugzilla* is a web browser. *Evolution* is more comparable to something like *Microsoft Outlook*, rather than *Bugzilla*. Of course, one can use both *Bugzilla* and *Evolution*, effectively replacing the need to have *Internet Explorer* and *Outlook* available.

'stellarium' installs into '/usr/local/stellarium', although you will need to add the appropriate lib directory to /etc/ld.so.conf before being able to use any dynamic libraries created by the build.

smbmount

Q First off, your mag is great! Now on to business. I've been using Linux for about two years now. I've settled on SuSE Linux and am using version 8.0 currently. I have a small home network which comprises of an HP server running SuSE 8.0, a Compaq Deskpro 4000 running SuSE 8.0, a PowerMac 7500 running MacOS 8.6, a Compaq Armada Windows NT 4 workstation, A Fujitsu ErgoPro E running SME Server 5.5 with

Blade 1 installed and a Windows 2000 Pro Dell Inspiron 7500. The Ergo Pro acts as DHCP, proxy (using dialup via USB modem) and file server for the network.

I have been having problems getting my Compaq Deskpro to connect to the SME Server to access my personal Home drive. When I first set up the SME Server and after installing the Blade, there was no problem. I update SUSE 8.0 on the Deskpro as well and there was no problem, not even after I rebooted. Then one day out of the blue I typed in

smb://smeserver/ggreaves in Konqueror and I get

An error occurred while loading smb://paratha/ggreaves/tars:

Internal Error

Please send a full bug report at <http://bugs.kde.org>

Unknown error condition

I thought I needed to reboot, so I did, but that didn't work. I can access my home drive from the Windows NT workstation, Windows 2000 Pro and Mac. I've logged a bug report with KDE but no response to the problem has been logged (bug id 48804). A few people have seen this bug but haven't offered any advice. Have you got any clues?

My next problem is that *KCron* does not work. I know that *crond* is started. What I want to do is have it start *xmms* with the *-p* switch. Using *KCron* or *VCron* I get an error message **command not found**. Entering */usr/X11R6/xmms -p* nothing happens and entering */usr/X11R6/xmms* just gives **** CRITICAL **: Unable to open display**. Using *KCron*, if I were to use the execute/run now *xmms -p* works. What am I doing wrong?

A It sounds like *Konqueror* is having problems with SMB filesystems, so it might be worth checking the KDE lists to see if anyone has posted about this. Things generally don't just break all on their own, and remembering if you performed a kernel upgrade or upgraded KDE packages should shed some light on the situation. Another option is to look at using *smbclient* at the command line to connect to your server, as this should throw up any connectivity or authentication problems that may have made *Konqueror* fall over. You can use

```
david@macha:~ (pts/13)
macha:~ 10:01pm Thu Dec 19
david:~ # smbclient //prendergast/Prendergast -U administrator -I 10.3.1.2
added interface ip=10.2.1.4 bcast=10.2.1.255 nmask=255.255.255.0
added interface ip=10.3.2.3 bcast=10.3.2.255 nmask=255.255.255.0
Password:
Domain=[COELAKE] OS=[Windows 5.0] Server=[Windows 2000 LAN Manager]
smb: > ls
arcldr.exe          AHSR  148992 Tue Dec  7 07:00:00 1999
arcsetup.exe       AHSR  162816 Tue Dec  7 07:00:00 1999
BJPr inter         DA    0 Mon Dec  2 18:09:57 2002
boot.ini           HS    191 Sun Dec  1 18:49:51 2002
Documents and Settings DA    0 Mon Dec  2 09:42:58 2002
FTM                0 Thu Dec 12 22:39:53 2002
NTDETECT.COM      AHSR  34468 Tue Dec  7 07:00:00 1999
ntldr             AHSR  214416 Tue Dec  7 07:00:00 1999
pagefile.sys      AHS  2145386996 Thu Dec 12 18:07:47 2002
PAPRPORT          0 Wed Dec 11 18:31:17 2002
Program Files     DAR   0 Wed Dec  4 19:01:08 2002
RECYCLER          DHS   0 Mon Dec  2 18:08:01 2002
System Volume Information DHS   0 Mon Dec  2 09:42:37 2002
WINNT             DA    0 Wed Dec 11 13:36:15 2002

39997 blocks of size 1048576, 36630 blocks available
smb: > |
```

smbclient will allow users to access Windows shares without needing to mount the filesystems locally which requires root access.

smbmount to mount the SMB share locally and then browse it with

Konqueror:

mount -t smbfs

//parathea/ggreaves/tars /mnt/smbfs

Your problem with *cron* is because the **DISPLAY** variable is not set. *cron* is not aware of your X server, so you need to write a small shell script to start *xmms*. Of course, one has to ask the question why you want to start *xmms* via the *cron* system, as all this will do is start *xmms* playing whatever is set in the playlist. To write the shell script, you would create a file, such as *xmms.sh* and put it in your home directory, which contains:

```
#!/bin/sh
```

```
export DISPLAY=:0.0
```

```
/usr/X11R6/bin/xmms -p
```

And then do:

```
chmod 755 xmms.sh
```

You can run this from the command line to check it works, then add the path to your *cron* entry instead of supplying the path to *xmms*. You didn't say why you wanted *xmms* to start playing frequently, so we would be interested to find out quite why you are using *cron* in this fashion.

jig don't

Q I have been using Linux for some time now in various forms, but my favourite distro is Mandrake. I was therefore pleased that you have continued to put this distro on you cover DVD. I have successfully installed it straight from the DVD, but I have a second computer that I wished to install it on that only has a CDROM. I have tried creating the ISO images following your instruction in the magazine using *jigdo*, but when I

run the script the program halts with error message

```
./jigdo-lite: line 1: jigdo-file:
command not found
```

```
./jigdo-lite: line 1: jigdo-file:
command not found
```

Error - template checksum mismatch!

The .template file does not belong to the *jigdo* file - the chances are high that the image generation process will break. I will abort now. If you know better than me and want this error to be ignored, enter the string "42" to proceed.

I have also tried running *jigdo* directly, asking for individual ISO image files to be created and I get the same message.

Can you please help? I would like to install Mandrake 9.0 onto my second computer.

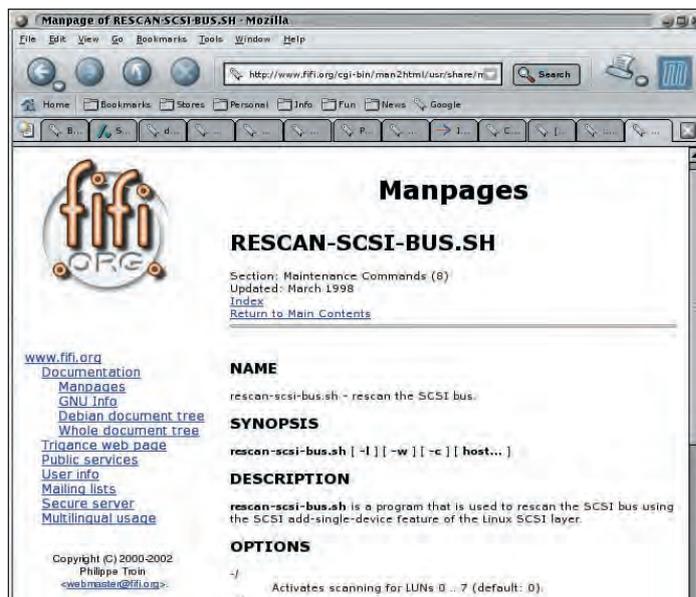
The Windows version also did not work, but I am running Windows ME and you did say people are having problems running *jigdo* under Windows ME.

I have also tried running Mandrake's *Makecd* script, which I managed to get to run but it aborted after about one hour saying there was a file missing. So I am having no luck using all three methods.

Paul Midgley

A It seems like the *jigdo-file* is not in a directory in your **\$PATH** environmental variable. As you're running the program with *.jigdo-lite* it would suggest that you've not actually installed the program and are simply running it from the build directory. You can get around this by *cding* into your build directory and doing:

```
export PATH=$PATH: `pwd`
```



Using `rescan-scsi-bus.sh`, it is easy to check for new SCSI devices on a hot-plug SCSI bus, so they don't all need to be done by hand.

◀ You should be then able to execute `jigdo-lite` from the command line and have it all work properly. Alternatively, you could run `make install` and install the binaries you have compiled, at which point it should function correctly. If you still experience problems, you may wish to check for an updated version of `jigdo` or look for a Mandrake binary RPM distribution of the program.

mount

Q I have just completed an apparently successful installation of Mandrake 9.0 from the Christmas 2002 issue DVD, but cannot successfully mount the CDROM, even when logged in as 'root'.

When I attempt this operation I get the following message:

```
Unable to enter file:/mnt/cdrom. You do not have access rights to this location.
```

Your help would be very much appreciated.

John Donaldson.

A This does not appear to be a regular error, so seems to be generated by KDE or whichever GUI you are using. You can mount `/mnt/cdrom` with the contents of your CDROM drive with:

```
mount /mnt/cdrom
```

If this produces an error, then you will need to find out which `/dev/hd*` device is for your CDROM drive. You can do this with `dmesg`, which you will want to pipe through `more` so you can

read it. You can then do:

```
# mount -t iso9660 /dev/hdc /mnt/cdrom
```

Assuming your CDROM drive is `/dev/hdc`. You may want to look at your permissions and configuration with your GUI to see if it is not setup correctly to mount your CDROM drive.

scsi bus

Q I was wondering if you guys knew of a better way to scan the scsi bus for newly attached devices on a Linux machine.

I have several Linux PCs running Red Hat 7.2 and 7.3. The are connect to a SAN. When I attach new devices to it via direct scsi or over the SAN via IP or Fibre Channel, I have to rescan the bus. What I do currently is issue the

following command.

```
echo "scsi add-single-device $a $c $s $l" >/proc/scsi/scsi
```

where `$a $c $s $l` is the ACSL (the scsi address of the new device(s)); then I `cat` the `/proc/scsi/scsi` to make sure its added and then continue with `fdisk`, etc.

Is there a better way to do this? What if I don't know the ACSL? I could make a script that increments the four dollar sign values and loop through it `cating` the results to the `/proc/scsi/scsi`. But I want another way.

In Solaris you can issue the commands `drvconfig`, `devlinks` and `disks` or the command `devfsadm` to rescan.

In HP you can issue the command `ioscan -fnC` to rescan.

In AIX you can issue the command `cfgmgr` to rescan and then `lsdev -c disk` to see the devices.

In Tru64 you can issue the command `hwmgrr -scan scsi` to rescan and then `hwmgrr -view devices` or `hwmgrr show scsi` to see the devices.

Why doesn't Linux have something like this? Or do they and I just don't know about it...

Any help would be appreciated.

Peter Zeppieri

A There is a great script called `rescan-scsi-bus.sh`, which can be found at www.garloff.de/kurt/linux. This actually `cats` the data into `/proc/scsi/scsi`, but searches the bus to find which devices are connected to the various hosts, channels and so forth. Anyone who's used `cdrecord` will probably be familiar with this utility. As it's a simple shell script, it's not too



Linux supports a wide variety of filesystems allowing partitions belonging to Windows installations to be shared with Linux.

difficult to modify it if it does not correctly pick up some devices, although as it works with the `/proc` interface, it is completely independent of the device actually being used.

DVD boot

Q Sorry if you have been inundated with mails about this but The DVD that came with LXF35 will not boot on either of my systems – both ABIT motherboards. I have booted SuSE 8.1 on this system – the one with a DVD and built 8.1 fine. I successfully built Mandrake 9 on the non-DVD (and slowest system) over my net with the help of the “boxout” – thanks.

Question – is the DVD faulty?

Ian

A It could be that the DVD is faulty. You can mail the DVD back to Future for a replacement and see if that fixes the problems you are having. As you've been able to boot DVDs and CDs in the past, and you're seeing the same on two systems, it would suggest by a process of elimination that the DVD is at fault. lf@futurenet.co.uk

Submission advice

We are happy to answer all sorts of Linux related questions. If we don't know the answer, we'll find out for you! But in order to give you the best service, it helps a lot if you read the following submission advice.

- Please be sure to include any relevant details of your system. "I can't get X to work" doesn't really mean anything to us if we don't know things like what version of X you are trying to run, what hardware you are running on.
- Be specific about your problem. Things like "it doesn't work" or "I get an error" aren't all that helpful. In what way does something not work? What were you expecting to happen? What does the error message actually say?
- Please remember that the people who write this magazine are NOT the authors or developers of Linux, any particular package or distro. Sometimes the people responsible for software have more information available on websites etc. Try reading the documentation!

We will try and answer all questions. If we don't answer yours specifically, you'll probably find we've answered one just like it. We can't really give personal replies to all your questions.

WRITE TO US AT:
Linux Format, Future Publishing, 30
Monmouth Street, Bath BA1 2BW or
email: lxformat@futurenet.co.uk

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).



January 2003

Product code:
LXFB0036(cd)
LXFD0036(dvd)

DVD HIGHLIGHTS:
 Evolution PIM/calendar/
 email client, Freevo
 personal video recorder,
 OpenZaurus, PilotLink

MAGAZINE FEATURING:
 Optimisation tips and tricks,
 Netbox Cubit review, build
 your own desktop, Xandros,
 Cubit, CORBA, Realsoft 3D,
 Instant Messenger roundup

CDs HIGHLIGHTS:
 Racer, Freedud, Linux
 From Scratch, GNOME
 Meeting, Opera 6.1, Film
 GIMP, phpOpenTracker,
 Tuxpaint, Snort



Christmas 2002

Product code:
LXFB0035(cd)
LXFD0035(dvd)

DVD HIGHLIGHTS:
 Mandrake 9.0 – 3 CD
 download edition, jEdit,
 BastilleLinux, Phoenix,
 VideoLAN, KDE 3.0.4

MAGAZINE FEATURING:
 Huge C++ IDE RoundUp,
 LXF Awards nominations,
 gadgets for Xmas, User
 Mode Linux, compiling code,
 process accounting

CDs HIGHLIGHTS:
 3CDs, including 2CD
 Mandrake 9.0, SpamX,
 Stellarium, Krename,
 ImagePress, Nogger,
 DigitalMusicCentre



December 2002

Product code:
LXFB0034(cd)
LXFD0034(dvd)

DVD HIGHLIGHTS:
 KDE 3.1 beta2, Scribus,
 Smoothwall2, Movix,
 GNUPrivacyGuard,
 GCompris, Drip

MAGAZINE FEATURING:
 Distro reviews of SuSE 8.1,
 Red Hat 8.0, Lycoris and a
 mini-distro RoundUp,
 scanning, PHP acceleration,
 Linux Expo UK 2002

CDs HIGHLIGHTS:
 Unreal Tournament
 2K3 (Demo),
 DemoLinux, Cardfile,
 RUTE, OoOictionaries,
 Mozilla binary packages



November 2002

Product code:
LXFB0033(cd)
LXFD0033(dvd)

DVD HIGHLIGHTS:
 Debian 3.0, mFighter,
 OpenOffice.org (bugfix
 release), DVDRip,
 Cyrus IMAP Server

MAGAZINE FEATURING:
 Sun's move into the Linux
 server market – with LX50
 review, The Liberty Alliance,
 Systems programming, using
 OpenOffice.org, Homebase

CDs HIGHLIGHTS:
 Kylix 3, GNUCash,
 BXPro, KDevelop,
 Opera, Vega Strike,
 Parted, AnjutaIDE,
 GTransferManager



October 2002

Product code:
LXFB0032(cd)
LXFD0032(dvd)

DVD HIGHLIGHTS:
 Knoppix, Drip, Squeak,
 extra FlightGear maps,
 Ogg Vorbis 1.0, Knoppix

MAGAZINE FEATURING:
 Building better databases,
 "Trusted Computing" –
 beware Palladium, USB 2.0,
 firewall roundup, Amiga
 emulation, Gentoo review

CDs HIGHLIGHTS:
 FlightGear (runs from
 disc), Aglaophone, UAE,
 Clam Antivirus, Perl
 5.8, Quanta Plus,
 Netclipboard, Mah-
 Jong, HTML-Mason,
 WebSuck, Epsutil



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

Other back issues also available, please call for further information, giving the issue number and/or date required

To order, call the back issues hotline on **0870 4448645**

or **+44 1458 271178** for overseas orders

Remember to quote the issue code shown above and
 have your credit card/switch card details ready.

Prices (inc post and packing): CD – UK£5.99,
 US & Eur £6.99, rest of world £7.99
 DVD – UK £6.99, US&Eur £7.99, rest of world £8.99.

LINUX FORMAT

www.linuxformat.co.uk

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!**

* or call the overseas number shown at the bottom of the coupon, below

RISK-FREE GUARANTEE!

You can cancel your subscription at any time and receive a full refund on all unmailed issues, no questions asked

Title _____ Initials _____ Surname _____
 Address _____
 _____ Postcode _____ Country _____
 Email address _____
 Daytime telephone number _____

Instruction to your Bank or Building Society to pay Direct Debits



Media with passion
 Future Publishing Ltd, Cary Court,
 Somerton, Somerset, TA11 6BR



Originator's Identification Number
7 6 8 1 9 5

Please send this form to address shown

1. Name and full postal address of your Bank or Building Society branch

To: The Manager _____ Bank/Building Society
 Address _____
 _____ Postcode _____

2. Name(s) of account holder(s) _____

3. Branch sort code _____
(from the top right-hand corner of your cheque)

4. Bank or Building Society account number _____

5. Instruction to your Bank or Building Society

Please pay Future Publishing Direct Debits from the account detailed on this instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Future Publishing and if so, details will be passed electronically to my bank or building society.

Signature(s) _____ Date _____

Ref No (Office use only) _____

Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

You may photocopy this blank form and then complete it. You cannot fax or e-mail Direct Debit forms.

QUOTE ORDER No. LXFP37

LINUX FORMAT

Please choose your method of payment **1** or **2**

- 1.** Direct Debit (UK only – double CD Edition) – I wish to pay £27.45 every 6 months
 or Direct Debit (UK only – DVD Edition) – I wish to pay £29.95 every 6 months

Please complete the form on the left

I understand that I will receive 13 issues during the next 12 months.

2. Cheque/credit card 13 issues

UK	<input type="checkbox"/> £54.90	<input type="checkbox"/> £59.90
Europe (including Eire)	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
North America	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
Rest of the World	<input type="checkbox"/> £91.90	<input type="checkbox"/> £96.90

Cheque (payable to Linux Format. Sterling cheques drawn on a UK account.)

Visa MasterCard American Express

Switch Issue No/Valid date _____

Card No _____ Expires _____

Signature _____ Date _____

UK READERS RETURN THIS COUPON TO:

Linux Format, Subscriptions, Future Publishing Ltd, Freepost BS4900, Somerton, Somerset, TA11 6BR.

OVERSEAS READERS RETURN (POSTAGE PAYABLE) TO:

Linux Format, Subscriptions, Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6TB, UK

WEB: www.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 +44 870 444 8645

OFFER ENDS: 27.02.03

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 9.0
● DemoLinux ● LinuxFromScratch ● Sorcerer Linux ● Knoppix
● Red Hat 7.3 ● Slackware 8.1 ● Debian 3.0 ● Gentoo ● Yoper
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-1.01.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.01.i386.deb – The same, but a debian package.

Someap-1.01.tar.gz – This is usually source code.

Someap-1.01.tgz – Same as the above, tgz is abbreviated form of tar.gz

Someap-1.01.tar.bz2 – Same, but uses bzip2 compression instead of zip

Someap-1.01.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.01.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.01.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7.x PPC Linux.

Someap-devel-1.01.i386.rpm – A development version.

Installing from tarballs

A tar ball is a two stage archive. First the files are archived into a single file with tar and then compressed with Gzip or Bzip2. To unpack, cd to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
tar xvzf --bzip2 /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in .tar.gz or .tgz, and the second for Bzipped files, ending in .tar.bz2 or .tbz2. Naturally, you change the paths to suit the location and name of the archive. and replace /mnt/cdrom with whatever is applicable to your system (e.g. /cdrom). This normally unpacks the archive into a directory of the same name, enter that directory with:

```
cd progname-2.1.0
```

To compile and install the software, type the following three commands:

```
./configure
make
su -c "make install"
```

The last line will prompt you for the root password, as this stage must be run as root. If you are already logged in as root, just type **make install**. This will give you a default installation. If you want to change any aspect of the install, type

./configure --help to see the options available. For example, you are usually able to change the default location with the PREFIX argument. When you have finished installing, you may remove the source files with:

```
cd ..
rm -fr progname-2.1.0
```

You should also log out as root, before you do anything you may later regret.

Defective CDs

In the unlikely event of your CD/DVD being physically damaged we'll send you a new, working version within 28 days. Send your defective disc – complete with your name, address, and a description of the fault – to:

Linux Format, Future Publishing Disc Department, 3B Athena Avenue, Elgin Industrial Estate, Swindon, SN2 8HF

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



Squadron Leader Neil Bothwick almost didn't have time to compile the Linux Format DVD this month, and it's all Flightgear's fault. Tally ho, chaps...

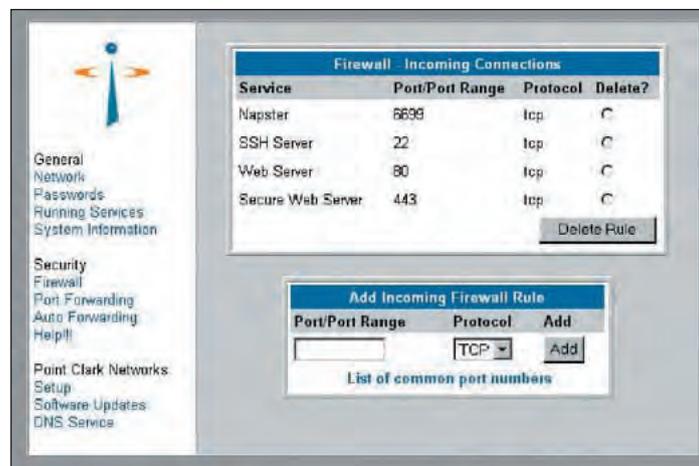
Lucky DVD users get the full contents of the coverdiscs, plus loads of bonus material every month. – that's nearly 2GB of material. So if you have the DVD issue, here's the second part of what's on your coverdisc. If you don't have a DVD drive, you should think very seriously about upgrading your hardware this month, as there's a double-sided DVD packed with a whopping 8GB of essential Linux programs planned for *Linux Format* 38 in March 2003!

DISTROS IPCOP

Using a redundant PC as an Internet gateway seems to be a popular project

for Linux developers. This month we have three such projects, with varying features. *SmoothWall* has been around for a while and deservedly gets a lot of good press. We had a beta of *SmoothWall 2.0* on the DVD a few months ago. 2.0 is still in beta but 1.0 has only just been released, replacing the previous 0.99 version. Naturally, a milestone release of such a popular project is on this month's cover DVD.

IPCop performs a similar function to *SmoothWall*, in fact it started life as a fork of the *SmoothWall* project. *IPCop* has now reached version 1.2 and adds support for the Speedtouch USB modems supplied by many UK ADSL ISPs. This project is also a pure router system. It takes an old PC and converts it into a router and firewall, connecting your computer or network to the Internet while protecting it from the hostile inhabitants of that domain. Using a dedicated machine for this task has at least two key advantages. The processing power required for this task is minimal, an old 486 or P100 can be dusted off and brought back into service. It also means that no sensitive services are running on the only computer that is directly connected to the Internet.



Use your old computer as your guardian to the Web – ClarkConnect combines an internet gateway, router, firewall and servers, all in an old PC.

If you want a machine that combines the features of a router/firewall and various server duties, you will need a somewhat more powerful computer than a 486, and a different package. *ClarkConnect* is just such a package. In addition to the routing and firewalling duties described about, *ClarkConnect* can also act as a web, email, file, print or proxy server, or any combination of these. Obviously, the more of these you run and the more traffic they handle, the

more powerful a box you'll need to run *ClarkConnect*. One other significant difference between *ClarkConnect* and the other two packages is that it only supports broadband connections to the Internet, while the others support modem and ISDN too.

DVD users get to choose between the three systems, try them all and see which suits your needs best. With the limited space on the CDs, we were only able to include *IPCop* there.

GAMES FLIGHTGEAR

Flightgear was included on our CD and DVD last year, runnable directly from the cover discs. DVD users also got more than a gigabyte of extra scenery files. Now the program has been updated. Along with various enhancements and improvements, this version has a new model of a very old aeroplane. This is the centenary year of powered flight, so the latest *Flightgear* has a model of the 1903 Wright Flyer. What would the Wright brothers have thought of the idea of us 'flying' their first aircraft around the World without leaving our homes?

This time we have not made *Flightgear* runnable from the cover



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.



An old plane for a new program. 100 years after the first powered flight, a model of the Wright Flyer is added to the new version of Flightgear.

disc, because that would have meant including the same scenery files again. So this is the basic installation, as tarball and Debian packages. There are two packages to install, using the standard `./configure && make && make install` process – one package contains the simulator and the other has a basic collection of textures, sounds, sample scenery, and other data files needed by it. If you have the CD or DVD from issue 32 (back issues available on page 97) you can copy over the scenery files once *Flightgear* is installed on your hard disk.

SYSTEM NVIDIAFREE86DRIVERS

NVIDIA have a history of providing support for Linux users of their hardware. Some may object to the fact that their drivers are not Open Source, but this does mean that they are able to use non-GPL technologies in their drivers, allowing their cards to run programs like *Unreal Tournament 2003*. Despite being closed source, NVIDIA's drivers still fit the "beer" definition of free, so here they are for you. There is a large number of packages to choose from, depending on your distribution, kernel version and CPU. To make picking the correct ones easier, NVIDIA have provided the NVchooser script. Run it with `sh /mnt/cdrom/System/nVidiaXFree86Drivers/NVchooser.sh` and it will advise you on the best packages to install. Whichever package format you end up using, you will need to install two packages, `NVIDIA_kernel` and `NVIDIA_GLX`. The included README and PDF files detail the installation process, which closely follows the standard methods for installing from tarball or RPM. However, there are a couple of catches, so read the documentation before installing. For example, when updating older drivers from RPM, the previous `NVIDIA_GLX` package must be uninstalled before the new version is installed, the normal `rpm -Uhv...` process doesn't work. You will also need to follow the instructions for editing your `XF86config` file if this is the first time that you have installed these drivers on your system.

It is a good idea to keep copies of the packages you install on your hard disk. If you upgrade your kernel the drivers will cease to work until you reinstall them, so having them where you can easily find them makes sense.



Create an interactive web site with minimum effort and choose from over 300 themes using *PostNuke*.

Always backup your existing software before attempting a new install.

SERVER POSTNUKE

More and more web sites are using some sort of system that generates web pages on the fly, instead of the old static, everyone gets the same page, HTML. There are many different ways of doing this, using a variety of technologies. One of the favourite combinations for Linux servers is LAMP (Linux, Apache, MySQL and PHP). Now you could follow our PHP tutorials and create your own system from scratch, or you could use an existing application that has had plenty of field testing. *PostNuke* is one such application, in fact it is the one used for the *Linux Format* website. *PostNuke* provides the tools you need to build a web site that encourages user interaction. It provides for posting news articles, discussion forums, surveys, themes, download areas, in fact most of the things you would want to do with such a site. Pop along to www.linuxformat.co.uk to get an idea of some of the capabilities, or install the prog and try it for yourself.

In addition to *PostNuke* itself, we have included a large number of themes. You can use these to change the appearance of the site, a *PostNuke* site doesn't have to look like other *PostNuke* sites. There are over three hundred themes on the DVD, but you

can create your own if you feel that doesn't give you enough choice or you relish the challenge.

SOUND LILYPOND

There are some things you expect to find on *Linux Format* cover discs, the latest version of *Mozilla*, kernel updates, popular games. That's not a bad thing of course, having expectations, and finding them met, is good. However, predictable is not good, so here is something rather different, something that shows the diversity of software written for Linux. *LilyPond* is a typesetter for music. In the same way that a program like *TeX* reads a file and produces professionally laid out

pages of text, images and formulae, *LilyPond* lays out sheet music. It takes its input from a plain ASCII description file and outputs beautifully typeset sheet music. The font used and the layout algorithms were inspired by engraved music, so you can expect that same clear and elegant look.

LilyPond was written for use with classical music, but will work with almost any musical genre. The input file format is specific to *LilyPond*, but several programs are capable of outputting in this format, and *LilyPond* comes with filters to convert popular data file formats. There are source and RPM packages of the current stable release, along with a tarball containing the latest development version. [LXF](http://www.lilypond.org)

Create professionally typeset music with *LilyPond*, using the output from most music composition programs.

CoverdiscDVD

DVD CONTENTS AT A GLANCE

Desktop

ConfiguratorForGnome	Edit advanced GNOME settings
FVWM	Classic highly-configurable virtual window manager
GWintree	Produces complete family trees from GEDCOM data
Kaspaliste	Literature database for Linux and KDE
KewlSessionManager	Session manager for X
KricketScoreboard	Keeps scores of cricket matches
PDFcrypt	Allows you to set permissions on a PDF file
PersonalFileManager	A personal file manager
Pipebench	Shows the status and a benchmark of piped commands
SpaceChart	Map of the stars in 3D
SquirrelSQLClient	Graphical Java database client
SVGIconThemes	A collection of icons in the SVG format for KDE and GNOME
Ted	A easy-to-use rich text processor for Linux
TightVNC	Enhanced VNC distribution
X-CD-Roast	Package dedicated to powerful and easy CD creation
XMLTV	Tools to fetch and process TV listings in an XML format
Xrmap	An X client for displaying the CIA World vector map
YASE	Text indexing and retrieval system

Development

Cocoon	XML processing framework
JavalDEforEmacs	Complete Java development environment for Emacs
KDevelop	Integrated development environment for Unix/X11
libsndfile	Library for reading and writing sound files
Nive3d	Portable 3D engine
OpenCOBOL	A COBOL compiler
PyPgSQL	Python DB-API 2.0 interface to PostgreSQL
PythonMilter	Python binding for the sendmail milter API
ScreenSaverDisabler	A screen saver disabling library
Source-Navigator	An IDE with source code comprehension features
SybaseModuleForPython	A Python DB-API-compliant interface to Sybase
XMLSecurityLibrary	XML Signature and XML Encryption standards

Distros

ClarkConnect	Transforms an old PC into an Internet gateway
CRUX	Lightweight, i686-optimised Linux distribution
MSC	A lean distribution delivering extreme performance
Sisela	Boot floppy for routers with PCMCIA and wireless support
SmoothWall	Fully functional Linux Firewalling application server
Trinux	Minimal Linux distribution that boots from floppy or CD-ROM
Yoper	Stable and optimised OS supporting various package systems

Games

AvoidTheRoid3D	A 3D asteroids-like multiplayer game
DroidBattles	A robot-programming game for Unix/X11
FlightGear	Flight simulator
nmaFPS	A simple, portable 3D first person shooter
OpenSteelball	A futuristic soccer game
Trackballs	An arcade game similar to Marble Madness

Graphics

Camsource	Grabs and processes frames from a v4l device
Fandango	An extensible CAD engine
gPhoto	GNU Digital Camera download software
KMediagrab	A DVD and TV grabber and a video translator
PanoramaViewer	An applet for displaying panoramic images
Ploticus	A graphics and plot presentation package
Video4LinuxGrab	Real-time capturing software for a Video4Linux device

Internet

AndromedelIRCD	An IRCd based on bahamut with host encryption
deStats	A cool IRC logfile analysis tool
MailStripper	An MTA-independent anti-spam, anti-virus mail scanner
Paranoy	Paranoid mail client, supports SSMTP, POP3+SSL and more
Rabbit	Caching webproxy to speed up surfing over slow links
TTY-Grin	A news and email client for Unix
WWWOFFLE	Proxy server with special features for use with dial-up

Mobile

HolidayCalc	HolidayCalc calculate movable holidays (eg. easter)
KAddressBook-sync	Convert KAddressBook files to Zaurus address book format
Leaflet	Leaflet is a auto-saving note-taking program
Mplayer	Mplayer 0.90 rc1 compiled for Zaurus
PyQt	Set of Python bindings for the Qt toolkit



Optimise your NVidia graphics card, whether for 2D use or for games like *Unreal Tournament 2003*, with the latest drivers direct from NVidia.

Python	Zaurus Python distribution as a set of .ipk packages
PythonIDE	Qt-based Python IDE for handhelds running Qtopia
TabManager	Move Qtopia applications from one tab to another
VolumeControlApplet	Controls Zaurus Mic and Speaker volumes independently
ZMoney	Small and fast program to manage accounts and budgets
Zsync	Copy PIM information from the Zaurus to Evolution 1.2

Office

AccWhizz	Multilingual business accounting application
AquaDataStudio	Database query tool
VetTux	Animal clinic management system

Server

Dnsmasq	Simple lightweight caching DNS forwarder
HTTP-WebTest	Test for remote URLs or local Web files
NextvendApplicationServer	Internet business application server
PgSQLHierarchicalQueries	Hierarchical queries support for PostgreSQL
phpImagick	ImageMagick extension for PHP4
PostNuke	An easy-to-use/install WebLog/portal system
Xerces2	The next generation Apache Xerces-J XML parser.
ZopeWeatherApplet	Zope Python applet to display weather information

Sound

AnMP3LEnder	Server for streaming MP3s over a network
GLAME	Audio processing tool and sound editor
GnomeALSAmixer	ALSA audio mixer for GNOME
KmusicDB	Music collection manager for KDE using PostgreSQL
LilyPond	The GNU Project music typesetter
OpenBeatBox	Virtual Drum Machine
Streamixer	Audio stream mixing and resampling toolset
StreamTranscoder	Converts MP3/Vorbis streams into MP3/Vorbis/WMA streams
ZynAddSubFX	Software synthesizer

System

CatweaselDeviceDriver	Driver for the Catweasel Advanced Floppy Controller
CatweaselFloppyTools	Catweasel ISA card driver
CommonUNIXPrintingSystem	An Internet printing system for Unix
CramFS	Compressed filesystem
Deborphan	Debian orphaned library finder
GRUB	GRand Unified Bootloader
LeptonsCrack	Generic password cracker
Linuxconf	Sophisticated administrative tool
LinuxTestProject	Test suite for Linux
MatroxMillenniumIIXv	An Xfree86 Xv extension for the Matrox Millennium II
Moosic	Powerful music queue manager
NetfilterLogsAnalyzer	A netfilter log analyzer with a Web interface
NVidiaXFree86Drivers	NVidia XFree86 4.0 Drivers
OpenwallLinuxKernelPatch	Security hardening patch for the Linux kernel
Samba	Tools to access to a server's filesystem and printers via SMB
SimpleRescueCD	Mid-size bootable CD (7MB) suitable for business card CDs
VirusUpdateServer	A virus update server for McAfee products

User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally natter about all things Linux. You can find lots more information online at: www.lug.org.uk

1 Hampshire

URL www.hants.lug.org.uk
Contact Hugo Mills

2 Bristol & Bath

URL www.bristol.lug.org.uk

3 Scottish

URL www.scottish.lug.org.uk

4 Oxford

URL www.oxford.lug.org.uk
Contact Alasdair G Kergon

5 Kent

URL www.kent.lug.org.uk
Contact John Mills

6 Brighton

URL www.brighton.lug.org.uk
Contact Johnathan Swan

7 Worcestershire

URL www.worcs.lug.org.uk

REVISED
DETAILS

8 Northants

URL www.northants.lug.org.uk
Contact Kevin Taylor

9 Anglian

URL www.anglian.lug.org.uk
Contact Martyn Drake

10 Milton Keynes

URL www.mk.lug.org.uk
Contact Denny De La Haye

11 Doncaster

URL www.doncaster.lug.org.uk
Contact Andy Smith

12 Moray

URL www.moray.lug.org.uk
Contact Stewart Watson

13 West Wales

URL www.westwales.lug.org.uk
Contact Dan Field

14 Wolves

URL www.wolveslug.org.uk
Contact Jono Bacon

15 Peterborough

URL www.peterboro.lug.org.uk
Contact Steve Gallagher

16 Edinburgh

URL www.edinburgh.lug.org.uk
Contact Alistair Murray

17 Tyneside

URL www.tyneside.lug.org.uk
Contact Brian Ronald

18 Leicester

URL www.leicester.lug.org.uk
Contact Clive Jones

19 Greater London

URL <http://gl.lug.linux.co.uk/>
Contact John Southern

20 Surrey

URL www.surrey.lug.org.uk
Contact Jay Bennie

21 Cambridge

URL www.cam-lug.org.uk

22 Devon & Cornwall

URL www.dclug.org.uk
Contact Simon Waters

23 Falkirk

URL www.falkirk.lug.org.uk

24 Manchester

URL www.manlug.mcc.ac.uk
Contact John Heaton, Owen Le Blanc

25 Hertfordshire

URL www.herts.lug.org.uk
Contact Nicolas Pike

26 West Yorkshire

URL www.wylug.lug.org.uk
Contact Jim Jackson

27 Sheffield

URL www.sheflug.co.uk
Contact Richard Ibbotson

28 Staffordshire

URL www.staffslug.org.uk

29 North East

URL www.shofar.uklinux.net/NELUG

30 London

URL www.lonix.org.uk

31 Thames Valley

URL www.sclug.org.uk

32 Liverpool OpenSource

URL http://linux.liv.ac.uk/_liv_linux_ug/
Contact Simon Hood

33 Deal Amiga Club

Email superhighwayman@hotmail.com
Contact John Worthington

34 Chesterfield

Email spirelug@yahoo.co.uk
Contact Robin Needham

35 South Derbyshire

URL www.sderbylug.org.uk
Contact Dominic Knight

36 Belfast (BLUG)

URL www.belfastlinux.cx
Email russell@belfastlinux.org

37 Wiltshire

URL www.wiltshire.lug.org.uk
Contact Jason Rudgard

38 South London

URL www.sl.lug.org.uk
Email ben@ilovephilosophy.com

39 Cheshire

URL www.sc.lug.org.uk
Contact Anthony Prime – enquiry@sc.lug.org.uk

40 North Wales

URL www.northwales.lug.org.uk
Contact Jonathan Cole

41 Midlands

URL <http://midlandslug.port5.com/>
Contact Pete Thompson

42 Cumbria

URL www.cumbria.lug.org.uk
Contact Jamie Dainton

43 Dorset

URL www.dorset.lug.org.uk
Contact John and Mat

44 Shropshire

URL www.shropshire.lug.org.uk
Email shropshire@lug.org.uk

45 South West

URL www.southwestlug.uklinux.net
Email southwest@lug.org.uk

46 South Wales

URL www.swlug.org.uk

47 North London – see 87

URL www.kemputing.net/lug/anlug-aims.html

48 Malvern

URL www.malvern.lug.org.uk
Contact Greg Wright

49 Huddersfield

URL www.hud.lug.org.uk
Contact Adam Brookes

50 Nottingham

URL www.nottingham.lug.org.uk
Contact Godfrey Nix

51 St Albans & Luton

URL www.lust.lug.org.uk
Contact Michael Culverhouse – mike@easily.co.uk

52 Wrexham

Contact Paul Kersey-Smith
Email paul@pkls.fsnet.co.uk

53 Preston & Lancs

URL www.preston.lug.org.uk
Contact Phil Robinson

54 Derry

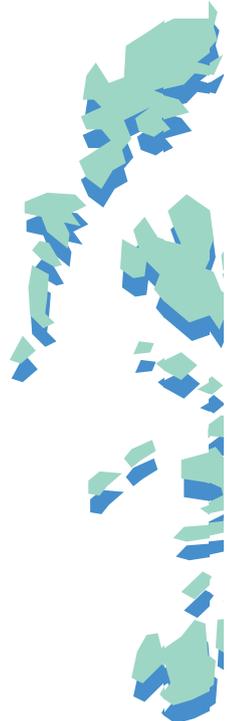
URL www.derry.lug.org.uk

55 Isle of Wight

URL www.iow.lug.org.uk
Contact David Groom – info@iow.lug.org.uk

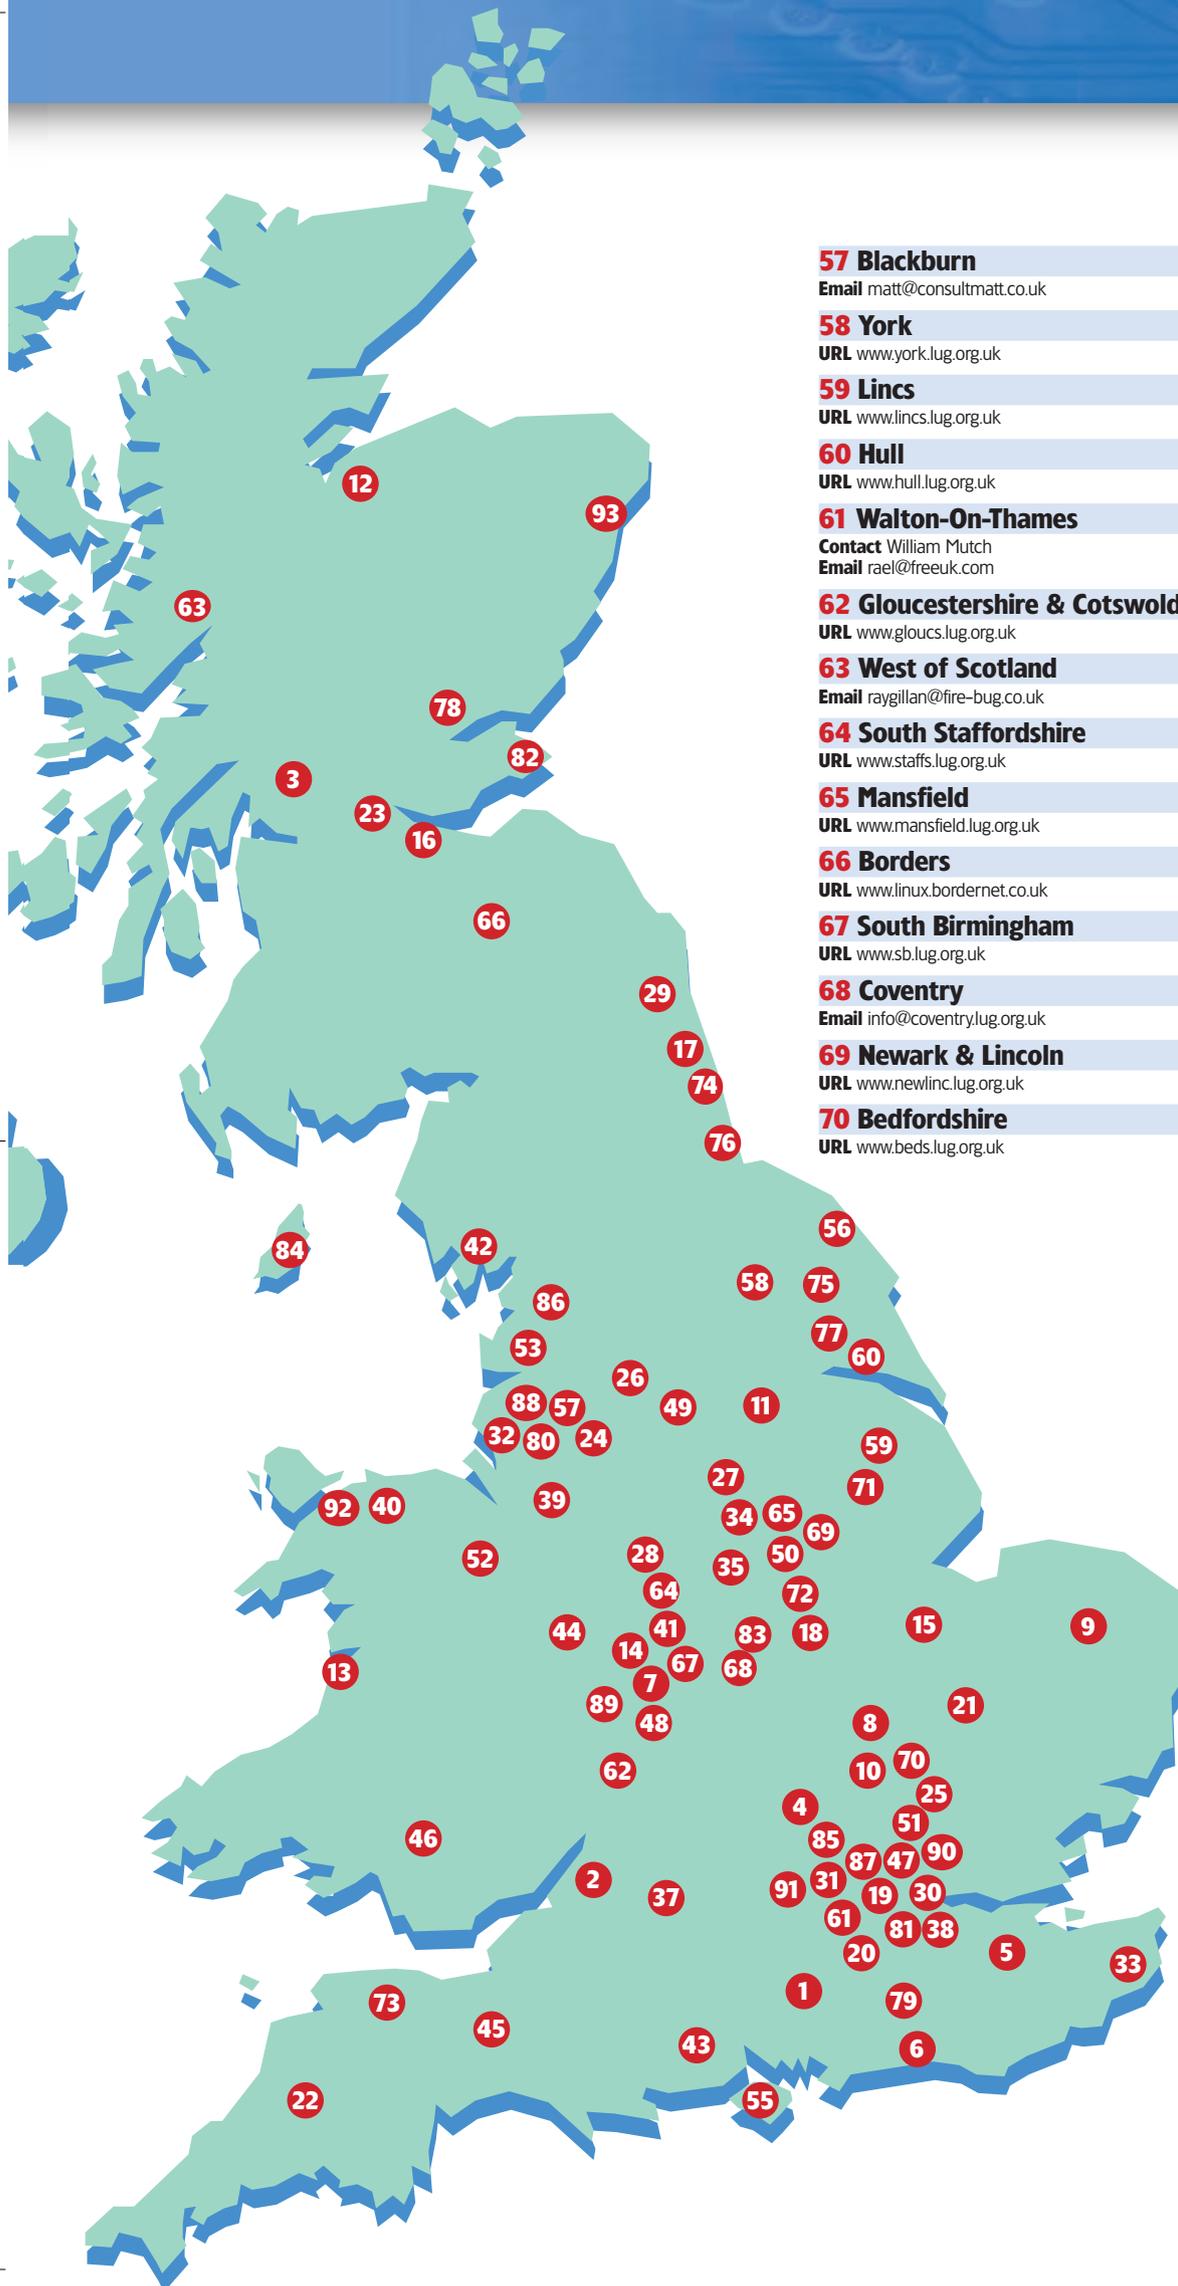
56 Scarborough

URL www.scarborough.lug.org.uk



54

36

**57 Blackburn**

Email matt@consultmatt.co.uk

58 York

URL www.york.lug.org.uk

59 Lincs

URL www.lincs.lug.org.uk

60 Hull

URL www.hull.lug.org.uk

61 Walton-On-Thames

Contact William Mutch

Email rael@freeuk.com

62 Gloucestershire & Cotswolds

URL www.gloucs.lug.org.uk

63 West of Scotland

Email raygillan@fire-bug.co.uk

64 South Staffordshire

URL www.staffs.lug.org.uk

65 Mansfield

URL www.mansfield.lug.org.uk

66 Borders

URL www.linux.bordnet.co.uk

67 South Birmingham

URL www.sb.lug.org.uk

68 Coventry

Email info@coventry.lug.org.uk

69 Newark & Lincoln

URL www.newinc.lug.org.uk

70 Bedfordshire

URL www.beds.lug.org.uk

71 Lincoln

URL www.lincoln.lug.org.uk

72 Loughborough

URL www.loughborough.lug.org.uk

73 Exeter University

Email N.J.Murison@exeter.ac.uk

74 Sunderland

Email thomas.croucher@sunderland.ac.uk

75 East Yorkshire

Email sharkonline@whatemail.com

76 Cleveland Open Source Group

Email openlug@digitalmedia.co.uk

77 Beverley

Email vladimir_lukyanov@hotmail.com

78 Dundee & Tayside

URL www.dundee.lug.org.uk

79 Sussex

URL www.phpworld.co.uk/~swlug

80 Wigan & St Helens

Email paul.f.johnson@ukonline.co.uk

81 Brixton

URL www.communitytechnology.org.uk/~linuxhome

82 St.Andrews, Fife

URL www.standrews.lug.org.uk

Email stuart@nx14.com

83 Nuneaton

URL www.nuneaton.lug.org.uk

84 Isle of Man

URL www.iom.lug.org.uk

Email helix@manx.net

85 Aylesbury

URL www.aylesbury.lug.org.uk

Email drbond@educational-computing.co.uk

86 Lancashire

URL www.lancasterlug.org.uk

87 North London – see 47

Email jason@voyagercomputers.co.uk

All the following are new details:**88 Ormskirk**

Email rob@northwestlinux.co.uk

89 Hereford

Email rbjh@good-news.fsnet.co.uk

90 East Herts

Email madtom1999@yahoo.com

91 Swindon

Email nick.trueman@ntword.com

92 Menai

URL www.menai.lug.org.uk

93 Aberdeen

Email mark@arricc.net

94 Shetland

Email c_s_s_butler@yahoo.com

**** Staffs ****

There will be a meeting on 18th February to kick start a North Staffs LUG. Details on <http://titanic.co.uk/staffslug>

LUG OF THE MONTH

South Wales

The South Wales LUG has been running since 2001 with its first few meetings being held in Cardiff. Due to the large area that the LUG covers it wasn't long before meetings were also being held in Swansea. We now regularly hold two meetings a month (one in each location) where we usually get between 10 and 20 members in attendance. The LUG is very informal and the meeting places vary from pubs to cafés and

occasionally restaurants.

Our plans for 2003 are to organise a few events such as getting the LUG involved in this year's Linux Install Day, talks and workshops. A few of our members are trying to set up projects that the LUG can get involved in, such as translating GNOME into Welsh.

We welcome anyone at our meetings and have occasionally had joint meetings with the West Wales LUG. There is a meeting

schedule on our website as well as mailing lists and an IRC channel where meeting announcements are made and anyone can ask for help or assistance. Here's another link – it's the Welsh Gnome Translation page that some of our members are currently working on – <http://tgb.org.uk/projects/gnome-cy/> irc: **#swlug** on irc.freenode.net www.swlug.org.uk



Worldwide Linux User Groups

Free Software users across the globe

Africa

EGYPT

URL www.linux-egypt.org
Contact Hesham Bahram

GAUTENG, SOUTH AFRICA

URL www.glug.org.za
Email glugmin@revolution.org.za

Australia

ADELAIDE

URL www.linuxsa.org.au
Email mtippet@anu.edu.au

ALICE SPRINGS

URL www.aslug.org.au

MELBOURNE, VICTORIA

URL www.luv.asn.au
Contact luv-committee@luv.asn.au

PERTH

URL <http://plug.linux.org.au/>

SYDNEY

URL www.slug.org.au

Europe

AUVERGNE

URL www.linux-arverne.org
Email Cyril.Hansen@wanadoo.fr

COSTA DEL SOL (English speaking)

URL www.fuengirola.lug.org.uk

DENMARK

AIssund www.alslug.dk
Esbjerg www.eslug.dk
Fyns www.flug.dk

Midt-og Vestjylland www.mvjlug.dk

Nordjylland www.njlug.dk

Skåne Sjælland www.sslug.dk

Trekantsområdet www.tlug.dk

Vest-fyn www.haarby-net.dk/vflug

Århus www.aalug.dk

EIRE

URL www.linux.ie
Email root@linux.ie
URL www.dilu.org
Email glossary@dilu.org

GOTHENBURG

<http://nain.oso.chalmers.se/LUGG/>

India

URL www.linux-india.org
Email newsmaster@linux-india.org

TRIVANDRUM

URL www.river-valley.com/tux
Email anil@river-valley.com

North America

ALASKA

URL www.aklug.org
Email deem@wdm.com

ALBERTA

URL <http://calgary.linux.ca/>

BATON ROUGE

URL www.br lug.net
Email dpuryear@usa.net

BAY AREA

URL www.balug.org
Email aftyde@balug.org

CLARKSVILLE, TN

URL www.cllug.org
Email tux@cllug.org

DENVER

URL <http://clue.denver.co.us/>

FLORIDA

URL www.flux.org

LOS ANGELES

URL www.lalugs.org
Email dank@alumni.caltech.edu

NORTH COLORADO

Email nclug@nclug.org

OTTAWA CANADA

URL www.oclug.on.ca

TAMPA

URL www.suncoastlug.org
Email president@suncoastlug.org

UHACC Normal, IL

URL www.uhacc.org
Email lug@uhacc.org

VIRGINIA TECH

URL www.vtlug.org
Email nega@vt.edu

South America

BUENOS AIRES

Email dcoletti@impost.com.ar

CHILE

URL www.linux-chile.org

MONTEVIDEO

URL www.linux.org.uy

PARAGUY/ ASUNCION

Email rolgiati@conexion.com.py

SAO PAULO

URL <http://gul.ime.usp.br/>
Email gul@ime.usp.br



In the latest in this series on advocating GNU/Linux, Jono Bacon looks at marshalling your arguments.

In the past months we have used

this little feature to discuss how to gather and present information that can be used to promote and advocate the use of Linux and free software to more people.

This month we start to examine the inevitable subject of Total Cost of Ownership (TCO).

TCO is a generic figure that shows how much a particular solution costs to run; this does not mean just the initial cost of the solution. As an example, a small business may download Linux from the Net and install it, but there is also the cost of technical support, upgrades, the time used to install and configure the machine and other issues. In this sense Linux is not free to own, but is free to obtain.

The issue about TCO is one that has been entrenched in the Linux world as cost is one of the major benefits of using Linux (in the sense it can be obtained freely). TCO is something that needs to be carefully discussed in the business community and all issues

need to be accounted for. With TCO being so important, it is essential it is given careful thought and consideration.

TCO can be a thorny matter in that it can be very objective in how it is calculated, and there is no clear rule of thumb. What is important however, is to take into account the right issues when adding up the TCO. For instance, technical support may be from a vendor who charges on an incident-by-incident basis, or it may be the *ad hoc* technical support from the Net. Which is cheaper? Initially, you may think the *ad hoc* style is, but bear in mind that it will almost invariably take you longer to fix an issue, and the time you spend may be depriving you of other income. This issue is largely down to the type of customer who is implementing the solution; a huge corporate will have very different TCO issues from a small startup or home user.

Next month we will take a look at some of the specific issues with TCO we need to consider, and how these issues compare to competing solutions.

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

NEXT MONTH

Issue 38 on sale Thursday 27 February

EDITORIAL

Editor Nick Veitch nick.veitch@futurenet.co.uk

Art Editor Julian Jefferson julian.jefferson@futurenet.co.uk

Production Editor Matt Nailon mattewnailon@futurenet.co.uk

Editorial Contributors

Jono Bacon, Neil Bothwick, Alain Buret, David Cartwright, Andy Channele, David Coulson, Chris Denton, Philip Downer, Richard Drummond, Hoyt Duff, Marco Fioretti, Professor Les Hatton, Paul Hudson, Jon Kent, Biagio Lucini, Neil Lucock, Michael Saunders, Richard Smedley, Charlie Stross, Martin Turner

COVER CD PRODUCTION

CD Editor Neil Bothwick

Senior CD Editor Kate Hadley

ART CONTRIBUTORS

Photography Louise Parker, James Wilson, Digital Vision, Corbis, Superstock, Getty Images

Illustration Paul Bateman, Chris Winn, Shane Collinge

ADVERTISING SALES

Portfolio advertising manager

Matt Stanley 01225 442244

matt.stanley@futurenet.co.uk

Portfolio deputy ad manager Diane Clydesdale

Business development manager Damian Hughes

Account managers Zoe Drew, Giles Crosthwaite Scott,

Ad sales executives Jo Doyle, Stephen Hall

MARKETING AND PROMOTIONS

Marketing Manager Fiona Tully

Subscriptions Executive Marie Spicer

PRODUCTION

Production Co-ordinator Diane Ross

Group Production Manager Clare Tovey

Advertising Production Co-ordinator Jo Crosby

MANAGEMENT

Publisher Dave Taylor

Group Advertising Manager Simon Moss

Group Publisher John Weir

Circulation Director Sue Hartley

Managing Director Colin Morrison

DISTRIBUTION AND CIRCULATION

Circulation Manager Jamie Malley

jamie.malley@futurenet.co.uk

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

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

Subscriptions and Mail Order

Phone +44 (0)1458 271178. See page 98

Email linuxformat.subs@futurenet.co.uk

Copyright No part of this publication may be reproduced without written permission from our publisher. We assume all letters sent – by email, fax or post – are for publication unless otherwise stated, and reserve the right to edit contributions. All contributions to *Linux Format* are submitted and accepted on the basis of non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing recognises all copyrights and trademarks. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight. All CD-ROM demos and reader submissions are supplied to us on the assumption they can be incorporated into a future covermounted CD-ROM, unless expressly stated to the contrary. We cannot be held responsible for mistakes or misprints. *Linux Format* recognises all copyrights in this issue. Where possible we have acknowledged the copyright holder. Please contact us if we have failed to credit copyright.

Disclaimer All tips in this magazine are used at your own risk. We accept no liability for any loss of data or damage to your computer, peripherals or software through the use of any tips or advice.

Printed in the UK by Midway Clark (Holt) and © Future Publishing Ltd 2003

LINUX is a trademark of Linus Torvalds, GNU/Linux is abbreviated to Linux throughout for brevity.

All other trademarks are the property of their respective owners

Future Publishing Ltd. is part of The Future Network plc.

The Future Network produces carefully targeted specialist magazines and websites for groups of people who share a passion. We aim to satisfy their passion by creating titles that offer superb value for money, trustworthy information, multiple ways to save time and money, and are a pleasure to read or visit. Today we publish more than 80 magazines and over 30 magazine websites and networks from offices in five countries. The company also licenses 32 of its titles resulting in over 60 local editions in a further 23 countries. The Future Network plc is a public company quoted on the London Stock Exchange (symbol: FNET).

Non-executive Chairman: Roger Parry

Chief Executive: Greg Ingham

Chief Operating Officer & Managing Director, UK: Colin Morrison

Group Finance Director: John Bowman

Tel +44 (0)1225 442244

BATH • LONDON • MILAN •

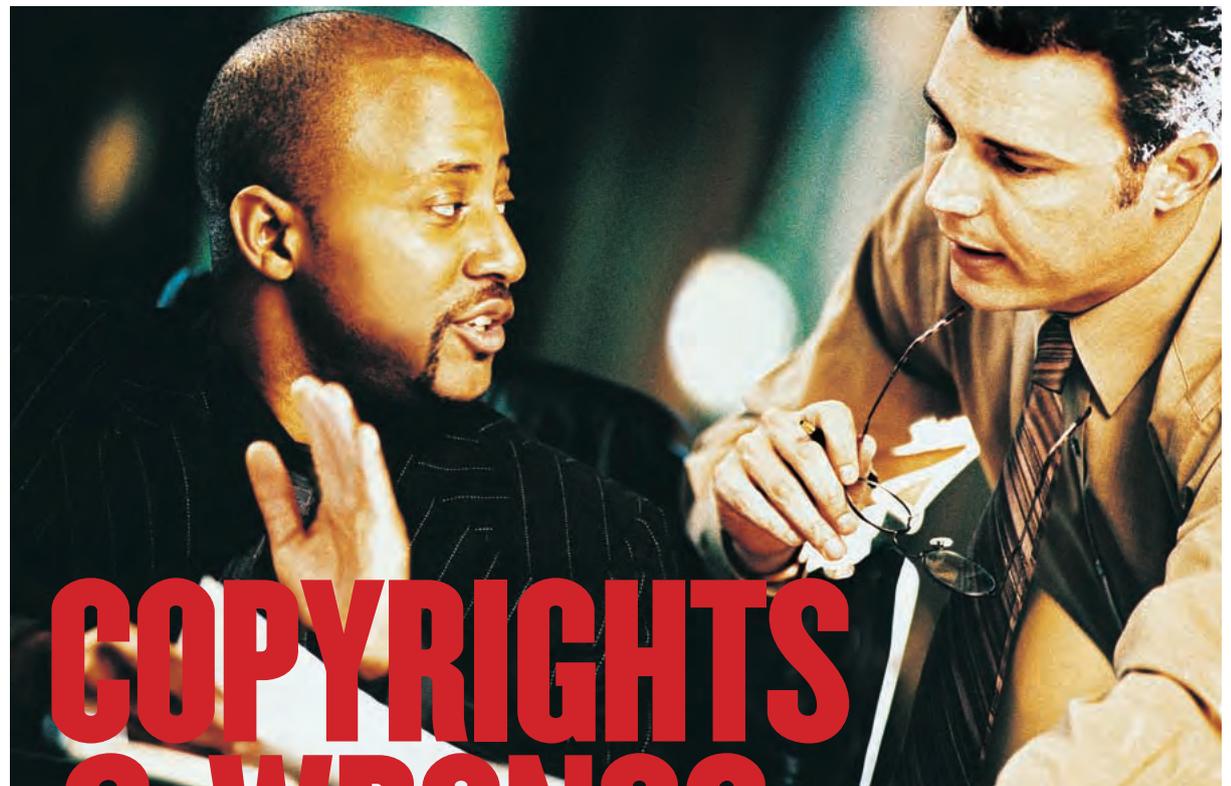
NEW YORK • PARIS •

SAN FRANCISCO • WROCLAW

www.thefuturenetwork.plc.uk



Media with passion



COPYRIGHTS & WRONGS

Who says it's illegal for you to watch DVDs on Linux? Why does your printer only accept expensive ink cartridges? Should using a screen-reader make blind people liable for prosecution? The wonderful, colourful world of Intellectual Property law & what it means to the average computer user, explained in depth...

FINK

The awesome power of Linux on your OS-X box – find out how the folks at Fink are getting on with porting your favourite Free Software to work on your Apple Mac.

Get Gimping

Part one of our new series *Getting To Grips With GIMP* starts next Issue. Learn all the special tricks and tips of professional image editing and unleash the power of this superb software with our easy-to-follow series.

LINUXPRO

Every Month, *Linux Pro* brings you the best of real-world Linux for IT professionals. Next month: firewalls explained, advanced *Apache*, Backup solutions and plenty more.

DVD BONANZA!



In another World first, *Linux Format* brings you a **DOUBLE-SIDED Linux DVD!** Two sides = twice the number of updates, apps, utilities and distributions at **NO EXTRA COST TO YOU!**

The exact contents of future issues are subject to change

DON'T MISS YOUR COPY OF LINUX FORMAT!

Get it delivered to your door every month – subscribe on page 98

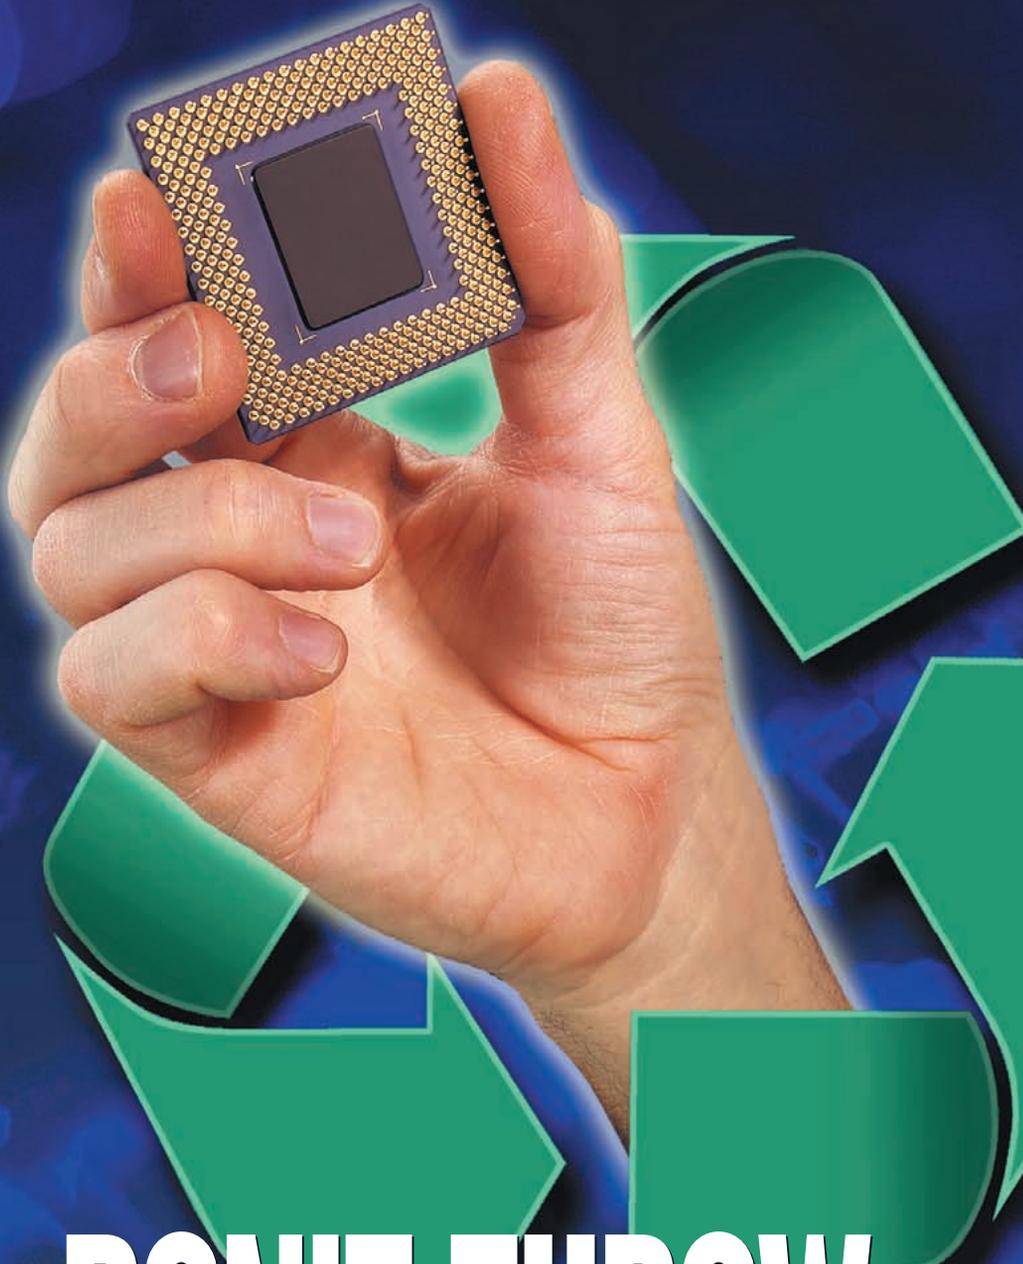
LINUX PRO

REAL-WORLD LINUX FOR I.T. PROFESSIONALS

FEBRUARY 2003

**NEW
BIGGER
SIZE!**

ONLY AVAILABLE WITH
LINUX
FORMAT



Taken into account

Finance solutions for Linux in our ACCPAC case study

Business Continuity

Could your organisation survive an IT disaster?

GNOME Accessibility

The GNOME project addresses your accessibility needs

Parallel Computing

Advanced parallelism explored with OpenMP

DON'T THROW IT AWAY!

As IT inventories groan with excess, outdated hardware, we look at how most equipment can be recycled and repurposed with the power of Linux **PAGE 9**

Welcome

Twenty-four pages of real-world Linux for IT professionals

Well, your support for Linux Pro has been so overwhelming that as from this issue, we have 'supersized' it to the same format as the main magazine. There may be the same number of pages at the moment, but obviously we'll be able to get twice as much on them, which means more case studies, interviews, articles and tutorials about the real world implementation of Linux!

Our cover story this issue is a look at the recycling of hardware. The recent 'downsizing' of many businesses has resulted in a glut of surplus IT hardware, but can a two year old desktop be put to any good use? It certainly can according to the organisation Access to Recycled technology, who breath new life into otherwise redundant hardware. The organisation repurpose machines for charitable purposes, but the principles can be used just as easily to recycle your corporate junk into more useable systems.

Accountancy software has always been a bit of a bugbear for those migrating to Linux, which is why we were glad to here about ACCPAC's support of Linux with their scalable finance software. We spoke to customers Globecast to see just how good a solution it is.

We also have the concluding parts of two of our series this issue. The accessibility series ends with a look at GNOME accessibility features and how they can address your organisations needs. meanwhile our clustering expert finishes his look at parallel computing with an advanced guide to OpenMP.

That's what we have for this issue, but remember that this magazine is here for your benefit. Write to us and let us know what you'd like to see in Linux Pro, whether that be specific articles or tutorials, or just general indications of what you'd like to see more or less of. We'll be back next month with a look at firealling, advanced Apache configuration and lots more!

Nick Veitch Editor
nick.veitch@futurenet.co.uk



“the recent downsizing of many businesses led to a hardware surplus. Can an old desktop be of any practical use? With Linux, the answer is Yes!”

Contacts

EDITORIAL

Editor

Nick Veitch nick.veitch@futurenet.co.uk

Art editor

Julian Jefferson julian.jefferson@futurenet.co.uk

Production editor

Richard Smedley
richardsmedley@futurenet.co.uk

Contributors

Richard Smedley, Biagio Lucini, Jono Bacon, Martin Turner

Photography

Digital Vision, Getty Images, Superstock

Publisher **Dave Taylor**

ADVERTISING

Portfolio ad manager

Matt Stanley matt.stanley@futurenet.co.uk

Portfolio deputy ad manager

Diane Clydesdale diane.clydesdale@futurenet.co.uk

Business development manager

Damian Hughes damian.hughes@futurenet.co.uk

Account managers

Zoe Drew, Giles Crosthwaite Scott

Ad sales executives

Jo Doyle, Stephen Hall

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

CONTENTS

GLOBECAST

Real world accountancy solutions from ACCPAC **p4**

RECYCLING

Putting old hardware to good use with Linux **p7**

BUSINESS CONTINUITY

Could your business survive a data disaster? **p12**

ACCESSIBILITY

Accessibility with GNOME **p14**

OPEN MP

Compting to extremes **p18**



7

18

ACCOUNTANCY



Taking Linux into Account

When talking about migrating business to Linux, there are often one or two software solutions that prove difficult to replicate. One of these has traditionally been enterprise level accountancy software, but this is no longer the case.

ACCPAC, a Computer Associates company, now provide their ACCPAC Advantage series software on Linux. In fact, far from merely supporting Linux, they barely stop short of rabid evangelism on the ACCPAC website!

ACCPAC is a world-class business management solution that offers the kind of features required for modern, dynamic businesses. Flexible data analysis tools combined with real-time queries and reports provide the data required

NICK VEITCH investigates a success story for Linux in the almost uncharted waters of accountancy.

The Linux ACCPAC client uses a friendly GUI driven interface.

to make informed business decisions. With multicurrency and multicompany capabilities it's also suitable for global companies or those based in multiple offices.

But one of the key features of ACCPAC that the developers are keen to underline is the flexibility when it comes to the choice of the systems you run it on. ACCPAC Advantage software will run on a number of database solutions and on a range of server solutions including Microsoft, but also Linux.

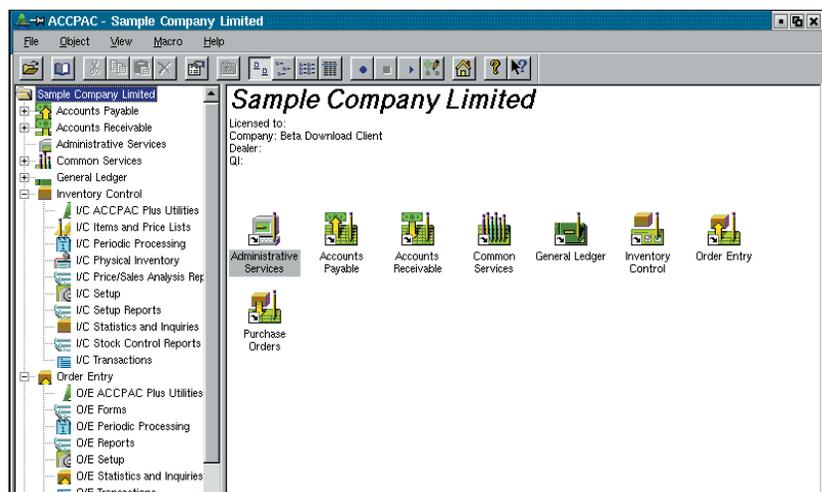
The Linux option

The ACCPAC website states that many people choose Linux because it has a reputation for availability, reliability and stability, and the company wants to support Linux as an option for its customers. This extends not only to the server software, but the clients too. The clients and server software interoperate completely, no matter which platform they are running on, so you can serve data from a Linux server to Windows and Linux desktops for example.

Of course, this is all very well in theory, but how does it work in practice? To find out we got in contact with Marco Tinirello, IT manager of communications company Globecast, who were one of the first companies to implement ACCPAC's Linux solution.

LXF We understand that you were running Linux before deploying ACCPAC?

MARCO TINIRELLO Globecast is a company owned by France Telecom, and we do satellite distribution for their



customers. This means that mainly we do video broadcasting for the likes of Viacom and MTV – big permanent distributions to Sky channels – that sort of thing. We also do newsfeeds and so on for various agencies and journalists, so we do a lot of *ad hoc* work too. That's the main part of our business.

Not so long ago we started doing data delivery via satellite and, because of the expertise we had in house, we started creating our own solutions. Because of the expertise we had in house, the techies here were comfortable with it and knew how it worked, this is mostly on Linux. They knew they could build solutions that would work. That worked for us too from a management point of view because it was obviously cost-effective.

That's how we got started in Linux, probably four or five years ago.

LXF So you were no strangers to Linux. But why did you choose it to run your accounts system on?

MT We had a history with ACCPAC. We used to run Novell servers with Btrieve running ACCPAC. The Novell servers were quite happy, but I don't think the Btrieve solution was one that was particularly suited to the way we were using it. We had countless problems – it was quite an old version of ACCPAC to be fair to them, but we had weekly, perhaps daily problems with people being crashed out from the system and we had to restore data from backups – we found it extremely unreliable.

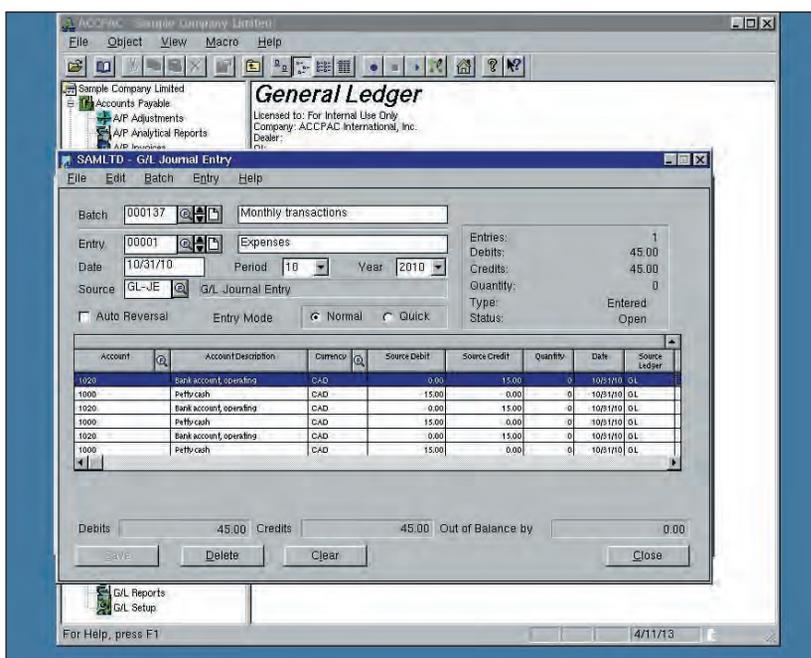
On the other hand we found that Linux was particularly reliable, in that we were happy with the performance on our customer solutions, so we mentioned it to ACCPAC, almost jokingly, and said "Could we not just run this on Linux, it would be so much easier". They said "well, we might be able to do something for you there". And that's how it all came about really.

So I wasn't a strategic plan as such, it more fell into place, but it's worked very well I have to say. I don't think in the year or so since we installed it we have had a single problem, which isn't bad. The main reason to go away from the old version was the reliability – I was confident that at least the host operating system was going to be fine. We had the in-house ability to support the OS.

LXF What about the end users of the system?

MT The end users are still Windows users. They're completely transparent to the back-end OS. All of their desktops are a mixture of versions of Windows. All of them, providing they can have an ACCPAC client (and all the versions of Windows we run are supported), can access the accounts system, and in truth they are oblivious to the fact that it runs on Linux, apart from the fact that it now runs without any problems. In terms of support, it's probably our biggest success story at Globecast because it was causing us quite a lot of grief, in that it was failing regularly. The finance department have quite tight deadlines reporting to Paris month in month out, and I was often the cause of them missing that deadline, so there was plenty of pressure.

LXF I suppose you have some more flexibility that you can run the clients on any system.



Ledger functions benefit from fast and easy data entry.

MT That's more an ACCPAC feature. As it stands we have several backend servers which are Linux based. We replaced our Novell servers with Linux also. We do have some Windows technology at the backend for some very specialised software we use for some of our satellite business, which I would replace if I had the chance. Our Novell servers were replaced with Linux simply because they came to the point where they were an inconvenience to us. Using IPX was making the routing between sites a bit of a pain and more expensive than it needed to be. We chose to scrap the Novell servers and run Linux. That's another success story for us, they now run quite happily with little intervention.

LXF It seems to be quite common to replace old Novell servers with Linux

“The ACCPAC website states that many people choose Linux because it has a reputation for availability, reliability and stability”

MT I think that's probably because the people who were using Novell in the first place were probably anti-Microsoft, so to be able to jump ship to another non-MS platform is probably a good option. In truth Globecast's primary reason for not going the Microsoft route is really just the cost of software. It seems ludicrous to spend huge amounts on Microsoft software when you can do the same if not better on the Linux platform – it just doesn't make financial sense. This isn't about techies not liking Microsoft or me not liking Microsoft, it comes down to money at the end of the day – if you can do it on a Linux

ABOUT GLOBECAST

Created in 1997, GlobeCast is a global leader in professional broadcasting and Internet delivery via satellite, and the satellite services subsidiary of leading telecommunications group, France Telecom. GlobeCast uses France Telecom's worldwide backbone and benefits from France Telecom R&D, featuring Europe's premier telecommunications research laboratory and research facilities in California's Silicon Valley, to constantly innovate and upgrade its solutions.

GlobeCast is now pioneering the development of total Internet via satellite. By focusing all of its worldwide Internet activities in one dedicated department, GlobeCast is developing and marketing the very latest and cutting-edge Internet via satellite services.

ACCOUNTANCY

◀ box why on earth are you going to invest in however many NT servers.

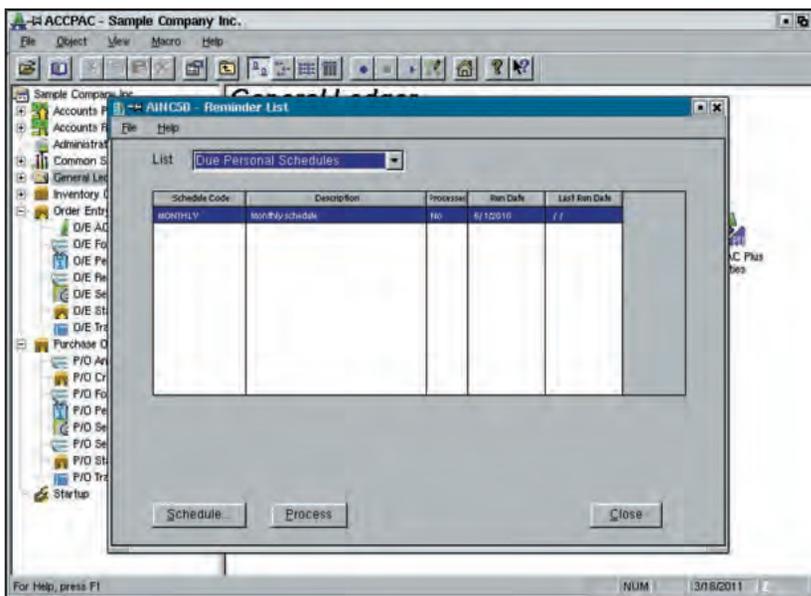
LXF We often hear from people who would like to use Linux for an accounting solution but there aren't many solutions running on Linux. What do you think you would have ended up doing if there hadn't been a version of ACCPAC for Linux.

MT Ultimately our problems were caused by the inefficiencies of the old Btrieve software. I think we would have been forced onto another database solution, which almost certainly would have been *SQL Server* on NT or Windows 2000. Or perhaps to one of the other bigger, well supported relational databases. Which would have been a big pain for us because we don't have anyone with the skills required to support that. We were quite fortunate to have the skill to go down the Linux route, and we caught ACCPAC at the right time I think – they didn't have many Linux customers at the time and weren't really geared up to support Linux at that time, but I was prepared to take a gamble that we would be able to do the Linux support at our end.

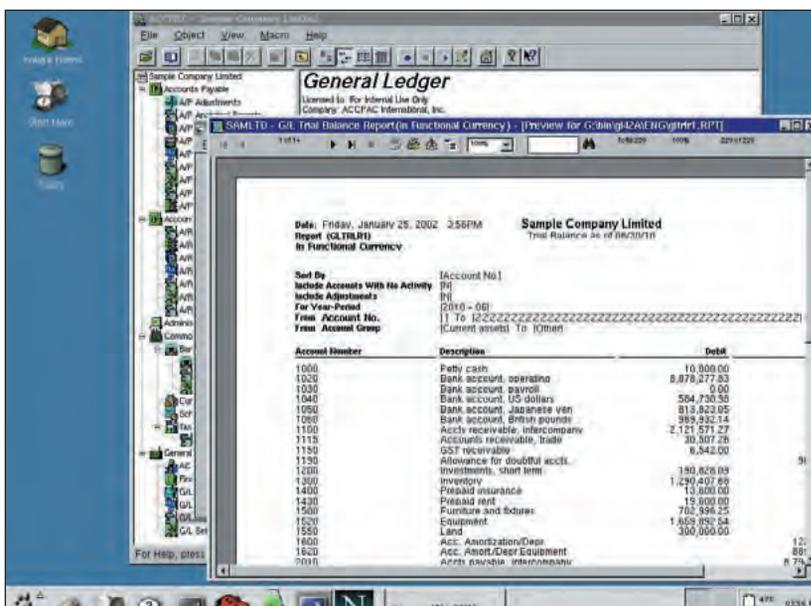
LXF So when you switched over to the Linux server, was this problematic at all?

MT No, absolutely not. It really is a bit of a joke how easy it was. We bought a new server for it, just a simple 1U HP server. Installed Linux on it – it was one of the standard distributions, I think Red Hat 6.2 at the time. That installed in about an hour. ACCPAC turned up and said, "we need to install this" and produced a CD with two or three RPMs

"The Windows users can access the accounts system, oblivious to the fact that it runs on Linux, apart from the fact that it runs without problems"



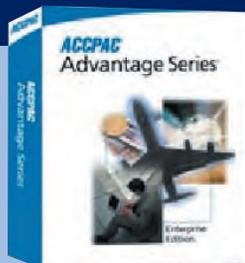
Recurring transactions can be automated with schedules and reminders.



Customisable reporting is flexible enough to process complex data requests.

ABOUT ACCPAC

ACCPAC International, Inc., a subsidiary of Computer Associates International, Inc., empowers small and medium-sized enterprises through a broad range of end-to-end e-business solutions designed to enhance competitive advantage and increase profitability. Award-winning product lines include ACCPAC Advantage Series (Enterprise Edition, Corporate Edition, Small Business Edition and Discovery Edition), ACCPAC Executive Series, ACCPAC Pro Series, ACCPAC eTransact,



ACCPAC eCRM, ACCPAC Warehouse Management System, ACCPAC HR Series, ACCPAC Business Analysis Suite, Simply Accounting, FAXserve and Simply BitWare.

Based in Pleasanton, California, USA, with offices in Asia, Australia, Canada, Mexico, the Middle East, South Africa, and the United Kingdom, ACCPAC has more than 500,000 registered clients and 5,000 business partners in more than 100 countries worldwide. Visit ACCPAC at: www.accpac.com

on it, which maybe took about four minutes.

Then they wanted to see some files over the network, so we configured *Samba*, that's maybe another five minutes, and the rest of the day was spent by ACCPAC porting the data over. So in terms of setting it up ready for them it maybe took about two hours if you aggregate everything together. Of course we had to run around and install the clients, and log everyone into the new database. The whole migration was done in less than a day.

Who do you credit for that? The simple installation of a Linux box – well, partly. I guess the real credit goes to ACCPAC who were able to port the data across to a new database format quite easily.

It was blissful. If only it was always that simple.

LXF It seems so. Things don't often happen like that in the world of IT!

MT No, in fact, quite regularly it doesn't, as we often find! ■

FEATURE NAME

Ensuring Business Continuity



If your organisation believes that Business Continuity is all about having back-up computers in the event of a disaster, it is wrong! Martin Turner, a senior consultant with systems integrator GFI Informatics and a specialist in Business Continuity, considers what turns an IT-centred Disaster Recovery plan into a true Business Continuity plan.

This article is about ensuring the continuance of your business. Mention Business Continuity and many organisations will say they are covered as they have back-up computers. The odd thing is that you will usually get this answer even if you ask business people rather than members of the IT department. The term Disaster Recovery is often confused with Business Continuity. Having 'back-up computers' will not provide Business Continuity; in fact it will probably not even provide real Disaster Recovery.

What does the term 'disaster' mean? This is something that disrupts the smooth running of your business. It can be anything from a loss of an entire set of offices, losing the people who run the business or losing data from a computer system as a result of a virus or malicious hacking.

How long can your organisation survive in the case of a disaster? If you are a manufacturing company and you lost your premises, where would you relocate to, and how long would it take to relocate? If you are a public body and you lost your major systems, what impact would this have on your ability to provide service? If you are a financial institution and you lost access to your corporate data how long could you survive without it? Likewise, if you are a dot.com company how long could you survive if your web servers were unavailable?

Notice that of all these questions, only the final one relates solely to IT. For a dot.com company, the loss of the website (by losing the web server) is a disaster as anyone attempting to access the website would find it unobtainable and would go somewhere else, possibly never to return.

FEBRUARY 2003

For all other questions, a good Business Continuity plan could actually stop a company from going bust in the case of a disaster.

Obviously any plan will include information on Disaster Recovery – *i.e.* the technical side of Business Continuity. This will cover a full plan as to what will happen to voice and data communications, and how computer applications will be provided in the case of a disaster. The full Business Continuity plan will also cover other items relating to the business requirements, such as plans for the duplication or restoration of paper records and also details of where the business will relocate to if a disaster made the premises inaccessible.

A unique plan

A Business Continuity plan is not something you can buy in a box! Each one is different, depending on the type of business of the organisation, and perhaps more often, depending on the available budget. As usual, money rules! No board of directors will authorise investing a (potentially) large amount of money without a very, very good reason.

Therefore, the start of any Business Continuity plan is a risk assessment. This does not need to be expensive to carry out. As a starting point, the head of each of the main departments of an organisation needs to be asked simple questions such as 'What would it cost the company if you were unable to access your offices or computers?' and 'What would you need to do first to re-establish your department's functionality subsequent to a disaster?'

A risk matrix can be produced showing a list of potential disasters that could befall a company together with the probability of that disaster happening, graded as red, amber or green. The impact of the disaster upon the organisation would also be graded as red, amber or green.

Clearly, the probability of a specific disaster occurring and its impact will depend largely on the particular circumstances of the organisation, such as location and type of business. For example, premises located near a river potentially may be susceptible to flooding, whereas those located on a hill probably are not. Similarly, a manufacturing plant suffering a flood may well experience a greater impact to its business than a service business with a largely field-based workforce experiencing the same 'disaster'.

These are simplistic examples and obviously there will be other considerations, but they do serve to demonstrate how significantly this assessment of risk can vary from organisation to organisation or even within the different business activities and sites of a single company.

Losses

A final column in the matrix shows the estimated potential loss of finance per day in the case of the specific disaster occurring. Look closely at any item which shows red/red – *i.e.* a high probability, and your organisation would suffer greatly. This matrix is a powerful way of convincing the board that a full Business Continuity plan is essential.

Once there is an agreement to continue with the production of the plan, then each department in the

company must be examined in some depth, noting the 'things' that are used in the day-to-day running of the business. This would include any paper records such as contracts, client instructions, correspondence, etc., together with details of the computer systems that are essential to the business. Other points to note would include the number of people required to run the business and who these should be as it may be necessary to run for some time on a skeleton staff. For each of the essential items there must be a plan to restore it in the case of a disaster.

Avoidance

All this talk of what to do in the case of a disaster is certainly not wasted, but the cheapest and most effective way is to avoid the disaster occurring in the first place! Stupid? Not at all, we are talking about protection against a disaster.

We are all aware of some aspects of physical protection

“A Disaster Recovery site may range from a duplicate set of premises to a reciprocal arrangement with a friendly company”

in the form of fire alarms, burglar alarms, locks on doors, CCTV, etc. Technical protection can be split by hardware and software; with RAID drives to provide redundancy for disk drives, dual power supplies, clustered machines; with software protection including firewalls, anti-virus software, and ensuring the security aspects of any operating system are used correctly.

Prevention will always be more cost-effective than the creation of a full Business Continuity environment. However, regardless of the effectiveness of preventative measures, a Disaster Recovery site is usually required. There are many ways in which this can be provided, ranging from having a duplicate set of premises well away from the main premises (most expensive), through to having a reciprocal arrangement with a 'friendly' company to provide a certain amount of space in the case of a disaster.

Larger organisations tend to be spread across multiple sites and therefore may be able to make use of some space in a remote site to create a disaster site. A similar approach may be taken with regard to IT. For example, rather than having a set of computers set up and waiting on a remote site for a disaster to happen, it is more cost effective to have a test machine placed off-site which could be used for production systems in the case of a disaster at the main premises.

Whatever your situation; whatever size your organisation, you need a Business Continuity plan. Don't employ consultants to write it for you, though. It's your plan and it is unique to your organisation. So, do employ experienced consultants to help and guide you in creating it! ■

For more information e-mail: swakefield@gfi-informatics.com www.gfi-informatics.com

INFOSECURITY EUROPE

GFI Informatics limited are exhibiting at Infosecurity Europe, Europe's largest and most important information security event. Now in its 8th year, the show features Europe's largest free education programme, and over 200 exhibitors at the Grand Hall at Olympia from 29th April - 1st May 2003. www.infosec.co.uk

Accessibility

Everyone's desktop



Within the last two articles we've examined some of the problems faced by various disabled users, and the solutions which make computers accessible to all users. However, to a certain extent, we've saved the best till last – GNOME2.

GNOME2 was designed almost from the start with accessibility in mind. With the next point release, 2.2, scheduled for release at the end of January, we have the potential for assistive technologies working right out of the box on the *nix desktop.

GNOME2.2 will see improvements to the theming and keyboard navigation. Anyone who cannot use a mouse will now find that all official GNOME applications respond to keyboard navigation. There are further refinements over GNOME2.0, which already implemented good keyboard navigation, as well as bounce keys, sticky keys, and all the other aids to people who are restricted in the control of their hands.

Themes continue to improve, to help users with low vision: high and low contrast, inverse themes, and better support for large fonts.

Assistive technologies

In addition to the themability – useful for those with low vision – and keyboard navigation – for those who cannot

As GNOME 2.2 approaches release, RICHARD SMEDLEY investigates the features that will make GNU/Linux the most accessible of desktop operating systems.

use the mouse, but also invaluable for those who do not wish to use one – GNOME2.2 will work with the beta bolt-on assistive technologies: *GOK* and *Gnopernicus*. *GOK* – the *GNOME Onscreen Keyboard* – is far more than just a basic on-screen keyboard. The ability to do on-the-fly word and command prediction points to its potential in mobile and embedded applications as well.

Gnopernicus is a project for users with low or no vision. Automated focus tracking and fullscreen magnification aid the computer user with low vision. For the totally blind, *Gnopernicus* provides a screen reader for the GNOME2.x desktop, as well as access to *GTK+2* and Java-based GUI applications, through speech and braille output.

(R)evolution

All of the many and diverse GNOME developers have become aware of accessibility issues in writing their programs, and all of the apps bundled in the “official” GNOME release will (soon) work with the built-in accessibility features (such as better support for large fonts), as well as with the add-on assistive technologies – *GOK* and *Gnopernicus* – which are approaching their 1.0 release. This includes the mail and calendar app, *Evolution*; the spreadsheet, *gnumeric*; the word processor, *abiword*; in short, most of the applications needed by the average

desktop user. Over the course of the coming year these apps will become more complete solutions for those with different disabilities, as the accessibility features and assistive technologies continue to be refined, and bugs are fixed in the various applications themselves.

If, like the author, you don't need an all-encompassing desktop like GNOME and KDE, you will be relieved to find that all the assistive technology works without having to run the GNOME panel on top of your window manager. You can still keep your Next-like environment, or whatever you prefer. All modern window managers, which respond to ICCM extensions, should operate with the GNOME2.2 apps and their assistive layer.

OOo, Moz and Java

However it is not just GNOME apps that benefit from the work done on ATK and AT-SPI (see *KDE box*). The next release of *OpenOffice.org* (and, naturally, *StarOffice*), and the next of *Mozilla*, will provide the same support for the themes, keyboard navigation, and work with *GOK* and *Gnopernicus*. These features are already in the current developer releases for those who wish to test out a fully accessible desktop now.

Another area of apps is opened up by Java. Anything using Java Swing classes on a 1.4.0-compatible JVM (such as Sun's own Java Virtual Machine) will also benefit from the assistive technologies discussed here.

Assistive technologies

GNOME2 features excellent built-in theming and keyboard navigation – but the assistive technologies – *GOK* and *Gnopernicus* – will not officially feature until they reach their milestone 1.0 release – sometime in the first half of 2003. They will most likely make the official GNOME release with 2.4 (in either core, or the so-called “fifth-toe” of extra applications). Let us, however, emphasise that if you can install them now (yourself, or with support) you will find that they are not beta quality software.

Gnopernicus currently interfaces with *ViaVoice* or *Festival*, as well as various common braille terminals. The magnifier component may be used as a standalone app.

GOK is now at the stage where it will provide all the key-entry requirements of those who are reliant solely on pointing devices. However people in this position are not likely to be able to **Ctrl-Alt-F2** to a shell to fix things if there are problems – so extensive testing is important here. As Bill Haneman, GNOME Accessibility Architect, says of *GOK* in the run-up to the 1.0 release, the “robustness requirement is extremely high. It will be ready for testing, but we will hold off on 1.0 until we've had more opportunity for user testing. It will be feature complete.”

In the distro

While *GOK* and *Gnopernicus* may not be bundled into GNOME2.2's official release, have no doubt that many distros will bundle up their own “complete” versions of GNOME2.2, with as full an accessibility suite as possible. Many distributions of GNU/Linux have provided some support of their own for a more accessible desktop distro and it would certainly be surprising if Sun Microsystems

were not to bundle up their accessibility work on GNOME with their own version of Linux, as well as Solaris.

SuSE, Red Hat and Mandrake, to name three of our readers' favourite distros, are sure to package the software soon, and Debian, to name the fourth, almost never fails to bring interesting Free Software to its easy packaging system. Something to be appreciated by anyone who has struggled

“Anybody needing assistive technologies on the Mac will be able to run a complete desktop: GNOME, Mozilla, Evolution, OOo”

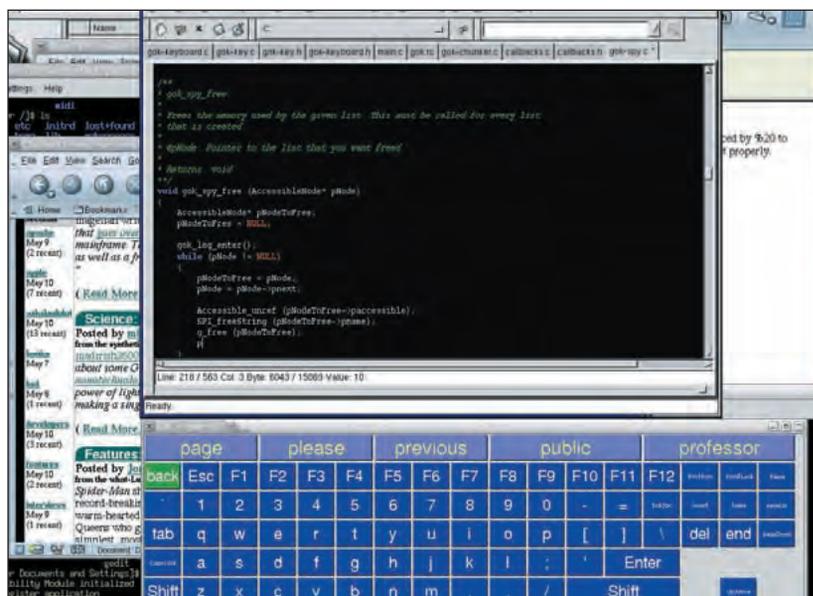
with the many dependencies of cutting-edge GNOME releases in the past.

The reverse side of this is that some distros may add in GNOME 1.4 apps, as well as their own bugfixes, branding and modifications, which may mean that not everything bundled will be compatible with the assistive technologies.

Companies and communities

Sun made large headlines when they paid 50 developers to help develop GNOME. It even made headlines with its work on accessibility for the *nix desktop. >>>

GOK – the GNOME Onscreen Keyboard – gives text manipulation from a pointing device.



PICTURE: DAVID BOLTER, UNIVERSITY OF TORONTO, CANADA

A BITE AT THE APPLE?

The other desktop Unix

GNOME2 compiles and runs on top of Darwin/OS-X – thanks to Fink. So, anybody needing assistive technologies on the Mac will be able this year to run a complete desktop of GNOME, Mozilla, evolution and OpenOffice.org.

There was a time when Apple were among the leaders in computer accessibility. Sadly that time is long passed. Perhaps the Free Software solutions reaching maturity in the coming months will spur them into their own solutions for OS-X – or even, if this is not too much to hope, producing something compatible with GNOME2 and thus expanding standards and compatibility for disabled computer users.

In light of the legislative requirements in Apple's key market – education – doing nothing is not an option for the self-proclaimed “leader of the desktop Unix market.”

ACCESSIBILITY

« However GNOME remains largely a collaborative effort, with many to thank for their help.

Working to provide accessibility has been very rewarding for those involved, with lots of positive feedback – which has given the GNOME accessibility coders a great deal of confidence in their approach to the desktop, which will provide “a substantial improvement in the quality of computing experience” and provide a “clear value proposition” that no-one else can match, according to Haneman. It has also brought them the prestigious Helen Keller Achievement Award for GNOME2’s Accessibility Framework.

Proper resourcing has been necessary to make it work, and to ship it and bundle it. Sun have provided the in-

“The next 12-18 months will be most interesting for equal access to computing”

house accessibility experts, and the experience of users and coders with disabilities, to help the project along. Other partners have played a part at all levels. Organisations such as the American Federation for the Blind (AFB) have given valuable feedback. The university of Toronto (which has a long history of research and development in accessibility for the Windows platform) has been a valuable partner, as has the Trace Center, mentioned in the first of these articles.

Much coding for *Gnopernicus* has been done by Baum,

KDE

Community efforts

GNOME may be the first choice for Sun Microsystems, and for previous releases of Red Hat, but in Europe – where Mandrake and SuSE are favoured distros – their first choice of KDE is more widely used. Competition on the desktop is good – but has this translated into a spur for accessibility on Linux’s other Free desktop? After some problems – including the disappearance of the KDE Accessibility website for several weeks – the project is back on track with more developers involved in the discussion of how to implement full Accessibility, à la GNOME, into KDE. KDE is dependent on TrollTech’s Qt – and that company are happy to respond to the needs of KDE developers.

Unfortunately this is a huge undertaking, unlikely to succeed without substantial backing in the way that the GNOME foundation found in the form of Sun engineers, accessibility expertise and company-wide resources.

CROSS PLATFORM

There is an alternative to going it alone: collaboration. Assistive technologies are a great area for cross-platform compatibility. Imagine if you could only use vital bits of hardware – such as your keyboard and mouse – with *GTK* apps alone, not with those which use *Qt*. This is effectively the position in which many disabled users will find themselves when a complete accessibility solution

works under GNOME 2.x, but not with KDE apps. There is no *nix desktop equivalent of the Win32 API or MFC class, for which developers on the dominant desktop are developing, and the case for collaboration and working to a common standard is quite compelling.

The KDE project now has the option of using the same accessibility layers (AT-SPI – the Assistive Technology Service Provider Interface, and ATK – the Accessibility ToolKit) as *GTK* and *Mozilla* use, via an intermediate wrapper. Practical ways of implementing this have been proposed by the pragmatic developers on the KDE accessibility mailing list, now a decision must be made. AT-SPI is a (cross-process) CORBA based API that lets different applications communicate to let Assistive Technology work for the users. However KDE does not use CORBA, for very good reasons. ATK is an in-process GObject-based API that does what AT-SPI does mostly, but “in-process”, it doesn’t cross process boundaries. ATK doesn’t use CORBA, so is much simpler to use.

PROKLAM

Much work has been done on one important accessibility aid, *Proklam*. *Proklam* is a plug-in based KDE project aiming to provide both a standard text-to-speech interface for KDE applications and a simple-

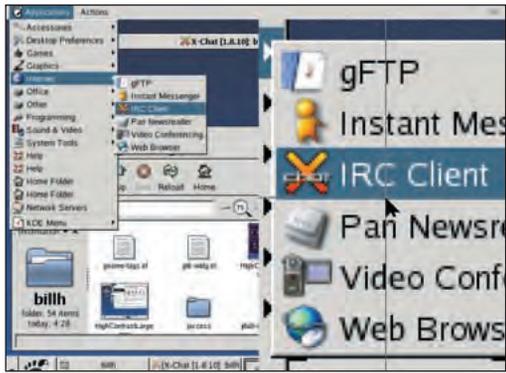


KMagnifier makes KDE more accessible to those with poor vision.

to-use configuration dialogue for the user. One of the plug-ins of *Proklam* enables it to use the **gnome-speech** class of the GNOME Accessibility Project. *KMouth* allows the computer to speak on behalf of those who have no speech. Phrases can be typed, or selected from a history. It will use the system’s configured speech synthesiser, but future versions of *KMouth* will use *Proklam*’s text-to-speech interface.

kmag, the *KMagnifier*, was a KDE 1.0 app that has been resurrected for duty in KDE 3.0, and will be useful for image analysis, as well as for the visually impaired. *Speaker* is a text to speech plugin for the *Konqueror* Web browser and *Kate* text editor, using the Festival speech engine, and currently supporting the English and German languages.

The final link in the current chain of KDE accessibility aids is *KMouseTool*. Many users can move a mouse, but find pressing the mouse buttons *extremely* painful. This is the case not only for those with repetitive strain injuries such as carpal tunnel syndrome and tendonitis, but also those with more general conditions such as arthritis, fibromyalgia, and chronic pain syndrome. *KMouseTool* sends a click whenever the mouse pauses briefly – this pause is, of course, configurable. In drag mode the mouse will pause after “clicking down” – move the mouse during this pause and it will drag until you pause again. The tool knows when to left-, right- or double-click, depending upon the window it is in. A version of *MouseTools* is also available for the MS Windows platform.



Now I see it! *Gnopernicus* offers a better view of the desktop for the sight-impaired..

Jimian have been major contributors of code, and Red Hat have also been involved, but away from all the company and institutional support – including the W3C – community support is also growing.

Although most GNOME developers have now written one or more pieces of accessibility code into their apps, it is still quite hard to incorporate newcomers into assistive technology programming, as there is quite a steep learning curve. The coming challenge for the developer community is how to build on their successes to bring in more volunteers. A challenge which faces other successful Free Software projects of a certain size.

Voice recognition

An advantage of Free Software is not having to wait to deal with problems. The source is available to download and test, before and between major releases. The developers are easy to find and contact by email, or IRC. Patches and bug reports are readily accepted.

Voice recognition is the one area where Free Software lags behind the competition. However proprietary software has little to offer here beyond *ViaVoice*, since the semi-disappearance of *Dragon Dictate*. A Free Software solution would not suffer through the troubles of one company in this way, but first we need something more functional than *Sphinx*, Carnegie Mellon University's (CMU) academic attempt at a speech input system.

The future

GNOME 2.4 (release date "some time in 2003") will feature incremental support in theming and should integrate the "visual bell" help for hearing impaired users. This latter app should also be available as a patch for 2.2.

GNOME *Mag*, used by *Gnopernicus* for screen magnification and mouse training, should be available in 2.4 as a stand-alone app. GNOME *Speech*, the text-to-speech engine, should also make an appearance. Full voice recognition is a long-term aim. As is internationalisation in speech recognition, and other areas. Cross-compatibility with other toolkits, such as KDE, is another aim.

Meanwhile, in the shorter term, it is only with the 2.2 and 2.4 releases of GNOME that the accessibility features and, in particular, the assistive technologies (*GOK* and *Gnopernicus*), become widely available for extensive, real-world user testing. The next few months will give the GNOME coders a

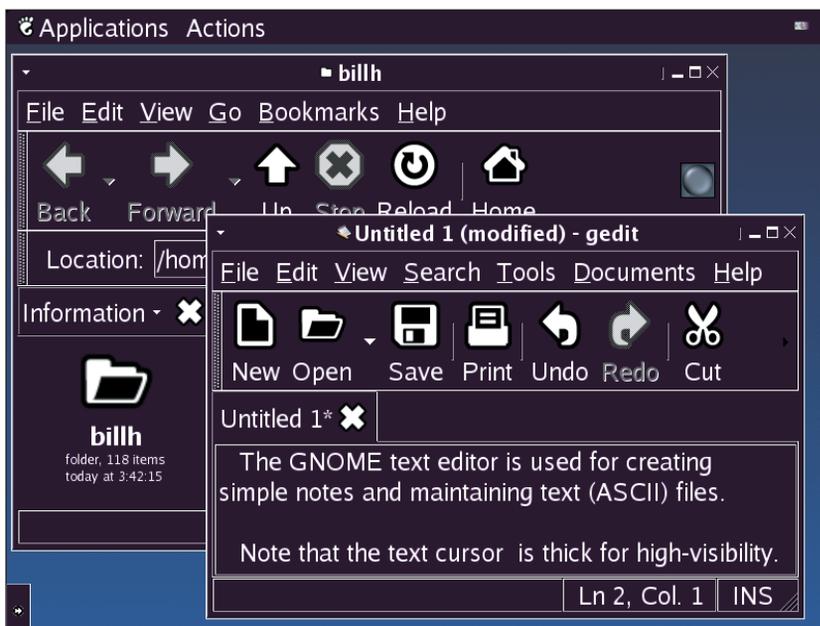
INSTALLATION

Coming to a distro near you

When GNOME 2.2 is released, shortly after the publication of this magazine, the accessibility features will be available easily as various distros, and concerned individuals, make packages available. This will be followed shortly by the appearance of GNOME2.2 in new releases of GNU/Linux distros. However if you are not prepared to wait a single day for the appearance of GNOME2.2 you can simply upgrade your GNOME2 install with GARNOME, from www.gnome.org/%7Ejdub/

[gnome/download](http://gnome.org/download)

No more painful messing about with cvs or individually updating dozens of libraries :-). Full instructions at www.gnome.org/%7Ejdub/garnome. Slackware and Debian packages are also made available by the community. The assistive technologies – *GOK* and *Gnopernicus* – will be available, in pre 1.0 form, at the same time as GNOME2.2. There will be a link for users from the developers' site at: <http://developer.gnome.org/projects/gap/>



Theming is an accessibility feature which makes the Linux desktop available to many computer users with vision problems.

chance to refine their accessibility project for real-world conditions – but will also enable packagers with an interest in accessibility issues, such as SuSE and Red Hat, to ensure that the installation process is right for the end user.

LINKS

GNOME Accessibility
<http://developer.gnome.org/projects/gap/>
KDE Accessibility
<http://accessibility.kde.org/>
Free Desktop Accessibility
www.speechinfo.org/fdawg
Sun
www.sun.com/access
GNOME roadmap
<http://developer.gnome.org/dotplan/>

LSB

GNOME accessibility may be the biggest thing in computer accessibility for a number of years, but it is not there yet as a "click-and-run" product. Those with short term needs, and the necessary I.T. support, can dive in and try it now. Others will need to watch this space for a few months.

The next 12 to 18 months will be most interesting to those concerned with ensuring equal and Free access to computing power. The work of the GNOME project and its supporters is something for which we all have grounds to be grateful.

Now that the Linux Standards Base (LSB) are taking an interest in accessibility standards and implementations – and talking to the GNOME Accessibility project – GNU/Linux finds itself teetering on the verge of being the undoubted best choice for accessibility solutions. ■

PARALLEL COMPUTING

PARALLEL COMPUTING

Extreme OpenMP

Welcome back to our mini-series on parallel programming. In this issue we are going to complete our tour of OpenMP. The material presented here is mainly meant as a quick overview of OpenMP features, more than as a complete reference. In order not to keep the discussion too abstract, we shall also show a typical example of parallelising an existing application that is not based on an “embarrassingly parallel” algorithm. We will see how in some cases, with a little modification, an algorithm that looks at first sight not suitable for parallelisation can be turned into an efficient parallel program. But your curiosity must wait a couple of sections: before doing that, we have to prepare the ground by discussing a few more features of OpenMP.

Last time we introduced the concept of a directive as an explicit instruction for parallelisation. The main directive is

In the last installment of this series, BIAGIO LUCINI introduces more OpenMP directives and shows us a few useful tips for getting tomorrow’s weather forecast in time.

#pragma omp parallel, and all other directives have to be included in the scope of it (the only two exceptions will be discussed in the next section).

The last statement requires further comments. We distinguish a lexical or static extent and a dynamic extent of a directive. The static extent can be seen as the structured block on which that directive takes effect; it does not span multiple files and does not extend to procedures called from within the structured block. It is possible to have OpenMP directives outside the lexical scope of a “parent” directive. Those directives are called orphaned directives. The dynamical extent of a directive is formed by its lexical extent and the extend of the orphaned directives referred to from within that lexical extent.

For instance, in the example

```
1 #include <stdio.h>
```

```

2
3 void printgreetings()
4 {
5 int j;
6 #pragma omp for
7 for (j = 0 ; j < 10 ; j++) {
8 printf("Hello, World! \n");
9 }
10 return;
11 }
12
13 int main()
14 {
15 int i;
16 #pragma omp parallel
17 {
18 #pragma omp for
19 for (i = 0 ; i < 10 ; i++) {
20 printf("Hello, World! \n");
21 }
22 printgreetings();
23 }
24 }

```

the **for** directive at line 19 is in the lexical extent of the **parallel** directive, while the **for** directive at line 6 is an orphaned directive that makes sense just when inserted in the dynamic extent of a "parent" directive, as in this case.

This classification might seem academic, but as a matter of fact there are specific rules for binding and nesting directives: a program that ignores them is illegal and when compiled can be executed incorrectly. For those rules we refer to the specialised literature.

There are also rules that state how parallel and serial regions must be related. We won't go into much detail here, since often they are pretty obvious. For instance, it is reasonable that it is not good practice to conditionally jump from within a parallel region to a serial region. In fact, this is illegal in OpenMP. As usual the specialised literature contains more information.

More on work sharing

As we have pointed out in the previous part of this tutorial, probably **#pragma omp for** is the most used construct to distribute work among different processors. Another possibility for parallelisation is to let different processors execute different parts of the code. This is accomplished via the "sections" directive, as in the example below:

```

#pragma omp sections
{
#pragma omp section
{
/* first structured block */
}
...
#pragma omp section
{
/* last structured block */
}

```

```

}

```

The **sections** directive encloses one or more **section** directives. Each of them is executed by only one thread in the team and different sections can be executed by different threads. This gives an effective way of parallelising part of the code other than **for** loops that do not depend upon each other. A "sections" directive implies synchronisation at the end of the referred structured block (unless a **nowait** clause is specified).

The clauses accepted by **sections** are **nowait**, **private**, **firstprivate**, **lastprivate** and **reduction**, which have all been introduced in the previous part.

Many programs will use the **parallel** directive just as a container for a **for** or **sections** directive. When this is the

“It is good programming practice to avoid synchronisation wherever possible – it’s bound to slow down the resulting app”

case, there are shortcuts, respectively the **parallel for** and the **parallel sections** directives. Those directives accept both the clauses of the **parallel** directive and of the **for** or **sections** directives but **nowait**.

The last work sharing construct is the **single** directive. When it is specified (as usual with **#pragma omp single**), the following structured block is executed by only one thread of the team. This is useful when executing parts of the code that aren't parallelisable or for which parallelisation is not required, like *e.g.* reading some inputs. **single** accepts the clauses **private**, **firstprivate** and **nowait**. The latter means that the default behaviour is that the threads resynchronise at the end of a **single** structured block.

Synchronisation

When the code contains dependencies, parallelisation (when possible) will require some extra work. In those cases it is often necessary to synchronise the threads at some point or at least to communicate some variables among threads. A few OpenMP directives allow to deal relatively easily with those cases.

The most immediate synchronisation directive is **barrier**. This directive gives an explicit point where all threads of the teams have to synchronise. The execution will restart only when all the threads have reached the barrier. The "barrier" statement must be enclosed by curly braces:

```
{#pragma omp barrier}
```

There might be regions in the code that are not thread-safe. If those regions are in the scope of a **for** or **parallel for** directives, a way out may be the **ordered** directive. This directive specifies that the following structured block has to be executed in the same order as in the serial case. Being applicable for loops for which the **ordered** clause has been specified, this directive allows only one thread at a given time in the ordered region.

Similar to the **single** directive is the **master** directive, the differences being that in the latter case it is specifically the master thread that executes the instructions and there is no synchronisation at the end of the region. 

MORE RESOURCES

Compunity (www.compunity.org) is the **OpenMP community website**. This site provides **good information on hot topics and events**. The **link collection is also remarkable**.

ERRATA FOR PART 1

The following "errata" apply to part 1 of this series: **1) Page 7, first column, top: "topology" should have been "typology"; 2) Page 7, first column, bottom: the correct definition of the efficiency is $e = t(1)/(p*t(p))$; 3) Page 9, second column, top: the opening sentence of the penultimate paragraph should have read "On the opposite site we have distributed memory machines".**

PARALLEL COMPUTING

When the value of a shared variable must be refreshed, the **flush** directive can be called. The construct is **#pragma omp flush (list)**, where **list** is a list of all variables that have to be reread. If **list** is omitted, all shared variables are flushed back to memory. The **flush** directive finds its best application when a processor uses a shared variable that in principle could have been updated

“For the library routines the names are the same in Fortran as in C, and the returned variable is often of the same type”

by another processor. **flush** is implied by most synchronisation and work sharing directives.

The last two synchronisation directives are concerned with code that has to be executed only by one thread at a time. A typical example is when a variable can be updated by more than one process. For a single instruction, there is the **atomic** directive. The statement following a **#pragma omp atomic** can be evaluated by more process at a time, but the implied writing in memory is executed atomically. The statement must have the form **x operator = expression**, where **operator** is one of **+, -, *, /, &, ^, >>** or **<<** and **expression** is a scalar expression not referencing the variable **x**. If the critical region contains more than one single statement, we must use the **critical** directive. The syntax is **#pragma omp critical name**, where **name** is an optional label. Separate regions delimited by **critical** directives with the same **name** are treated as the same region. Critical regions for which the label has been omitted are treated as the same critical region. Only one thread is allowed inside a critical region at a given time.

Before concluding this section, it may be worth remarking explicitly that synchronisation directives do not accept clauses.

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>
#define N 100
#define PREC 0.000001
#define ITER 1000000

unsigned long int dim=N*N;
float spacing = 1.0 / (float) N;

typedef struct neighbours
{
    int up;
    int down;
    int right;
    int left;
} NEIGHBOURS ;

unsigned long int calculate_point(int icoord1, int icoord2)
{
    int ic1, ic2;
    unsigned long int lpoint;
```

OTHER CLAUSES FOR THE “PARALLEL” DIRECTIVE

Persistent threads

The last two clauses for **#pragma omp parallel** we have not yet covered are **copyin** and **if**. **if** is followed by a scalar expression enclosed by (); if that expression evaluates to zero, the following structured block is not executed in parallel.

This gives a way to parallelise a chunk of the program conditionally. The **copyin** clause refers (as well as all the other clauses but **if** and **nowait**) to data scoping. It accept a list of variables enclosed by (). This clause is strictly related to the **threadprivate** directive. This directive gives a way to make variables of a thread persisting in different parallel instances of the same thread.

threadprivate variables depend (as the name suggests) on the particular thread.

For the variables specified as arguments of **copyin**, **threadprivate** variables (that otherwise would have been uninitialised on entering a parallel region) are initialised with the value they have in the master thread.

```
ic1 = icoord1;
if(ic1 < 0) ic1 = N - 1;
if(ic1 == N) ic1 = 0;

ic2 = icoord2;
if(ic2 < 0) ic2 = N - 1;
if(ic2 == N) ic2 = 0;

lpoint = ic1 + ic2*N;

return lpoint;
}

void define_grid(NEIGHBOURS *neigh)
{
    int i, j, ip1, jp1, im1, jm1;
    unsigned long int point;

    for ( i = 0 ; i < N ; i++ ) {
        ip1 = i + 1;
        im1 = i - 1;
        for ( j = 0 ; j < N ; j++ ) {
            jp1 = j + 1;
            jm1 = j - 1;
            point = calculate_point(j,i);
            (neigh + point)->right = calculate_point(ip1,i);
            (neigh + point)->left = calculate_point(jm1,i);
            (neigh + point)->up = calculate_point(j,ip1);
            (neigh + point)->down = calculate_point(j,jm1);
        }
    }

    return;
}

double update_grid(float *potential, float *function,
NEIGHBOURS *neigh)
{
    unsigned long int point;
    float new_pot;
    double diff, diff_max=0.0;
    extern unsigned long int dim;

    for ( point = 0 ; point < dim ; point++ ) {

        new_pot = 0.25*(potential[(neigh+point)->right] +
            potential[(neigh+point)->left] +
            potential[(neigh+point)->up] +
            potential[(neigh+point)->down] -
            spacing*spacing*function[point]);

        diff = new_pot - potential[point];
        if ( diff < 0 ) diff = - diff;
        if ( diff > diff_max ) diff_max = diff;

        potential[point] = new_pot ;
    }
}
```



Earth Simulator – you can't get much more massively parallel than the world's largest supercomputer.

```

return diff_max;
}

void define_function(float *function)
{
    int icoord1, icoord2;
    unsigned long int point;

    icoord2 = N/2;

    icoord1 = N/4;
    point = calculate_point(icoord1, icoord2);
    *(function + point) = -30.0;

    icoord1 = (3*N)/4;
    point = calculate_point(icoord1, icoord2);
    *(function + point) = +30.0;

    return;
}

int main()
{
    NEIGHBOURS *neigh;
    float *potential, *density;
    int finished;
    int count, point;
    double error;
    extern unsigned long int dim;

    neigh = calloc(dim, sizeof(NEIGHBOURS));

    potential = calloc(dim, sizeof(float));
    density = calloc(dim, sizeof(float));

```

```

define_grid(neigh);
define_function(density);

/* main iteration */
finished = 0;
for ( count = 0 ; count < ITER ; count++) {
    error = update_grid(potential, density, neigh);
    point = calculate_point(N/4,N/2);
    printf("%f \t %f \n", potential[point], error);
    if ( error <= PREC ) {
        finished = 1;
        break;
    }
}

free(neigh);
free(potential);
free(density);

count++;

if ( finished == 1 ) {
    printf ("The computation required %d iterations \n",
count);
}
else {
    printf ("The required accuracy was not reached after
%d iterations \n",    count);
}
return;
}

```

PHOTO: JAPAN MARINE SCIENCE AND TECHNOLOGY CENTER THE EARTH-SIMULATOR CENTER

Avoiding synchronisation

While it is a good thing to have directives for dealing with cases where syncing is required, it is a good programming practice to avoid synchronisation as much as possible: synchronisation implies communication among threads or pauses in all threads but one, and this is bound to slow down the resulting application. There are cases in which synchronisation cannot be completely avoided, but it is anyway recommendable to reduce the use of synchronisation directives to a bearable minimum. Often this requires rewriting a considerable part of the serial app. We are going to discuss in this section a typical case.

Our prototype application is given in listing 1. This is an example of computation on a grid of points approximating a continuum space region. There are countless applications of this technique, an obvious one being weather forecast. For those who like the details, in our particular case the code implements the computation of the electromagnetic potential due to two point-like particles with opposite charge located at $(1/4, 1/2)$ and $(3/4, 1/2)$ on the compact bidimensional domain $[0,1] \times [0,1]$ with periodic boundary conditions (i.e. with the conditions $f[0,i] = f[1,i]$ and $f[i,0] = f[i,1]$ for any function f and integer i in $[0,1]$). This is done by iteratively solving the discrete Poisson equation on a grid with points spaced $1/N$ apart. The iterative process starts from the trivial value of the field and stops when a given precision is reached. The continuum limit is recovered at large N .



PARALLEL COMPUTING

◀ The goal is to determine the value of the array “potential” on each point of the grid. The problem is that each entry of the array is related to other entries of the same array, and this results in a spiral of dependencies that will make your life harder if your boss has asked you to parallelise that code.

Before giving up, let us give a closer look to the serial code. The source of our headache is the iterative computation for determining **new_pot** (routine **update_grid**). This value depends on the value of **potential** in nearest neighbour sites and the problem arises when for instance one process is trying to update the value in the point (i,j) using the stored value at (i+1,j) while another process is trying to do the reverse. This could be thought as a case for using synchronisation directive, e.g. by locking **new_pot** inside a critical region, but since the algorithm is relatively simple, the resulting parallel application will run slower than its serial counterpart, since latencies due to the syncing process will account for a substantial part of the execution time.

The key for parallelisation is the observation that, due to the fact that only the values of “potential” at nearest neighbour points determine its value at a given point, we could update at the same time two non-neighbour points. We need a way to go through this procedure systematically. The serial algorithm works by mapping sequentially the points of the grid onto a set of integers used as indices of array functions and then cycling over the argument of “potential”. If we don’t want to introduce extra overhead due to the way arrays are laid into memory, the best thing to do is to think of a different mapping that allocates sequentially points with a non-competing update. That map

is given in the new **calculate_point** below:

```
unsigned long int calculate_point(int icoord1, int icoord2)
{
    int ic1, ic2;

    unsigned long int lpoint;
    int poffset, doffset, sumcoord, parity;

    poffset = dim/2 - 1;

    ic1 = icoord1;
    if(ic1 < 0) ic1 = N - 1;
    if(ic1 == N) ic1 = 0;

    ic2 = icoord2;
    if(ic2 < 0) ic2 = N - 1;
    if(ic2 == N) ic2 = 0;

    sumcoord = ic1 + ic2;
    parity = sumcoord - (sumcoord/2)*2;

    if (parity == 0) {
        doffset = 0;
    }
    else {
        doffset = poffset;
    }

    lpoint = (ic1 + ic2*N)/2 + doffset;

    return lpoint;
}
```

Now an index less than $N^2/2$ points to a point with even sum of coordinates (“even” points) while the second part of a vector refers to “odd” points. It is now easy to see that the value of “potential” on a “even” point is determined only by values on “odd” points. Said differently, all “even” points can be updated at the same time. The same is true for “odd” points.

Now it is just matter of rewriting the loop iterations in **upgrade_grid** as a double **for** loop and parallelising the inner one:

```
double update_grid(float *potential, float *function,
    NEIGHBOURS *neigh)
{
    int point;
    int par, start, end;
    float new_pot;
    double diff, diff_max=0.0;
    extern unsigned long int dim;

    for ( par = 0 ; par < 2 ; par++ ) {
        start = (par*dim)/2;
        end = (par+1)*dim/2;
        #pragma omp parallel for default(shared) private(point,
        new_pot)
        for ( point = start ; point < end ; point++ ) {

            new_pot = 0.25*(potential[(neigh+point)->right] +
                potential[(neigh+point)->left] +
                potential[(neigh+point)->up] +
```

OPENMP FOR THE FORTRAN PROGRAMMER

Some minor variations

There is not much difference between how OpenMP is implemented in Fortran or C/C++: all directives, clauses and library functions available to C are available also to Fortran. A Fortran directive starts with a sentinel, which can be one of **!\$OMP**, **!\$OMP** and **!\$OMP** for Fortran 77, while the last one is the only valid sentinel for F90 in free form. Corresponding directives have the same logical function, so in Fortran there is **!\$OMP PARALLEL**, **!\$OMP DO**, **!\$OMP SECTIONS**, etc., the rule of thumb being to replace the **#pragma omp** you find for a C/C++ directive with **!\$OMP**. In Fortran directives have also a closure, so a **!\$OMP PARALLEL** is ended by a **!\$OMP END PARALLEL**, a **!\$OMP DO** by a **!\$OMP END DO**, etc. In this case, the rule is that every time in C you need curly braces for delimiting the structured block to which a directive applies in Fortran you need a **!\$OMP END <directive_name>**.

The clauses for directives are almost the same, with minor variations (for instance, in Fortran **MAX** is supported as an operand for reduction, while in C is not). Also in Fortran, as usual, words are case insensitive, so **!\$OMP** is the same as **!\$omp**. For the library routines, the names are the same as in C and also the returned variable (when any) is often of the same type (the exception being that flag variables in Fortran are logical expressions, not integer values as in C). Note that a function that returns void in C translates into a subroutine in Fortran. As for compilation, the Intel Fortran Compiler behaves exactly in the same way as the C/C++ one, *ifc* being the name of the compiler executable. *pi.f90* on our CD/DVD gives an implementation of OpenMP in a very simple case. A somewhat more complicated example can be downloaded from the OpenMP site.

```

        potential[(neigh+point)->down] -
        spacing*spacing*function[point]);

    diff = new_pot - potential[point];
#pragma omp critical
    {
        if ( diff < 0 ) diff = - diff;
        if ( diff > diff_max ) diff_max = diff;
    }
    potential[point] = new_pot ;
}
}
return diff_max;
}

```

The full parallel code is on the coverdisc, but the above procedure is the only one that has been modified for parallelisation: our efforts have been focused on the most computationally intensive part of the code. We could also have parallelised the loop in **define_grid**, but this would have saved just little time (if any at all: try it for yourself). The solution given above is the simplest possible one, but not the optimal one: for instance, if we had extended our parallel region to the iteration loop in the main program (not as a parallel **for**), which in the present form is forbidden because of the **break** instruction, we would have eliminated the latencies due to the frequent fork/join mechanism in the present version of the parallel code. Also, the critical region could have been avoided by defining a “reduction” by hand (the set of allowed reduction operators in C does not include the maximum among some numbers), e.g. by defining a shared vector that contains **diff_max** corresponding to thread **k** in the (k-1)th entry and then taking the maximum of its components outside the parallel for. All these improvements (note also the restriction that N is even) are left to the reader as an exercise.

Lib functions & enviro variables

As we already know, directives (and supported clauses) are only one component of the OpenMP specifications. The other parts are library functions and environmental variables, often closely related to each other.

OpenMP library functions and routines offer a high-level and system-independent interface for managing parallelisation. Among the provided functions, we have already seen **omp_set_num_threads()**,

omp_get_num_threads() and **omp_get_thread_num()**.

Those functions are concerned with thread management.

Other functions used for managing threads are

omp_get_max_threads(), which returns the maximum possible value of the thread identifier, and

omp_get_num_procs(), returning the number of

processors that the program can use (we stress that this does not have to coincide with the number of threads).

The number of threads available for executing a program can be adjusted dynamically. This is obtained by setting to **TRUE** the environment variable **OMP_DYNAMIC**. A setting of **FALSE** disables dynamical thread adjustment. Similar results can be obtained by using the function **void omp_set_dynamic(int dynamic_threads):** if **dynamic_threads** is non-zero, the dynamic mechanism is

switched on, otherwise the number of threads will be fixed. This function, which as usual has priority over the corresponding environment variable, has to be called from a serial region of the program. The behaviour of **omp_set_num_threads()** is modified when the dynamic thread mechanism is enabled: in this case, the specified number of threads will be the maximum allowed number of threads. **int omp_get_dynamic()** is used instead to determine if the dynamic thread mechanism is enabled or not.

omp_in_parallel() (called from within the dynamic extent of a parallel region) returns zero if that part of the

“Parallel regions inside other parallel regions, with slave threads creating new teams for which they act as masters”

code is being executed serially, a non-zero integer value otherwise. This function can be used in parallel regions subject to an **if** clause to determine whether the condition has prevented parallelisation or for setting a switch on the basis of whether a given section is executed in parallel or not.

It is also possible to have parallel regions inside other parallel regions, with slave threads of the team creating new teams for which they act as masters. This feature is called “nested parallelism”. Not all the vendors implement it. Where implemented (as in the case of Intel), it is usually disabled by default. It can be enabled by calling **void omp_set_nested(int nest)** with a non-zero value of **nest** (**nest** equal to zero is used to disable nested parallelism) or by setting the environment variable **OMP_NESTED** to **TRUE** (the other possible value being **FALSE**). **int omp_get_nested()** can be used to check nested parallelism.

Finally, OpenMP defines library functions for managing locks. They are **omp_init_lock**, **omp_destroy_lock**, **omp_set_lock**, **omp_unset_lock** and **omp_test_lock**. For usage and syntax of those functions we refer to the specialised literature.

This concludes our overview of a standard OpenMP implementation. Different vendors can offer different extensions to the standard. Those (unportable, but useful in some cases) extensions are documented in the compiler reference material.

Conclusions

Our mini-series on parallel programming ends here. As anticipated in the introductory part, the main goal was not quite to teach parallel programming, but to help you to decide whether parallel programming is, in a given specific case, an effective way of doing things. In our tutorials we have discussed OpenMP, which to our knowledge is the most cost-effective solution, since it can be implemented also with cheap off-the-shelf hardware. OpenMP is also a good solution from a technical point of view, since it often allows parallelisation of existing apps with little modification of the source code. The material we have presented here is by no means exhaustive, parallel programming being a huge and complex field. However, it should be enough at least to get you started if you've decided to follow the OpenMP path. ■

