

Movie magic...

When we were discussing this month's cover feature, which, as you will have gathered, is about using DV equipment with Linux, I got to thinking about why it was necessary and the different experience you would have on say a Mac or a PC. Under Windows XP, when you connected your Firewire camera to the PC, it would pop up one of those messages extolling you to create a movie with *Windows Movie Maker*. On your average Linux box, absolutely nothing will happen at all. Even if you run editing software like *Kino*, you'll probably be told that no devices are available – you'll have to load the relevant modules manually more often than not, and get bogged down in where *V4L* puts the devices. It isn't a question of technology – the technology is all there. But perhaps because of the disparate nature and distributed development effort, sometimes things don't tie together as you might want them to. It is eminently possible to create a hotplug script that recognised a camera being added to the system,

loaded the modules and then started *Kino* if you wanted. But would it work globally – maybe the user doesn't even have *Kino* installed...?

With an OS controlled by one organisation, such things are a lot easier, but these are perhaps some of the issues that need to be worked out to make Linux truly ready for the modern desktop generation, who don't have the first clue about the technologies involved or how systems actually work. In the meantime, it is still necessary for us to publish articles like our main feature.

On a slightly related topic, we should also celebrate what the Linux community has done well. The achievements in the last few years have been spectacular, and the rate of improvement in many areas is probably accelerating. Our awards, voted for by the readers, are there to offer some small recognition for the effort that has gone into improvements in software, in usability and in promoting and developing Linux as a whole. Do try to find time to become part of the experience: find out how on page 46.



Nick Veitch EDITOR

Make movies the Linux way – get your DV camera talking to your computer **p48**

Your vote counts! Decide who most deserves a *LXFi* **p46**

Write your own game! And discover the secret of *Trout Wars* and *SDL*! **p70**



MEET LINUX FORMAT'S TEAM OF WRITERS...



Andrew Channelle
LXF's newshound and the Beginners' best friend, Andy organises his entire life with *KOrganizer* this month.



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



Mike Saunders
Scouring the Internet for new software for *Hot Picks*, he's a bit of a compression software expert too.



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



Paul Hudson
At the end of the day, Paul is trying to prune root-and-branch his use of tedious cliches that leave us all sick as parrots.

Biagio Lucini
Digital video expert and clustering know-all? If he's seen the movie *Weird Science*, it's time to be afraid...

Hoyt Duff
LXF lifer and leading Linux author, Hoyt is a master of finding easier and better ways to do things with Linux.

Chris Brown
When it comes to packets and protocols, Dr Chris knows what's what, and isn't afraid to draw a diagram.

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

Hans Huberland
Real-life sysadmin who lives in the basement at Rackspace, Hans is here to answer your sysadmin queries.

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

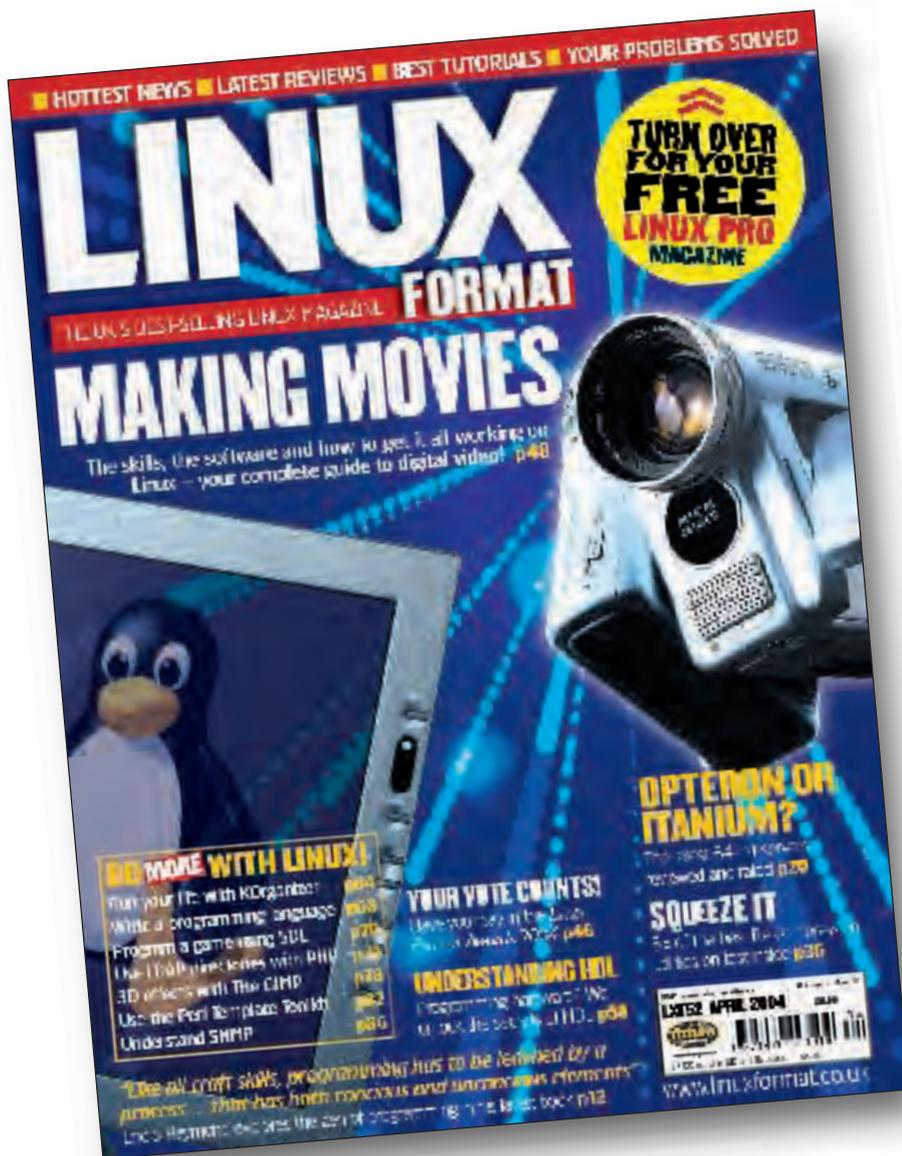
Or send your letters to:
LINUX Format, Future Publishing,
30 Monmouth Street, Bath, BA1 2BW
Phone: 01225 442244
Fax: 01225 732295

More contact info on p114

LXF52 April 2004

Contents

LINUX
 FORMAT

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

COVER HIGHLIGHTS

Hollywood 48 calling Linux!

Your complete guide to making and storing successful movies with Linux

20 Opteron or Itanium?

The argument settled: two servers go head-to-head

36 SQUEEZE IT!

Six of the best compression utilities rounded up

46 Your vote counts

 How to have your say in the *LXF Awards 2004*

58 What On Earth is... HDL?

Unlock the secrets of programming hardware



'Independent studies have shown that Linux is more expensive than Windows in four out of five business situations!'

08 If "independent" means "funded by MS..."

'But what's a Zen story without a little mystery to tease the drunken monkey of the mind and provoke enlightenment?'

12 Eric S Raymond on his brand-new book.

LINUXPRO

IT Law Changes in UK copyright affect *your* business **Grid** Back To The Future with Oracle **Virtual machines** Save *your* business time and money **Embedded Linux** MontaVista's CEO speaks exclusively to us



20 Saviour of Itanium
or exercise in futility?



06 News

Analysis of the latest Linux stories

12 Mailserver & Reader Tips

Got an opinion? Share it with us!

24 Linux Megapack review

Pile of popular distros for pence

26 Intel C++ Compiler 8 review

Lightning-quick coding system

27 Portland C/C++ Compiler 5.1

Features as far as the eye can see...

28 KDevelop 3.0.1 review

Free IDE sets the bar higher

30 Arkeia 5.2 review

Flexible network backup solution

32 KOffice 1.3 review

More powerful productivity today

33 Book reviews

Geek humour in *The Joy of Tech*

42 Hot Picks

Latest, greatest Open Source apps

64 Beginners' Guide to KOrganizer

Sort your life out!

68 Compiler writing tutorial

Functions can return useful values

70 Game Programming tutorial

NEW SERIES First steps with SDL

74 PHP and LDAP tutorial

Get directory info from inside scripts

78 The GIMP tutorial

Executing 3D design in a 2D app

82 Perl Template Toolkit tutorial

NEW SERIES The basics explained

86 SNMP tutorial

NEW SERIES Network essentials

90 Answers

Got a problem? Ask our experts

98 User Groups

The Great 2004 LUG Explosion!

114 Next Month

Gazing into the LXF crystal ball

12



58



24



70

COVERDISCS

A DVD or 3 CDs packed full of
the latest Linux goodies **105**



CDS A, B AND C

KDE 3.2 Latest version of desktop environment; **ClarkConnect** and **SME Server** Everything for your network; **LMsensors** Keep track of what's happening inside your PC; **hdparm** Tune up your hard drive ; **SpamAssassin** If only it *really* did maim those spamming morons! **LiarLiar** Don't let your boss get hold of this... **QTVision** Watch TV on your PC **Video apps** everything you need to follow our main feature



DVD

8 DISTROS! includes a rescue that fits on a credit-card CD; **Gcompris** kids love Linux; **MyPasswordSafe** store and protect your logins; **SystemRescueCD** does what it says on the tin

Please **CAREFULLY** read the coverdisc instructions on pages 105 onwards before installing from the coverdiscs!



SAVE MONEY!

SUBSCRIBE TO LINUX FORMAT

See **page 102** or phone 0870 4448645



Newsdesk

- Kernel 2.6 help ● SCO News ● Lionstracs Mediastation X-76 ● ● Linux TCO lower than Windows
- Lindows and Lin---s ● MS Office for Linux? ● Visual Web Development ● Chessbrain vs human

DEVELOPMENT GATHERS PACE

Getting ready for the next-generation distros

Red Hat's 'community' offshoot (Fedora) and MandrakeSoft are both ramping up development effort on their latest distributions, and both are expected to feature a large selection of cutting edge software including kernel 2.6, the latest releases from both GNOME and KDE and the most recent edition of the Mozilla suite.

Fedora Core 2 will be the second major release from Red Hat's recently divested community development project and, as we go to press, the first beta/release candidate (Yarrow) is undergoing testing. The 'official' release of Fedora Core 2 (FC2) is scheduled for April 6 2004.

Most significant additions to FC2 include KDE3.2, GNOME 2.6 (the test version includes a development snapshot of GNOME 2.5, but the final version should be available on March 8, in time for integration into the final release) and Mozilla 1.6. The main reason for the test release is to allow early adopters to put the latest kernel (2.6.2) through its paces before it is released into Red Hat's revenue-spinning Advanced Server line.

Amid the usual bug reports and random crashes, testers have been finding a lot to admire about the new kernel, including the long-awaited inclusion of the *Advanced Linux Sound Architecture (ALSA)* and a general

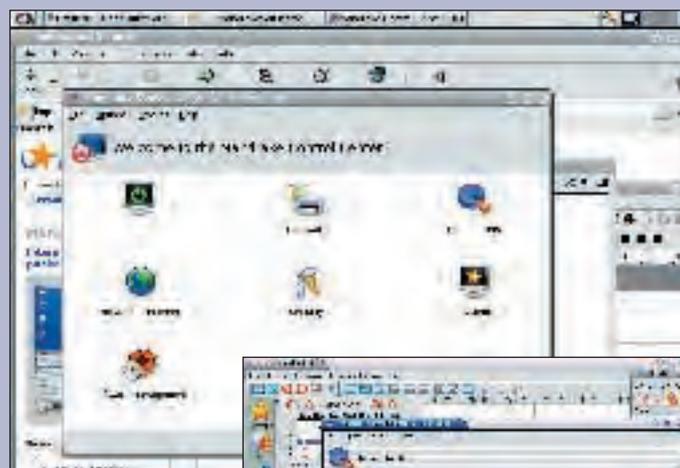
MANDRAKE LEADS THE RETREAT FROM XFREE86 4.4 GPL incompatibilities?

Mandrake's move to drop *Xfree86 4.4* in favour of a previous release has been followed by the developers of Debian, Gentoo and OpenBSD. Red Hat has also indicated that it won't ship *Xfree86 4.4*. The problem has arisen due to changes in the *Xfree86* license decided on, according to David Dawes, to protect the work of contributors in the distribution of binaries. The so-called *Advertising Clause* inserted into the *Xfree86 4.4* license has come, according to Mandrake, too late for inclusion into Mandrake 10.0; and is also likely to be incompatible with the GNU General Public License.

Red Hat's Mike Harris, discussing the change on freedesktop.org's mailing lists said Red Hat said there were no plans to use the software in future releases. "The new *XFree86* license seems to be intended to restricting existing freedom for no real-world, technical or other gains. At least no gains that are beneficial to the community." He said it was time for Open Source developers to band together and give *X11* back to the community.

"I very much look forward to working together in collaboration with the

improvement in speed. The binary download consists of four CD ISOs at about 640MB each, while the sources will consume a further quartet of CDs.



Above: **Macromedia Flash is fully supported: now you can see all those web pages where style eclipses content...**



Above: **Stepping back to *Xfree86 4.3* makes things more stable overall, with little perceivable loss of performance for the GUI user.**

Debian project, FreeBSD, Mandrake, SUSE, X.org foundation, the other BSDs, and any/all other interested parties on a true Open Source solution," he stated.

Mandrake has also been touting its 2.6 flavoured wares in the form of a brand-new release candidate for Mandrake Linux 10. The main change

in this release candidate is a drop back to *Xfree86 4.3* (see *Mandrake* box above) due to licensing issues. However, MandrakeSoft has also

LICENSING

Let Java go

Sun Microsystems has answered

criticisms from some sectors of the Open Source software(OSS) community about its control of Java by announcing a third OSS licensee.

Following in the footsteps of the Apache Software

Foundation and JBoss,

ObjectWeb has licensed the *Java 2 Enterprise Edition (J2EE) 1.4* specification and *Compatibility Test Suite*. This Open Source-friendly approach, announced by Sun in March 2002, is designed to improve the profile of Java in OSS development environments.

Christophe Ney, President of ObjectWeb's Executive Committee said that the move by Sun was vital to the project's community of developers and users. "It will consolidate the position of *JOnAS* as a keystone of our open-source ecosystem," he said.

Sun's Mark Bauhaus said the *J2EE 1.4 Application Server* was making a significant impact on Open Source development. "The rocketing momentum is being maintained with more than 400,000 downloads of our free *J2EE 1.4 Application Server* and broad adoption of *J2EE* compatibility in the Open Source community."

However, Eric S Raymond, President of the Open Source Initiative (see page 12) recently published an open letter to Sun claiming that the company's "insistence on continuing tight control of the Java code has damaged Sun's long-term interests by throttling acceptance of the language in the Open Source community." Java is now challenged, Raymond said, by scripting languages such as Python and Perl, and with the choice between control and ubiquity, Sun is choosing to remain in control.

He pointed out that Sun's terms were so restrictive that Linux distros cannot even include Java binaries for use as a browser plugin, let alone as a standalone development tool.

"You have millions of potential allies in the Open Source community who would love to become Java developers and users, if it didn't mean ceding control of their future to Sun," he wrote.



LINKS

www.mandrakelinux.com
<http://fedora.redhat.com/>
www.knoppix.org/

announced major changes in the way it distributes Linux, enabling the company to keep its reputation for bleeding-edge technology for the hard core Linux user, while placating more conservative home and business users who would rather have "rock-solid stability" than pretty transparent windows and the occasional crash.

MandrakeSoft says that it will now split development into two areas, called *Mandrake Linux Community* and *Mandrake Linux Official*. The former will be based around periodic releases from the *Mandrake Cooker* – much as happened before – while the latter will be released a few months after the *Community* product and will feature all of the errata releases making it "particularly solid and problem-free."

It is though the development changes are a result of Red Hat's successful Fedora project, which works along the same lines, and criticisms of the quality assurance problems experienced with Mandrake 9.2 which 'featured' a number of bugs, including one which had the potential to destroy certain CD-ROM drives.

And finally, Knoppix hackers have been beaver away on a 2.6-based live distribution which, beyond the obvious advantages of being able to carry a fully-fledged OS in your pocket, should also allow nervous upgraders to run comprehensive hardware checks before diving headlong into an upgrade.



NEWSBYTES

■ As you're aware, a bumpy month for **Microsoft** was topped when a serious security breach led to the release of portions of Windows 2000 and Windows NT source code 'into the wild'. While the breach is not thought to be as serious for MS as first thought, Open Source developers have pointed out the danger of future 'IP infringement' cases if a hacker concerned with interoperability issues can be shown to have seen the source code.

■ The process of migration from Windows to Linux in **Munich** is not going as smoothly as initially hoped. In addition to a few financial and technical difficulties, the body charged with managing the migration has said it has encountered some resistance to change within city hall. SUSE and IBM are expected to offer technical support especially in the area of application testing.

■ One city authority keeping a careful eye on developments in Munich is **Paris**. French local government officials are said to be keen on finding an alternative to the vendor lock-in that Microsoft's new licensing engenders.

■ An enterprising IT trainer in the German city of Schwäbisch Hall has told how he used women to demonstrate the advantages of Linux and handed out **stuffed penguins** to make workers feel warm and fuzzy. "We put the chairwoman of our workers' council on stage in front of all the municipal workers, and showed her using the new system. After that, we found that no man would say that he couldn't use his PC now that everyone knew a woman could do it."



■ **IBM** has prepared a strategy guide to aid developers in their transition from Windows to Linux. Written by the company's e-business architect Chris Walden, the guide covers nine stages of migration: from learning how to 'think Linux', and a crash course in console basics, to finally installing and using software on a day-to-day basis. There's also an excellent introduction to *Webmin*. www-106.ibm.com/developerworks/linux/library/l-roadmap.html

Jono Bacon

The founder of UK Linux, KDE developer and all-round nice guy, Jono is also a musician who's tunes have been featured on Slashdot.



OPINION

Cult hero or big hitter?

“ I remember when Linux really started making an impact on the world. Around this period we were starting to see commercial adoption, software availability, and of course, the birth of *Linux Format*. Back then, I remember reading a comment from a member of an IRC chatroom. This individual was expressing his view that Linux was losing its edge and cult status, and was becoming "too commercial."

Many people are cautious of popularity. A possible reason for this trepidation is that popularity can often be perceived as 'selling out' and low quality burgeoning as focus turns to profit. You only have to look at the amount of rubbish cluttering up the music charts to see an element of reality in this view. The issue here, however, is that these typical references to popularity are usually constructs of the commercial world. The charts are after all, filled with products of huge record companies.

Linux is of course different. The Linux community is *not* run by commercial vendors, it is simply supplemented by them. Red Hat, IBM and HP have no more weight in how our kernel gets developed than a bleary-eyed hacker slumped over a computer in his or her bedroom. Surely there is no possibility of selling out? Surely the concept of popularity with regards to Linux is one that can only do our community good?

Linux can only get better with popularity. Linux is here as a result of us working together, and an elitist culture would be detrimental to our progression. I for one welcome the increasing adoption of our favourite OS and the future looks bright. ”

WHEN A FACT IS NOT A FACT

Face the facts says MS

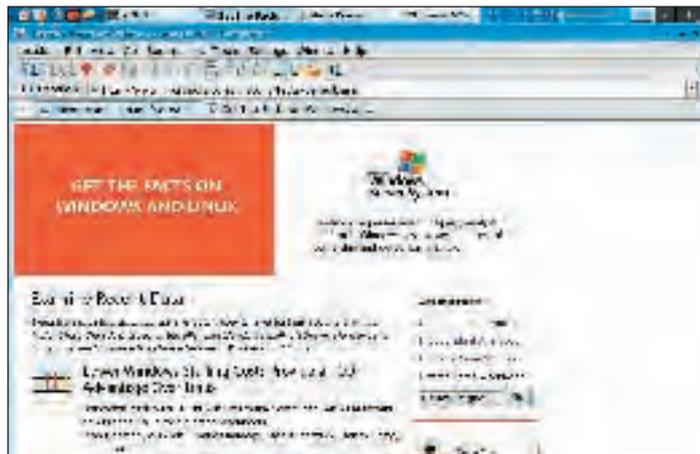
Linux is more expensive than Windows in almost every situation, Microsoft claims. A new advertising campaign backed by a website called *Get the Facts on Windows and Linux* aims to counter the argument that Linux is free, highlighting differences in 'price' and 'Total Cost of Ownership' in enterprise use, and the support savings potential that the ubiquity of the MSCE programme can bring.

Some tech pundits have suggested the campaign may backfire, by actually raising the profile of Linux as a genuine alternative to other systems. In terms of the Gandhi principle of competition – First they ignore you, then they laugh at you, then they fight you, then you win – Microsoft's decision to launch such a campaign would perhaps suggest that the 'fight' stage is now well underway.

The site is bulging with supposedly authoritative 'independent' reports and

case studies that prove the economic sense of ditching Open Source solutions – including Apache and Linux – in favour of Windows Server 2003, .NET and *Internet Information Services (IIS)*. However, a quick scan reveals that almost all of these independent efforts were in fact sponsored, supported or commissioned by one organisation. Which particular company might that be? Er, Microsoft.

In one study, an analyst contrasts the cost of a Windows-based database system – Server 2000 (\$4,000) plus SQL Server (\$9,000) – with that of a Linux system featuring Red Hat Enterprise Server (\$3,500) and Oracle DB (\$15,000), while completely ignoring the more cost-effective solutions which are one of the key features of Open Source software. How different the numbers would have been if Metagroup had used the initial costs of setting up a Debian/MySQL combination, which



An example of the post-modern definition of facts? Sadly, poorly specified analysis only contribute to greater cynicism about such reports.

could be acquired for anything between the cost of a few gigabytes of bandwidth and some blank CDs for an unsupported setup, and a fully supported installation costing in the region of \$4,000.

While no-one would impugn the integrity of the report compilers – which include some of the world's largest IT analyst group – the whole feeling that the site engenders is: "Well, they would say that wouldn't they?"

Linux Web Watch/

Kernel 2.6: when it's time to do it, it's time to do it!

Developers extend a helping hand to unsteady Linux-lovers

With Linux distribution developers concentrating hard on preparing their next-generation releases and users everywhere bracing themselves for the

big move, a number of websites have sprung up to take the terror out of upgrading for all levels of Linux user.

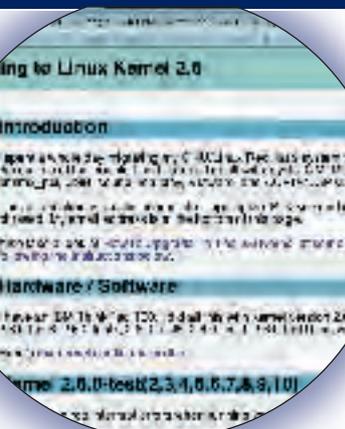
A developer-centric document on the process of porting device drivers from 2.4 has been posted by sometime LXF scribe William von Hagen at the TimeSys 2.6 Linux Resource Centre (www.timesys.com/index.cfm?body=26program.cfm). Von Hagen's strategy is to create a five-part migration plan, and make it available to registered users of the site.

If you're more interested in the possibilities of the new kernel (or

« Long-time readers will remember Bill Von Hagen's excellent articles in *Linux Format* on using Linux with Apple Macs.

you want something to read while FC2/Mandrake 10 is downloading) Joseph Pranevich has compiled a fairly readable overview of the new features at www.kniggit.net/wwol26.html. The document has also been translated into nine other languages.

If your geek supremacy have an overwhelming need to be more cutting edge than the rest of the world, the Linux Kernel Archives (www.kernel.org/) always has the scoop on what's coming and when it's likely to be released. It's always a good place to hang out if you want to name-drop ("...and as I said to Linus last night"). However, it is very technical and often very boring following the mailing lists.



» Sharing your experiences is often a better route to success than wading through reams of technical documents.

If you want some reassurance and handholding, visit <http://thomer.com/linux/migrate-to-2.6.html> to relive the experience of one user's successful attempts at upgrading Red Hat 9. Thomer Gil hit a number of snags, and the workarounds that he shares may make your life easier.

NEWSBYTES

■ Educational software and services company **Epcorn** has announced the availability of its web-based eLearning module for SUSE 9.0. The module makes user personalisation and results tracking much easier across a variety of operating systems. See www.epcorn.com/elearn for more information.



■ **Intel**, in an attempt to make up for lost time, is ramping up its efforts to support Linux, initially in efforts to bring native Centrino drivers to the OS. Intel is widely perceived to have missed the Linux boat and underestimated the potential of the market, something a spokesman said "would not happen again."

■ **Microsoft** has applied for and received a patent for the "systems, methods and data structures for encompassing scripts written in one or more scripting languages in one file." The patent for 'XML script automation' was applied for in December 2000. One developer said the patent simply stated a process that has been going on in the industry for years (with regards to, for instance, Javascript) and that the patent should not have been granted. Some experts fear that, while MS is unable to patent XML technology itself, it may intend to patent so many ways of using it, that real-world use becomes all but impossible without infringement.

■ As well as announcing a septet of major new customers, **Oracle India** has revealed that Linux has grown to become the second most popular platform among its developers. A spokesman for the country's central bank, which has successfully migrated its IT systems to Linux, said the move would mean savings in the region of 30-40 per cent, representing at least US\$4 million so far. One of the key factors, said Oracle India's MD Shekhar Dasgupta, was the one-stop support the company was able to offer customers. "We provide technical support for Linux at no additional charge to customers," he said.

■ **Passport software** has announced the availability on Linux of its flagship RealWorld Accounts package. RW32 version 10 is built on the foundation of the company's RealWorld Classic Accounting Software and features a complete range of accounting tools.



CrossOver Office already supports MS Office and Lotus Notes. Could it eventually become the basis of an IBM office solution?



COMPATIBILITY

MS Office for Linux?

No sooner had OSDL CEO Stuart Cohen suggested that "Microsoft will port some of its applications to Linux" as the platform grows, IBM says its long-term desktop Linux strategy involves bringing *MS Office* to Linux. Despite a few rumours, Microsoft remains adamant it is not involved in the project, countering with the claim that Big Blue may be planning on some kind of emulation technology.

Stefan Pettersson, an IBM Sweden spokesman told infoworld.com that the company had embarked on a long-term desktop strategy that would first involve the release of a Java-based *Lotus Notes* client later this year. Going forward, he said that *MS Office* would become increasingly important for desktop Linux.

"It will be possible to run the *MS Office* package on the Linux platform. Exactly how it's done, I can't reveal right now. But we're working together with Microsoft, who have provided us with part of their code." This working method, Pettersson said, was similar to

the efforts made on the creation of an *Outlook* client for IBM's Domino Server.

Cohen, the head of the organisation that is quickly establishing itself as a hub of enterprise Linux activity, said that there would come a point when it would make business sense for MS support a rival OS, as it has in the past with the Mac. His comments come as – according to International Data Corporation (IDC) – the Linux desktop marketshare had risen to 3.2 percent (up from 2.8 per cent last year) overtaking the Macintosh for the first time. IDC predicts that it will achieve a desktop penetration of six per cent by the middle of 2007.

A further study by Merrill Lynch suggested that 43 per cent of CIOs would consider a switch from Windows to Linux. Merrill Lynch analyst Steven Milunovich said the figure was something of a surprise.

"I had expected governments to be interested, but now Linux is on the radar of corporate chief information officers," he said.

Hoyt Duff

The co-author of *Red Hat Linux 9 Unleashed* runs a fishing pier when he's not shouting lovingly at his computers.



OPINION

MyDoom:
author found

“ We may have found the MyDoom author! It could very well be Stephen Evans, BBC North America Business Correspondent. Supposedly computer-savvy, yet disliking Linux, could he have engineered the virus to thwart the rising popularity of the Open Source movement?

In a shabbily written and poorly edited piece for the BBC website, Mr. Evans uses similar innuendo to link the infamous MyDoom virus to a zealous Linux community. This pitiful puffery is a sad example of the current state of the Linux community. By now several things should be clear:

- Microsoft mail clients are excellent and efficient transmitters of malicious code.
- The email infrastructure is just as bad, with headers easily forged and mail relays easily co-opted.
- The shameful SCO business has gotten out of hand with endless opinions and speculation substituting for rational thought and the due process of law.
- The anonymity of the Internet makes zealotry a sport without penalty; such mindless zealotry is detrimental to the widespread adoption of Linux.

Let's finish development of a Free and Open drop-in replacement to *MS Exchange Server* (not just a work-alike), the final 'killer app' of the Linux movement. Let's create a better email system to deal with open relays and spam, unclogging our mailboxes and our bandwidth. Let's let the legal system deal with SCO; it has dug its own grave by now. Finally, let's kick the zealotry back several notches and find a more positive way to advocate the advantages of Linux and FOSS. Let's get a better life. ”

LINUX MARKETING

Lindows: dashing about

After winning a provisional victory in the US courts, in which the term 'Windows' was judged to be a generic term and, thus, not suitable for use as a trademark, Lindows boss Michael Robertson has used his marketing and publicity genius to create a new Linux operating system: Lin---s (pronounced Lin-Dash).

Lin---s is a version of the Debian-based desktop distribution aimed at Dutch resellers and users who have been prevented from using the Lindows name due to its apparent 'similarity' to Windows. The fact that Microsoft does not actually have a product called 'Windows' seems to be lost on the Dutch judiciary.

As ever, the new PR stunt comes complete with the regulation welter of new domain names (www.lindash.nl, www.lin---s.com and a less predictable www.choicepc.com), and also a retooled Dutch edition of LindowsOS with all references to the Lindows name replaced with the new 'Hangman' Lin---s logo as shown on the right.

In typically bullish mood, Robertson said effusively, "Dutch citizens deserve the same choices that are currently available to the citizens of more than one hundred countries around the world. LinDash ensures that the Netherlands will have affordable, virus-free options instead of just expensive Microsoft software."



It can be dangerous to be right when the law is wrong...

Embedded Linux News

- Michael O'Brien, author of the *GoAhead* web server, has launched a new company dedicated to developing an embedded web server with the security features vital to success in today's "more hostile" world. The first application from the company, **MBedThis**, is *Appweb*, an Open Source web server O'Brien has characterised as a mini-*Apache* built on modular principles. With a minimal footprint of just 110K and 400K for a full server, O'Brien says *Appweb* will outperform *Apache* on standard hardware, but that the real benefits will be felt on embedded devices where such exceptional performance translates into low power consumption. www.mbedthis.com/

- Linux could finally begin to make a mark on the music business with the launch of the very exciting **Mediastation X-76** from Lionstracs. This portable digital keyboard/studio packs everything the modern musician needs into a sexy instrument bursting with ports, screens and knobs. Available with either AMD



Athlon or 64-bit Opteron processors and an enormous hard disk, the keyboard is capable of streaming samples directly from disk, which means multi-gigabyte instrument samples. Very exciting. www.lionstracs.com

- Sharp** has announced another addition to the Zaurus family. The SL-C860 is based on the same core design as SL-C750/760 but features a faster processor, more user accessible memory and, if rumours prove to be correct, English to Japanese and Japanese-to-English translation software. Early examples should be available soon.

- After divesting itself of the **Symbian** PDA/smartphone operating system (sold to Nokia), Psion is rumoured to be expanding its Netbook product line with Linux as the focus. "We see Linux as being very interesting, not only in terms of technology, but also in market dynamics; lots of companies want to move in that sort of area when they buy equipment these days," the company said.

DISTRIBUTED COMPUTING

Man vs machine ends in a draw... for now!



The makers of Chessbrain, a nascent distributed chess application, have attempted to get their project into the *Guinness Book of Records* by becoming the first distributed computer to play chess against a human player.

Grandmaster Peter Heine Nielsen faced the might of 2070 networked computers across the world running a variety of operating systems and managed, after 34 moves, to hold *Chessbrain* to a draw. Heine Nielsen said he'd set a number of traps during the game that computers usually fall for, but was surprised when *Chessbrain* didn't fall for them.

Despite various problems occurring during the setup phase, and some difficulty accommodating everyone who wanted to take part, the game itself went smoothly.

Teams linked into the network from 64 countries, and while the US took the top position in the throughput league, closely followed by the UK, Denmark managed to supply the most discrete machines thanks to use of a 250-processor cluster and the fact that the event took place there.

The attempt was accepted by the Guinness organisers and should appear in the next update of the book.



By Pamela Jones, www.groklaw.net/

■ In response to IBM's success in compelling SCO to reveal the 'stolen code,' the company has detailed approximately 3,700 lines of source code allegedly dumped from AIX and Dynix into Linux. As expected, none of the code constitutes "line-by-line code copying" that has been "copied and pasted from Unix into Linux." (Darl McBride, 16/06/2003) This number is also far short of the "millions of lines of code and not a few dozen" highlighted by McBride on 22 October 2003.

■ The big news this month was Novell suiting up for battle and entering the fray. First, it released correspondence with SCO (www.novell.com/licensing/indemnity/legal.htm) disputing SCO's copyright claims. The letter also claims authority to overrule SCO's "termination" of IBM's AIX and Sequent's Dynix licenses. Novell also registered copyrights on Unix, the same code SCO registered a copyright for in July, which means the ownership issue must be decided by a judge. Then, for the cherry on top, Novell put up on its website copies of two AT&T newsletters from 1985, one of which showed that AT&T sent a clarification to all licensees

stating that it made no claim to derivative works, the equivalent to a bomb as far as SCO's claims to IBM's AIX as a derivative work.

SCO promptly sued Novell in Utah State court for slander of title. Novell moved the case to federal court, the same court where the IBM case is being heard, and then promptly filed a motion to dismiss on the grounds that SCO failed to adequately plead two elements necessary in any slander of title action. If, like most of the human race, you haven't a clue what slander of title is, Groklaw has gathered some information that explains it.

(www.groklaw.net/article.php?story=2004021116125699)

■ The other big news as we go to press is that SCO has dropped its trade secret claim against IBM, and is asking the judge to allow it to amend its complaint to add a copyright claim instead.

However, IBM says the copyright allegations relate not to System V and not to any line-by-line copying, but rather to SCO's "legal theory" of what constitutes a derivative work. Under this theory, all of AIX is a derivative work – even the parts that were not written by SCO, but does not include any System V code, which is copyrighted to IBM.

At a hearing on February 6th, IBM told the judge that SCO had failed to provide full discovery, as required in her order. Groklaw has a transcript of the complete hearing. (www.groklaw.net/article.php?story=20040209231214944) SCO then again asked the judge to compel IBM to turn over all versions of AIX. The judge, in an unusual move, said she would take both matters under advisement and would render a judgement in about a week. This has yet to happen, which may have something to do with the fact that SCO turned over a pile of supplemental answers and documents to IBM on the eve of the hearing. Significantly, she asked IBM what relief they would like if she were to find that SCO had failed to obey her order, and IBM did not take the opportunity to ask for a summary judgement at this time, which some took as an indication that IBM wishes the matter to be decisively settled on the merits once and for all.

■ Red Hat also filed a Motion in Delaware court, apparently seeking to bring the judge there up-to-date with new information. Considering how long the judge in that case has been waiting to rule on SCO's motion to dismiss, no doubt Red Hat had plenty to tell her.

David Cartwright

An IT consultant since the phrase was respectable, David specialises in Linux systems and solutions.



OPINION

Micr[Open Source]oft?

“ Well, deary me – Microsoft has had a bad time of it this last few weeks, haven't they? First the MyDoom virus, next the supposedly critical security holes, and then finally the leaking of some of its source code onto the Internet. Is it just me, or has the world gone bonkers about this source code leak? Everywhere I turn, I find people in the media (both the IT and mainstream varieties) screaming about what a disaster it is to have your source code on the Internet for all to see (and, they point out, for all to probe for bugs).

Has nobody (except us Linux folk, of course) spotted that making one's source code accessible can actually be a *good* thing? After all, the Open Source movement spend most of their days pointing out that the only way to be sure something's secure is to be able to examine its source code. So how can it be good for Linux/BSD *et al*'s Open Source code to be accessible to the masses, but bad for Microsoft code to be available in the same way?

I actually think the leak will do Microsoft some good, because despite all the bashing it gets, it does actually have some rather good technologists over there in Redmond. I reckon that there's every chance that, once the world and his dog has examined this leaked code, no significant bugs will have been found, and instead of Microsoft being shown up yet again for its lousy security, the opposite may well be the case.

I can see the headlines now: "Microsoft accused of deliberate leaking of source code in shock adoption of Open Source verification techniques"... ”

MORE LINDOWS SUCCESS

Visual web development – WYSIWYG

The project initiated by Lindows boss Michael Robertson to bring usable visual web page development to Linux has borne its first fruit. *Nvu* version 0.1 is built on the *Mozilla Composer* codebase and counts former *Composer* luminary Daniel Glazner as lead developer and maintainer. Fitting in with the new *Mozilla* focus on discrete applications, *Nvu* provides complete what you see is what you get (WYSIWYG) page editing, a tabbed interface and, for the first time, a (sometimes) usable FTP pane, which makes working with remote servers a lot easier than the previous versions of *Composer* and appears to be modelled on the similar tool within *Dreamweaver* and *Adobe GoLive*.

The current version is available in both binary, source and Click-'N'-Run form, though *LXF* had some difficulty



Nvu is built on the often-neglected *Mozilla Composer*.

installing on a non-Debian machine. However, if you really must try *Nvu* out, visit www.nvu.com/landing_page.html for details on how to acquire a free edition of LindowsOS

Developers' Edition. Downloaders can get it for nothing, but if you're on dial-up you might want to consider the US\$10 (plus shipping) CD version.

A Windows version is also available.

Mailserver

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

★ Letter of the month

This month's winner receives a copy of **Dave Baum's *Definitive Guide to LEGO MINDSTORMS™***

Beagle 2 and Linux

I have been intrigued to read your little pieces in both your *LXF50* and *LXF51* about the role that was played in our ill-fated space mission by our favourite Open Source operating system.

Your *LXF50* piece mentioned the GNU/Linux workstations at the National Space Centre at Leicester (the Beagle 2 Lander Operations Control Centre), so I thought I would just get a mention for the role played by Linux systems at the Lander Operations Planning Centre at the Open University, where we would have planned the science operations, and worked on the returned data, had we got any.

We have two DNUK workstations running the Space Craft Operating System (SCOS) on top of the "ESA preferred" distro – which is SUSE – as well as a DNUK/SUSE server as our 'gateway', which would have allowed the Science teams to collect their data and update the central archive; and of course our file server in the Control Centre was also a SUSE box, with *Samba* connectivity for all those pesky Windows workstations that we were required to supply. We have nothing but good things to say about SUSE, and the kit all worked well. Interested readers can still view our information online at www.beagle2.com/ and www.open.ac.uk/science

I'd like to take this opportunity to offer many thanks to Duncan Carter, resident über-geek at the OU and to Dr Ed Chester (who is

also very über, but probably would not like to be labelled a geek!) and all the Flight Operations boys at Leicester; Stewart Hall of SciSys (godlike scripter), and the NASA team at JPL, serious "hackers" all...

Jed Cawthorne, *AMBCS Ground Systems IT Engineer & Data Archiving Team Leader Beagle 2 Lander Operations Planning Center, The Open University*

Many thanks for the further detail on Linux's involvement with the Beagle Project. It's unfortunate that it never

got a proper chance to be used, but we certainly hope that Beagle 2 will not be the last in the series. Irrespective of the outcome of the mission, I'm quite sure that the actual design and construction of the lander in the first place was very beneficial in many ways, and that plenty has been learnt.

I hope you'll accept our slightly tongue-in-cheek star letter prize for this month, the entertaining and informative work on *Lego Mindstorms*, as published by Apress.



Current theories as to what actually happened to Beagle 2 include airbag failure and theft by curious Martian hackers interested in Linux...

Printer types

You state in your Hardware feature in *LXF51* that there are two types of printer. My recollection – a bit hazy after 20 years – is that in the 1970s and 80s there were *three* main sorts of printer.

The simplest was the IBM 'golfball', incorporated in the Diabolo printer. The ball spun and dipped to create a single character at a time.

Then there was the 'daisywheel' or 'starwheel', which had all the characters around the rim but in the plane of the disk. It rotated about its axis and, again, printed one character at a time.



Quickest – and most expensive of all, by a long way – was the 'lineprinter' which had around 132 (?) wheels with all the characters around the edge of the rim, facing out. All the wheels rotated and stopped momentarily at the required character. When they were all in the right place, there was a bang and the whole line was printed at once – hence the name. Eventually, printing was delegated from IT centres to small groups of users who acquired cheap dot-matrix, and later laser and bubble-jet printers.

I haven't seen a line printer myself for decades, though I expect they do still exist; but card-readers of the same age are no more than museum pieces, along with teletypes which were the original

universal input/output machines. I remember seeing a big Creed round about 1963 that projected pink punched tape (memory says about 20 feet) into a reserved playpen, just like a baby projectile vomiting: a striking sight when in full flow. 'Fast' magnetic memories at that time were rotating drums that had displaced mercury delay lines and CRTs. Those were the days when one programmed in Mercury autocode and CHLF3, and drove the computer from its console oneself.

RS Clymo, *Robertsbridge*

There were indeed many types of printers in years gone by. I had a particular fondness for the daisywheel in the old computer room at the University of Surrey. Being serial-driven from the main server, it would speed up or slow down depending on the load on the server; and if you were really bored, you could write scripts to vary the processor load and get it to hum a little tune. I've never seen a Creed, but it sounds fantastic.

Where is the LUG?

I've been reading your mag now for about a year and must have dropped enough hints because was given a subscription for Christmas! I was very much a newbie 18 months ago, but now use Linux as the main operating system on my box. I've tried many distros but have settled for the moment on Mandrake 9.2.

The one piece of advice I never see published is to get involved in a LUG. Doing so has been invaluable, not only in finding different solutions to the problems I've had, but also making new like-minded friends.

One thing I still have problems with is DVD video: until support gets better I'm having to use Windows applications through Wine.

Arden King, *York*

We list all the existing UK LUGs at the back of every issue of LXF, to make it straightforward to get in touch with a group in your area. More information about LUGs, including suggestions about starting your own can be had from www.lug.org.uk/. There are quite a few new LUGs that have started up in the first few months of 2004 shown on our map – check it out to see if there's one near you.

As for DVD and Linux – your wish is our command: see this month's cover feature: there are plenty of video resources for both viewing pre-

READER TIPS

UNFAIR IPCOP!

Having used *IPCop* as my firewall at home for a while, I thought that your *Roundup* review in LXF51 was slightly unfair, to say the least.

IPCop as a main version may not have been updated recently, but smaller fixes are available in short order once a problem appears. There are 8 fixes currently available for 1.3.0 at present, fixing amongst other things the kernel exploits you said were not fixed. These fixes can be easily downloaded and applied by users, often without even requiring a reboot.

In addition *IPCop* does support a caching proxy, in the form of *Squid*. This is operable from the interface as is the VPN support. I currently run *IPCop* on an AST 486DX4/66 with 24MB of RAM and 200MB of hard disc space, and it provides great service for a home network, supporting four machines over a 600Kbps broadband connection.

I am also looking forward to the release of *IPCop1.4.0* which will support amongst other things four network ports (RED – Internet, GREEN – internal



IPCop has limited capabilities, but as a fork of *Smoothwall*, it has a web interface for management.

network, ORANGE – DMZ and BLUE – wireless).

Carl Scarlett, *via email*

FEDORA FETTLING

In Fedora Core 1 (on the LXF49 coverdiscs) the menu item "System-settings Add/Remove Applications" appears to be broken! After reading packages, it will ask for fedora-core disk x, but when disk is inserted it will exit with – "ERROR installing packages - OK."

SOLUTION

To update / change packages
 1 Insert installation disk 1 (or cover disc 1) on the installed desktop
 2 At the prompt "Auto Run



Packet filtering rules are easily added through a simple form.

Yes/No" select YES; you will now be in the *Package Management* dialog make your selections by ticking the boxes as normal.

3 Click Update – the installer will start updating the system, then ask for fedora-core disk (1,2 or 3). With disks 2 and 3, just insert and click OK. If at any time disk 1 is asked for, click NO on the auto run prompt: if you click YES, you will run another package manager! Also, the apt package is apt-0.5.15cnc3-0.1.fri386.rpm

Tony Windebank, *via email*

LINUX PCS

You may already be aware of this, but I've not noticed it appear anywhere in your magazine: Scan computers (www.scan.co.uk) is offering its custom 3XS systems with a choice of MS Windows, Libranet GNU/Linux, or nothing, and even more shockingly, they're offering *OpenOffice.org1.1* as the default software package! Obviously, I can't vouch for the quality of the systems, since I don't have one, but it's certainly nice to see.

On another matter, I wrote in a while ago about my USB mouse not working with *Lilo*. It turns out that I had to tweak my bios settings and enable "usb functionality for DOS". Odd!

Jim Rutherford, *via email*



Visit www.scan.co.uk to buy a PC with Linux or without an OS.

recorded discs and making your own DVDs of home movies or picture archives to view on your television, and all the necessary applications are on the coverdisc too.

Loving the legacy

I'll make a small bet that Biagio Lucini (*What on Earth is Fortran, LXF50*) first learned to program in Fortran. Well, I learned Algol 60

first and several of those improvements he quotes for Fortran 90 are actually just the 'blasted Yanks' finally adopting features of a proper programming



MAILSERVER

« language. Incidentally, have you ever wondered why Ada, Pascal and C look vaguely related? You got it: all descended from the first really well thought out high-level programming language; whereas Fortran has no descendants.

As for being ideal for numerical applications: has Biagio ever had to cope with a dynamic, that is runtime defined, two dimensional data array? In Algol 60 (as in Basic actually) it's a doddle, whereas in C it's very slightly tricky; then try it in Pascal – it's impossible (apparently deliberately!) and in Fortran – ugh! That's in 'IV of course, although I gather that '77 is no better.

languages then in use. Fortran was designed with one eye on implementation, which caused some terrible decisions but fast binaries; Algol was designed to the requirements of numerical analysts and to hell with how to implement them, which made it a delight for such work. (And the compiler writers rose to the challenge.)

So why did Fortran beat Algol? The Algol Committee made two fundamental mistakes: firstly, they completely ignored the input/output and file access side because that wasn't relevant to publishing algorithms (which killed the portability); and secondly, they



Whenever LXF mentions legacy code, we get impassioned letters...

“We tend to take things that already exist for granted, but just think about what it would be like if our Linux boxes lacked Miguel de Icaza’s tools...”

The differences go back to their origins. Fortran was invented to make programming numerical problems easier than in *Assembler*; Algol was invented to make describing numerical algorithms for publication clearer than the mess of

weren't IBM, so America simply ignored Algol except for later language design.

There's one last wrinkle to this tale: the C language is derived almost directly from Algol (via BCPL) and retains Algol's structure

of declaring procedures first, followed by the main body. In spite of this, C is invariably taught as if it is Fortran, which requires the main body to come first followed by subroutines. Very messy.

Funny lot, the Yanks!

Thomas Edward Groves, *via email*

I remember Algol68 when I went to University. I learned Algol68, Fortran, Lisp, COBOL and a whole host of other languages that I've never used since – in fact, not many people were using them at the time, either! It's interesting to look back, and also to wonder what languages will be used

20 years from now. C certainly seems to have solved the portability issues, and is an economical and elegant language, but who knows?

Hit for a six

Hello from a regular GNU/Linux user and *Linux Format* reader from Southern Europe.

Considering your *LXF50* article *Biggest Hitters*, I wouldn't object at all that Linus Torvalds, Richard Stallman and Mattias Ettrich are (among others) personalities that changed the world for our benefit, and your comments about them are on target.

Helpdex

shane_collinge@yahoo.com



What annoys me a little though is this: The co-founder of the GNU Network Object Model Environment (what a name anyway?!) GNOME Desktop, programmer of *The GNU Image Manipulation Program (The GIMP)*, *Midnight Commander*; and nowadays a developer of the Mono Project – a free implementation of the .NET platform for Unix-Compatible OSes – is listed only in the honourable mentions box at the end of the piece.

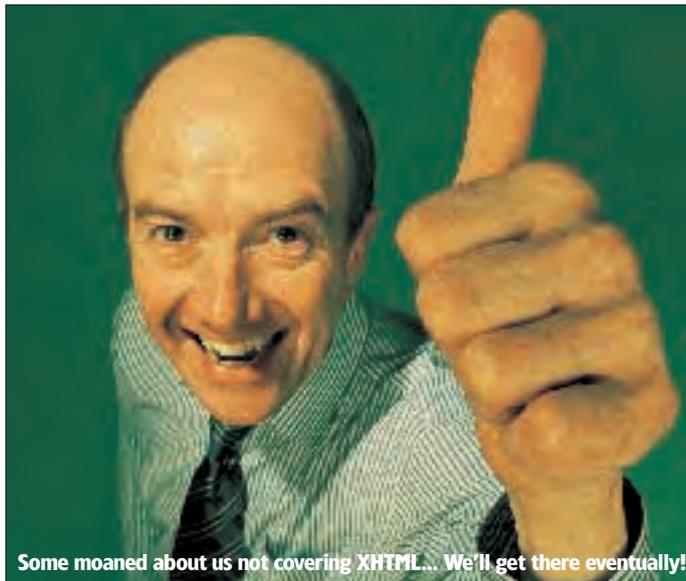
Miguel de Icaza gave us more than we would normally expect. Just think about it: a classic file manager with the same feeling we old DOS command-liners needed for a soft transition to the console, the only serious Free image editor that exists for almost every platform (including Windows) and an astounding Desktop based on a toolkit made from scratch!

So, when I need a quick file management tool or perform some quick image editing, I know exactly what to use. And when I want to make any Wintel users jealous, guess what desktop I show to them besides KDE? That's right...

We tend to take things that already exist for granted, but just think about what it would be like if our Linux boxes lacked these tools. My personal taste has nothing to do with my opinion; in fact, I am a big fan of KDE as a Desktop, and *Blackbox* as a Window Manager, and would gladly accept *The GIMP* could boast a slightly more practical workspace. But at the same time I believe that Miguel de Icaza is – apart from being an excellent programmer – one of the *Biggest Hitters*; thus deserves a paragraph of his own and a photo, too.

George, *via email*

As you might imagine for such an emotive subject, we've had quite a few emails on our selection: from well-reasoned ones like George's here, to some ridiculous rants questioning the parenthood of the LXF staff; and even a couple that try supporting their points of view with quotes from various flavours of religious and radical political texts. On reflection, George is probably right: Miguel de Icaza *does* deserve to be part of *LXF's 10 Biggest Hitters*, but who would we drop from the existing list to include his entry? You have to draw the line somewhere, and we had already cheated with the



Some moaned about us not covering XHTML... We'll get there eventually!

"Women in Linux" entry, so we could include more than 10. Perhaps more valid is the suggestion that, for example Bruce Perens, Stallman *et al* are merely celebrities these days, and the real influencers of Linux are elsewhere. At least the feature served it's purpose – to introduce some people you may not know and get you thinking about the rest!

If you want to know more about Miguel de Icaza's impressive input to many crucial areas of Linux, visit <http://primates.ximian.com/~miguel/>

HTM-well pleased!

I just want to say a big THANK YOU to Jono Bacon for his excellent series of articles on HTML/CSS. I have rebuilt our community website: www.balneario.co.uk in CSS and it looks great. As always when learning a new subject, you have some problems, I found myself leaving out the pixels (px) after each size attribute, and wondering why it wasn't working!

I found *Quanta Plus* very easy to use: I had been using a Windows-based HTML editor for years, and prefer hand coding, as you get exactly what you want. The CSS dialogue window in *Quanta*, is really easy to use, and I found that you could produce a style sheet in a matter of minutes. One tip with *Quanta* is to use the Project facility: it seems to start up by looking for a project to load. So, start off with a new project, and it will quickly load the last used project for you.

Thanks for an excellent magazine, I am patiently waiting for

the next distro release with the full 2.6 kernel, it sounds very promising. Derek Smith, *via email*

Jono Bacon writes: What a month – first someone on the forum suggests that I should get a pay rise (I didn't know your mother was registered on the forum – Ed), and now this! Thank you for your kind comments – it is always great to get positive feedback, and I agree that much of this positivity needs to be levelled towards *Quanta*; it really is a great tool.

HTML and CSS is such an involved topic that we unfortunately could not explore *Quanta* further due to space limitations in that particular tutorial; *Quanta* is something we are considering for more in-depth examination in a future issue. Linux is a truly viable web development system, with a stack of tools to help you get the most out of your web presence. Good luck with your site and keep us updated. [LXF](#)

SUBMISSION ADVICE

WHAT WE WANT:

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

WHAT WE DON'T WANT:

- Technical questions – direct those to our *Answers* 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

www.linuxformat.co.uk

PROTEUS COMPUTER SYSTEMS

www.proteus-systems.co.uk



Gentoo Linux Systems

Proteus Systems from £865.00
Apple iBook from £1039.00
Apple P/Book from £1289.00

We donate £20 to the Gentoo team for each Linux System sold configured with Gentoo Linux



Modems

ReWAN ADSL Modems
PCI Version £56.00
USB Version £56.00
MultiTech 56K Modem £90.00

Network Cards

Belkin F5D500 10/100 £9.54
Cometa 10/100 £8.99

Books

Linux in a Nutshell £28.50
Using Samba £28.50

Linux Based Distributions

SUSE Linux 9.0 £20.00
SUSE Linux 9.1 £25.00
SUSE Linux 9.0 £24.00
SUSE Linux 9.0 AMD64 £29.00

Linux/BSD CD Sets

Gentoo 1.4 £4.00
Debian GNU/Linux 3.0r2 £14.00
SUSE Linux 9.1 GPL £8.00
FreeBSD 5.2 £4.00
Mandrake 9.2 GPL £6.00

More Computer Systems,
Software, Components &
Services available via
our Web Site!!

www.proteus-systems.co.uk

Tel: (01332) 518559

Fax: (01332) 510349



PROTEUS SYSTEMS LTD. Registered in England No. 2721188

Reviews

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

LXF VERDICT EXPLAINED

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

Features: Does it provide the functions you need? Is it innovative?
Performance: How well does it do its job? Is it fast and reliable?
Ease-of-use: Is the interface well designed? Is the documentation well written, helpful?

Value for money/Documentation: Whichever is most appropriate!

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



10 The close-to-perfect product



8-9 Good, but has a few niggles



6-7 Does the job, but needs work



4-5 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...

Boston Superserver 6113

Will adding SCSI drives to servers powered by Itaniums bring them up to speed? **p20**

Systemax Mission 3602 server >>

Big performance, titchy price – one to add to the top of your wish-list **p22**

Linux Megapack

How to get a bunch of distros without overloading your local cable exchange **p24**

Intel C++ Compiler

Version 8 of the speedy system for X86 coding **p26**

Portland Group C/C++ compiler 5.1

More of a complete software development suite than just a compiler **p27**



KDevelop 3.0.1

IDE that proves the maxim “the best things in life are free...” **p28**

Arkeia 5.2

Network backup solutions need to be both simple to install and maintain, yet powerful enough to cope with preserving data from a range of environments **p30**

KOffice 1.3

Productivity improvements mean that *OpenOffice.org* should not be resting on its laurels **p32**

Book reviews

Fedora Desktop for Dummies Kit, *J2EE Developer's Handbook*, *Hacking the Tivo*, and the sublime *Joy of Tech* **p33**

LXFBENCH 2004 EXPLAINED

To comprehensively test the capabilities of machines we review, we have developed *LXFBench 2004*: a new benchmark suite designed to push hardware of all shapes and sizes to its limit.

The test is broken down into four distinct parts: multiprocessing, uniprocessing, RAM, and hard disk, of which the first two are largely similar. The multiprocessing test creates four child processes in order to take advantage of SMP hardware, then proceeds to run *oggenc* to encode a large audio file to Ogg format, uses the *GD* image library to resample a complex image several times, and also runs an external C program to calculate

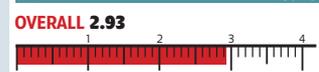
the hashes of random numbers. The *uniprocessing* test is identical except that it runs on just one CPU.

Both the *RAM* and *HD* tests use the *SQLite* database library to manipulate database information in RAM and on the hard disk respectively.

The overall score is an average of all four tests, and is presented as a bar graph for ease of reading. A score of 1 means that the machine has equalled our yardstick machine – a 1.8GHz Pentium 4 with 512MB RAM and an IDE hard disk. A score of 2, therefore, means that a machine has completed our tests twice as fast as the benchmark. The majority of the code was written using PHP 5, with the CPU-intensive tests written in C.

BENCHMARKS

CPU	6.07
SINGLE	3
RAM	2.17
HD	0.46



All our benchmarks, unless specifically noted otherwise, are run on a fresh installation of Red Hat Enterprise Linux 3 AS for the specific platform. All source code, including PHP itself, is compiled using *GCC* unless otherwise noted. The *mhash* library, created by Nikos Mavroyanopoulos and Sascha Schumann, is used for data hashing.



1U ITANIUM SERVER

Boston SuperServer 6113

The beleaguered Itanium fights on, but Paul Hudson asks: is it worth it?

BUYER INFO

1U rackmountable Itanium server with power SCSI drives. Also consider RX2600 reviewed in LXF48.

- **SUPPLIER** Boston
- **PRICE** £6,249 + VAT
- **CONTACT** 08707 515950
- **WEB** www.boston.co.uk/

With the second 64-bit chip from Intel hovering over the horizon, the Itanium chip is looking more and more niche. For ultra-high performance computing, Intel maintains that the chip is unbeatable, and that may well be so. However, few organisations have need to calculate nuclear explosions down to a molecular level, which makes us wonder quite how well the Itanium is suited to more general HPC tasks. Of course just the phrase 'general HPC tasks' is a bit of an oxymoron – a chip should really perform well irrespective of the tasks it is put to, whether that be reproducing explosions or just calculating Pi to the trillionth place.

Although Itanium was launched several years ago, it has yet to be declared a success by anyone but Intel. In the words of Scott McNealy of Sun, "That's a dog that just won't hunt". However, many manufacturers continue to back the architecture, not least of which are SuperMicro and Boston, which means there's life in the

old dog yet. This 1U server clocks in at £6249 before VAT, and for that you get two Itanium 1.4GHz chips, each with 1.5MB of L3 cache, four 73GB Ultra-320 SCSI drives, and 4GB of ECC PC2100 RAM. Yes, if you're sharp-eyed you will notice that this is largely the same machine as Systemax sent us in LXF50, and this is because Boston supplies Systemax with its Itanium servers. Having read our comments in the past about the hard drives, Boston switched this machine completely over to SCSI to give a disk

boost, and, of course, this issue we are able to put the machine adequately through its paces using the brand-new *LXFBench 2004*.

Boston was also kind enough to pre-install the Itanium build of Red Hat Enterprise Linux 3 AS, which would have been ideal if it weren't for the fact that it seemed to be some sort of pre-release of v3 that was incompatible with our CDs. As a result, we had to re-install the OS, but this was trouble-free as you might expect. RHEL has supported IA64 for some time, and

also happens to be the platform we use for our benchmarks, so it's good to see that Boston is shipping it as standard. Having said that, we keep nagging suppliers to send CDs with their servers, and only Armari seem to have woken up to this – is it really so hard?

A noise annoys!

The last time we looked at this machine it was incredibly loud, leading us to say, "the Armari Opteron server we reviewed back in LXF45 was extraordinarily loud, and, by some feat



It's only when you plug it in that you'll be able to tell the difference between this and the Systemax Mission 6507.



of mechanical engineering, this server is louder". That much has not changed with this machine – within seconds of turning it on, people across the other side of the building were looking over to see whether were testing a server or an industrial-strength vacuum cleaner. Boston did inform us that it is using the most expensive fans available to cool this machine, which is at least some comfort – it may be loud, but it won't ever overheat.

As with the previous model, there were no PS/2 ports around. We don't know about you, but we don't actually have a spare USB keyboard here at *Linux Format*, which always makes this a mite problematic. Luckily we are only a stone's throw from *PC Format*, who happily loaned us such a wondrous device. Again, Boston took the time to explain this to us, and it seems that none of the Itanium boards have PS/2 simply because it is considered an unnecessary legacy port. If Boston were to supply its units with a free USB keyboard and mouse set, we'd be happy to agree!

Making it work

The previous two Itanium servers we have reviewed here were both left unbenchmarked, because at the time we had no means to properly test such a high-end CPU. However, our new benchmark was specifically written with the Itanium in mind – it was even optimised by engineers at Intel – and, in order to really give the Itanium the best chance at performing, we also compiled the benchmark using Intel's own *C++ Compiler v8*. Compiling code

for the Itanium is a notoriously tricky task, and Intel's compiler does a much better job of it than *GCC*. Although this technically gives the Itanium quite an advantage – the Intel compiler will produce faster x86 executables than *GCC* also – we were surprised to find that it didn't really alter the outcome.

The first test in *LXFBench 2004* is multiprocessing, which runs three types of test. First, lots of number crunching is done using large floating-point calculations and also a variety of hash functions. Then, both *oggenc* and the *GD* image-processing library are used to recreate real-world scenarios. Combined, this is a very thorough test of compute power, and we expected the Itanium to perform particularly well. Sadly, it returned a score of just 2.6, which is particularly poor when you remember that the yardstick machine was not dual-processor. The single-processor score was 1.31, which is equally poor, and we have a feeling that if we had compiled the benchmark using Intel's *C++ Compiler* on the yardstick machine the difference would have been even less.

Although Intel is the first to cite the "MegaHertz myth" when discussing its Itanium systems, it is hard to ignore the fact that these CPUs are only clocked at 1.4GHz. Although there is a faster (and more expensive) 1.5GHz Itanium with 6MB of cache available, that's the fastest chip available at this time. As a result, even with the fastest chips available, a dual-248 Opteron is going to have a 1.4GHz lead over the dual Itanium, which is almost the equivalent of another chip. If you

adjust for the MegaHertz difference, the Opteron chip comes out about 50% faster clock-for-clock, which is still quite a lead. The Itanium has been at 1.5GHz for some time now, and subsequently we don't expect it to scale much further.

The *SQLite* database system is used for both our RAM and our hard disk tests, with the former running a variety of queries against a database stored in RAM, and the latter inserting several thousand rows into a disk-based table. This machine scored 1.47 and 0.5 for RAM and hard disk respectively, which is a mixed bag of results. A RAM score of 1.47 is quite respectable, and inline with the Itanium's performance in other tests, however RAM speed is one area where AMD continues to have dominance: DDR PC3200 RAM is commonplace on the Opteron.

The hard disk test result is quite disappointing, however, particularly because this system is powered by SCSI drives. There are three mitigations here that lessen the impact of this score. First, the machine was sent to us in a rush, so it is quite possible there are some hardware configuration issues that will be worked out when the machine hits the streets. Second, we installed the OS ourself, removing Boston's preconfigured copy of RHEL, which may have had optimised drivers installed for the RAID card. Finally, Itanium machines are not likely to be put to use as general servers, simply because their price is so high; so disk-based performance is not likely to be an issue. In HPC environments, data is mostly sent across the network and stored in RAM.

Unimpressive results

At a price of £6,249 before tax, this server will put a dent in anyone's budget whether or not they are Intel fans. Furthermore, Intel switching their Xeon chip over to the AMD64 architecture should send shivers down the spines of anyone who has already committed to IA64, as it means that Intel's chips are no longer the bellwether for HPC. The Inquirer, the reputable news source that it is, has taken to naming the Intel 64-bit extensions "iAMD64", and quite rightly – even Linus Torvalds said, "just because Intel didn't care about their customers and has been playing with some other 64-bit architecture that nobody wanted to use is no excuse for not giving credit to AMD for what they did with x86-64"

If you have £7000 burning a hole in your pocket, you could spend it on this server. Alternatively, you could use it to buy three of the Systemax Opteron servers reviewed on the following pages and never look back – no prizes for guessing which option we would choose!

Other options

With the newly announced Xeon 64 chip effectively forming the core of Intel's commodity 64-bit initiative, it is likely that the Itanium will be pushed even further upmarket, making it less likely to appear in devices like this, and more likely to be found in larger HPC solutions. If Itanium was the best all-round, there would be little need for the Xeon 64 at all.

An overall score of 1.34 puts this machine 33% faster than our yardstick, which, while nothing to be sniffed at, is not so impressive when you look at the price. If there were some redeeming feature to this system, believe us when we say that we would cite it and make the most of it. There is an argument out there that "people who want Itanium will buy Itanium", essentially saying that people are committed to a solution rather than to a performance point. This might have been the case before the dot.com crash, and is perhaps still the case in some operations. While in the long term the Itanium architecture may have more leg-room for better performance, the current technology doesn't seem to deliver the necessary 'bang per buck' at this end of the market. **LXF**

BENCHMARKS

CPU	2.6
SINGLE	1.31
RAM	1.47
HD	0.5

OVERALL 1.34



LINUX FORMAT VERDICT

FEATURES	7/10
PERFORMANCE	6/10
DOCUMENTATION	8/10
VALUE FOR MONEY	4/10

Underperforms and is expensive to boot: the SCSI drives don't do enough to rescue the benchmarks either, sadly.

RATING 6/10





1U OPTERON SERVER

Systemax Mission 3602



It's the anniversary of the Opteron this month, and Systemax is celebrating in style, says Paul Hudson...

BUYER INFO

1U Rackmounted Opteron server with surprisingly low price-tag. Also consider Armari reviewed in *LXF45*.

- **SUPPLIER** Systemax
- **PRICE** £2299+VAT
- **CONTACT** 08707 297644
- **WEB** www.systemaxpc.co.uk

Since AMD launched its AMD64 family with the Opteron last April, we have seen a steady flow of machines passing through *LXF Towers*. But do we get bored of them? Not at all – Opterons continue their conquest of the hearts and minds of Linux users both here and in the enterprise, simply because they are gosh-darned fast and reliable to boot. AMD's original goal of breaking out of its value market segment and into the high-margin server industry has been met, and AMD are now back in the black financially.

The last Systemax machine we reviewed was a 1U Itanium server in *LXF50* that aimed to get performance on a budget, but did so by chopping out everything but the IA64 CPUs. This time Systemax has got wise, and sent over an Opteron-based unit with two 248 Opterons (2.2GHz each), 1GB DDR266 RAM, 320GB IDE hard disks, two Gigabit Ethernet NICs, and

well-cooled by a slew of fans. What's more, this machine comes in at just £2299 before tax, compared to the £4999+VAT of the Systemax Itanium from *LXF50*, which shows that Systemax is competing aggressively against, er, itself. While sitting on the architectural fence might seem like an odd move to many, we think it makes sense – Systemax seems to be playing for both teams right now, and will probably continue to do so until one architecture wins conclusively.

This issue we also debut the new *LXFBench 2004* benchmarks detailed on the reviews intro page. It is our hope this will allow us to push this machine and others of its calibre right to the edge to get the most accurate performance figures possible.

Inside the box

Although we installed our own copy of Red Hat Enterprise Linux 3 AS for the tests, this machine came with a 32-bit version of Windows Server 2003 pre-installed. We still find it odd that 64-bit manufacturers would ship a 32-bit OS, but we had no problems at all when installing RHEL.

The unit itself is a slim 1U box with a heavy reliance on USB. Irritatingly enough, again there was no PS/2 port onboard, but this is becoming increasingly common in servers, so we

will have to ignore this. As is also becoming increasingly common in Opteron servers, the fan noise is actually quite low, which reflects well on both Systemax and AMD; and looking up the manufacturer's numbers it seems the Opteron is now a lot cooler than the Itanium. In the long-term, this is a big thing for AMD: if it can keep the Opteron cool enough, it will be very easy to utilise the chip in very large clusters and blades.

Still, there was little sight of much documentation in there, with the motherboard documentation being the extent of it. This is a very depressing situation, and is one that we really think we should start a campaign to correct. Yes, adding a copy of Red Hat Enterprise Linux to the box isn't cheap, but it should at least be optional so that customers can deploy Linux more easily.

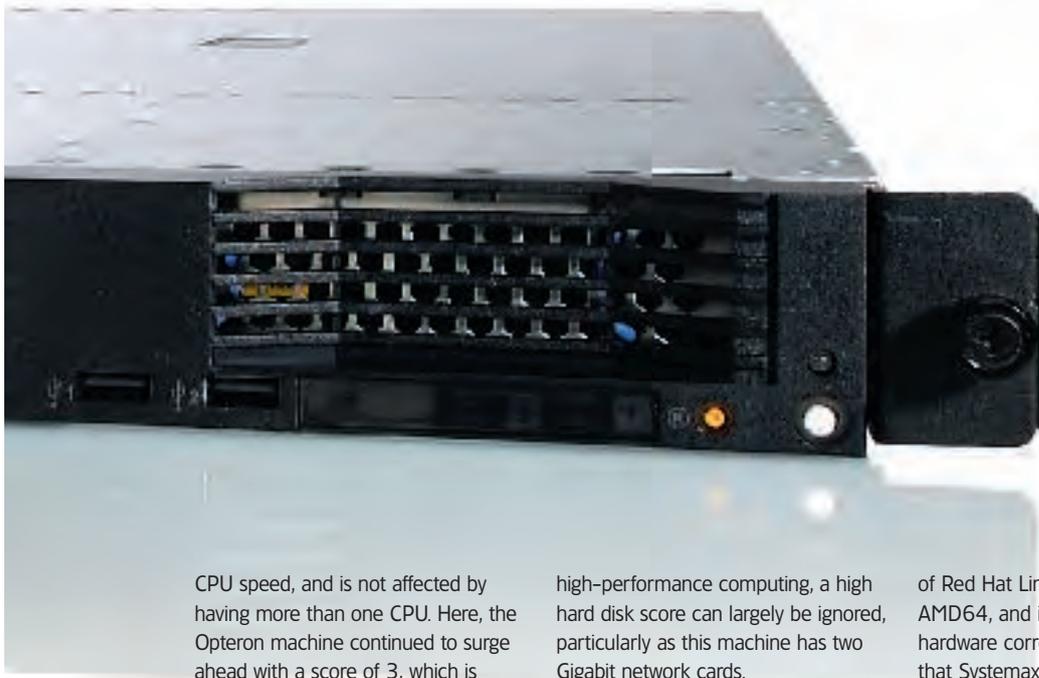
Performance

After being a long time in development, the brand spanking new *LXFBench 2004* benchmarking system is now available for use, and this was the very first machine that we tried it on. Given that this has a combined raw MHz of 4.4GHz, the scores were destined to be quite fast to begin with – keep in mind that our previous tests have found the Opteron chip to run about twice as fast as an equivalently clocked Xeon.

The first part of the benchmark tests multiprocessor performance, and does so by having four copies of the benchmark running simultaneously. Each one of these runs a variety of bit-shifting mathematical functions such as MD5 and SHA1, as well as running thousands of very large floating-point multiplications and divides. Each process then runs *oggenc* to encode a large WAV file to the Ogg format, and also performs some image resampling – both of which are real-world applications already being used around the world.

In this test, the Opteron scored 6.07, making it well over five times faster than our benchmark machine. Yes, there are two CPUs in this machine, which makes each chip about 3.04 times faster than the benchmark. As the benchmark is a 1.8GHz Pentium 4, a score of 3.04 is very good – keep in mind that single-CPU machines don't actually handle multiprocessing at half the speed of a two-way machine, as process switching and on-die cache limitations invariably takes their toll. As result, the score of 6.07 is actually higher than what we would expect.

We also run a uniprocessing test that does largely the same thing as the multiprocessing test, just without multiple processes. This is a raw test of



CPU speed, and is not affected by having more than one CPU. Here, the Opteron machine continued to surge ahead with a score of 3, which is another higher-than-expected score. It's important to keep in mind that we already know the Athlon XP chip is faster than the equivalently clocked Pentium, however the Opteron extends this lead by a large amount. To give you an idea quite how surprised we were at this lead, we sent our benchmark code off to Intel, and its engineers optimised it for us and sent it back. Despite that, we still found this machine performing above and beyond our expectations.

The last two tests are *RAM* and *HD*, both of which use the *SQLite* database system to create a database for reading and writing. The *RAM* test creates the database entirely in *RAM*, adds 25,000 entries to the system, then searches through. Naturally, this has some overlap with the *CPU* test in that a faster CPU will be able to execute the queries faster, but largely this is down to *RAM*. Here this server scored 2.17, which is excellent. For the hard drive test, where the test is about how fast data can be written, the machine returned its only low result: 0.46. As this machine is destined for

high-performance computing, a high hard disk score can largely be ignored, particularly as this machine has two Gigabit network cards.

An overall *LXFBench 2004* score of 2.55 is excellent, and shows this machine to be a great all-round performer. We still get emails from readers with questions about switching to Opteron – “Is it worth the money?”; “Is a safe investment?”; “Can it be relied on?”. The answer we always provide is “yes”; “yes”; and “yes”; and this is highlighted particularly well by this server. Systemax is a huge manufacturer that has a big backing in the industry – we can't think of a more reliable company to purchase from, particularly as their relationship with AMD is as strong as ever. If you really want to get the price further down, try asking for cheaper chips – the 248 CPUs are the fastest available right now and therefore cost the most. Asking for 246s or even 244s will still give excellent performance, but at an even lower price.

64-bit Linux

Having supplied Windows 2003 with the server, there is a lot of scope to question Systemax's support for Linux. Having said that, we did a fresh install

of Red Hat Linux Enterprise 3 AS for AMD64, and it detected all the hardware correctly – at least we know that Systemax made the effort to have Linux-compatible hardware. Furthermore, the performance of this machine shows that Systemax knows exactly what it is doing, and again proves that the Opteron is a processor leaving the competition in the dust for these configurations.

As for price, you have to admit that £2299 is very, very tempting. Yes, 1 GB of *RAM* might not be enough for you, and it's possible you will want *SCSI* drives rather than *IDE*, but everything else about this system screams “Buy me now”. While it is true that we have yet to see a system where Opteron underperforms in our expectations, this is certainly the fastest machine we have reviewed to date and at a cost that again shows Opteron as a chip to be proud of.

By summer 2004, we should be seeing the first samples of Intel's 64-bit Xeon chips, and it will be interesting to see how they face up against the Opteron. From what we can see, the new Xeon64s implement all of the AMD64 instruction set and add to the mix SSE3 and second-generation hyperthreading. Purely from a

performance point of view, we think the Xeon64 should revive the current Xeon line and maybe even get it neck-and-neck with the Opteron. However, Intel will need to make some deep cuts to its pricing, because right now a 2.8GHz Xeon MP chip costs \$3692, which means two of those chips alone costs more than this entire Opteron system.

As Intel's 64-bit Xeon chips are unproven, not yet available, and unlikely to be any cheaper than existing Xeons, we think that the Opteron will continue its popularity for the foreseeable future. Either way, we think the extra competition from Intel will help drive prices even lower and this can only benefit the end-user.

In the meantime, we recommend you take advantage of these low prices while you can. Systemax provides a three-year return-to-base warranty with this server and you can upgrade that to on-site for very little extra. If you can ignore the fact that you need to get the operating system yourself, this is probably the best low-end HPC machine out there at a price that would be criminal to miss. **LXF**

BENCHMARKS

CPU	6.07
SINGLE	3
RAM	2.17
HD	0.46
OVERALL 2.93	

LINUX FORMAT VERDICT

FEATURES	8/10
PERFORMANCE	10/10
DOCUMENTATION	7/10
VALUE FOR MONEY	10/10

Very fast and suited to even the lowest budget, we expect these to be flying off the shelves at Systemax.

RATING	9/10



PS/2 ports are a dying breed on servers of all kinds – we still pine for the old days though!

VALUE DISTRO BUNDLE

CheepLinux Megapack

Want to try out all the popular Linux distributions without getting a cap placed on your broadband?

Paul Hudson shows you how...

BUYER INFO

Distro bundle that gives you several top distros all in one box. Also consider: downloading them yourself and being the most hated user in your local cable exchange.

- **SUPPLIER** CheepLinux
- **PRICE** £29.99
- **WEB** www.cheeplinux.com/

Were you to take a demographic of LXF readers, you'd be likely to find that the majority of people buy the magazine for one of two reasons: to read the excellent *Flex & Bison* compiler-writing tutorial, or to get their hands on our coverdiscs. Alright, so perhaps I lied about *Flex & Bison*, but the coverdiscs remain popular as a source for the newest (and often some of the weirdest) free software available for Linux.

However, although our hard-working disc editor, Neil, does his best, space is always a strict limiting factor: we have just three CDs on which to cram all our content on, which means we have to choose very carefully which distros to supply each month. So, how can eager Linux users get their mitts on a selection of the top Linux distros all at once?

CheepLinux

As a distributor of Linux, CheepLinux, through its parent company Definite Software, is pretty much unparalleled – Definite runs CheepLinux, UKLINUX.NET, and DebianShop.com; and thus supports the backbone of the UK home Linux movement. CheepLinux has offered a huge range of Linux distros and software for some years now, even donating some of the proceeds to various Free Software projects. The latest and greatest package it has produced is

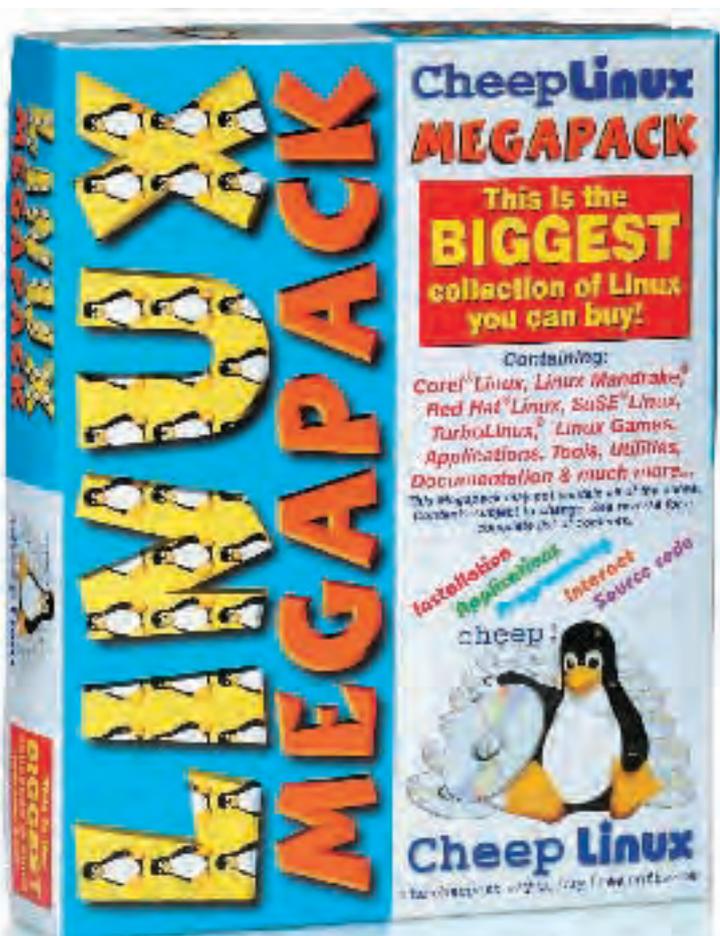
the CheepLinux Megapack, a set that at the time of writing comes with a bumper *fourteen* CDs. These are: Debian 3.0r2, Fedora Core 1 discs 1, 2, and 3, Mandrake 9.2 discs 1, 2, and 3, Slackware 9.1, Knoppix 3.3, *OpenOffice.org 1.1*, and a copy of the *Linux Documentation Project* on CD. Yes, that's only eleven CDs, but the box also comes with a special voucher inside that allows you to order three more CDs from CheepLinux free of charge (although not free of delivery costs; that costs another £2.99) from a choice of several.

Each of the CDs come nicely printed with its contents, and wrapped up in plastic CD books for dust and scratch protection. Beyond that, however, the boxes are empty: no manuals, no quick start guide, no FAQs, and no "if you have problems, why not check out <friendly-linux-people.com> where you can ask all the questions you want". This is what holds this boxed set from greatness: there's no help at all for newbies to Linux, not even a single sheet of A5 paper with a list of books that they can buy to help them out with their new distros.

It's arguable whether a newbie would go 'all out' on Linux by buying several distros at once, but just a little documentation wouldn't be too much effort on CheepLinux's behalf. Although the *Linux Documentation Project* CD is very welcome, and does mean users can at least get some online documentation when they need it. Of course, how a user trying out their shiny new Knoppix CD will be able to read the documentation CD is beyond us – perhaps if their system has two CD drives...

Other options

Is there much more to say about this? Well, no, not really – it's a box of the



most popular distros, all downloaded and burnt to CD so you don't have to do it yourself. It is worth mentioning the rest of the CheepLinux range, though – if the Megapack isn't for you, you can buy any of the individual distros direct from CheepLinux, either in multi-CD packs or on DVD. They still come on CheepLinux CDs and with no documentation, but you do get a much more detailed pack for your chosen distro. The Mandrake 9.2 boxed set, for example, is £24.99, and includes nine CD ROMs dedicated to Mandrake – three install discs, three source discs, a disc full of updates, and also an *OpenOffice.org* CD and Linux Documentation Project CD.

So, to conclude, if you want a Linux 'trial' pack that gives you the ability to try out several different distros easily, this is your best bet. The price is very reasonable, the CDs are all quality printed, and you're helping fund a very good cause to boot – give it a shot and see what you think! [LXF](http://www.linuxformat.co.uk)

CHECK OUR REVIEWS...

Before you try installing the distros provided, check our reviews to see what we thought of them. For example, we reviewed *OpenOffice.org 1.1* in LXF46, Debian 3.0 was reviewed in LXF32, and we had a special feature on Mandrake 9.2 in LXF47.

LINUX FORMAT VERDICT

FEATURES	10/10
DOCUMENTATION	N/A
EASE OF USE	8/10
VALUE FOR MONEY	10/10

The only docs are on CD, but you get so much value for money with this pack you won't mind printing it off yourself. Maybe not ideal for complete novices though...

RATING 9/10



COMPILER

Intel C++ Compiler for Linux 8.0

Biagio Lucini assesses the progress of the compiler that has gained the reputation of producing the fastest code for Linux on x86 systems.

BUYER INFO

Standards-compliant C/C++ compiler that produces optimal-performing code. Also consider *GCC/G++* and the *Portland Group C/C++ Compiler*.

- **PUBLISHER** Intel
- **LICENSE** Proprietary
- **PRICE** \$399 with a year support, Free non-commercial version with limited support also available
- **WEB** www.intel.com/

When considering C/C++ compilers in Linux, the first name that springs to mind is *GCC*. This is a great multi-purpose compiler, but fortunately it is not the only choice. In fact, some users are better catered for by compilers with more specific focus. The *Intel C/C++ compiler (ICC)* has reached version 8, and aims to provide users with a unified tool for the x86 and Itanium platforms. Since its appearance, *ICC* has got a lot of attention, since by and large it produces faster code than *GCC*.

Installation

Thanks to the vagaries of the Swiss Post Office, we review the unsupported free version of the compiler here, which is downloadable as a 65MB gzipped tar file. It's fully functional and shares many steps of the install procedure with the commercial release.

The compiler package contains several RPMs for x86 and Itanium architectures (respectively IA32 and IA64 in Intel jargon), an install script and a few other scripts used by the latter. Our review was performed on a 1.7 GHz

Name	Function	Hits
main.o	main	10

The code coverage tool (here at work on the benchmark *Scimark2*) comes very handy when debugging a project.

Pentium 4 system with 768MB RAM. The relevant packages to be installed in our case were the compiler itself, the required substitute headers (mostly a modification of GPL header files) and the optional debugger IDB. The latter is a symbolic source code debugger similar to GDB (but with an optional GUI) that supports parallel code.

The installation is supported only on RPM systems. The RPMs should not be installed with the `rpm` command though, since they need to be refined. This process is performed by the `install.sh` script, which should be executed in lieu of `rpm`. The Intel C++ compiler uses the *FLEX* license manager, and the installation script will check for a valid license in `/opt/intel/licenses`. However, it is possible to enter a different location when the script is executed.

As this is a major upgrade, users of previous releases should uninstall the

old version before installing 8.0. The installation went flawlessly on our test machine, which is set as a dual boot Mandrake 9.2/Knoppix 3.3. In the former case, we used the supported installation procedure, while in the latter we followed the procedure that is described at www.intel.com/software/products/compilers/techtopics/LinuxPlatformNotes80.pdf. That procedure is a bit involved and can be difficult to execute by a novice user, but we managed to perform successful installations on non-default locations and as non-root users.

What's new

The compiler comes with thorough docs and a good tutorial to help beginners to get started. Starting with this release, the Intel compiler supports the latest version of the Pentium 4 processor (codename Prescott). More generally, *ICC 8.0* inherits all the features from previous releases (like automatic vectorisation, autoparallelisation, *OpenMP* support and compliance with ANSI/ISO standard C/C++) and adds some more.

ICC 8.0 now boasts binary compatibility with *G++ 3.2*. Moreover, it now supports *GCC*-style inline assembly and several other extensions and macros typical of the GNU compiler. Also, some compiler options have been modified for better interoperability with *GCC*. New features welcomed by

developers are the code coverage tool – which shows what percentage of the code has been covered during execution; and the test prioritisation tool – which helps in understanding how to change the test suite after modifying the project).

How much will code performance improve with this new product? We tested alongside version 7.1 and used the C version of the benchmark suite *Scimark2* (see <http://math.nist.gov/scimark>). In the total score, we noticed an increment of performance of about 5 per cent in 8.0's favour, but looking carefully at the single benchmarks, we observed that there was a slight regression in all but one test, and that the much-improved performance on that particular one (around 20 per cent) overturned the result in favour of 8.0. For this reason, we concluded that *ICC 8.0* does not bring much gain (if any) in terms of speed. However, one has to bear in mind that the 7.x release of the compiler was already excellent from this point of view.

Conclusions

By and large, this is a great product; but it does not improve on the performance of the generated code. The new features of this compiler can make it a worthy upgrade for the users of previous releases. Together with the historical strong selling points of the Intel compiler (including the performance boost that many applications will get), those features should convince many others to give it a try. The only problem we envisage is the unsupported tricky installation on non-RPM systems, which might cause some problems to the inexperienced user. **LXF**

SYSTEM REQUIREMENTS

Hardware and software setup on x86 systems

- A system based on a x86-compatible processor, eg 450MHz Intel Pentium II processor or greater. Pentium 4 recommended
 - 128MB RAM (256MB recommended)
 - 100MB HD space, plus an additional 200MB during installation for the download and temporary files
 - *glibc 2.2.5, 2.2.93 or 2.3.2* and a 2.4.X kernel
- For a trouble-free install, use an RPM4-based distribution (installation of non-RPM system is possible, but unsupported)

LINUX FORMAT VERDICT

FEATURES	8/10
PERFORMANCE	9/10
DOCUMENTATION	9/10
VALUE FOR MONEY	8/10

A feature-packed product with excellent performance. It's a pity that installation is not supported on non-RPM systems.

RATING 8/10



SOFTWARE DEVELOPMENT SUITE

Portland Group C/C++ Compiler for Linux 5.1

In his eternal quest for the definitive compiler, **Biagio Lucini** gets to grips with the latest release of this well-known product.

BUYER INFO

C/C++ compiler with native support for parallel apps. Also consider *GCC/G++* and Intel C/C++ Linux compiler.

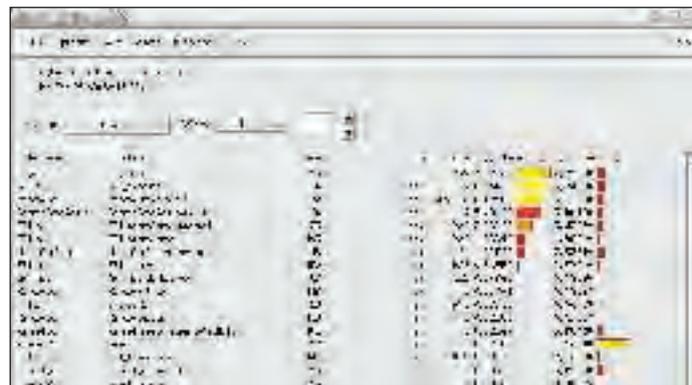
- **PUBLISHER** The Portland Group
- **LICENSE** Proprietary
- **PRICE** From \$239, free trial
- **WEB** www.pggroup.com/

The Portland Group is one of the longest players in the world of High Performance Computing (HPC): and the quality of its compilers has made it the *de facto* standard for Beowulf clusters. A fifteen-day trial of *PGCC 5.1* can be downloaded directly from the Portland Group website after registration, which is what we used to write this review.

Installation

It's pretty straightforward. All the vital components come packaged in a 37MB *tar.gz* file; it is recommended that you create a dedicated directory for decompressing it. After unpacking, it is just matter of reading the *INSTALL* file and to install the compiler by executing the *install* shellscrip. This will ask a few questions, among which whether one wants to install optimised libraries for AMD64 systems, where the compiler has to be installed and personal details of the user that will be needed when/if the full license will be installed. The package contains the C/C++ compiler, the Fortran 77 and 90 compilers, the profiler *PGPROF* and debugger *PGDB* for both x86 and AMD64 architectures. What components will be usable after the trial period will depend on the user license you buy. The license for the C/C++ compiler only includes the debugger and the profiler.

The installation of the license file is a complex step: *PGCC 5.1* can use a Portland Group-style or a *FLEXLM*-style license. The first option (the default) is



PGPROF (here at work on the Scimark2 benchmark) is a great graphical utility for finding bottlenecks in the code.

the easiest and allow an user to run concurrently as many executables generated with the compiler as he wishes. However, use of the compiler will be tied to a given user account. If more users need to access the compiler, then it is mandatory to install a *FLEXLM* license. This is slightly more complicated. Full details are given in the *INSTALL* file.

The electronic docs installed with the compiler includes the *PGI User's Guide*, which explains in detail and with an easy language the various possibilities offered by the compiler; the *PGI Tools Guide*, which addresses the profiler and the debugger; and the *PGI Workstation 5.1 Releases Notes*, mainly targeting users of previous versions.

In use

As a preliminary step before using the compiler, the shell variables *PATH* and *MANPATH* need to be updated to include `<INSTALLDIR>/linux86/5.1/bin` and `<INSTALLDIR>/linux86/5.1/man`. By default, `<INSTALLDIR>` is `/usr/pgi`, but this can be changed at installation time.

When starting to use a new tool, it is always a good idea to browse the documentation: these PDF docs are very detailed and easy to understand for anyone with some background in software development. In contrast, the

man page is quite involved, with a plethora of nested options. Generally speaking, having many options is a good thing, since this allows for more flexibility. However, those options should also be easy to manage, otherwise the user will tend not to look at many of them. In our opinion, this is a weakness of this compiler that is partly amended by the *User Guide*, which points out the most needed options.

Getting started with the compiler is very easy. A feature of *PGCC 5.1* is the high number of optimisation options, many of which are processor-specific. The *PGCC* compiler also supports many *pragma*-based vectorisation options, which allow for a fine-grained control over different parts of the code. The debugger and the profiler are two useful tools for developing software apps: with its intuitive interface, the former will save time in the bug-hunting process, while the latter's graphical report tool is great for finding bottlenecks.

The *PGI* compiler comes with native support for parallel applications, and the *OpenMP* support (including autoparallelisation) is integrated, while *MPI* support requires the *MPICH* libraries. All the built-in functions and routines are thread-safe, as they must be in a parallel environment. Both the

SYSTEM REQUIREMENTS

- Processor: x86 or AMD64 (multi-processor systems supported)
- Memory: 16 MB or more
- Hard Disk: 250 MB during installation, 100 MB for the installed software
- Kernel: 2.2.10 or higher

debugger and profiler are multithread-aware: a bonus, since other compilers don't come with concurrency-aware utilities, forcing the developer to buy them separately. Native support for parallel apps explains why *PGCC 5.1* is popular among Beowulf cluster users.

PGCC 5.1 is explicitly geared towards improving support on the AMD64, which is starting to get a sizeable share in the HPC world, and mainly concerns the companion Fortran front-end, which is claimed to have benefited from an increased performance of around 20% with respect to the 4.x release. The general expectation is that some of these improvements will reside in the so-called middle-end and back-end, shared by the Fortran and C compilers.

To test performance, we used the Scimark2 suite (<http://math.nist.gov/scimark>), which consists of a good sample of real-world HPC applications. To our surprise, we noticed that *GCC* overperforms *PGCC 5.1* by around 30 per cent, but we should remember that things may change on multiprocessor systems, where *PGCC*'s native support for parallelism can be exploited.

Maybe *PGCC 5.1* should not be thought of as a compiler, more a software development suite? The integrated graphical debugger and profiler with *OpenMP* and *MPI* support are a welcome addition to the compiler itself. If only the compiler is considered though, performance of the generated code might be disappointing. **LXF**

LINUX FORMAT VERDICT

FEATURES	8/10
DOCUMENTATION	7/10
PERFORMANCE	5/10
VALUE FOR MONEY	7/10

Great as a software development suite, disappointing for the performance of the generated code.

RATING 7/10



INTEGRATED DEVELOPMENT ENVIRONMENT

KDevelop 3.0.1



After a shiplot of expensive IDEs, Jono Bacon reckons the best things in life are Free...

BUYER INFO

Powerful IDE with a range of plugins to support GCC, G++, f77, pgmpf (Portland Group High Performance Fortran Compiler), pgf77 (Portland Group Fortran 77 Compiler), dcc (Borland Delphi Compiler) and fpc (Free Pascal Compiler).

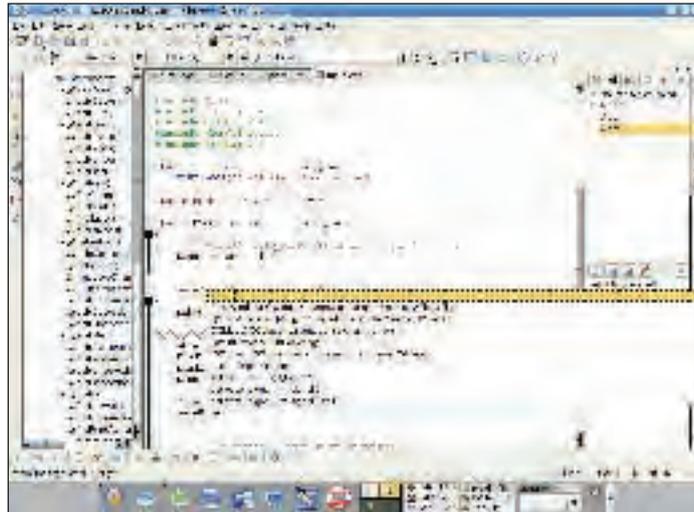
- **SUPPLIER** KDevelop
- **PRICE** free download
- **WEB** www.kdevelop.org/

Many years ago, I stumbled across the *KDevelop* project in its very early stages. The project was seeking to create a powerful IDE that could be used for all forms of development and various types of projects. Many years later, the project has put out the third-generation of releases. The true question that needs to be answered is – does *KDevelop* make you develop better?

This 3.x series of *KDevelop* is the first to be based on the completely rewritten codebase known as *Gideon*. Although *KDevelop* was in many ways the leading Free IDE, many parts needed to be rewritten to allow the sensible inclusion of additional features and plugins. This work has taken a few years to mature, but we can now explore the culmination of these efforts.

Features

KDevelop 3.0.2 provides a huge range of project templates. These templates give you a 'bare bones' application and build setup for a specific type of project or platform. This includes KDE, GNOME, Console, SDL, OpenGL, Ada, Database, Java, PHP, Ruby, Shell and others. This expansive range of templates is itself useful for the beginning developer; a template will get you up and running straight away so you can concentrate on the code and not the tool. The only downside to the portfolio of projects available is the support for PHP. Although you could admittedly use *KDevelop* for PHP development, a tool such as *Quanta* or *Bluefish* is far better suited to this kind



Install plenty of other tools, and access them from inside *KDevelop*.

of development. I see *KDevelop* as predominantly useful for compiled code/scripting projects.

When you have selected a project and clicked your way through the New Project wizard, you are then ready to begin development. *KDevelop 3.0.1* has included a number of new retractable panes with different information such as the class viewer, bookmarks, file information, new file quickstart buttons and more. In addition to this there are retractable panes on the right of the screen that are project-specific, so you can have the Automake manager present for a KDE project for instance. These panes are useful when retracted but can also be pinned to stick out.

In use

KDevelop is a strikingly flexible environment compared to some of the more rigid proprietary IDEs. Virtually everything can be configured: including the editor, the build system and how your project is managed. Code completion is one notable feature in this series of releases and not only shows methods available in the class that you have written but also classes in various toolkits. As an example, if I had a Qt object called `button` that is a `QPushButton` class, when I type `button->` I will get a list of `QPushButton` methods available. This makes development quicker, and in

many ways easier; it gives you an idea of which options you have available when using a particular toolkit.

KDevelop also includes a huge number of other tools and additions that are part of the environment. This includes documentation generation with *doxygen*, CVS support, patch difference viewing, memory leak checking, integrated debugging, full-screen mode, line numbering, different editor support, different output information and many other features. There really is a lot of functionality packed inside. One thing to be aware of though is that much of this functionality may require some other tools to be installed; as an example, you will need the *valgrind* tool to make use of memory leak checking. This is a good choice – it lets those programs do what they do best, and simply provides an interface to them all in *KDevelop*.

For everyday usage, *KDevelop* is a great IDE to use. Naturally *KDevelop* is not the easiest of tools in use, but this is a criticism levelled at most IDEs and development tools. You should expect to spend a little bit of time playing around with it before you can be fully conversant in its features and be productive. The only real criticism that I have in terms of usage of *KDevelop* is regarding how all of this functionality has been packaged within the menus

and dialog boxes of the application; it all seems a little too much sometimes. There is a feeling that there is so much that this tool can do that it is a little off-putting to have all of these features accessible from the menu all of the time. This is not a huge concern however, and a bit of exploration and experimentation will soon get you familiar with the system.

Conclusion

KDevelop is a great IDE. Out of all the free software IDEs that I have played with, I think *KDevelop* stands out as the most flexible and productive tool. Many of the other IDEs are developing well, but at the moment, *KDevelop* takes the crown as the cream of the IDE crop.

Whichever way you look at the 3.x series of *KDevelop*, I think it is difficult to see how it is not a vast improvement on earlier versions. *KDevelop* has matured into an incredible application with an absolutely vast amount of functionality included. This functionality is not only part of the *KDevelop* application itself, but augmented by a strongly integrated plugin system that inspires developers to craft their own *KDevelop* functionality.

If you have an ambition to get involved with software development for Linux, *KDevelop* is a wise choice that will gain you respect in the community. Aside from the expansive user interface, *KDevelop* is a great tool for beginners and veterans alike. The single issue that keeps this latest edition of *KDevelop* from a *Top Stuff* award is a similar lack of quality documentation that was also suffered by earlier iterations... **LXF**

LINUX FORMAT VERDICT

FEATURES	9/10
PERFORMANCE	8/10
EASE OF USE	8/10
DOCUMENTATION	6/10

If you do use this for development of a publicly available project, please make sure you make full use of those documentation generation tools!

RATING 9/10



NETWORK BACKUP SOFTWARE

Arkeia 5.2

Chris Denton tries out a backup server designed to make data protection a lot less trouble.

BUYER INFO

Veritas and Legato are big players in the server backup market, also consider an Open Source alternative such as Amanda.

- **DEVELOPER** Arkeia
- **PRICE** From \$290
- **WEB** www.arkeia.co.uk/

Were we living in an ideal world, users would never delete files they actually wanted, nor would hard drives ever fail. To guard against these and other eventualities, a backup solution is absolutely vital. It is perfectly possible to use standard Linux commands such as **tar** or **cpio** to achieve this, but to have any hope of being able to do so effectively some pretty hefty scripting would be called for. *Arkeia 5.2* promises to take the strain of managing backups away from the sysadmin, and hopefully make backups faster and more reliable too.

Arkeia is a networked backup solution based on the server-client model. It is compatible with a variety of OSes, including most significant Linux distros. The server obviously requires at least one backup device, which can be either a hard disk or a SCSI-based device such as a tape drive or an autoloader. Other kinds of hardware, such as an IDE CD-writer, for example,

are not supported and are therefore unlikely to work correctly with *Arkeia*. For the purposes of the review we decided to use Fedora as the server's OS, and selected a machine that handily contained an internal DLT 7000 tape device. Several backup clients were also earmarked, including a SUSE 9.0 desktop and a technical workstation running Windows XP.

The installation can be performed from either a RPM package or a gzipped tar archive. Although it is listed as supported on *Arkeia's* website, there was no RPM for Fedora provided on the software CD. We could have installed using the alternative method, but elected to try the package meant for Fedora's close relative, Red Hat 9. Fortunately this worked perfectly well, and appeared not to cause any problems either during installation or at any time afterwards. This was perhaps just as well, since when it came to install the clients the relevant versions were again not available. Emboldened by our first success, we installed the RPM intended for SUSE 8 on SUSE 9.0, and the zipfile marked for Windows 2000 on Windows XP. Again this seemed to be absolutely fine, although it is rather off-putting when the current versions of all these major operating systems are apparently overlooked.

With all the software in place it was time to fire up the graphical



The incredibly stylish GUI has a real dashboard feel throughout *Arkeia*.

management console and start to configure *Arkeia*. This proved to be a rather exciting experience because the GUI is truly exceptional. It has been elegantly constructed and manages the difficult feat of being both packed with information and relatively clutter-free. Even command-line junkies will be hard-pressed to keep things text-based when this interface is available and crying out to be used. The only drawback is that the layout is rather unconventional, so it may take a bit of getting used to. Still, the design is very intuitive and a product such as *Arkeia*, which is not as feature-packed as, say *Legato Networker*, lends itself very well to this type of thing.

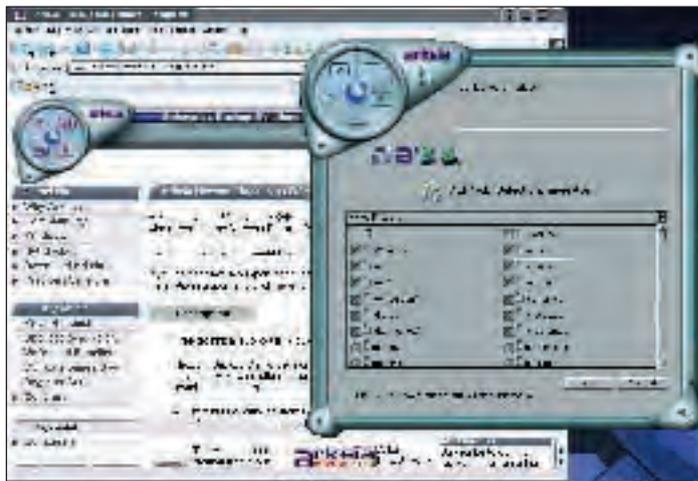
Restoration

The process of configuring backups is fairly simple, and is greatly aided by some excellent documentation provided on the CD. The terminology can be a little strange, especially the references to "savepacks" and "drivepacks", but these concepts are quite simple to understand, and the Quick-Start Guide makes things clear. Basically, a savepack is a group of clients to be backed up; for instance, you could have a savepack called 'Windowspack' and another called 'Linuxpack' and put all the relevant clients in each group. A drivepack is just a collection of backup devices, and on large machines you could put several of these into the same group to potentially improve backup performance.

Backups can either be kicked off interactively or periodically using a sophisticated and friendly scheduler. If

you decide to start one manually a status display appears that astutely visualises the data throughput by use of a speedometer. It is possible to view this screen for scheduled jobs, although of course they would normally run at the dead of night. *Arkeia* is very good at managing the capacity of your media so it is always clear how much has been written to a tape and crucially how much space is still available. The act of restoring data is also pretty much pain free. The GUI eases the process, which in truth comprises no more than a few clicks of a mouse. The attractive status display appears again here, equally as good for restores as it is for backups.

So it seems that *Arkeia* lives up to its promises. It may not be the most comprehensive backup solution, but it is efficient, friendly and above all, reliable. The cost, though modest, may put some people off; they should try *Arkeia Light*, a Freeware version that supports one Linux server and up to two desktop clients. For larger setups, the full version of *Arkeia* is well worth considering. **LXF**



Navigating the *Arkeia* client's file structure is simplicity itself.

LINUX FORMAT VERDICT

FEATURES	8/10
PERFORMANCE	9/10
EASE OF USE	9/10
VALUE FOR MONEY	9/10

Simple to install and maintain, yet powerful enough to cope with most environments, *Arkeia* is a first-rate backup management product with broad appeal.

RATING **9/10**



OFFICE PRODUCTIVITY

KOffice 1.3

Marco Fioretti examines the most serious contender to *OpenOffice.org*.

BUYER INFO

OpenOffice.org may rule the productivity roost at present, but it shouldn't get too complacent...

- **DEVELOPER** The KOffice team
- **PRICE** Freeware
- **WEB** <http://koffice.kde.org/>

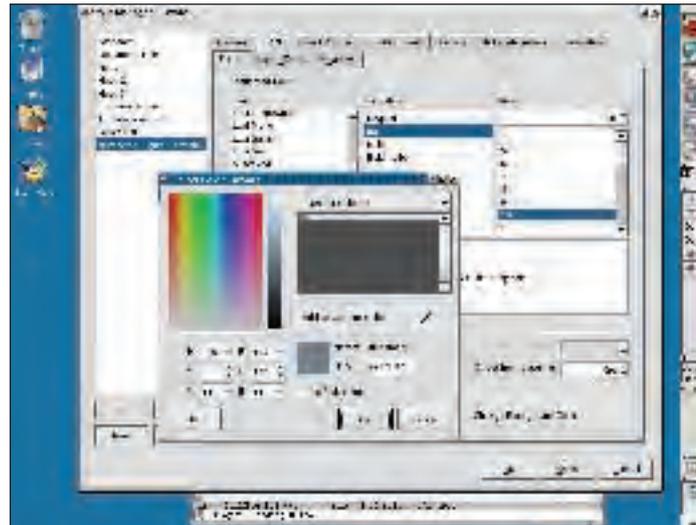
As many readers will know already, *KOffice* is the integrated office productivity suite of the KDE desktop environment, with its own release schedule, independent from that of the parent project. Version 1.3 of *KOffice* was released in late January 2004.

The main components are *KWord* (a word processor inspired by *Adobe FrameMaker*), *KSpread* and *KPresenter*, respectively a spreadsheet and a *PowerPoint*-like presentation program, as you might expect from the names. Other parts – arguably more specialised, but highly valuable for many users – take care of generating mathematical equations (*KFormula*), flowcharts (*Kivio*), charts (*KChart*) and business reports (*Kugar*).

Two separate applications – *Krita* and *Karbon 14* – are devoted to drawing. The first one is a bitmap image program like *The GIMP*, the second is for vector images.

What's new in 1.3

All the programs included in *KOffice 1.3* have a lot of new features with respect to previous versions. *KWord* gains automatic hyphenation, text sorting and



The interface to create new *KWord* paragraph styles is packed with options.

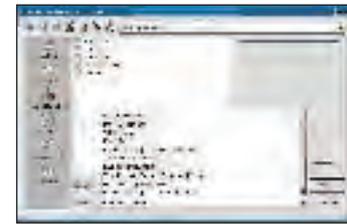
better spellchecking (through *Aspell*). The presentation tool has more customisable printing and preview of the slide transition effects. The spreadsheet offers more than 100 new formulas. The most important improvements, however, are in the area of import and export filters. Significant progress has been made, starting with compatibility with *OpenOffice.org*. This is a prelude to a very important event: the next version of *KOffice* will switch to the same native format as *OOo*, the OASIS standard (www.oasis-open.org/): this will guarantee the maximum interoperability with all *OpenOffice.org* users on all platforms. A complete report on the current status of *KOffice* filters is available online at <http://koffice.kde.org/filters>.

KOffice can be controlled remotely through Desktop Communication Protocol (DCOP0; just type `dcop --help` at the prompt to find out more) can be used for this purpose even inside shell scripts. Bindings exist for C, C++, Perl, Python, Java and other languages. DCOP is also suggested as a (safer) substitute for *KOffice* macros.

Using KOffice

KOffice is fast (once KDE is already running), whereas *OOo* is often slow, even after it starts up. *KOffice* is a pleasant change from this point of view. On computers where *Open Writer* gives you enough time to check the weather outside while opening a menu, *KWord* answers almost immediately. Another nice surprise is outline numbering. Unlike its more popular cousin, *KWord* offers three nested levels of numbered headers in the default style catalog: what's more, they actually work without further fiddling, incrementing exactly as they are supposed to.

We also found the procedure to create new custom paragraph formats (shown above) quite fast and intuitive. Generally speaking, the user interface of the whole suite is quite pleasant, visually, simple and well organised. The integration among different programs is another strength of *KOffice*:



Interoperability is much improved in *KOffice 1.3*: there are some of export filters available in *KWord*.

embedding different kinds of documents (charts, pictures...) into others (text, presentations...) is not a problem. *KPresenter* and *Karbon 14* also did their jobs quickly and without any problems.

Unfortunately, other parts of *KOffice* seem to be less mature than the word processor. *KSpread*, for example, didn't print an embedded chart properly. Installation from sources – especially in non-standard locations – presented several little traps, all documented in the February 2004 archives of the *KOffice* mailing list. All the *KOffice* developers, starting with David Faure, were really helpful in sorting issues out. Source installation is not a big deal however: by the time you read this, tested binary packages for all the latest distributions will be certainly available.

Conclusion

If your main requirement is the greatest compatibility with *MS Office's* file formats and/or user interface, *KOffice* is not the best choice. Conversely, if you are looking for a Linux office suite which is reasonably complete, different and faster, then give *KOffice* a serious try. *OOo* may get the most coverage, but this is its most serious contender. *Abiword*, *Gnumeric*, and the other bits of the *GNOME Office* initiative are as fast as *KOffice* or more, but not equally well integrated. Maybe *KOffice* is not quite ready for prime time yet, but it is a big step in the right direction. **LXF**

WHAT, NO DATABASES?

Where is something like *Microsoft Access*?

Writing memos or generating attractive pie charts is not enough. Small and medium businesses – as well as individuals who need to keep their tape collections and recipes under control – can't live without some simple database management interface.

The *KOffice* project does offer such a tool, called *Kexi*: since it wasn't included in version 1.3, but scheduled for a later standalone release, we decided that it

wasn't appropriate to review it this time. The developers' plan is to eventually merge *Kexi* back in the next major *KOffice* version. The program is supposed to do just what one would expect: to replace tools like *MS Access* or *FoxPro* in creating, updating and querying databases. *Kexi* will interact with several engines, including *SQLite* (www.sqlite.org), which doesn't need a separate server process to be used.

LINUX FORMAT VERDICT

FEATURES	7/10
PERFORMANCE	9/10
EASE OF USE	9/10
DOCUMENTATION	8/10

More of a stop-gap solution while waiting for OASIS formatting than a full version leap, but the new tweaks bode very well indeed for full *OOo* compatibility.

RATING **8/10**



Red Hat Linux Fedora Desktop Kit for Dummies

Getting to grips with Fedora couldn't be easier, could it? Andy Hudson tries to find out...

BUYER INFO

- **AUTHOR** Jon 'maddog' Hall, Paul G. Sery, Susan Douglas, Korry Douglas
- **PUBLISHER** Wiley
- **ISBN** 0-764-54263-X
- **PRICE** £33.99

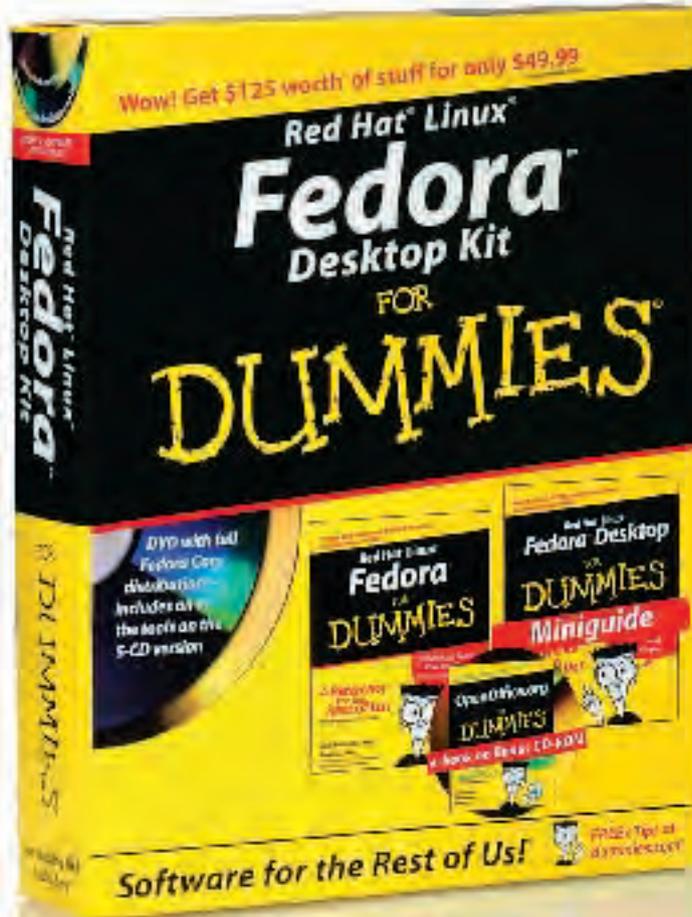
For quite some time, Wiley has been bringing us various *Dummies* books on the Red Hat distributions, so it was no surprise that when Red Hat Linux diverged into The Fedora Project and RHEL that Wiley would match this.

The *Fedora Desktop Kit* aims to be an all-in-one resource for the beginner to Red Hat/Fedora who wants to learn in an easy way. There are other books available, but *Dummies* has a long history of presenting their topics in a lighthearted and, dare I say, fun way that gets the message across in a form that is easy to digest and actually quite helpful. There are three separate parts to the kit – *The Fedora Project for Dummies* (with The Fedora Project on DVD), *The Fedora Desktop for Dummies* and the companion CD containing *Fedora Desktop for Dummies*, *OpenOffice.org for Dummies* and a number of RPMS for you to install.

The basics

The Fedora Project for Dummies is an update of *Red Hat Linux 9 for Dummies*, also written by the same team of Hall and Sery. Reading through the book, you can't help but wonder what the split was between the two authors. Nearly all of the examples have *lidia* as the localhost, which apparently is the name of Sery's wife. True, any book with Jon 'maddog' Hall as a co-author is worthy of some attention, it's just that I almost felt let down reading through it.

The book covers a surprisingly wide range of things to do, but instead of going really in-depth, it only shows you the basics of how to get things up and running. This is actually quite a good thing, considering that most people who pick up this book will be wanting to use Linux, giving them an adequate amount of knowledge to build on and doing so with minimum fuss. What was quite good was the



inclusion of a chapter discussing how to set up a streaming media server, which would appeal to many people eager to do 'cool things' with their kit. However – and this is a criticism of both this book and the desktop book – the editors haven't done a very good job of updating them to replace references like Red Hat Linux 10 or Red Hat Linux. Granted, there are notes at the beginning explaining that at press time Red Hat changed to The

Fedora Project, but the amount of times you come across references to Red Hat Linux 10 almost give you a sense of picking up the wrong book.

Second half

The second book, *Fedora Desktop for Dummies*, starts off in a similar vein, aiming to give you a number of useful tips on how to customise every aspect of KDE. Yes, KDE – the authors seem to be under the impression that this is the

default desktop for Red Hat Linux 10, sorry, *Fedora*. What's more confusing is the fact that they claim that when you see a Red Hat icon on the task bar, then you will be in GNOME, so switch to KDE and you will see the K icon – sadly, in Fedora, this simply isn't true. As default, the Red Hat icon is displayed on both desktops, so it's not a good start for a book that is aimed at beginners!

Thankfully, this is pretty much the only error in the book, as the rest of it reads very easily and provides some useful tips for customising your desktop experience. Perhaps surprisingly, they devote a whole chapter to working with remote desktops; in this day and age of homes often having more than one computer, this is definitely a welcome addition, as is the chapter on *Wine*. It's good to see that the authors give the user the knowledge of how to access existing Windows applications such as *MS Office* and *Photoshop* – this is an important area for the take up of Linux on the desktop, and this kind of thing can only help. Another welcome addition is a section on Linux gaming: long an underdog in the Linux world, and rounding the gaming section off is a nice intro to the ubiquitous *Tuxracer*.

Money-off coupons

The Companion CD includes *OpenOffice.org* info in eBook format – which is all very well, but you can't print from it. Adept Linux users will know where to find the info in other formats on the Internet, but for the novice, this condemns you to hours of nightmarish window-swapping while you work with OOo – unless you're lucky enough to have a PDA, laptop or another computer close to hand. The CD also contains two coupons for the Fedora Project on CD and money off the *OOo for Dummies* book. However these are in dollars, so aren't much use if you're wanting to buy elsewhere!

LINUX FORMAT VERDICT

Dummies is a generally good series, but this volume is let down by poor attention to detail in the updating, and an American bias that may annoy some global readers.

RATING **6/10**



J2EE Developer's Handbook

Employing the services of a fork-lift truck, **Andy Hudson** succeeds in conquering J2EE.

BUYER INFO

■ **AUTHORS** Paul J. Perrone, Venkata S. R. "Krishna" R. Chaganti, Tom Schwenk
 ■ **PUBLISHER** SAMS
 ■ **ISBN** 0-672-32348-6
 ■ **PRICE** £43.99
 ■ **PAGES** 1490

This book is big – there's no other way to describe it. You will need a good understanding of programming techniques before you start: it is such a complex topic to get into. The authors make it a little less daunting with code snippets throughout, and also flow charts that detail how the code works and the processes that it follows. However, it does get very in-depth very quickly, making it almost impenetrable for the fledgling developer.

The book deals with a wide range of technologies, covering everything you would realistically need to know about in Corporate/Enterprise



environments – like messaging, Javamail, CORBA, XML, SOAP and SAAJ, Java Naming and Directory Interface. The best thing is that it is pretty much platform-independent, as you would expect a Java book to be. Where a particular technology is relevant, like *Windows Active Directory* or Novell's *NDS*, the book explains clearly how to manipulate these services using J2EE.

The authors themselves state that the book is a bit of a tome, yet that at the end of reading it, you should have a comprehensive yet practical knowledge of the ins and outs of J2EE development. They're not wrong! But to get there you will need plenty of focus and caffeine.

Unfortunately, there are a few minor grumbles about this book. For one, the proofreaders missed a few typos (glass houses, stones!) – the

most obvious one being on the outside back cover ('Enterprise'). Another is the inconsistency of the diagrams that are used throughout. This might seem like nit-picking, but it does get rather annoying when they use different representations for the same ideas.

When you reach the end of the book, you can't help but feel that you have completed something of an odyssey. This is an entirely accurate feeling and this is a volume that will give you all the knowledge you will ever need to get the best from the J2EE platform. Though the term 'handbook' might be a little misleading, you would be wise to keep this book close to hand for its sheer reference value.

LINUX FORMAT VERDICT

A heavy-weight book that is not for the faint-hearted. Approach with respect and you will richly rewarded.

RATING **9/10**



Sharp Zaurus SL-C750/C860 From Japan, Now Available In The UK



Amazing 640 x 480 Screen, Full Keyboard,
Word/Excel Compatible, Presentation,
PIM, Web Browser, Video, MP3 ...
Linux In Your ShirtPocket !

ShirtPocket
LIMITED

Tel : 0116 2201777
Email: salesdept@shirt-pocket.co.uk
http://www.shirt-pocket.co.uk



Hacking The TiVo

The only mystery of the TiVo is how it hasn't caught on yet in the telly-addicted UK, says Alan Smithee...

BUYER INFO

- **AUTHOR** William von Hagen
- **PUBLISHER** Premier Press
- **ISBN** 1-592-00111-4
- **PRICE** \$29.99
- **PAGES** 402

These days, it's common knowledge that the last person you want to be trapped with in a lift is a TiVo owner.

TiVo may have floundered in the UK, but that won't stop anyone who bought one telling you in excruciating detail how TiVo has changed their life. And as Linux-heads, we can do so much more – my TiVo lives in the US with me, but on a visit to the UK, I set it up to (almost certainly illegally) re-encode the latest episode of



Enterprise into MPEG4 and send it to me via FTP.

I didn't get to do this because I am a super-genius, I got to do it because – at its heart – TiVo is a Linux-based computer and cleverer people than

me have worked out how to let me abuse it.

Hacking The TiVo is written by one of these clever people, and covers everything you need to know about different TiVo models, some useful tips and tricks, how to upgrade the hard

drive(s), and an entire chapter regarding hooking it up to a Linux system. An included bootable CD-ROM provides a lot of useful utilities, and makes DIY hard drive upgrades as about as easy as they could be. The only shortcoming is the chickening away from decoding recorded video programs, which is, sadly, understandable.

It's up-to-date and makes a great guidebook for all TiVo owners.

I think I'll be boring people in lifts about two things now: my TiVo, and this book on hacking TiVo.

LINUX FORMAT VERDICT

Hack your TiVo now: this book gives you the methods AND the inclination...

RATING **9/10**



The Best of The Joy of Tech

Paul Hudson has finally figured out how to get paid to read comics...

BUYER INFO

- **AUTHOR** Nitrozac and Snaggy
- **PUBLISHER** O'Reilly
- **ISBN** 0-596-00578-4
- **PRICE** £9.99
- **PAGES** 192

Alongside other classics such as *Dilbert*, *The Joy of Tech* stands out as an entertainment icon read and enjoyed by geeks around the world. Nitrozac and Snaggy (Liza and Bruce to their offline friends) have been working on the online edition of *Joy* for some time now, but this is the very first edition to appear in print. If you're a long-time fan, you will need no introduction to their sharp



and topical humour, but if this is your first toe-dip in the hilarious and satirical geek universe that's been so lovingly hand-crafted by its two creators, let us give you a taste.

Whether you get your kicks trying out the tips in *How to lose a geek in 10 seconds* ("Surprise him by re-organising all those icons on his computer's desktop for him"), reading all about the adventures of *Linux*

Lass ("Super Linux Evangelist! Software Freedom Fighter! Defender of the GNU and the Open Source Way!"), or just wending your way through the myriad of references to Mac OS X and *Star Wars* (such as Luke and Leia buying a Father's Day card titled, "You're sick, twisted, and evil, but you're still our Dad"), there you are sure to find *something* to amuse in here.

Each of the comic strips have been carefully scaled up to much higher resolutions than previously available, and they benefit greatly from the touching-up. Some of the colours used by the artists often looked a little too brash on screen, but this isn't the

case here at all – the pictures nearly always look better in print than they do if you look them up online, which makes the pictures live up to the humour in every way. Artistically speaking, this is as good as it gets, and it's good to see that O'Reilly has pulled out all the stops to make the book high quality.

Perhaps the most amazing thing about this book is its price: it's just £9.99. If this book were £19.99, we would still consider it a bargain because there are over 100 pages of high-quality and hilarious comics that will keep sides splitting in any geek environment. There are some in-jokes that will take a little getting used to, but that's no bad thing – watch out for digs at the 'Switch to Mac' campaign, relentless jibes about Microsoft (all in good humour, of course!), and guest-star appearances by George Lucas, Steven Spielberg, Steve Wozniak, and Cowboy Neal. This is all part of the fun, though, and the sooner you start reading the better! **LXF**

LINUX FORMAT VERDICT

This is geek culture at its finest – some of which will even be accessible for non-geeks – and we can't get enough of it! Three cheers for Nitrozac and Snaggy!

RATING **10/10**



Roundup

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

Compression software

Trying to cram the kernel source onto a floppy, Mike Saunders looks at Linux's best compression tools.

Although hard disk capacities are spiralling upwards and broadband Internet connections are becoming ever-more commonplace, compression tools and algorithms still play an enormously important role in the computing world. Compressed files save time when downloading, provide more space on CD-Rs, and are a necessity for archiving large amounts of data, particularly if that data is rarely used. Compressing larger files so they can be successfully sent by email is an essential task for most users.

Under Linux and most other modern Unix-like operating systems, the traditional tools for creating compressed archives are *tar* and *gzip* – the former bundles together files into one neat package, while the latter applies compression. *gzip* and *WinZip* are the *de facto* standards for compression, but various other formats are popular and – depending on the original file types – can often lead to even smaller results (such as *bzip2*, for instance).

Essentially, compression tools work by searching a file for redundancies: they find blocks of bytes or strings of text that are repeated throughout and replace them with smaller equivalents. Take this article, for example – the word “compression” appears very frequently, so substituting it with a single number or letter would make the text much shorter (but a bit difficult for human readers to understand). The compression



OUR SELECTION AT A GLANCE

- Ark
- Bloat
- File Roller
- GNOChive
- KArchiver
- TkZip

“Compression tools work by searching for redundancies: they find repeated blocks of bytes or strings of text and then replace them with smaller equivalents.”

application builds a table of these patterns for substitution, and then fills them in when expanding.

For this reason, compression works extremely well on text documents, but it can apply to many other file formats as well. Large blocks of static colour in BMP files can be replaced with a few bytes, and software binaries with repeating sections of machine code can be shrunk down well too. Some files which incorporate their own compression (eg GIF and MP3) don't see much of a reduction. Naturally, it

gets more complex than that, with many different algorithms and techniques in use, but the principles are still the same, whatever filetype is being compressed.

Perhaps the most well-known compression tool is *WinZip* on Microsoft Windows, which has grown over time to support a multitude of formats other than its native Zip. This is the archiving application that most home users will be familiar with – its interface and design have been imitated by many and graphical

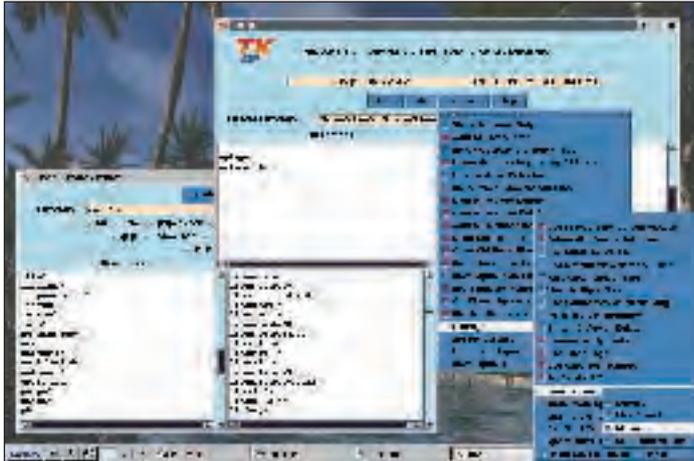
equivalents in Linux and other OSes tend to follow a similar path.

This month, we're looking at a range of the best known and most popular applications which assist in creating, managing and extracting compressed files, weighing up their interfaces, stability, performance, and range of features. Undoubtedly, support for a wide variety of compression formats is one of the most crucial things to look for, but layout design, documentation and reliability are all of equal importance as secondary considerations.

TkZip

Tcl/Tk mainstream comeback in a lightweight app.

■ **VERSION** 1.1.5 ■ **WEB** <http://woodsway.com/TkZip/>



The interface feels messy at first, but it's not too difficult to navigate.

Ark and *File Roller*, without doubt the two most prominent apps in our selection, aren't much good on low-spec machines: GNOME and KDE are weighty beasts, and the disk and

memory requirements mean they're not all that suitable for the old Pentium in the loft; nevertheless, one of the best aspects of Linux is its flexibility, and with a lighter desktop and smaller

apps like *TkZip*, that dusty machine can become a decent little workstation.

TkZip (Shareware, but free for non-commercial use) is written in Tcl/Tk, a scripting language and associated toolkit for quickfire development of graphical apps. While Python, Perl and *Gtk/Qt* bindings are much more popular these days, many coders still use Tcl/Tk and consequently it should be available somewhere in your distribution's package repositories. If not, check out www.scriptics.com for the full works.

There's a slightly unusual installation procedure (run the main program, and it'll ask you for a directory into which the files are copied), and for those fairly new to Linux the 'old-school' interface may be somewhat weird at first glance. Leaving aside the slightly garish stock colour scheme, *TkZip*'s interface is manageable after a few minutes of use and gets the job done adequately.

TkZip will work with *gzip*, *bzip2*, old UNIX compress, *ar* and *lzh* formats; it's not the broadest range, then, but still covers most essentials. Annoyingly, the interface freezes when opening large files, and the default viewing program

links are almost quaintly old (*Mosaic* and *XLess!*). There's no way to extract files selectively, nor can files be sorted by name and attributes *etc*, but the abundance of options to tweak and tune remedies the situation slightly.

Featurewise, *TkZip* naturally doesn't compare to the likes of *Ark* and *File Roller*, but its main strength lies in its desktop independence and skimpy resource requirements. And even with its little idiosyncrasies and seemingly strobe-lit colour scheme, it's usable for general archive manipulating chores. If you don't have KDE or GNOME installed, but need a simple and speedy compressor, it's worth a look.

LINUX FORMAT VERDICT

FEATURES	5/10
EASE OF USE	4/10
DOCUMENTATION	8/10
PERFORMANCE	9/10

Desktop independence makes this ideally suited for low-spec systems, and will also be loved by those people who glory in retina-burning colour schemes.

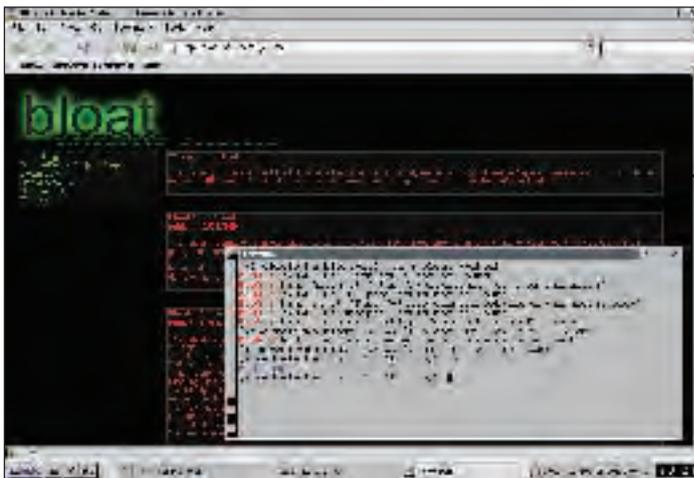
RATING **7/10**



Bloat

Only the version number is bloated!

■ **VERSION** 1.3.0.1 ■ **WEB** <http://bloat.sourceforge.net/>



Not visually stunning – thankfully, the red bits of info text can be disabled.

Apart from *Bloat*, all of the tools we've looked at here depend on a GUI to operate. For desktop users that's no problem, but on servers (or old boxes

without the muscle to manage X) the command-line tools come into play. The problem here, though, is that remembering every tool and command-

line option for all the different file formats is worthy of a Mensa application test. *Bloat*, from LXF reader Dave Wickham, aims to solve this by providing a single command for all types, making life for newcomers and admins easier.

Installing *Bloat* Perl script is a doddle, as it only depends on Perl to run. However, it doesn't contain any specific extraction code itself – you'll need the correct binaries for specific formats, and we've provided packages for all common ones on our coverdisc. The only other setup step is to copy over the sample *bloatrc* file into */etc* (or in home at *~/.bloatrc*).

Basic use of *Bloat* involves passing an archive file as an argument to the binary. For instance,

```
bloat foo.tar.gz
```

will throw up a couple of lines indicating the file format; the script looks for an extension first to determine the type, or uses the **file(1)** command otherwise. This is where *Bloat* shines – you don't need to know what the format is, what the extension means or what commands are required to extract it.

Bloat's supported format list is impressively thorough, although the extent of support does depend on the

utilities already installed. All of the regular formats can be managed, along with some intriguing additions including *Mozilla .xpis*, *Webmin* modules and *Unreal Tournament* packs. Curiously, *Bloat* and *File Roller* are the only programs in our *Roundup* that can extract RPMs easily. The config file allows for various output alterations to me made, along with the location of the dependant utilities, and it's well commented too. *Bloat* may be a petite and simple, but the author has executed it well and it's unquestionably handy when you've come across an utterly esoteric archive format and don't want to spend hours poring over man pages.

LINUX FORMAT VERDICT

FEATURES	4/10
EASE OF USE	8/10
DOCUMENTATION	8/10
PERFORMANCE	10/10

Not as featureful as the graphical tools boasted by other apps, but a great help for users who prefer the command-line to GUI fripperies.

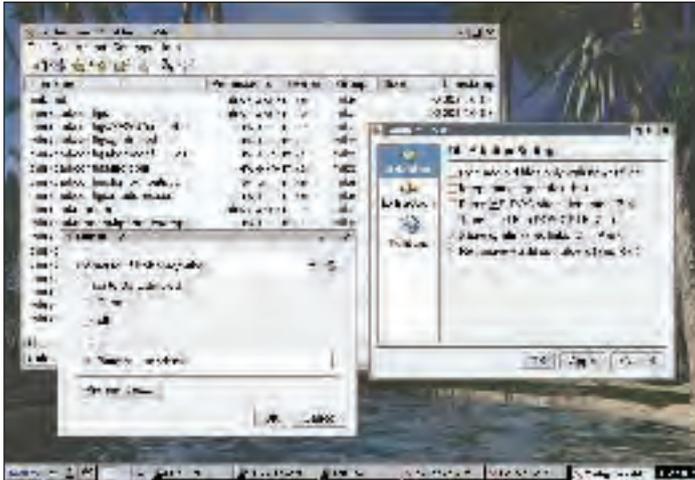
RATING **8/10**



KDE Ark

The animals went in two by two...

■ **VERSION** 2.19 ■ **WEB** www.kde.org/



Ark preparing to extract some files, with Settings box popped up.

Right from the early days of the project, KDE has always been equipped with a healthy bunch of small desktop applications and utilities. *Konqueror* allows for seamless

navigation inside many compressed archive formats, and another pleasing touch is the ability to create new archives via a right-click. Still, Linux's most popular desktop environment

wouldn't be complete without a dedicated compression tool, and *Ark* does the job very well indeed.

Ark is supplied as part of the *kdeutils* package, so if you're running KDE (or have it on your system) it should already be installed. Failing that, you'll need the appropriate *kdeutils* package, and if you're compiling from source then the development header files and related bits are required.

A simple and sparse layout makes up *Ark*'s front end; it's not overloaded with window furniture and superficial frills, thankfully. *Ark*'s initial layout doesn't scream out "Use me!" though, and an initial dialog offering to create or open an archive (if called without any filename arguments) would make it a little more approachable for newcomers. Also, when calling

`ark <filename>` with a large archive, the UI freezes as it's opened – a progress bar would surely be a better option?

New archives can be created either via the menu or by dragging and dropping files from other KDE applications. *Ark* handles a satisfyingly comprehensive range of formats: along with *.tar.gz/bz2*, *Zip* and *rar*, more

exotic formats including *Zoo*, *Lha* and Java's *.jar* are also catered for. As we expect from KDE apps, individual files can be plucked out of archives and dropped anywhere, and on the whole it works commendably well with nifty pattern matching for extracting files.

There's a handful of settings to fine tune aspects of certain formats and default folders locations, and the documentation is suitably thorough (but could do with a sprinkling of screenshots here and there). *Ark* is definitely the best choice for archiving and compression jobs under KDE – it may not have the most flashy UI, but it's reliable, a decent performer and supports all the common file formats.

LINUX FORMAT VERDICT

FEATURES	8/10
EASE OF USE	9/10
DOCUMENTATION	8/10
PERFORMANCE	7/10

KDE's top-notch answer to *File Roller* supports a respectable range of file formats, is straightforward, accessible and pleasant to use, despite the odd UI glitch.

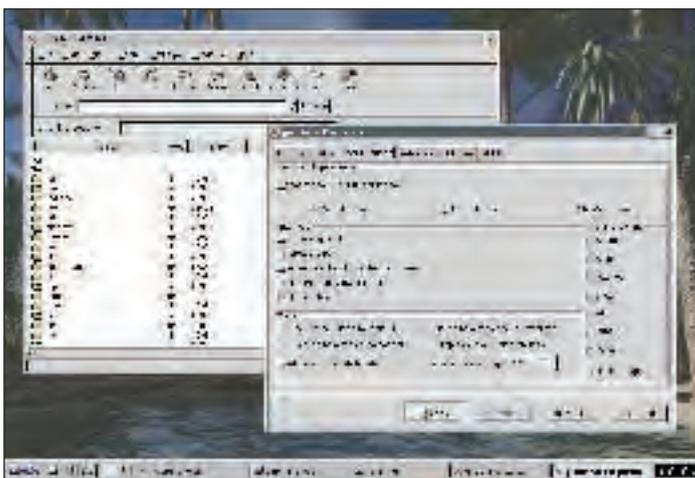
RATING **8/10**



Gnoarchive

Ideal for older installations, with plenty of potential.

■ **VERSION** 0.7.0 ■ **WEB** <http://gnoarchive.sourceforge.net/>



The configuration window is packed to the gills with settings to play with.

Thinking up a name for an Open Source application is always great fun, particularly when you want to indicate that it's part of a larger project. We have

Gnumeric and *Kontakt* as two of the most striking examples, and *Gnoarchive* (previously known as *Gnomerarc*) continues in the quest to prefix

everything with K- or Gno-. Recursive acronyms seem so passé these days...

Building *Gnoarchive* from source could be tricky with modern distros, as it's coded around GNOME 1.4 libraries. Many distros include legacy packages for supporting older GNOME apps – both versions of the desktop can live side-by-side – but if you're out of luck you can always pluck out the relevant packages from an older distro release and try to retrofit them in.

Gnoarchive's cluttered main window sports the standard GUI upholstery with a toolbar, location entries and status line. The retro-look of the file browsing component feels fairly dated when compared to the snazzier GNOME 2 and KDE 3 equivalents; appearance isn't everything though, and at least a great deal of file information is readily at hand. Happily, *Gnoarchive* uses tabs when working on multiple archives – a feature we'd like to see implemented in the others.

Some menu entries and dialogs are not totally functional in *Gnoarchive* at present, with toolbar buttons working much more effectively than the menus. Yet it stands strong in its format range,

with *Zoo*, *Lha* and *Rar* available alongside the regulars (depending on external utilities already installed). Default paths for these utils are mostly in `/usr/local/bin`, but can be changed through the detailed configuration window. Many aspects of the UI and plugin system (seemingly unused) can be modified, and as new code is rolled in, its versatility will please power users.

Ultimately, *Gnoarchive* only loses points as a result of its unfinished state – it's certainly usable, and has some lovely touches; but the occasional bug pops up and its interface style doesn't meld well with modern desktops.

LINUX FORMAT VERDICT

FEATURES	6/10
EASE OF USE	5/10
DOCUMENTATION	8/10
PERFORMANCE	7/10

Worth trying if you're an experienced user running older GNOME 1x installations, but suffers from some incomplete features. The plugin system will be great when finished.

RATING **5/10**



COMPRESSION SOFTWARE THE VERDICT

Compression tools aren't the most complicated programs: there's often very little to pick between their featuresets. They're not 'killer apps' in the same way that word processors and Web browsers are, but nonetheless they're still a vital part of any modern desktop setup. Being able to create and extract archives in a broad range of formats is a must for anyone who spends a lot of time downloading – it's not uncommon to find different formats being used for different file types and Internet sites.

Bearing this in mind, *File Roller* and *Ark* come up trumps in providing inbuilt support for an exhaustive range of formats; it's true that *Bloat* deserves to be up there too but it needs other extras to be installed as well. Your choice of the polished apps, be it *File Roller* or *Ark*, will be influenced by your use of GNOME or KDE – while both will run fine under either desktop, they work best in their native environments.



File Roller is bundled with well-written and informative documentation.

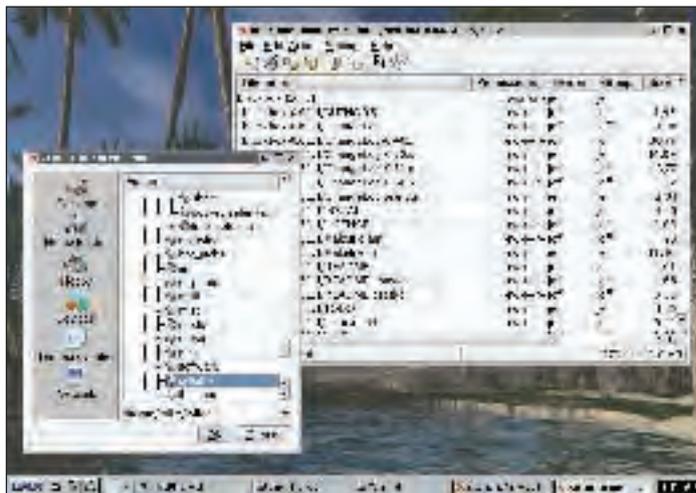
File Roller just pips *Ark* to the post, although the differences are nearly negligible, as they're both approachable, cleanly designed and backed up with good format support. However, *File*

Roller has a slight edge in featureset and interface design (eg not freezing when opening large archives), and for a GNOME app it's a zippy (pun intended!) performer to boot. Power users may

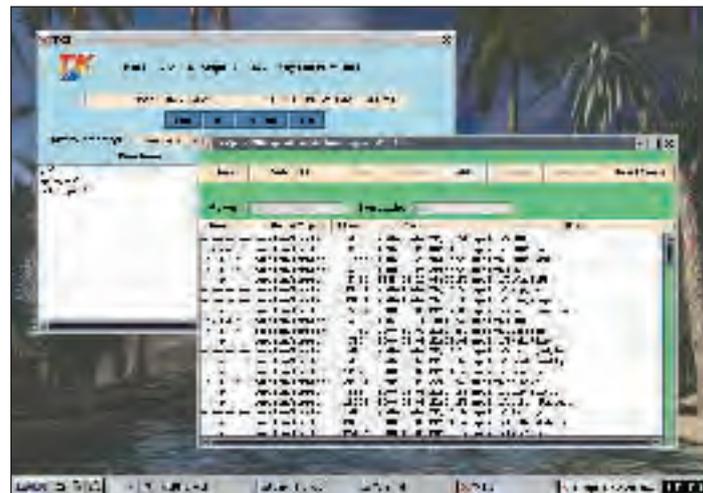
prefer *Ark*'s greater number of settings and options, though.

TkZip and *Bloat* are still worthy of attention. The former is quite weak in terms of features and user interface, but it's orders of magnitude lighter and on older machines it makes a world of difference. *Bloat* isn't ideal for novices, but those who want to explore the shell prompt will find it a lifesaver at times – there's no need to remember individual programs and command-line options.

In the table below, we've listed the number of file formats each program supports, their memory usage (gauged when started, as loading files would skew the results) and other goodies that may be available. *Filtering* refers to the ability to select certain files for extraction based on a pattern, *Password* indicates whether or not the software will open passworded files (which in turn depends on file format) and *Testing* denotes applications which allow an archive's integrity to be confirmed. [LXF](#)



Ark strutting its stuff with KDE's outrageously smooth Plastik theme.



Argh! Luckily, you can change *TkZip*'s colour scheme to something saner...

TABLE OF FEATURES

NOTE: Figures for Slackware 9.1 on a 1.2GHz IBM ThinkPad, 384MB RAM. Memory usage varies from system to system, and is affected by program configuration and other factors.

Name	License	Toolkit	Package size	Memory usage	Formats	Filtering	Password	Testing
<i>Ark</i>	GPL	Qt3	In <i>kdebase</i>	15MB	16	Yes	No	No
<i>Bloat</i>	GPL	Text	37KB	N/A	20	No	External	External
<i>File Roller</i>	GPL	GTK2	1.2MB	11MB	14	Yes	Yes	No
<i>Gnouchive</i>	GPL	GTK1	620KB	5MB	8	No	External	Yes
<i>KArchiver</i>	GPL	Qt3	870KB	16MB	8	No	No	Yes
<i>TkZip</i>	Shareware	Tcl/Tk	175KB	4MB	7	No	No	No

HotPicks

The best new Open Source software on the planet!



Mike Saunders

A coder since Amiga times, Mike's a Linux and BSD guru.

This is the place where we get to profile some of the hottest software around. Each month we trawl through the hundreds of Open Source projects which are released or updated, and select the newest, most inventive and best for your perusal. Most of the *Hot Picks* are available on our coverdiscs, but we've provided web links if you want to make sure you have the very latest version.

If you have any suggestions for software that you think we should cover, email us at linuxformat@futurenet.co.uk with "Hot Picks Request" as the subject-line, or contact us by post through the address on the *Mailservers* pages.

HOT PICKS AT A GLANCE

MagicPoint	42
LiarLiar	43
Colorize	43
SquashFS	44
DOSEMU	44
Squirrel	45
Gtkboard	45

HOT PICKS 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!



PRESENTATION SOFTWARE

MagicPoint



■ VERSION 1.10a ■ WEB www.mew.org/mgp/



MagicPoint presentation demonstrates an embedded XTerm/Pico in a slide.

When you're trying to explain a new business strategy to an assembled group of 200 people, or give an easy-going tutorial in step-by-step format, presentation software can work wonders. Of course, anyone who's been subjected to lengthy meetings and lectures will be all too aware of mindless abuse of these tools – buzzwords, clichés and laughably redundant image effects all make their way into many presentations. Still, Linux is doing pretty well on this front, with *OpenOffice.org Impress* leading the pack, with *KPresenter* and *Agnubis* following closely behind.

MagicPoint is similar to those tools in terms of the end result, but goes about it in a different way. Most significantly, instead of relying on a graphical interface to create the files, *MagicPoint* works with its own plain text file format. As author Yoshifumi Nishida states: "it is designed to make simple presentations easy while making complicated presentations

possible." Naturally this won't be attractive to everyone, but constructing a presentation in this manner isn't as hard as it seems...

We had a brief look at *MagicPoint* way back in issue 16, and as it's moving along well and gaining popularity: *LXF* reader Dan Thomas suggested that we give it some more coverage. Compiling and installing *MagicPoint* is a doddle, with no esoteric dependencies. It's best to sort out some high quality fonts beforehand though, and if you're running an older distro you'll need to make sure *Freetype* is up to scratch too

Fancy features

MagicPoint's coherently designed markup language, demonstrated in the supplied samples, consists mainly of directives preceded by a percentage (%) sign and followed by a value. For instance, an `%image "tux.jpg"` line will display a picture in the current slide. The most immediate way to learn the tags is by scanning through

the sample presentation; it's all self-explanatory and the commands are easy to remember. Among *MagicPoint's* catalogue of presentation effects are the usual text alignment and colour attributes, bullet points (created with tabs in the file), cut-ins (new slide moving across the previous, although way too fast on modern boxes), pauses, colour gradients for backgrounds and various other bits and bobs. Busy slides can be pre-cached too for better performance. It may not have all the snazzy special effects of *MS PowerPoint* and its ilk, but there's plenty here to stop your audience from dozing off.

Perhaps *MagicPoint's* most prominent and celebrated feature, though, is its ability to embed the output of command line operations and standalone X applications in a slide. This is enormously useful for folding animations and other external goodies into a presentation to give it extra spice, and with the `-o` flag (which cancels fullscreen mode and lets the window manager do its work) there's nothing better for demonstrating hands-on usage of a program. Note: if you have trouble getting this feature to work, enable the `-U` flag and switch to a different WM, as they don't all work correctly here.

Other neat features include the ability to toggle on a virtual laser pointer with the `X` key, allowing the mouse to be used for drawing over the current slide (especially useful for highlighting vital points), and you can produce HTML output for viewing on the Web. See <http://puchol.com/cpg/software/mgp> for supplemental templates are available at making a quick slideshow.

MagicPoint is never going to have a huge mainstream impact as long as it involves hacking files by hand, but just 10 minutes of experimentation make it a walk in the park to build up simple presentations in a less stress-inducing than clicking your way through graphical apps. Hopefully at some point we'll see a few GUI front-ends to the file format, but right now it works like a charm and is well worth investigating if you produce presentations frequently and are sick of the traditional tools.

LIE DETECTOR

LiarLiar

■ VERSION 0.5.1 ■ WEB <http://liarliar.sourceforge.net/>

Lies are a never-ending and often disastrous part of our social fabric, and yet they can be life-savers at times. Without them, we'd have a harder time escaping from awkward chores, attracting the most desirable life partners and initiating bizarre legal actions against communities of volunteer programmers. Being able to detect mendacity is a subject that's constantly explored; while there are few obvious signs to look out for, proper lie detector software is perhaps the best bet.

LiarLiar is a 'voice stress analysis (VSA) tool' which examines sound data and attempts to spot signs of stress – hopefully letting you know that your subject is telling a whopper. It's not too difficult to build from source, providing you have *GTK*, *Gtkmm* (the C++

bindings for that toolkit) and other common development files. As it uses the *GStreamer* framework, you'll need to make sure that side's all set up correctly before first use.

You can configure *LiarLiar* to read directly from the microphone port, but by default it scrutinises MP3 files for fibs. Loading an MP3 and hitting play will go through the file, constantly monitoring stress levels and giving an indication via the lower bar. On our main test machine, an 800MHz Pentium 3, the visual updates stuttered after a while – it's a known problem, but we found that removing the debug ifdefs around line 76 of the *simple.cpp* plugin gave a textual update. When we tested *LiarLiar* on one of our faster boxes, it worked fine with no problems at all.



100 per cent obvious lie? Cripes. Won't harm my career in politics though...

So, the big question: can it accurately detect lies? Unfortunately, unless you have access to all kinds of stealthy recording equipment, your subject will be all too aware of the proceedings. We coerced three people (two men and a woman) to speak truths and lies into the mic; in some cases the tool appeared to recognise

real corkers but on the whole the results weren't that conclusive.

As they say, your mileage may vary – nevertheless, you could find it to work better with quality microphones. It's an intriguing little project, but it's wise to equip yourself with asbestos underwear because your pants may, quite literally, catch on fire.

LOG FILE TOOL

Colorize

■ VERSION 0.3.4 ■ WEB <http://colorize.raszi.hu/>

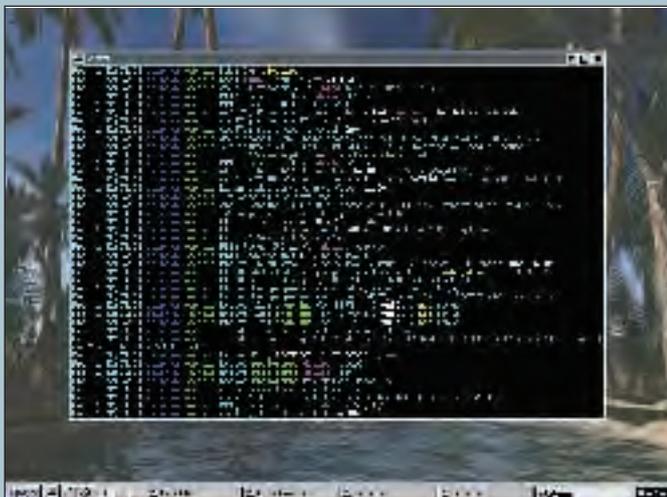
All competent server admins spend a fair bit of their time poring over log files, looking out for any unusual behaviour and making sure everything's running smoothly. However, the plain text format of most logs can tend to be rather headache-inducing after a while – additionally, sometimes it's all too easy to miss a crucial line when you're scanning through large amounts of data. As a result, various tools exist to improve readability of log files, and *Colorize* is a clean and simple option.

You'll need Perl and its *Term::ANSIColor* module to run the program; these are installed by default on all major distros so you should have few trouble getting it running. The installation process is equally simple – just copy the *colorizerc* file (which defines colours in a friendly format) into

`/etc` and then the main script anywhere in your path.

For basic usage, piping a log file to the main script will churn out a more attractive rendition without too many bold or flashing retina-burning attributes, as is happens all too commonly in many text editors that have a colourise option. In a web server log, for example, each field is coloured uniquely and makes it appreciably easy to scan through referers, return codes or anything else. *Colorize* also works pleasingly well with various other formats, including secure and messages etc. in `/var/log`, and tips are provided for properly integrating the tool with **syslog** and making a prettier **tail(1)**.

As an added bonus, *Colorize* can generate HTML output with the **-h** flag – it works well and renders decently in *Links* and *Gecko*, but the files that are produced with the **-h** flag tend to



The `/var/log/messages` output in lovely multicoloured form.

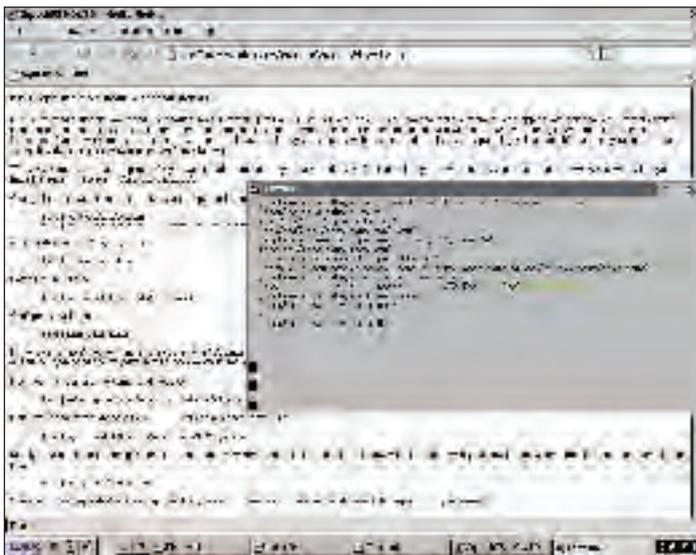
be on the huge side, with a copious amount of `<font color=` tags. Editing the script to add new words or change its behaviour isn't too difficult providing you have a smattering of Perl knowledge, and the code is fairly clean and well-commented for the most part.

Even though it's essentially just a single-purpose and simple program at heart, *Colorize* has certainly been implemented well (there's also a version C for improved speed) and any program that makes an admin's life just a little bit more pleasant is always A Good Thing™.

COMPRESSED FILESYSTEM

SquashFS

■ **VERSION** 1.3-r3 ■ **WEB** <http://squashfs.sourceforge.net/>



No messing around with *tar* and *gzip* – *SquashFS* does it all automatically.

The most popular Linux filesystems (*ext2*, *ext3* and *ReiserFS* et al.) are engineered with performance and stability in mind. This is fine for desktop PCs, but poses problems on embedded and low-spec kit, and is wasteful if you're archiving large amounts of data. *SquashFS* aims to provide a transparently accessible compressed filesystem for Linux using *Zlib*.

There are two main components to *SquashFS*: a patch for the kernel to enable support either built in or as a module; and the toolset used for creating and manipulating *SquashFS* mounts. Supplied in the tarball are a handful of patches for various kernels, but note that you'll probably need the vanilla kernel source to patch correctly, as most distro-supplied kernels have already been hacked and tweaked.

With support compiled in and the *mksquashfs* binary built and installed, setting up a new compressed filesystem isn't too taxing. *SquashFS* uses the loopback driver to mount filesystem images onto mount points; this is similar to browsing an ISO image without burning it to disc, and to create

a new filesystem image you simply:

```
# mksquashfs directory/ name.sqsh
```

This produces a single file representing the directory, and the compression ratio will depend on the type of files (eg MP3s and JPEGs won't shrink a great deal, unlike plain text and bitmaps). Directories, data and inodes are compressed, and to remount use

```
-t squashfs -o loop
```

on the command-line. On our test machine (800 MHz, 320MB RAM), *SquashFS* took 16 seconds to write an 80MB directory as an image, bringing the size down to 25MB – not too shabby at all.

Although the mounted filesystems are read-only, there's an option to append data to existing filesystems, improving the tool's flexibility. *SquashFS*'s coders hope to have the code rolled into the 2.7 kernel development tree – we hope so, because it proved to be robust and reliable in our testing. *SquashFS* is hugely useful if you're running an archive server, creating a trimmed down floppy/USB key distro and need as much space as possible, or just wanting more room to store old files.

DOS EMULATOR

DOSEMU

■ **VERSION** 1.2.0 ■ **WEB** <http://dosemu.sourceforge.net/>

Eagle-eyed readers will recall our look at *DOSBox* in issue 48's *Hot Picks*. The DOS emulation scene is going strong, fuelled by nostalgia for classic games and the occasional need to run legacy apps, and as PC performance improves it's becoming much more pleasant. *DOSEMU* is much older than *DOSBox* (early releases dating back to 1993) and has one outstanding benefit: it uses the CPU's virtualisation features rather than emulating every instruction. In practice, this leads to much better performance, but there is a downside to this method: it only runs on x86 boxes.

While the build and installation process doesn't require any special somersaults, configuring *DOSEMU* to work correctly can be horrendously complicated. There are user config files, system config files, built-in config files... We hoped this new release

would bring along a saner setup procedure, but at least these files are well-commented and easy enough to navigate. One solution if you're close to punching the monitor, though, is to remove `~/dosemu` and run as root – this avoids any restrictions.

You'll need a standalone version of DOS to run the emulator; the developers suggest *FreeDOS*, which does a sterling job, but for increased game compatibility you may want to use a full MS-DOS release (or Windows 95-made boot floppy). `dosemu.conf` allows you to tweak various speed, memory, display and sound settings (check out the *hogthreshold* option), and a carefully tuned setup will work wonders.

As with any emulator, compatibility is the key issue and *DOSEMU* does well on this front – with support for *DPML*, *EMS* and *XMS* it can cope with a good range of titles, and *Doom*, *Duke*



Extreme violence in a *Microprose F1GP* session (also showing config file).

Nukem 3D and similar action games work fine. Two of our faves, *Worms* and *Microprose F1GP*, ran without any major glitches.

Current *DOSEMU* users will find a plethora of new additions since the last stable release, including fullscreen

and joystick support, improvements to sound and performance, among others. In all, *DOSEMU* isn't quite as polished as *DOSBox* but runs orders of magnitude faster, so it's best to try your program here first and fall back on the latter if it fails.

PROGRAMMING LANGUAGE

Squirrel

■ **VERSION** 0.9-alpha ■ **WEB** <http://squirrel.sourceforge.net/>

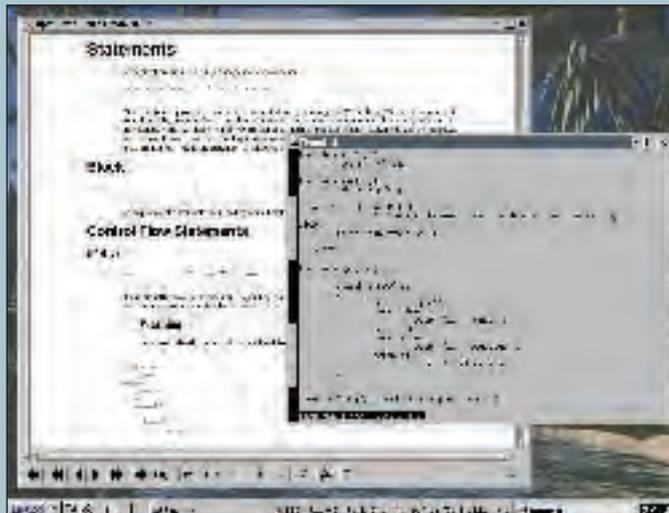
Programming languages are up there amongst text editors and desktop environments as constant fuel for flamewars, with zealots nailing their flags to particular masts despite the obvious futility of the whole situation. C vs C++ and Perl vs Python battles are often severely aggressive for instance, and when you've been coding in a language for many years, it's easy to develop an emotional attachment. Conversely, you could well be utterly sick of the thing, so new languages are always worth investigating to check out their technical merits and capabilities.

Squirrel is an object-oriented scripting language with a strong focus on fast running speed and, in particular, games. Nevertheless, it's capable of a lot more, and has been influenced by Python, JavaScript and

Lua. *Squirrel* runs on Linux and Windows – there's not much in the way of install documentation that would be helpful to newcomers, but a single **make** will produce the main *sq* binary for installation in your \$PATH.

Source files for *Squirrel* generally use the .nut extension (can you see a theme emerging here?) and at first glance the language shares many superficial similarities with C/C++ (and a touch of Perl into the bargain). Blocks are contained with in curly braces, statements and variable definitions end with a semicolon, and on the whole it shouldn't be too challenging for intermediate C coders to get started with.

With support for dynamic typing, exception handling, tail recursion, auto memory management and other goodies, *Squirrel* is undoubtedly up-to-steam for many 'serious' programming



Squirrel's PDF manual sitting alongside an example of some code.

tasks, and its standard library includes a decent range of math, system and IO functions. In terms of documentation, the supplied PDF file provides a solid overview of the language's workings with details on embedding *Squirrel* in another application.

Naturally, at this stage you wouldn't want to attempt a major project in *Squirrel* – the developers have yet to finalise the featureset but there

shouldn't be any drastic changes from here onwards. When 1.0 arrives, we'll no doubt start to see more software built with the language. Right now, *Squirrel* will be of most interest to coders who enjoy the ease and rapid development of high-level scripting languages like Perl and Python, but feel more comfortable with the syntax of C, and fancy exploring new pastures.

GAME COMPENDIUM

Gtkboard

■ **VERSION** 0.11pre0 ■ **WEB** <http://gtkboard.sourceforge.net/>

While Linux is starting to make slow but steady progress on the 3D action game front, it's always been well endowed with desktop distractions and coffee break puzzlers. All the usual frill-free desktop games have been implemented in one way or another – including Chess, cards, *Breakout*, *Tetris et al* – usually in the GNOME and KDE bundles. The imaginatively named *Gtkboard* goes one further, though, and attempts to wrap 31 games into a single application (with an eventual goal of 100!)

You'll need to build *Gtkboard* from source, as no binary packages have yet been created, but it's a breeze with the typical **./configure, make** and **make install** routine. It only needs the GTK toolkit to compile (typically *gtk-devel* in most distros) although it'll find SDL for sound and the GNOME

libraries to add other extras if they are deemed necessary.

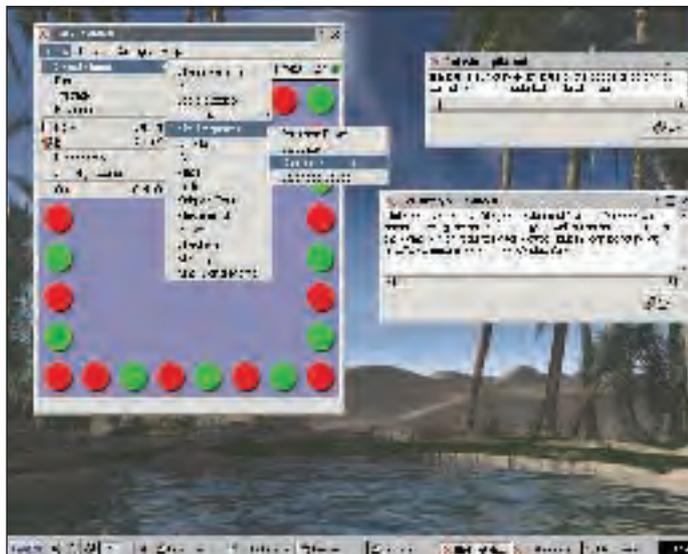
Thankfully, *Gtkboard's* coders have avoided any jazzed-up gimmicks, instead letting the player get straight into his/her favourite game; these are neatly categorised in the main menu, while the last three played are instantly accessible. Joining the well-known pastimes of Chess, Othello, Checkers and *Mastermind* are some more action oriented offerings including *Pacman* and *Tetris*. Most have their own high-score boards, and although a few haven't been totally completed; yet the majority are fully playable.

Before you start wondering about the bloat of packing all these games together, the developers point out that the AI (search algorithms) for many two player board games is shared – separate engines don't need to be

written. Also, *Gtkboard* has been crafted so that additional games can be added via a plugin/API system which results in less wasted code.

One major feature planned for the 1.0 release is networked play, which will really add a whopping layer of crunchy icing on an already sweet (but slightly unfinished) cake, and while

some of the implementations aren't the most lavish ever they still do the trick. If you have a penchant for mini desktop diversions and fancy trying out a few more esoteric games without the hassle of having to compile endless packages and the associated dependencies, *Gtkboard* is well worth trying out. [LXF](#)



More games than you could need in a year's worth of months of Sundays.

AWARDS 2004

It's time to have your say in recognising tremendous achievements over the last year.

Once again, the time has come to reward those associated with the linux community who have given their best with the recognition that can only come from receiving an *LXFi*. Our awards are highly regarded, mostly because they are completely chosen by you, the readers, and therefore by Linux users in general. The website (www.linuxformat.co.uk/awards/) is now open for nominations, but first why not peruse the categories on offer that are listed here, and accept a few tips on what you should be considering.



BEST DESKTOP APPLICATION

This is a hotly contested category that really is open to all sorts of Linux projects. Previous winners include *The GIMP*, but in reality, virtually anything that you can install on a Linux desktop could win this category. We suggest that you should be considering a brilliant new project that has appeared over the last 12 months or so, has recently been open-sourced, or one that has made dramatic improvements over its previous incarnations. Please note that Office and Internet Software both have their own categories.



BEST DISTRO

This is one that always causes a stir. The award usually goes to one of the big three, but with the distro market in flux so much, perhaps anyone could take it. What we've certainly seen is a lot more interest in distros such as Knoppix and Gentoo. Perhaps this year it could all be different.

BEST DEVELOPMENT TOOL

This is a fairly obvious category, and one not limited to open source software. Previous nominations include *JBuilder*, *Kylix* and Intel's *C Compiler* for example. While we've seen new versions of *Anjuta*, *KDevelop* and *Zend Studio* – as well as improvements to the venerable *GCC* – many others have appeared or had an upgrade during the last year, so there should be a fair amount of competition.

BEST HARDWARE SUPPORT

This award is supposed to go to the company or organisation who has provided the best support for using hardware under Linux, whether that concerns printers, processors, graphics cards or whatever. You might think that IBM deserves the award for supporting a huge range of servers on Linux, or maybe nVIDIA or ATI for their graphics card drivers. There's a whole world of choices...

FREE SOFTWARE PROJECT OF THE YEAR

Could it be GNOME, or KDE? Will it be *Mozilla/Firefox/Thunderbird/Waterkitten* or whatever its various bits end up being called? Maybe it

will be the Kernel itself? Or perhaps it will be something that you wouldn't first think of? Either way, get your thinking caps on!

EMBEDDED LINUX AWARD

Won last year by Trolltech for *Qtopia*, this category, it should be remembered, can also include appliances as well as developer tools. So as well as



MontaVista, think about Ayrsoft, Neoteris, Navaho...

BEST LINUX GAME



Another straightforward category, but also one that is often the source of more than a few surprises. Previous winners and shortlist entries have included *Return to Castle Wolfenstein* and *Unreal Tournament 2003*. The demise of Loki Games has left a bit of a gap in the commercial side of the market, but there are plenty of great Open Source games that deserve a mention, and some cracking new releases from Linux Game Publishing for instance.

BEST INTERNET SOFTWARE

A more specific category that has produced some predictable and some more surprising winners in the past. *Mozilla* won last time, but this time you'd have to wonder which bit of *Mozilla* to choose (and try and remember the name, which may include some elemental force and some sort of animal). Nevertheless, you might want to reward *BitTorrent*, *KMail*, *Evolution*, *Konqueror*, *Opera* or a host of others.

BEST SERVER SOFTWARE

Won last year by a resurgent *Apache*, there is no end to the potential winners in this category either. As well as databases such as *Oracle*, *MySQL* and others, there are technologies such as the recently updated *Samba*.



BEST SUPPORT RESOURCE

Previously won by the *Linux Documentation Project* and Google, there are plenty of resources, not just online – like book publishers, for instance – that could be in the running for the award this year. Once again, we're sorry but we won't accept any nominations for our own website and forums, flattering though it may be.

BEST OFFICE SOFTWARE

OpenOffice.org romped home with the top prize last year, but there's plenty of scope for competition here. Perhaps the newly revamped *KOffice* can be a contender, or individual apps such as the excellent *Gnumeric*?

HOW TO VOTE

The process (for our readers anyway) is fairly simple. Firstly we will be collecting nominations through our website at www.futurenet.co.uk/awards/. From the nominations we receive we will construct a shortlist, published on 1 May 2004, and then put a voting slip on the website. All the votes will be verified and counted, and hopefully we'll deliver the results towards the end of Summer. Unlike many other awards, the *LXF Awards* are completely decided by you, the readers, which makes them more prestigious for the winners, if a little more difficult to manage. We do take this responsibility seriously, and would like to remind you that you can nominate whomever you like, but the actual voting process will be monitored quite strictly. Anyone entering multiple votes will have ALL their votes discounted. **LXF**

BEST COFFEE/TEA

A new category for last year, we thought it time that caffeine vendors were given recognition for their tremendous efforts in supporting developers (and magazine teams) through lengthy hours of hard slog. Nominate your favourite fixes now!



BEST WEB HOST OR ISP

Plenty of contenders – including last year's winners UKLINUX.NET – but the result here should be straightforward: which ISP/host has given the world the best Linux service and pricing in the last year?



ENTERPRISE AWARD

The usual suspects line up again to receive the award for best use or promotion/advocacy of Linux in the Enterprise. Will IBM's scary adverts do it for them, or maybe HP's volume of sales? Or Novell's recent purchase frenzy of anything to do with Linux? Or what about the Munich governmental migration?

MAKING MOVIES

Editing video on your Linux home desktop

MAKING MOVIES

Whether you need to add some sparkle to a video of your typically British sunshine-free holiday, or transfer your existing home-movie archives to digital format, **Biagio Lucini** is the man with the megaphone in the Director's chair...

cover feature



One of Linux's most impressive success stories is the recent trend for many of the biggest players in the movie industry to switch from proprietary and expensive solutions to Linux. The backing of major companies like Pixar and LucasArts underlines how easily Linux can bring concrete benefits, as they have found Linux to be at least as efficient as closed-source apps, but also cheaper and

more customisable as well. But what does this mean for the more modest home user, whose greatest ambition is just to clean up the movie of their summer holiday a bit; and maybe store it on DVD? Although both scenarios might seem similar on the surface, the way they are tackled is not: commercial movie production mostly concerns the server and high-end workstation side – where speed is critical above all else, while home video-editing is more desktop-focused, and speed is secondary to usability.

For many, desktop in Linux continues to be a minefield. While we constantly experience improvements and innovations, even the most zealous Linux advocates acknowledge – albeit somewhat grudgingly – that some things are still missing. Rather than getting bogged down here in debating whether Linux is a viable desktop system for home video-

editing, we will be more pragmatic and look at what is currently available and how you can get started using it.

First, we will briefly review how videos from different sources can be imported, discussing both hardware and software issues. Then, we will analyse in detail how to connect your Firewire devices up to your system so you can transfer your creations quickly and without fuss. Next, we will touch on clip editing and outline the features of *Kino*, a popular video-editing application. Finally, we will be looking at how to render your finished results so that you can watch it both on personal computers and on playback devices like DVD players. If you're new to video-editing on Linux, you might find it best to work through this feature step-by-step. Otherwise we recommend you just jump in at the section that interests you most and work from there.

MINIMUM HARDWARE SPECIFICATIONS

Don't resign yourself to watching a busy cursor!

video-editing – or to be more accurate, video rendering – is a very CPU-intensive process, and is really second only to 3D gaming when it comes to high system requirements. For this reason, we recommend that you use a modern computer, preferably with a recent Pentium 4 or AMD64 processor, and 512MB of RAM at the minimum. You *may* be able to get away with less, but you have been warned!

Owing to the need for various scratch files during a processing job, a big hard disk is essential. As a rule of thumb, one hour of video takes up about 12GB of space, and the same again is needed for processing. The hard disk should also be reasonably fast, since Firewire can reach a transfer rate of 400 Mbits/sec.

ANALOGUE VIDEO-CAPTURE

When you start to edit your clips, the first problem you will be faced with is how to get the video from your capture device on to your PC. Video can come from different sources: a camcorder, a digital camera, a webcam, etc. Sadly, the techniques for grabbing images from these sources is different in each case.

Usually, Linux will treat digital cameras as nothing more than USB storage media, which means that to get movie files from them is a simple matter of executing this as root

```
# mkdir /mnt/camera
# modprobe usb-storage
# mount -t auto /dev/sda1
/mnt/camera
```

If you have other SCSI drives installed, or even USB storage, you may need to use **sdb1**, **sdc1** etc, instead of **sda1**. Unless you keep a mind-map of your device structure, this may be down to trial-and-error. To make your life a bit easier in future, you may want to add an appropriate entry in `/etc/fstab` and

also have the `usb-storage` module loading automatically at boot time, so that access to the camera is possible for normal users. If you're using a modern distribution where storage media is dynamically mounted when detected, you shouldn't need to do any of this – Mandrake, for example, will do it all for you. If you have problems getting your system to recognise the camera as a USB storage drive, you may have better luck if you opt to use *gPhoto2* (www.gphoto.org) to retrieve the files from it.

Digital camcorders can usually be connected to a PC via a USB or a Firewire (IEEE 1394) port. When a choice is given, we recommend Firewire, as it is an Open standard and hence is very well supported in Linux. USB, on the other hand, while often eminently suitable for hotplugging, is usually based on proprietary transmission protocols, so you will likely find it more problematic to get working.

Some digital camcorders also have

an S-Video port, and this gives you the additional option to connect the camcorder to a capture device. Although this might be more familiar to some, it isn't recommended because better results are usually obtained using Firewire. Among

“Pixar and LucasArts have found Linux to be at least as efficient as closed-source apps, and cheaper and more efficient too.”

capture devices, the most popular ones are entry-level TV-Tuners, such as Hauppauge's *WinTV*. Many PCI TV-Tuners, and these are supported by Linux via the *bttv* driver, which is compiled by default in most modern distributions. If you're using Mandrake, you can use *DrakxTV* to configure these devices, but other distros will have something similar.



VIDEO FILE FORMATS

XviD, DivX, MPEG, AVI *et al*

A PC works with files. So, if you want to modify movies and clips, you must first get them into the computer in the form of a file. Moreover, this file must be in a standard format so that the editing and viewing software can easily read both the video and audio streams.

As raw video would require huge storage resources, a number of techniques exist that allow you to compress video images without significant loss of quality. The most popular of these are based on the natural limitations of the human eye, which perceives a continuous sequence of still images to be 'moving'. This is most convincing at a rate equal or greater than 25-30 frames per second. The first video compression technique was MPEG-1 (1993), which has survived as the official format for VCDs. The most modern successors of MPEG-1 are MPEG-2 (used for DVD) and MPEG-4, of which DivX (closed source) and XviD (Open Source) are popular implementations.

The audio that is associated with moving images is less of a problem, since the information does not require as much space. Hence, it would be possible to

deal with raw audio (in fact, the DV format used in camcorders stores the raw audio). Audio compression techniques have a less perceivable loss of quality too, and both Ogg and MP3 formats are very popular.

It's a wrap!

Having sorted out the problem of storing audio and video in a comprehensible format, the following step would be to package them in such a way that the images and corresponding audio are associated. A *de facto* standard in this area (because it is widely used in the Windows world) is AVI. AVI (Audio Video Interleave) is not actually a format, but a container – it specifies how video and audio are packaged together, not what their formats have to be. This in turn means that there are no restrictions on the video and audio formats you can use. The AVI format is interleaved, which means that a segment of audio data follows immediately a segment of video data. Currently, AVI is very well supported in Linux, thanks to the efforts of contributors to the *Linux AVI file library* (<http://avifile.sourceforge.net/>); and it supports a comprehensive list of

compression formats for audio and video. The most popular competing formats are Apple's Quicktime (native for *Cineerra*), Real's Real Media and Microsoft's Windows Media.

As long as one can easily convert movies from one format to another, the actual format of the file does not really matter that much. What is needed for this operation is an encoder, *ie* a program that accepts as an input a file in the first format and gives in output a file in the other format. Programs that convert a file back into images are called instead decoders. A good encoding tool supporting the most common file formats is *Transcode*, with playback being done in both *MPlayer* and *Xine* – both of which allow reproduction of almost any video file.

Usually the DV format that digital camcorders output is packaged as an AVI file, with the video encoded in an almost lossless format (DV). DV can be of type 1

and 2 (DV1 and DV2, respectively). The difference between the two is that that in DV2's case the audio is stored together with the video, whereas in DV1 they are stored separately. *Kino* can understand both DV1 and DV2, while *Transcode* and *MPlayer* will understand only DV2 and *MainActor* will save only in DV1. A close relative of DV is *M-JPEG* (<http://mjpeg.sourceforge.net/>), mainly used in semi-professional capture devices.



MAKING MOVIES



In any case, it should be enough to execute (as root)

```
# modprobe btvtv
```

Things are different for the ATI All-in-Wonder cards, which are supported by the Gatos project (<http://gatos.sourceforge.net>). Part of this project is also *AVview*, a TV-watching and movie-capture application specifically for those cards.

“Most modern distros will dynamically mount your storage media when detected, saving you a lot of hassle.”

If you're a laptop owner, you will probably be more interested in USB TV-Tuners. These are mostly supported by the *usbvision* driver (<http://usbvision.sourceforge.net>), though the configuration in this case is slightly more involved.

Capture devices

In capture devices, the hardware issues are kept separate from the software issues by the *Video4Linux* application layer (more information at www.thedirks.org/v4l). The *Video4Linux* API is a thin layer between the two that allows software applications to capture and use images from any device controlled by a *Video4Linux* compatible driver. Popular applications of this type are *XawTV* (<http://bytesex.org/xawtv/index.html>) and *TVtime* (<http://tvtime.sourceforge.net>).

Low-specification TV tuner devices usually have software in place to convert the signal from analogue to digital. Conversely, high-end capture cards regularly include special hardware to encode the analogue signal directly to MPEG1 and MPEG2 format, which makes them a great deal more efficient, even allowing you to use them as part of a PVR device. Through projects like *Freevo* (<http://freevo.sourceforge.net>) and or *MythTV* (www.mythtv.org), lower-end cards can be turned into powerful PVR systems with just a little tweaking. Even without these two, it is still possible to record the images to the hard disk with video-editing



TVtime is a very capable application for displaying images from Video4Linux devices, but there are plenty of alternatives.

applications like *Cinelerra* or simple TV-watching applications like *XawTV*, both of which can capture and record images from a *Video4Linux* device.

Many webcams can also be treated as *Video4Linux* devices, which means that pretty much all you've read so far about capturing images from a *Video4Linux* device also applies to webcams. Owing to the relatively high number of chipsets and corresponding kernel modules, it is difficult to give

general guidelines on how to set up a webcam in this limited space – see our special *Get All Your Hardware Working* feature in *LXF51* for more on this. As usual, you can often find quick hints and tips from other users by looking up the manufacturer and model on Google. Although if you like tweaking with your system internals and *really* want to get your hands dirty, you can root around in the USB section of the kernel configuration.

LINUX VIDEO-EDITING APPS

There is a bias towards *Kino* in this feature since it's generally regarded as having a gentle learning curve for newcomers to video-editing. As you scale the ranks of video mastery, or if you find you have particular needs that *Kino* doesn't fulfil, there is a veritable feast of other applications available for you to try out, some of which are included in the Magazine section of this month's coverdiscs. There's a brief introduction to the more mature and stable apps in the boxes in this feature that have a red top and clapper-board like this one.

As the field of video-editing in Linux is evolving rapidly, it's possible you might find some applications a little hard to install and use. However, two programs not examined here that you might want to look at are:

Kdenlive
www.uchian.pwp.blueyonder.co.uk/kdenlive.html

A video editor that aims at integration with the KDE desktop.

LVE
<http://lvmpeg.sourceforge.net>

An application with an innovative look and great ambitions.



IEEE 1394 AND LINUX

Digital camcorders are increasingly popular, so the chances are you will need to know how to use them at some point. The most popular method for transferring data to and from camcorders is over an IEEE 1394 port, most commonly known as Firewire. If you have a Sony machine, the chances are you it will be branded as 'iLink', but internally it's identical. Irrespective of what it's called on your system, IEEE 1394 is a standard data transmission protocol that allows to reach a transfer speed of 400 Mbits/sec.

Not many entry-level PCs come with an IEEE 1394 port as standard, since it is often only needed by professionals. As a result, IEEE 1394 ports are more common in high-profile workstations and laptops. However, a PCI expansion card can be fitted in any desktop system with a free PCI slot, and only costs around £30. Laptop owners can opt for a PCMCIA expansion card, but do make sure that the card you choose is Linux-compatible. The IEEE 1394 project for Linux (www.linux1394.org) ensures support for Texas Instruments PCILynx/PCILynx2 and OHCI-compatible chipsets. As the latter is by far the most common, it is very simple



to shop for a supported IEEE 1394 board. We will assume from here onwards that the IEEE 1394 chipset is OHCI-compatible.

With the hardware in place, the first thing to do is to check whether it has been recognised correctly. By issuing the command

```
lspci
```

from the root account, the output should contain a line like

```
02:01.2 FireWire (IEEE 1394): Texas Instruments PCI4451 IEEE-1394 Controller
```

Once the hardware has been correctly identified, you are just one step away from making it do useful things. The only thing holding you back is if you have compiled your own

custom kernel. Modern Linux distros usually come with plenty of modules precompiled, and the IEEE 1394 ones are no exception. However, if you made your own custom kernel, it's possible you skipped over the IEEE1394 configuration tab without setting it up. The box *Kernel Configuration* below shows the settings that are required for connecting a digital camcorder to a Linux-powered system.

The next step is to load the required modules. Again, modern distributions do a great job at hardware discovery. For instance, with Mandrake, our OHCI-compliant controller was found and configured, with the corresponding modules loaded.

Though often not easy to identify, some laptops do have a built-in IEEE 1394 port.



KERNEL CONFIGURATION

Getting your hardware working

The IEEE 1394 subsystem has been a part of the Linux kernel since the 2.3.40 release, and a backport exists for the 2.2.x release. Hence, it is very likely that your kernel don't need patching; but if it does, it is possible to download the most recent version from the new (no longer SourceForge) Linux IEEE1394 project site at www.linux1394.org/.

Here, we will focus only on the IEEE 1394 subsystem configuration, making the explicit assumption that the reader knows how to compile a Linux kernel, so please refer to the official documentation for more general information.

Although it does not cause any major problems, the subsystem IEEE 1394 is still tagged as experimental. For this reason, to be able to access the appropriate configuration page, we have to select **Y** for the option "Prompt for development and/or incomplete code/drivers". Then, we can move to the

IEEE 1394 part, accessible directly from the main configuration window (if you are configuring with **make xconfig** or **make menuconfig**).

Next, we must make sure we include the general IEEE 1394 support in the kernel (statically or as a module) and support for the specific chipset. Assuming that the chipset is OHCI-compatible and that we want to build the subsystem as a set of modules, we will select

```
[m] IEEE 1394 (FireWire) support (EXPERIMENTAL)
```

```
[m] OHCI-1394 support
```

For the specific use that we want to make of IEEE 1394 (ie connecting a camcorder), the following options must be selected as modules as well:

```
[m] OHCI-1394 Video support
```

```
[m] OHCI-DV I/O support
```

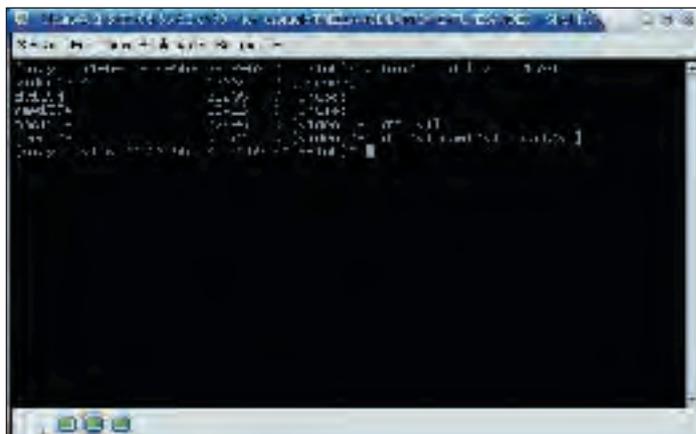
```
[m] Raw IEEE1394 I/O support
```

This setup described here is shown in the screenshot on the right.



The configuration window for the IEEE 1394 subsystem.

MAKING MOVIES



The output of the command `lsmod | grep 1394` shows that in this case all the modules needed for capturing images from a camcorder have been correctly loaded.



If your distro has auto-detection, running the command

```
lsmod | grep 1394
```

as root will show the modules `ieee1394` and the chipset-specific module, which in our case was `ohci1394`. If the modules are missing from the output, they can be loaded manually with this command:

```
# modprobe ohci1394
```

TRANSCODE

The Swiss Army knife of video-editing



Transcode (www.theorie.physik.uni-goettingen.de/~ostreich/transcode) is an excellent audio- and video-editing tool that utilises all the power of the command-line. If correctly configured, *Transcode* supports virtually all the video/audio formats on the market, both in input (decoding) and output (encoding). For this reason, it is particularly useful when you want to work with a file with two different tools that do not support a common format.

Suppose for instance, that we have the file `viddv2.avi` imported with *Dvgrab* directly from our camcorder, and that we want to convert it to `DivX5` for easy sharing with our friends. With *Transcode*, this conversion can be performed by executing this somewhat tortuous command:

```
# transcode -i viddv2.avi -x dv.avi -I 1 -C 1 -z -k -o viddivx5.avi -y divx5
```

By way of explanation, the options we have used above can be broken down as follows:

- `-i viddv2.avi` specifies the input file
- `-x dv.avi` specifies the format of the input file – this is often not needed, since *Transcode* is usually able to automatically detect the format for you, but sometimes may not.
- `-o viddivx5.avi` specifies the name of the output file
- `-y divx5` specifies the format of the output file

The other options are more technical, but no less useful:

- `-I 1` de-interlaces the video, for smoother transitions among images
 - `-C 1` enables antialiasing;
 - `-z` reverts the video upside down;
 - `-k` inverts red and blue.
- In particular, the first two options guarantee optimal vision of the image, particularly on TV screens and non-TFT monitors, while the last two are there to avoid incompatibilities between *Transcode* and some other programs. These should only be used in instances where the video looks weird as a result of these inconsistencies – that is, you'll know when you need them!

The simple example here only gives a brief glimpse of what *Transcode* can do: with its modular design, the ability to manage and convert many clips, and the possibility of separating audio from video and to take the audio from an external source (eg an MP3 file), *Transcode* has something for everything. As it also has the ability to resize images among its many other useful features, *Transcode* is simply indispensable for manipulating movies of all kinds. As usual, the man page explains in details all the supported features.

If you don't fancy more console work, you can try *Gtranscode*, available from (www.fuzzymonkey.org/newfuzzy/software/gtranscode).

To get the camcorder working, you will have to load the modules `raw1394`, `video1394` and `dv1394` in the same way. These operations can be automated by creating a simple initialisation script like this one:

```
#!/bin/sh
echo " "
echo "Loading IEEE 1394 modules: "
for ieee1394mod in ieee1394
ohci1394 raw1394 video1394
dv1394; do
    modprobe $ieee1394mod &&
    echo "$ieee1394mod loaded"
done
echo " "
```

You may want to have that script run at boot time to save you running it yourself manually each time that you need to use it.

The interaction with the camcorder is managed by the following libraries:

- *libraw1394* – provides an abstraction layer of the IEEE 1394 protocol to various applications;
- *libdvc* – supports camcorders;
- *libdv* – handles the DV format used by camcorders;
- *libavc* – controls playback functions of the camcorder from the PC.

The corresponding packages should be included in your distribution. If not, all these projects are hosted at SourceForge, so getting them is just a matter of paying a visit to the site. For instance, the latest version of *libdv* is available to download from <http://libdv.sourceforge.net>.

At long last, you can now plug your camcorder in! You will of course need a IEEE 1394 cable, which, unless you have bought a top-of-the-line camcorder, is probably sold separately. Before you rush out and grab the first one you see, you should know that IEEE 1394 cables exist in three types:

- four-pin to four-pin
 - four-pin to six-pin
 - six-pin to six-pin
- For connecting a camcorder to a computer, usually the four-pin to four pin variant is needed, although you should read your camcorder's manual first. As this is *Linux Format* and not *Plugging-In Firewire Cables Monthly*, we'll assume you can handle the task of connecting your camcorder to your PC. However, do be careful to make sure that your camcorder is switched to playback mode. To double-check whether your setup is in fact correct, you can use *Gscanbus* (<http://gscanbus.berlios.de>) by executing the command

```
# gscanbus &
```

If all is well, the IEEE 1394 connection tree ending with our camcorder should be displayed as shown in the example screenshot below right.

Gscanbus should work if executed as root and, with a correct setup, will also work when executed with a normal user account. If problems arise, one of the very first things to check are the permissions of `/dev/raw1394`. A quick way of giving the necessary permissions on the device to all users is to execute this command as root:



Confusingly, the IEEE 1394 port of camcorders is sometimes labelled by manufacturers as 'DV out'. Here it is possible to also see the S-Video port.

“Although it does not cause any major problems, the subsystem IEEE 1394 is still tagged as experimental – so take care!”

```
# chmod o+rw /dev/raw1394
```

Once you are able to interact with your camcorder through *Gscanbus* as a normal user, you are ready to get started capturing video from your device. A simple command-line utility to do this is *Dvgrab* (<http://kino.schirmacher.de>). Keeping your camcorder switched on in playback mode, try executing:

```
# dvgrab --autosplit --format dv2 myclip
```

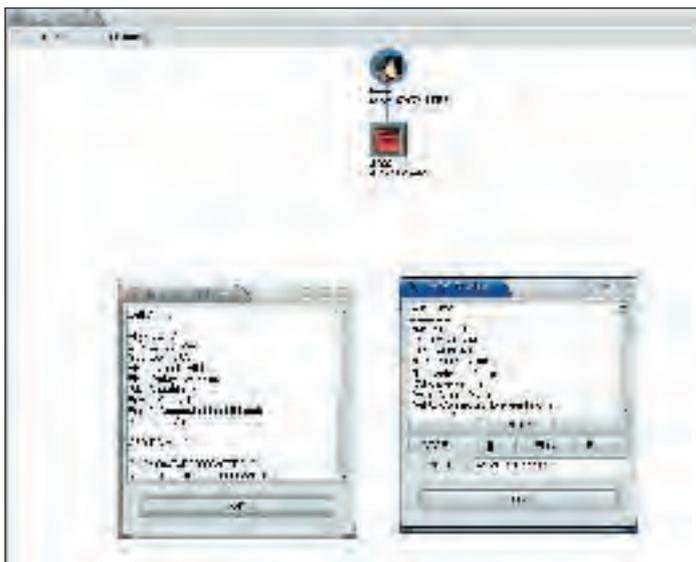
The `--autosplit` instructs *dvgrab* to create a file for every sequence, which means that each file will be named `myclipxxx.avi`, where `xxx` is a three-digit number. Also, as requested by the `--format dv2` parameter, they will all be in the DV2 AVI format. This format was chosen because it works well with *MPlayer*. When you're finished capturing video, the capture process can be halted by pressing **Ctrl+Z** to background the capture task, followed by

```
killall -9 dvgrab
```

Dvgrab is a very useful application, whose possibilities go far beyond this

simple test – you should try experimenting with the extra options to see what suits you best. With its built-in support for all the most popular video file types and the flexibility granted by the slew of options it offers, *Dvgrab* is an excellent and easy-to-learn alternative to the capture utilities integrated into graphical video-editing programs. It is also a fundamental tool whenever video-editing software does not support direct capture from camcorders, so it's certainly worth keeping to hand. For advanced usage, you should refer to the *Dvgrab* man page for the latest information.

If the test with *Dvgrab* has been completed without any problems, and your output files can be played back using *MPlayer*, you're all clear to move onto the next step and fire up a full-featured video editor – this is where the real fun begins. As long as you have followed our instructions carefully, you shouldn't need to play around with *chmod* or device mounting any more – it's all much easier from now on! >>>



Gscanbus can show information on your IEEE 1394 chipset and the connected camcorder. It also offers a practical playback interface that allows you to test whether you can interact with your camcorder.



CINELERRA

www.heroinewarrior.com



Cinella is one of the most complete video-editing apps for Linux.

Cinella is the recent successor to *Broadcast2000*, which was the first fully featured Open Source video editor for Linux. *Cinella* is aimed at users with a technical understanding of video-editing, and as such the developers haven't made many efforts to make the interface intuitive or attractive. On the other hand, users experienced at editing can jump straight in and take full advantage of its advanced features. In its default view, *Cinella* displays a main window, a resource window, a compositor window, and a preview window. Although it's clunky at first, the interface does grow on you and the end results will repay your initial efforts.

Cinella supports capture from camcorders and from *Video4Linux*

devices. If you find the capture interface unstable, try updating to the latest release. Alternatively, try using *Kino* or *Dvgrab* to import the image. On the plus side, there's a first-class selection of effects, a timeline with separate audio and video tracks and good drag-and-drop support. As always, there are caveats too: *Cinella* is infamous for a pronounced tendency to crash. While there is little chance of losing your work (*Cinella* comes with a good crash-handling function), it is annoying to get frequent freezes. Hopefully this instability will be fixed in the near future, which will really help turn *Cinella* into a killer application. Even in its present less-than stable form, *Cinella* still deserves kudos.

« EDITING VIDEOS WITH KINO

One of the advantages of digital video is that you can import your videos to your PC and modify them in a quick and easy way. By editing the video on PCs, all the fuss of analogue playback and video tape is firmly banished to last century. More importantly, the time needed to get great results from your work is greatly reduced, and you get to try out some of the many filters and effects available to you. Sure, you'll be hard-pressed to make a film to rival *Star Wars* for special effects, but few people would want action-packed laser shootouts in their wedding videos anyway. For some, the biggest advantage brought about through digital video-editing is the ability to backup and restore data easily: which means you can try out effects and – if you don't like them – just roll back to your previous copy. Of course, to benefit from all this, we need an appropriate application. The technical name for a program in this category is 'non-linear video editor', which is a fancy way of saying you can jump around your movie as you please and make changes. This might sound like

the most basic feature a video editor would expect, but keep in mind that video-editing used to be done entirely on celluloid, which meant that editors had to literally chop and change pieces of film with scalpels in order to the desired effect. When remixing was required, tapes had to be played back and recorded over and over again, a process that left the original master tape for Queen's hit song *Bohemian Rhapsody* so thin you could see daylight through it!

Although 'non-linear video-editing' sounds like a great term to impress your mates in the pub, we'll be using the more generic term of 'video editor' from now on. Generally, a video editor will have the following features:

- a graphical interface, for easy execution of the editing operations
- a capture function, to import the video directly from the camcorder
- a resource area, where different clips are accommodated
- a work area, to experiment with various modifications
- a timeline, for a quick glance at the whole video and for easy identification of different points in the clips
- a preview area, for visualising how the clip would appear with the actual modifications
- an effects section, for choosing special filters and transition effects
- an export function, to save your modified clips to files or back to your camcorder

Of course, this general scheme is implemented differently in the various video editors, and you may find your editor uses a weird and wonderful name for an otherwise-mundane task!

Choose Linux

Linux is about choice, and this applies to video-editing apps too. If you're after all-out breakneck speed, then *Kino* is the best place to start, but if you're more interested in power functions to take your videos to the next level, *Cinelerra* would be preferred. For pure ease-of-use, *MainActor* has a strong fan-base, and there's even a special application available for quick modifications:

Avidemux. They all do a good job at improving even the dullest clip, so there's no harm trying out each to see what suits your workflow best. There's no better way to find out which video editor is the right one for you than to experiment a bit with all of them. Many of these applications can be installed easily, and you will probably find that your distribution already has binary packages provided. Of the four, only *MainActor* is proprietary, but a fully featured trial version is available from the official website.

Kino represents a reasonable compromise between stability, usability and features, and this makes it attractive to the largest range of users. Furthermore, it also supports a high number of formats – both in input and in output – and has the best available capture function around. It is also the only video editor that allows you to write your video directly back to the camcorder, which is reason enough for many to choose it. It's not perfect, though – *Kino* has a limited range of special effects, fiddly audio management, and rather limited drag-and-drop support.

Kino has recently had a face-lift: and at the time of writing, version 0.7 has just been released. If you used the previous release, you'll notice the biggest change in 0.7 is that it has been ported to *GTK2* (GNOME 2). Although the functionality between 0.7 and its predecessor exhibits little difference, the interface now looks a lot cleaner and more polished than before. A noticeable exception to the last statement is the scriptable Export/DV pipe, which adds even more flexibility to an already highly customisable application. It's enough to open the preferences window in order to realise just how many aspects of *Kino* can be configured according to taste and need – *Kino* is quite the tinkerer's paradise.

One popular feature of *Kino* is that it is wholly extensible through plugins, however this is somewhat limited in v0.7, because the recent port to *GTK2* means that all the plugins for older versions have to be ported before they can be used with the new iteration.

TO UPGRADE OR NOT TO UPGRADE?

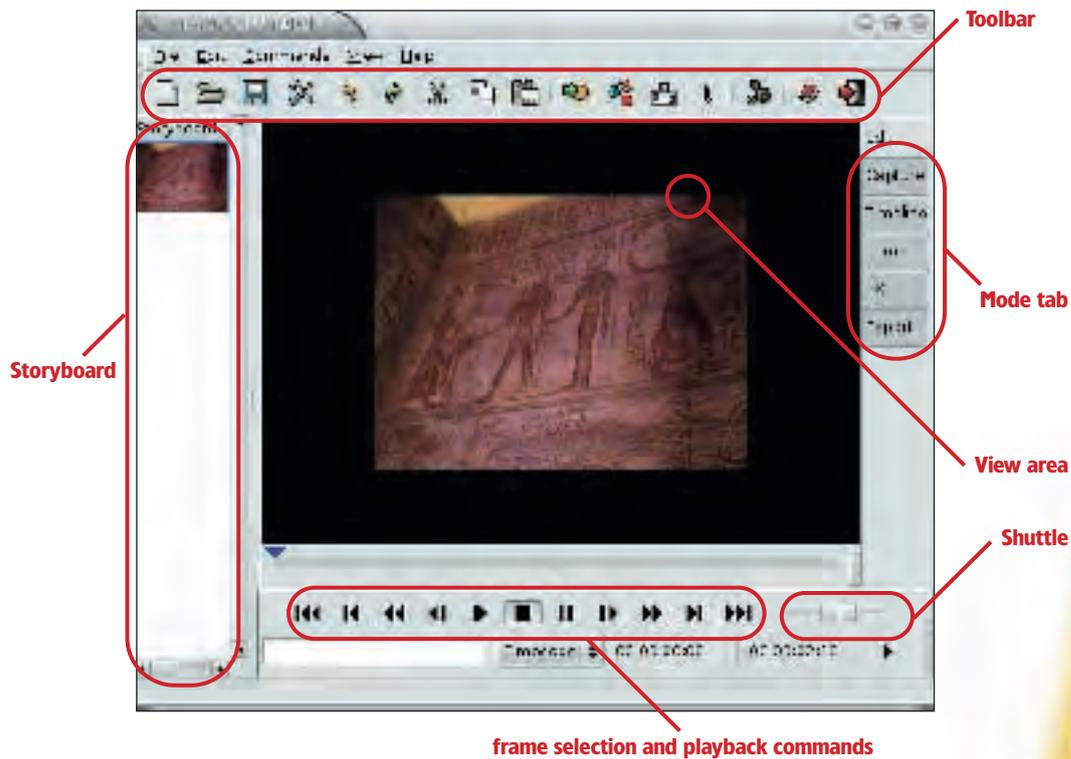
A couple of tips for a safe install of new releases.

Most of the video-editing applications for Linux are relatively new, and all of them are under heavy development. For this reason, releases usually appear at a very rapid rate. Since those releases often fix critical bugs or add new important functionality, it would be desirable to perform the upgrade. However, it's not always possible to get binary packages for your distribution of choice. As a result, users are obliged to upgrade with compilation from sources. Generally speaking, this is not difficult, and often just requires some familiarity with the standard build process. However, video-editing applications are an exception: in order to have so much functionality, they require a number of optional libraries. The result of this can be that even if those libraries are not installed, the configuration script will often appear to have been completed successfully, while as a matter of fact it has just silently disabled some features. The risk is then to trade a slightly

defective version for a more seriously defective (for you) update. Hence the question is: *should you upgrade straight away or should we wait for the availability of third-party binary packages?* Our answer is that a full upgrade is still worth the risk, since that risk can be highly reduced or even eliminated with little attention. The first recommendation is to run the configure script with a non-default location as a prefix parameter, for example:

```
# ./configure --prefix=/opt
```

This will configure the scripts in such a way that binaries are installed in `/opt/bin`, `/opt/usr/bin` etc., eliminating the risk of overriding the old executable. In this way, the old and new binaries can be run side-by-side. Then, if you decide that some essential functionality is missing, you can repeat the configure process and try to track down any notices about missing libraries until the executable that is produced works as you want it to.



The *Kino* interface in the Edit mode is more straightforward than some video apps.

Having said that, there is more than enough in the vanilla application to have a great time while touching up family movies.

The window of the program is shown above – as you can see, it has a simple but very functional design. The first thing that users should familiarise themselves with is the concept of views, or, in the *Kino* language, ‘modes’. These are subsections where similar operations are grouped together for easy access.

The default mode is the Edit mode, but you can change the mode by clicking on the tabs to the right. In the Edit mode, all basic functionality for video-editing such as cutting, pasting, scene insertion, and scene split splitting is enabled and can be accessed through the toolbar at the top of the screen. Imported clips can be ordered in the Storyboard (the column on the left-hand side) from top to bottom according to their desired chronological order in the movie. >>

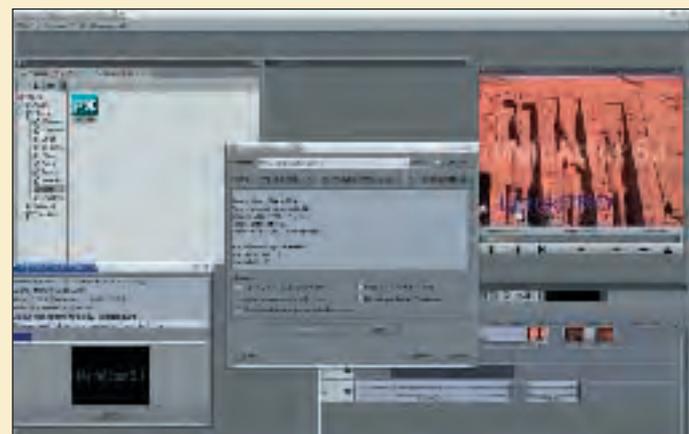


MAINACTOR

www.mainconcept.com

MainActor 3.x has been the yardstick for video-editing in Linux until very recently. Now, this venerable application has been withdrawn in order to make room for the forthcoming 5.x release, and the preview release shows that the jump in the version number is more than just a marketing move: there is a wealth of innovative new features available. At the time of writing, it is possible to try out the beta 2 version, which is distributed as an RPM package. Users of non-RPM-based distros can install it by converting to the appropriate format with *Alien* or extracting it with a pipe like `# rpm2cpio MainActor_V5-5-2.01.i686.rpm | cpio -id`. Though stability is not to be expected

from a beta release, we didn't experience any problems when we tried it out. The intuitive interface, the fact that almost any editing function can be executed via drag-and-drop, the long-awaited capture function, and also the wide range of effects could well make this proprietary app a reason to crack open your wallet and splash out. However, much will depend on the price, which is still unknown at the time of writing, and also whether the new features make it to the final release. Two aspects that definitely deserve attention before the final version is released are the rendering process (it's very slow) and also the restricted range of export formats. If these are fixed, the final version should be a masterpiece.



The new *MainActor* supports very few formats, but that may change.

MAKING MOVIES



There's a large space reserved for the viewer, where the clip selected in the storyboard is shown, and you can easily isolate a single frame using the fine-resolution navigation commands and the shuttle on the bottom-right.

Clips can be imported using the command-line or by clicking on the appropriate button in the toolbar. Movies can also be imported directly from the camcorder or indeed from a *Video4Linux* device in the Capture view, where images from the source can also be played.

Timeline

The Timeline view displays stills of the clip taken at regular intervals, to help the user quickly identify critical sections of the movie, while the Trim mode offers a practical interface for cutting sequences at the beginning and end of the scene.

“Just open the preferences window in order to see just how configurable *Kino* is – it's quite the tweaker's paradise.”

Although *Kino* does not have many effects, the few that it does have are offered in the FX mode – this is one of the least intuitive areas of the application, so it will be covered in more detail in a tutorial in *LXF53*.

Of course, the goal of editing a video is to be able to enjoy the results of all that hard work. To this end, all

INTERLACED VS DE-INTERLACED VIDEO

Smoother images on your screen

One of the most common techniques in video-editing is interlacing, which consists of dividing the image into lines and storing separate information for even and odd lines. On displays like TV monitors, reproduction of interlaced videos gives the impression of a smoother movement. This is because televisions draw frames as a series of two fields, so the result of interlacing is that the frame rate appears to be doubled. Video outputted by digital camcorders is interlaced, for better rendering on TV screens.

However, images from an interlaced source played on a non-interlaced display like a TFT screen have noticeably rough edges around fast-moving objects. In this case, better results are obtained if the video is de-interlaced with some tools. Most video editors have an option for de-interlacing images in exported clips, although it may be named differently.



video editors offer an export function, and *Kino*'s Export tab has these options:

- 1 Export back to Camcorder (IEEE1394 tab)
- 2 Export in a handful of high-quality formats (DV tab)
- 3 Export static images (Stills tab)
- 4 Export audio only (Audio tab)

5 Export in some formats with high compression rate (MPEG)

Since there are so many options and strategies involved in exporting your movies for others to view, we cover them separately in a little more depth in the following section below: *Sharing Your Masterpiece*.

SHARING YOUR MASTERPIECE

The key thing to keep in mind

when exporting your edited clip is your intended display format – how will others play back your movie? Essentially, the options are:

- 1 put it on some variety of videotape
- 2 save it as a file for playing on a computer
- 3 put it in a CD/DVD, for playback on a DVD player

In the first instance, you will need a program that allows you to put your creations back to your camcorder, which means *Kino* is the only option in this situation. If you want to export to a file, you have a lot more leeway

because players like *MPlayer* and *Xine* support a wide selection of formats. With files, the most important choice is whether you want to favour the quality of the images at the expense of larger file size, or *vice versa*.

If large files are OK, opt for a format like DV2 as it provides excellent quality. If you'd rather sacrifice quality in exchange for smaller files, use an MPEG-4 format such as Xvid or DivX. Both of these compression formats require fairly high-powered CPUs to play back smoothly – anything over 1.5GHz should be up to the task.

NTSC or PAL on DVD?

The last option brings us into the field of DVD authoring, which is a huge subject, and certainly one that requires its own feature. Unfortunately at present, there are no graphical tools that allow easy creation of a DVD, so be prepared to open a console.

First, we need to create a file in MPEG-2 format, which is the standard for DVDs. If not installed, it is now time to get the MJPEG-tools package from the CDs of your distribution or from <http://mjpeg.sourceforge.net>. We use **movie.avi** in our code

examples that follow. The specific formats of the audio and video component don't matter, since we are going to extract audio and video with Transcode, which works smoothly with all formats. At this stage, what we need to know is whether we want to use the PAL or NTSC video format. PAL is the standard format in most of Europe and Australia, while NTSC is mainly used in North America. If you live in another area, be sure to check! For PAL, you would do

```
# transcode -i movie.avi -V -y mpeg
-F d -Z 352x288 --export_fps 25
--export_asr 2 -E 48000 -b 224
-o moviedvd
```

while for NTSC

```
# transcode -i movie.avi -V -y mpeg
-F d -Z 352x240 --export_fps
29.970 --export_asr 2 -E 48000
-b 224 -o moviedvd
```

This will create the files **moviedvd.mpa** (audio) and **moviedvd.mv2** (video), which will need to be rejoined with

```
# mplex -f 8 -o moviedvd.mpg
moviedvd.m2v moviedvd.mpa
```

The file **moviedvd.jpg** is now in the correct format.

For the next step, we need to use *Dvdauthor*, which is available from (<http://dvdauthor.sourceforge.net>), to parse an XML file with information about the DVD and create the correct tree structure. In our next example that we show here, we created the file **dvdtree.xml** and then entered the following XML data:

```
<dvdauthor dest="DVD">
<vmgm>
<menus>
<pgc>
<vob file="moviedvd.mpg"/>
</pgc>
</menus>
</vmgm>
<titleset>
<titles>
<pgc>
</pgc>
</titles>
</titleset>
</dvdauthor>
```

From the directory where the file **moviedvd.mpg** resides, we then execute

```
# dvdauthor -x dvdtree.xml
```

This code will create the subdirectory 'DVD', which will be the root directory of the disc. To test the disc, we can now use *Xine*:

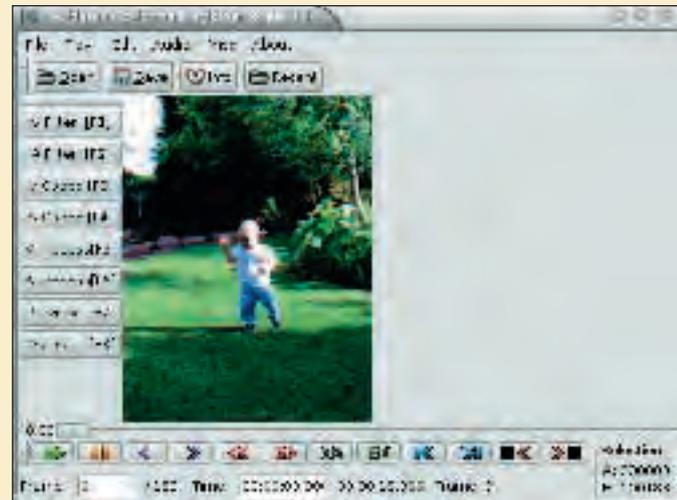
```
# xine
dvd:/PATH/TO/DVD/VIDEO_TS/
```

where **/PATH/TO** must be replaced by the correct path to the DVD directory.

It is now time to burn the finished DVD, and this is quite easy if you have a graphical application such as *K3b* (www.k3b.org). You can also do it from the command-line, but if the expression **dvd+rw-tools** is obscure, we recommend you stick to the GUI!

AVIDEMUX2

<http://fixounet.free.fr/avidemux>



The practical and intuitive interface of Avidemux2.

Avidemux has the cleanest and most functional interface of all the video-editing programs shown here. However, it isn't as feature-packed as its competitors, meaning you should steer clear if you're after video-capture functionality or if you want to use

uncompressed audio. On the other hand, it can work with virtually any kind of AVI file and offers a wide range of effects. In our opinion, *Avidemux* is probably best-suited to small, almost on-the-fly editing operations more than for complex works.

Of course a CD can be used instead of a DVD if the size of the directory tree is below 700MB.

The case that we used as a demonstration here was a very simple one. For a more complex example, see http://mightylegends.zapto.org/dvd/dvdauthor_howto.php.

CONCLUSIONS

Although the killer app for video-editing on Linux has yet to be produced, that needn't stop you from using what exists already. Fortunately, Linux offers a wide choice of good programs that do the job well enough, so you can pick-and-choose as you see fit. Perhaps the hardest part for most people will be the setting up their hardware to work with IEEE 1394 – if you find yourself struggling here, it could really be as much a fault of your distro as it is yours, so you might want to consider switching to a more friendly distro. As it stands, a user unfamiliar with the command-line and concepts like loading kernel modules may give up all hope before they get to running

Kino, which is a shame. Hopefully this guide will help get you through these teething troubles and allow you to concentrate on *actually making movies*. As you've seen, the process of capturing, converting, editing, and outputting a complete movie involves using many different tools. At first, this will probably confuse you, but it will become easier with experience.

Where possible, we have tried to provide links to Internet sites where you can learn more about the software we've been using here. Though we've included most of the apps on this month's coversdiscs, we strongly recommend you investigate the sites – as per usual, the easiest way to get

started is to learn from others who have been there before you.

In any case, once the installation step has been mastered we hope you'll agree that, through quality apps like *Kino* and *Cinelerra*, you can make your home videos look all the better. Who knows? This time next year you might even be challenging George Lucas at the box office... [LXF](#)

“Struggling to set up IEEE1394 could be as much your distro's fault as it is yours – so consider switching to a friendlier one.”

What on Earth are... HARDWARE DESCRIPTION LANGUAGES?

Marco Fioretti explains how to *make registers* – rather than just the code with which to fill them.

» You Linux types sure love your acronyms. What does HDL mean?

The acronym – which, for a change isn't recursive – stands for *Hardware Description Language*. It applies to any programming language that is designed to model and implement *hardware* blocks instead of software programs.

» Eh? Isn't hardware designed drawing rather complicated and obscure electrical diagrams?

At the end of the day: yes, all computer hardware can still be defined as an electrical circuit of some kind. HDL allows us to describe certain classes of circuits at a higher level.

» What kind of hardware can be modelled by HDLs?

Purely digital integrated circuits (ICs) where the chips work only in binary mode. Some HDLs can deal with analog designs too, but due to their complexity and relevance to Linux and computing, we will not be covering them in this article.

Digital ICs can be permanently wired (as in "If there is a bug, you just blew up some millions of dollars;") or field-programmable. The first variety is called ASIC (Application-Specific Integrated Circuits): CPUs fall in this category.

The second class includes FPGA (Field-Programmable Gate Array) and similar: this type of IC can be programmed to connect its internal transistors in different ways, in order to correspond to different functions. This blurs traditional differences between hardware and software, which can cause confusion.

» When are ASICs used?

ASICs are developed with huge initial investment, when one needs any combination of millions of units, very low power consumption, or very high speed. FPGAs are the way to go in all other scenarios, or for initial prototyping.

» Is there only one HDL?

It would be a dull life indeed if this were the case. There are several HDL languages, but the two industry standard today are Verilog and VHDL. The

first (www.verilog.org) was developed by Gateway and then by Cadence in the Eighties, released as Open Standard in 1990 and became IEEE Standard 1364 in 1995.

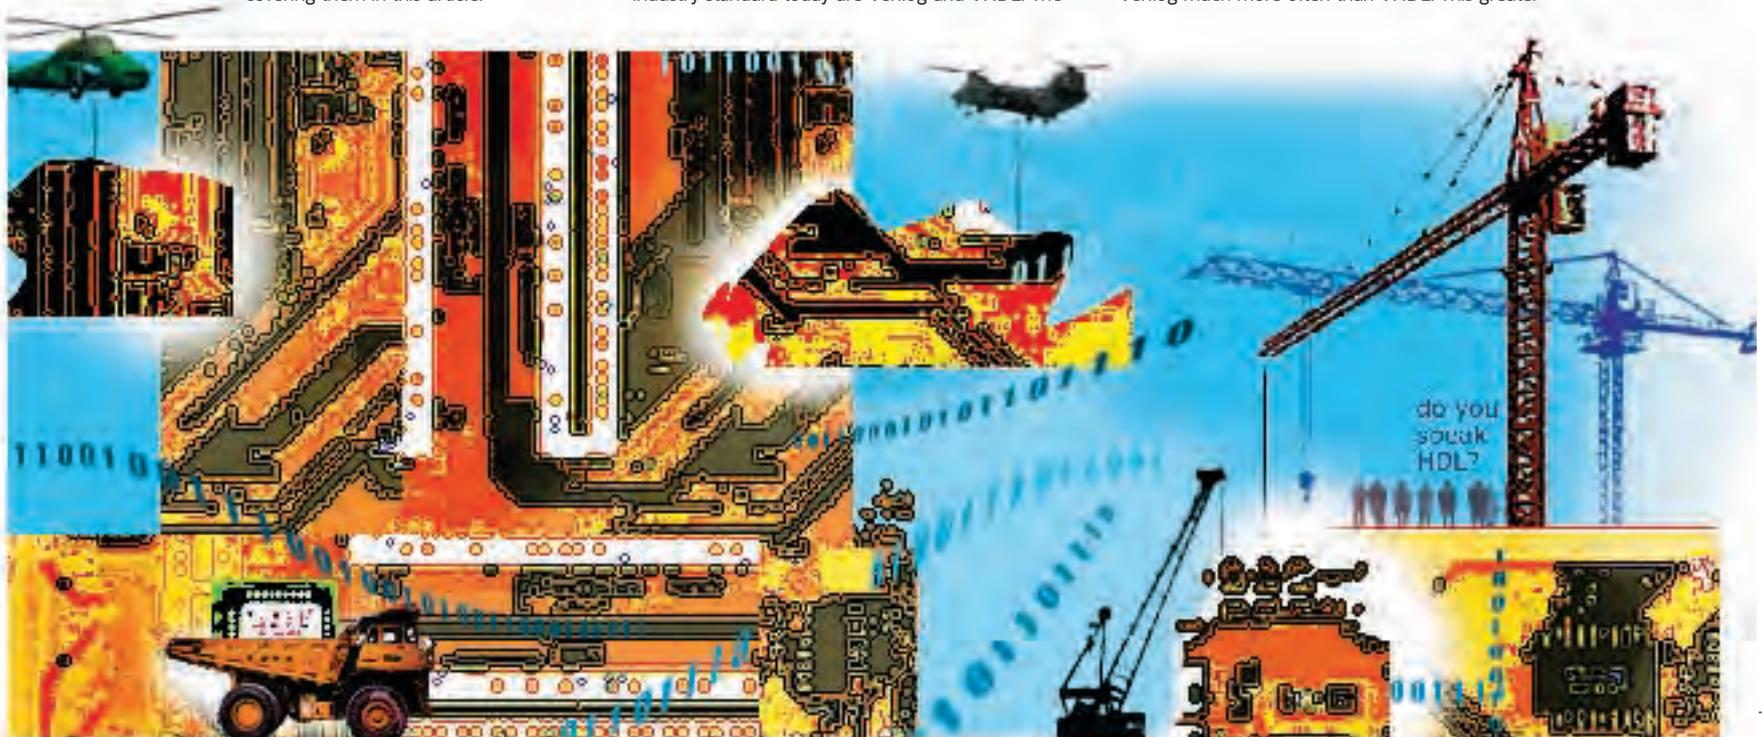
VHDL (Very High Speed HDL) came out of the American Defense Department's VHSIC (Very High Speed Integrated Circuits) program. Development was started in 1981 and then passed in 1986 to IEEE, becoming Standard 1076-1987.

» Can they be mixed?

Yes. You can't alternate Verilog and VHDL in the same "module" (more on this later), but modules written in the two different languages can be combined and nested in almost any conceivable way.

» Which one is better?

Which is better, *Vi* or *Emacs*? Debian or Red Hat (er, Fedora, sorry...)? Seriously, you didn't expect a straightforward answer to such an explosive question, did you? Verilog vs VHDL is a holy war of the first order. The author has been exposed to Verilog much more often than VHDL. This greater



OPEN SOURCE TOOLS FOR HDL

How to apply your scripting skills to hardware design

HDL languages make possible to design digital hardware writing good old ASCII text. Huge quantities of ASCII text, often scattered over thousands of files. The good news is: wherever there are giant amounts of text and files that need to be managed as fast as possible, all your traditional UNIX scripting experience is still valid. Really good HDL designers can very often teach *awk*, *sed*, *find*, *perl* and endless one-liners to veteran sysadmins.

As it often happens with Open Source, however, they don't need to do it every time. There are already plenty of GPL tools available to generate HDL files and templates, connect them to each other, automate simulation and synthesis, and so on. Perl and TCL are the most commonly used languages: TCL is also

embedded in several commercial simulators, so one can use or merge existing scripts right inside a simulation.

Before you ask, yes, *Emacs* does have both Verilog (www.verilog.com/verilog-mode.html or www.veripool.com/verilog-mode_veritedium.html) and VHDL (<http://opensource.ethz.ch/emacs/vhdl-mode.html>) modes, with the usual selection of increasingly tortuous multi-finger macros to automate code generation or launch simulations. Another editor with good HDL support is *NEdit*.

The main online portals for these kinds of resources are www.tcforeda.org and www.openeda.org (EDA stands for 'Electronic Design Automation'). A free Verilog simulator is *Icarus*, which can be found at <http://icarus.com/eda/verilog/>.

familiarity – and the need to save space – are the only reasons why the following examples are all in Verilog. Verilog and VHDL tutorials containing examples similar to those in this article, are available online at www.optimagic.com/tutorials.html.

»» Could I learn just one HDL?

As an introduction to the subject, or if you don't plan to do it professionally, yes: both languages will be around for a while. Regardless of one's personal preferences and company tradition, the reality is that code reuse is as good a practice in hardware as it is in software. Consequently, a good designer must be familiar with both languages, and mix them as each single project demands.

»» What does an HDL look like?

If you look at it while doing something else, you might think it sounds like C, or maybe C++. Since they are supposed to model the functionality of a circuit, HDLs have all the basic constructs of software languages: **for** and **while** cycles, **if** and **case** statements, functions and so on.

»» Great! This means that, being a good C programmer, I am already good at digital design too, am I not?

Maybe. Maybe not. As a matter of fact, there *are* 'Pointy Haired Bosses' who desperately want to believe this, reasoning that if the same designer can do software now and hardware after lunch, his company can sell all kinds of 'solutions' all day long.

»» What is the big deal then?

The reality is that – just as web design is deeply different from assembly coding – code and circuits remain very diverse. To begin with, hardware is real-time. *Really* real-time. Sometimes, just one nanosecond later can be *too* late, and it can be caused by slightly different coding styles. In other cases, a different implementation may dissipate twice the power of the other. It takes time and silicon knowledge to get used to all this.

»» So I must have a firm grounding of electronics knowledge, but why can't I just do it in C?

There are compilers around which try to convert C to hardware. Regardless of the chosen approach, any hardware modelling solution must be aware of (and be able to describe) two exquisitely hardware proprieties: concurrency and timing.

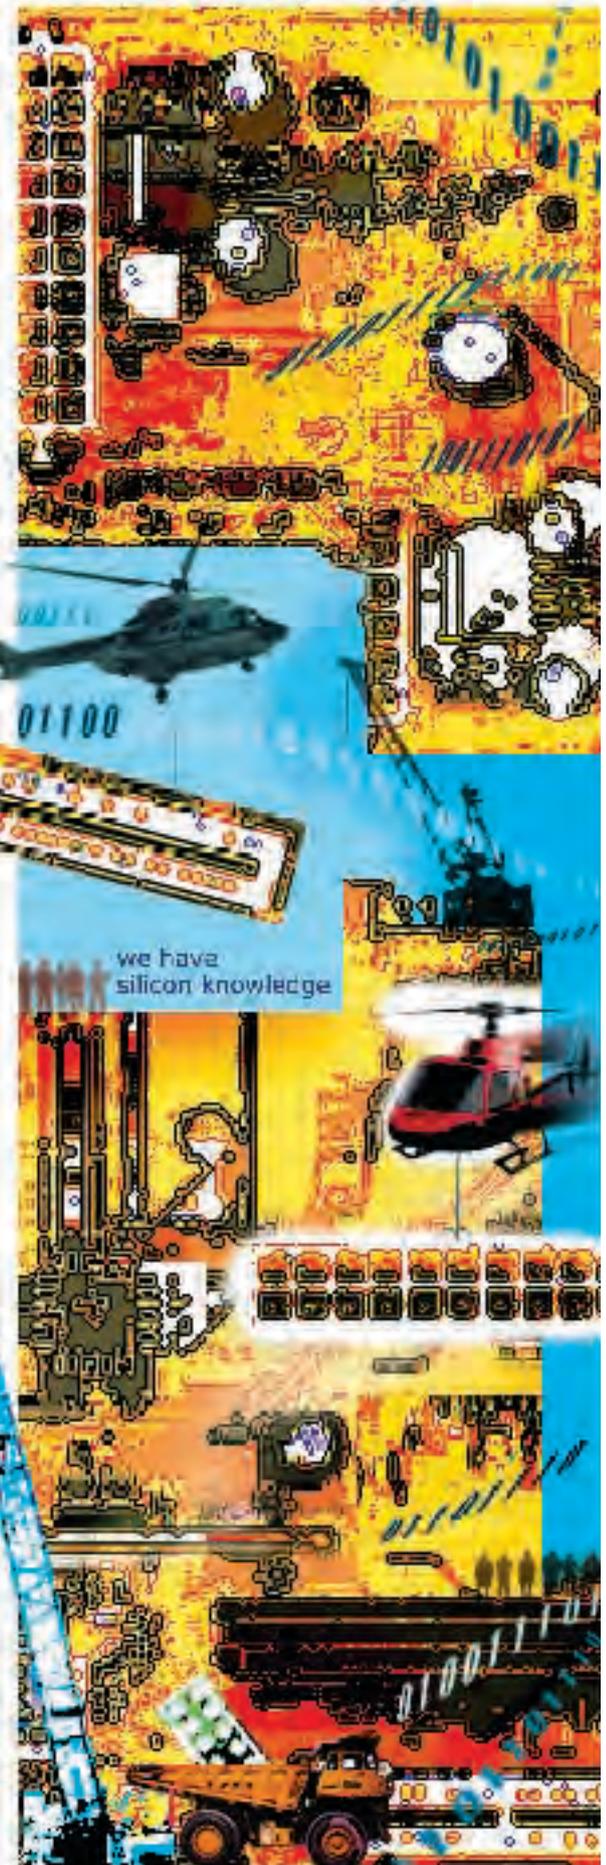
»» This sounds like the hard part, let's do them one at a time

Let's begin with concurrency. Software evolves all in one flow, even when it contains (conditional) loops, branches and such. Imagine starting from line one of a program with certain initial values, and following through the source what happens: at any given moment, only one line of code will be under execution. All the others will be patiently waiting their turn in cache, RAM or disk. A coarse simplification of this could be saying that software is **serial**.

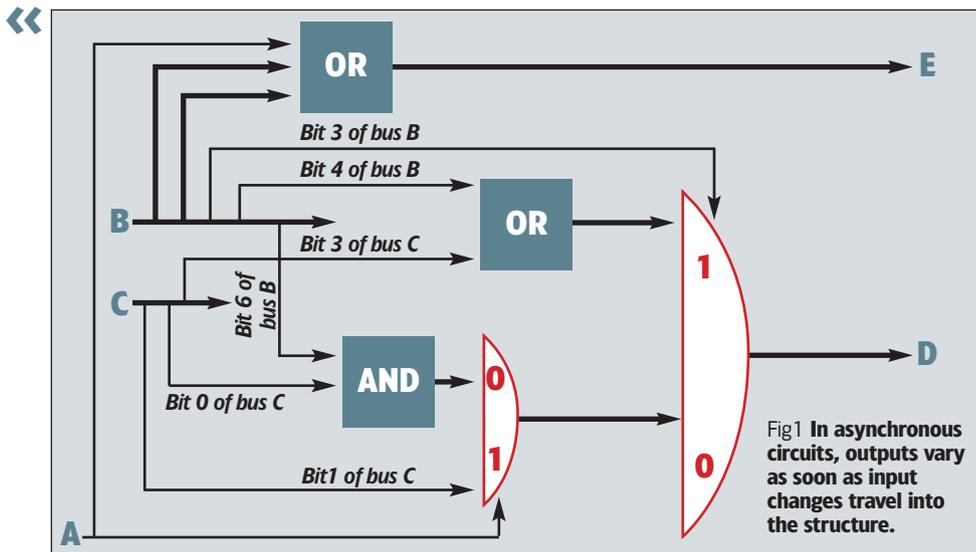
Following this approach, hardware can be described as **parallel**. Each different sub-block of the same chip (eg PCI controller and mixer in an ASIC for audio PC boards) will happily mind its own business, talking to the others only when some electric pulse forces it to do so. HDL languages have extra constructs to define concurrent processes, *ie* pieces of hardware that must do different things simultaneously.

»» Fine (more or less). What about timing?

The other thing that must be modelled with great accuracy is *when* things must happen, at which speed and in which order. The designer must have the possibility to say "this signal must be updated or sampled exactly 10 nanoseconds after that pin had a zero-to-one transition". And it must be done in such a way that it surely leads to the creation of a properly connected flip-flop, otherwise the chip will be useful only as a toaster.



WHAT ON EARTH HDL



>> That's enough about the theory... Show me the code!

Gladly! Let's start with a very simple combinatorial process. In digital hardware lingo, "combinatorial" means a circuit in which all outputs are recalculated instantaneously whenever any of the input changes. The only delays are those caused by the wires and by the actual times needed by the electrical pulses to travel through them and inside each transistor.

```
always @( A or B or C)
```

```
begin
```

```
E = A|B|C;
```

```
if (B[3] == 1'b1)
```

```
D = B[4] | C[3];
```

```
else
```

```
D = (A == 1) ? C[1]: B[6] & C[0];
```

```
end
```

We can already see that Verilog code looks almost identical to (very low-level) C, doesn't it? **if** statements, both in traditional and ternary form, and bit level operators (**B[4]** means the fifth bit of bus **B**, starting from **B[0]**, which is the least significant). Even the **begin** and **end** keywords are just another way to enclose a piece of code.

>> Uh, sure. But what does it actually mean?

The snippet of HDL language above means "I want a digital circuit like the one in **Fig1**! The most significantly HDL part of the example is the initial statement.

The **always** keyword marks the beginning of a process, *ie* the description of an independent part of the hardware block to be designed. The content of the parentheses (**A or B or C**) is called the 'sensitivity list' of the process. The first statement

means just what it says: *every time any* of the variables in the sensitivity list changes, the whole process must be executed. A complete design includes many multiple processes (*ie* simple circuits) interacting with each others. The inputs of our examples (**A**, **B** and **C**) will come from external physical pins or from other HDL code; I mean, circuits. See – it can get confusing!

>> Are different processes executed in sequence?

Not at all. This is what concurrency is all about. When you are simulating the design, or actually building it, the instructions inside each process are executed, or look at, sequentially, one at a time. But different processes happen *simultaneously*.

>> There is no clock in the example you used. Is this HDL of yours good for synchronous designs too?

Heavens, no! Clocked circuits make everybody happy: marketing can say that the next CPU is better just because it has a higher clock; people with a lot of spare money feel justified to buy it; and overclocking addicts have something to boast about. Clocking, *ie* synchrony, is good for the HDL designer too: registered signals move around the chip in a much more predictable way, and nasty loops (**A** depends on **B** which depends on **A**) are avoided. All HDL languages have built-in clocking support.

>> How do you make (or recognise!) a synchronous process?

By enclosing the related code in a different **always** statement, and using a different assignment syntax. A synchronous process in Verilog looks like this:

```
always @(posedge X)
```

```
begin
```

```
B <= D;
```

```
A <= C;
```

```
C <= B | A;
```

```
end
```

The **posedge** qualifier inside the sensitivity list means "I want the outputs of the following process to come straight out of flip-flops, switching on the rising (positive) edge of signal X, the clock." The resulting circuit is the one in **Fig2**. Modelling-wise, it is like saying "**A**, **B** and **C** must be recalculated *only* when there is a positive front of the clock signal".

>> Got you! Even though I'm new to the subject, it's clear that there is a loop in the code above!

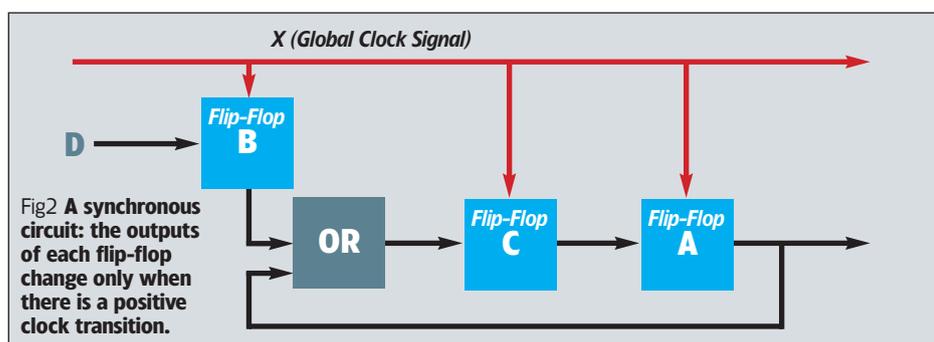
You are correct: **A** depends on **C**, which depends on **A**. But are you? Actually, no. The trick is in how the assignment of the next value is done, and in the use of the **posedge** list. To cut a long story short, the syntax above means:

"Whenever there is a positive transition of X, write down all the current values of all the variables to the right side of all the assignment operator (<=) without stopping at each statement to actually do the maths and assign the new value after that, use all the collected (ie frozen!) values to calculate the next value of each register."

This makes the transitions at each clock cycle independent from the order in which statements are written (demonstration of this is left as exercise for the reader). It also just happens to be the way in which real flip-flops behave. The **<=** assignment operator is said to be "non blocking".

>> How is this stuff partitioned? Do HDL 'objects' exist?

In broadest theoretical terms, there is nothing that would prevent a whole complex design from being written as one single bunch of HDL processes, all in one single file. In practice though, just how readable and maintainable do you think the whole Linux kernel would be if it was written in the same way?



In addition to this consideration is the fact that chips, unlike programs, must have their top-level signals connected to specific physical pins.

HDL languages handle all this with the concept of module: a module is declared in this way:

```
module DATA_SHIFTER (
  Clk, // all I/Os must be declared
  Reset,
  Data_in,
  Data_out
);
inputClk, Reset; // single bit input wires
input [7:0] Data_in; // 8 bit input bus
output [7:0] Data_out; // 8 bit output bus
reg [7:0] Data_out; // must be also declared as
// register, because we want to
// generate it

always @(posedge Clk or Reset)
if (Reset) Data_out <= 8'd0; // Yes, Reset can be
else // asynchronous

Data_out <= {Data_in[6:0], // and data buses can
Data_in[7]}; // be built concatenating
// and shuffling bits

endmodule //
```

An HDL module may be much more complicated than our simple `DATA_SHIFTER` example here: regardless of its complexity, it would still be a kind of black box, able to interact with the rest of the hardware only through its I/O ports.

»» How are modules connected together?

You *instantiate* them inside a module at a higher level of hierarchy, like this:

```
//top module already declared somewhere above

FIRST_DATASHIFTER DATA_SHIFTER (
  .Clk(Clk_100_Mhz),
  .Reset(Master_Reset),
  .Data_in(Downstream_Data),
  .Data_out(Data_shifted_once)
);

SECOND_DATASHIFTER DATA_SHIFTER (
  .Clk(Clk_100_Mhz),
  .Reset(Master_Reset),
  .Data_in(Data_shifted_once) // Output of one
  module
  .Data_out(Data_shifted_twice) // goes inside
  another
);
```

»» How many levels of HDL coding style exist?

Software code can be written in different ways for different purposes: from fully portable to only working on one

combination of CPU and OS. HDL designs are no exception to these bound or cross-platform ideas. In this case, however, what matters is the possibility (or the need) to transform a module in real silicon working as expected at the required clock speed.

High-level HDL modules, constructs or instructions that EDA tools cannot map to Carnot maps and connected transistors are said to be coded at *behavioural* level: they are only for proof of concepts, or to generate inputs to feed for verification to the real design. Behavioural HDL can be recognised also because it doesn't describe any hardware *architecture*, just the way something should work.

»» What kind of HDL can be mapped to silicon?

The next-lowest level is called RTL (Register Transfer Level). The code examples in this article belong to this category: they have structure. This kind of code describes which flip-flops have to be in the design, and which boolean functions connect them to each other. The designer can control perfectly how many clock cycles are needed to perform each task. In spite of this, everything is still generic enough to be portable to (almost) any kind of digital hardware: we speak of flip-flops without defining transistor width, depth, and so on.

»» Can HDL be even more detailed? If your answer is yes: why?

Sometimes, counting clock cycles is not enough: it is necessary to get even closer to the actual silicon that will be used in the final chip design. Imagine a designer with the assignment to implement a high-speed clock tree: this must be placed and balanced almost entirely by hand, to minimise clock skew.

In the real world, we would already know the silicon target. For example, we would know that the code will be mapped on a SomeSuperASIC Gate array, 01 micron technology, which offers as hardware macro a 2 output buffer, called `SA_2BUF_20`, with a 20 picosecond (yes, picoseconds: 10E-12 seconds!) internal delay. The clock tree would then be built using such macros directly:

```
N1 SA_2BUF_20 (
  .IN(EXTERNAL_CLOCK),
  .OUT_1(TMP_CLOCK_1),
  .OUT_2(TMP_CLOCK_2)
);

N2 SA_2BUF_20 (
  .IN(TMP_CLOCK_1),
  .OUT_1(TMP_CLOCK_3),
  .OUT_2(TMP_CLOCK_4)
);

N3 SA_2BUF_20 (
  .IN(TMP_CLOCK_2),
  .OUT_1(TMP_CLOCK_5),
  .OUT_2(TMP_CLOCK_6)
);
```

and so on an so forth, possibly for two thousand lines. This is called *gate-level* HDL. Each wire is routed manually to a basic boolean unit: you can't get any more accurate than that, unless you start designing the basic gates yourself.

»» If gate level is so optimised, why isn't it always used?

For the same reasons why software developers don't do everything in assembler. Gate level is disgustingly precise, but terribly slow to code, really hard to debug, and definitely not portable: if the boss announces triumphantly that he switched to some other silicon vendor, you must either resign or know your `awk` and `Perl` real well, to avoid weeks of manual editing and its attendant troubles.

»» Apart from the abstraction level, all designs of the same function lead to the same hardware, right?

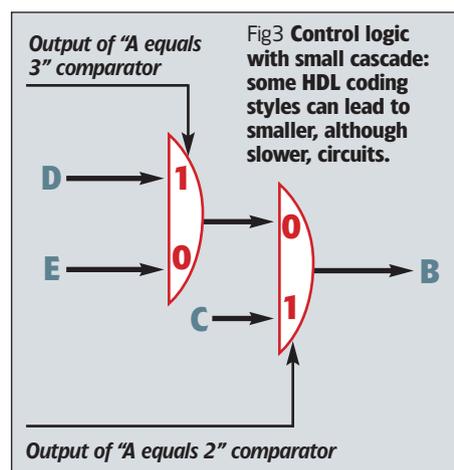
No. If that were true, we HDL guys should go find some honest work. Compare this first piece of pseudo code:

```
if (A == 2) B = C; //cascaded IF implementation
else if (A == 3) B = D;
else B = E;
```

with this second example:

```
case (A) // multiplexer implementation
2: B = C;
3: B = D;
default: B = E;
endcase
```

They are perfectly equivalent functionally, but give the two different hardware structures shown in **Fig3** and **Fig4** respectively. The first will be slower but smaller, the second faster but bigger. »»



WHAT ON EARTH HDL



>> Are HDL designs as portable as software can be?

You mean portable like Java? "Write once, run anywhere"? Yeah, sure; if you write only at RTL level; and partition your modules in the right way (what is fast enough on one silicon may not work on another, intrinsically slower one); and don't use vendor-specific RTL modules; and don't rearrange your RTL Boolean functions to take just half of the area... but only on this particular silicon family. See the problem? More seriously, HDL is just like C: if it is written and structured properly from the beginning, it will have less than maximum performance, but work with null or very little effort on many different silicon families. If not... back to the keyboard every time.

>> Software is tested running it on hardware. How is an HDL design tested? Do you have to make a chip?

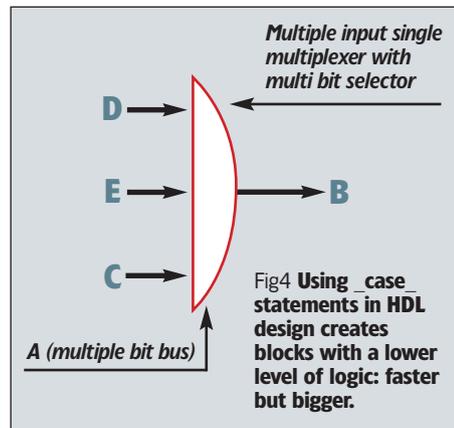
Writing a test bench: an HDL module which instantiates the actual design, generates the same inputs that the actual chip will receive, and dumps to an ASCII log or to some graphic databases everything that one needs to observe to decide if the design is correct. Writing and running good test benches can (and should) take much more time than writing the actual code. Otherwise, one could go to market with a super CPU which can't even handle divisions properly. Of course, this is just an hypothetical case, isn't it?

>> Digital hardware is made up of transistors and connecting wires drawn on silicon. How do we get there from ASCII models?

Ouch, right. We have to ship this wonderful model some day. Once the design has passed whatever simulation you could come up with (remember that broken division...) it is time to do logical synthesis. This process converts the HDL design at RTL or gate level to a **netlist**: kind of a huge Carnot map of the whole thing, which describes the functionality in terms of registers and basic boolean gates: AND, NAND, OR, XOR and so on. The cost per-seat of an industry standard ASIC synthesiser can easily match the annual budget of a small town.

Once the netlist has been simulated, layout happens: each kind of basic boolean gate is converted to an equivalent (predefined) cluster of transistors. All clusters are placed in a more-or-less random position on the silicon. The space left between them is used to draw the metal strips connecting each transistor to the others.

If and when all necessary wires have been routed successfully, timing is checked: if there is even one single wire which propagates its signal slower than the required clock speed, the process starts over; sometimes from the synthesis stage. Constraints can be defined to guarantee that certain critical parts will be laid out just as the designer wants.



>> Sounds big, like compiling the kernel or OpenOffice.org...

No, it's *much* worse than either of those. Every trial will most likely need Gigabytes of RAM and last from a few hours to several days (even more than one week): you don't want to be close to an HDL designer when the server breaks on the tenth layout day, just hours before the planned delivery...

>> Can HDL designs be Open Source?

Of course: code is code. The author can license it in any way that he or she prefers. The big difference with respect to software, of course, is that you can create and use Linux with just one relatively cheap PC.

"Manufacturing" an actual Pentium Clone out of your wonderful HDL design, however, could prove just *a bit* more difficult. New hardware can be Free as in Freedom, never as in free beer. You should sell your

whole county first (inhabitants included) to buy the required machinery and the personnel to run it. Alternatively, you could sell just your neighbours (gee, two birds with one stone!) to convince some existing silicon foundry to do the job for you.

Don't despair yet. Several hardware projects have indeed been done under Open Source licenses. Opencores (www.opencores.org) is a perfect example: they have a lot of complex macros, proved in real designs, freely available under the GPL or similar licenses.

>> Is HDL design possible on Linux?

Most certainly. Otherwise, why should we bother writing about it in this magazine? In the last few years, Linux has gained a lot of recognition among the big EDA players: companies like Synopsys, Cadence and Mentor all offer Linux versions of almost all their products; while several ASIC and FPGA vendors are following suit. Linux is very much appreciated for the same reasons which made it popular in other fields: rock-solid simulation farms can exist on commodity hardware, without the price tag of Solaris or other traditional Unix platforms. Linux arrived later in the HDL field, but is there to stay.

Tutorials >>>

Our experts offer help and opinions on a whole host of Linux applications

YOUR GUIDE TO GETTING THINGS DONE!

Whether you are just starting out in Linux, or an experienced veteran, there's always more to learn. Every issue of *Linux Format* is packed full of practical advice, and nowhere is it more concentrated than in our tutorials pages.

Here you'll find expert guides to all sorts of things, from Basic Linux usage to understanding and deploying network solutions, from simple script coding to the complexities of Perl regular expressions, Java server apps and more. We aim to bring a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact us, by email at linuxformat@futurenet.co.uk or by snail mail, or log on to www.linuxformat.co.uk and post your suggestions in our special forums? Hope to hear from you soon!

Nick Veitch EDITOR

HOW CODE IS REPRESENTED

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap
between lines:
begin
  mniWordWrap.Checked := false
end;
Usually, you'll find the code on
our CD/DVD too.
```

THIS MONTH TEACH YOURSELF...

Beginners' Guide to KOrganizer

Sort out your life and you'll have more time to drink Guinness. Er, and also do other things **p64**

Flex and Bison

Functions are cool, but they are better when they actually return useful values! **p68**

Game Programming

Discover the joys of the Simple Direct Media Layer, and get worried about large fish **p70 >>**

PHP and LDAP

Access directory information from within your PHP scripts **p74**

The GIMP

Create three-dimensional effects in the best 2D art package – complete walkthrough. **p78**



Perl Templates

Having written a new book on the topic, Dave Cross is best placed to explain the template toolkit **p82**

SNMP

In the first part of a new, short series, Chris Brown introduces the SNMP protocol **p86**

TIP OF THE MONTH!

LXFBENCH YOUR PC

Ever wondered how your machine would score compared to those featured? The *LXFBench 2004* benchmark suite is available for you to try out on your own machines. The benchmarks are online at www.linuxformat.co.uk/lxfbench2004.zip. You need PHP 5 to run parts of the suite, so snag it from www.php.net and extract it somewhere. Then run:

```
./configure --enable-pcntl --with-sqlite --with-gd --with-zlib
```

You will need to have *libz-devel* and *libpng-devel* installed on your machine already. Compile PHP with **make**, then **make install** and **make**

install-cli and you're ready to run the PHP part of the tests.

Next, download *mhash* from <http://mhash.sourceforge.net> and install it. You may find it puts its libs into */usr/local/lib* rather than */usr/lib* and its include files into */usr/local/include* rather than */usr/include*.

With the *lxfbench.c* file, run the following command:

```
gcc -O2 -lm -lmhash -o lxfbench lxfbench.c
```

PHP bundles its own versions of *SQLite* and the *GD* image library, so if you want to use *ICC* we suggest you use it on everything. Having said that, it does issue many warnings on both

mhash and PHP – nothing fatal, though. Once you have *lxfbench.php* and the *lxfbench* executable in the same directory, you should be good to go: run `php lxfbench.php` and sit back while the benchmark runs.

Once you've performed the tests, get involved! You can either go online to www.linuxformat.co.uk and post your benchmark results online (please include your system specs too) for others to view, or, if you're really eager, you're welcome to submit changes to the *LXFBench* code back to paul.hudson@futurenet.co.uk for possible inclusion in *LXFBench 2005*.

BEGINNERS' GUIDE TO LINUX APPLICATIONS

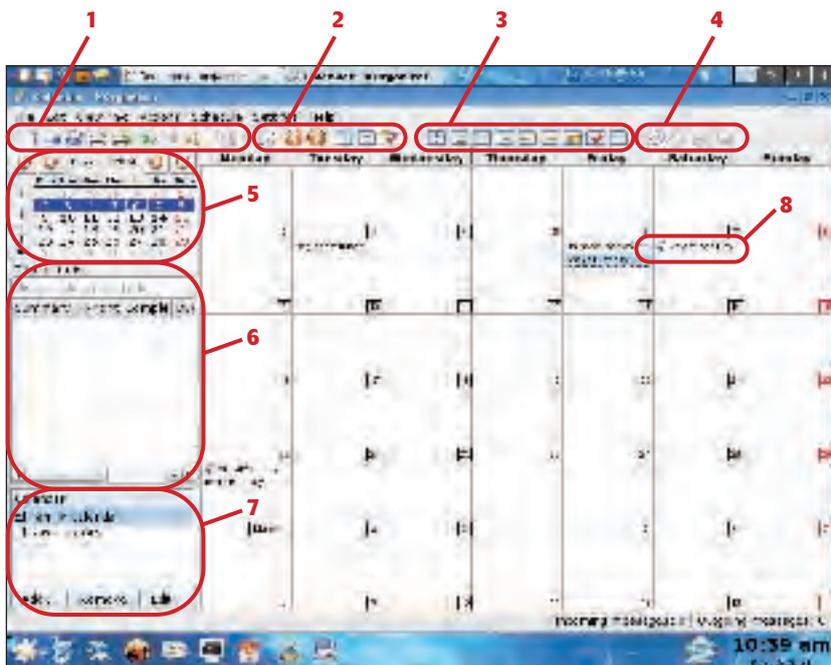
Sort your life out: KOrganizer

Since it's bound to bucket with rain over the four-day Easter Holiday, why not sit at your PC and follow **Andy Channelle's** example by resolving to be more organised using a mixture of old and new KDE technology?



Being more organised is probably the one ambition that most people aspire to at some point in their lives, usually three hours before a significant – and somewhat unprepared for – deadline or exam. For most of us, the tools we need to be more organised are literally at our fingertips, and whether you favour the GNOME or KDE desktop, there is a suite of applications that will integrate not just into your desktop, but also into your life. GNOME users can call

on the skills of Ximian's excellent *Evolution* application – which we will be examining next issue – while followers of all things beginning with 'K' will find much to admire in the new *Contact* collection, which is an integral part of the recently updated KDE 3.2. If you are still running a pre-3.2 version of KDE, don't worry, *Contact* is simply a 'meta-application' bringing together *KMail*, *KOrganizer*, *KAddress* and *KNotes*, and most of this tutorial will relate to the calendar portion, *KOrganizer*.



KOrganizer: please refer to the numbered parts in the text for explanation.

The Toolbar

KOrganizer has the normal KDE look with a toolbar ranged across the top of the main window, divided into general and application-specific sections.

1 Standard icons

These are the icons you can expect to see on any KDE applications, covering jobs such as opening and saving files, printing, undoing, cutting, copying and pasting.

2 Navigation and Search

As well as the standard tools, the 'Main Toolbar' also includes icons for searching events for a specific text string, a pair for moving through the calendar and a trio of buttons for zeroing in on the current date, adding a new event and adding a new 'To Do' task. The navigation duo inevitably behave differently, depending on your 'current view'. In the month view, for instance, it moves you back and forth a week at a time.

3 Views Toolbar

This has another selection of icons covering how the main calendar window displays data. We can see a single day, the working week (traditionally running from Monday to Friday), whole week, next three days or a calendar month. There are also options to see 'What's Next', which displays a detailed view of the next job on a To Do list, a list view of all upcoming appointments,

a full-screen To Do list and a Journal. The Journal entries are tied to individual calendar days and are excellent for using when notes need to be attached to certain dates. *LXF*, for example, has used the Journal mostly for note-taking during lectures and meetings, which saves the trouble of organising disparate *OpenOffice.org* files for later revision; we can simply open up the Journal window and click through the dates on the 'Calendar widget' (on the left-side pane) to find the right note.

4 Schedule Toolbar

A quartet of icons that is used for managing group meetings. These icons can be used, in conjunction with entries in the main calendar, to remind expected attendees of the time, date, agenda, etc. of a meeting or to respond to another user's meeting request. See below for a full description of setting up meetings.

5 Calendar Widget

You can see another example of it by clicking on KDE's clock. Across the top are navigation buttons which move through months (inner buttons) or years (outer), and a month/year label which, in *KOrganizer 3.2*, can be clicked on to open a drop-down to go to any month in any year. The body of the calendar is arranged so that dates with events recorded will be highlighted in bold, and the current day is shown by a black bounding box. Clicking on a date – in all but the month view – will take you to the 'day view' for that date. If your calendar has more recurring events than is rational and every day is displayed in bold, you can opt to hide events that occur daily or weekly from the Calendar Widget. Go into 'Settings>Configure KOrganizer..' and choose the 'Views' option. Now deselect 'Show events that recur daily in date navigator' and its weekly counterpart and hit 'OK'.

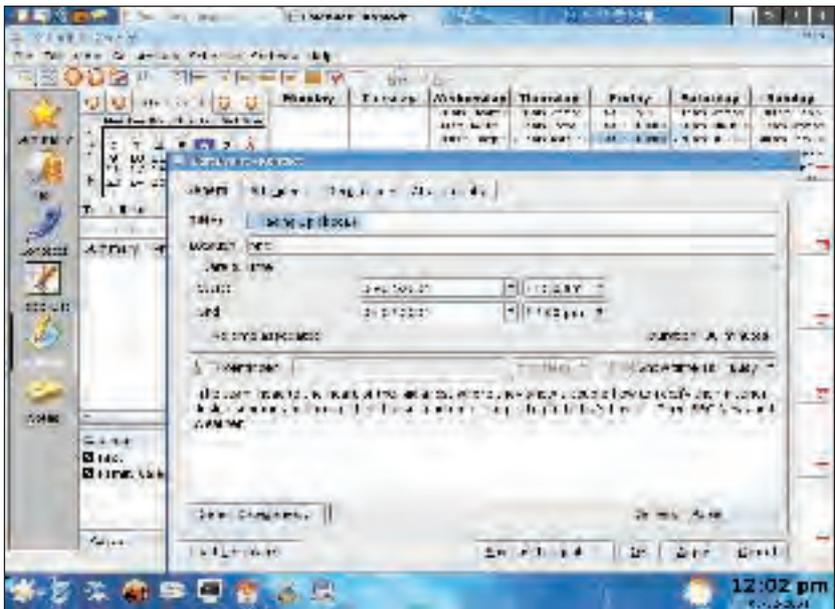
6 Making Entries

Entries in *KOrganizer* can be divided into two groups: Events are given a specific dates and time and go into the main calendar, while To Do entries have their own space in the side pane and on the desktop. The difference is they are not bound to a specific date, even though you can set a deadline for a project which does appear on the calendar. In most respects the To Do entry is the same as a normal one, but there are a couple of differences. A To Do entry can be given a priority, which helps you to (surprise) prioritise your workload, and also a completion designation shown as a percentage. While this may seem a little superfluous to some, being able to see when something is due, how important it is and what needs to be done to finish the job makes time management a less arduous task. To Dos can also be given 'sub-To Dos' if, for example, a project is to be done in well-defined stages.

7 – The Calendar section

There are times when manually adding a slew of dates to your calendar would take so much time as to be pointless. For instance, adding UK Saint's days and holidays by hand will take some time simply to find out each of these dates, let alone input them. Fortunately, someone thought of this when creating the iCalendar standard (with the .ics extension) which has been adopted by KDE, Ximian, Apple and others. What this means is that another user may have gone to the bother of creating a calendar featuring the events we want to include in our own efforts. We just need to locate the correct file, download and import it into our application.

A number of websites exist merely to serve up calendars. If your tastes are mainstream and, it must be said, dominated by events in the US, the first port of call should be



www.apple.com/ical/library/, which has a range of downloadable files including one of UK holidays. Apple's decision to standardise on the .ics format has inspired many sites such as www.icalshare.com, which is home to a collection of over 1,600 calendars arranged into sensible categories. There is still an American focus, but it's not so pronounced; for instance there are calendars featuring the fixture lists for the majority of Europe's big football clubs from Arsenal to York City.

There are a couple of ways you can add data to *KOrganizer 3.2*, which makes it particularly flexible when it comes to sharing data with family members, or keeping the default view clean. The first, and most permanent, way to add dates to a calendar is to use the Merge facility under 'File>Import>Merge Calendar..'. The advantage of this is that it integrates the selected .ics file into the currently highlighted calendar and is ideal for adding things like national holidays. The disadvantage (which is also a weakness in *Evolution*) is that once the data are imported, it is impossible to de-merge easily. So while importing the history of NASA may seem like a cool idea, it can make for a very busy calendar, and when quick, visual access to information is important, this is not ideal. Users of *KOrganizer* versions prior to 3.2 and *Evolution* can, of course, have multiple instances of the application open featuring different views. Fortunately the latest version of *KOrganizer* adds another way of importing calendars while keeping information discrete, even though it is displayed as though it were merged.

In the Calendar box select the 'Add' button, to launch the import dialog box. There are a number of options here, but we want to get a calendar from a local file, so select that and hit the OK button.

You can now use the browse button to specify the file to import, rename it if necessary and select whether the data is in iCalendar or vCalendar format. Finally hit 'OK' again and wait for the file to be loaded. You should now have two entries in the Calendar box and clicking the radio button to the left of the entry will display or hide the data in the main screen. This is a great feature for sharing individual calendar information among groups, but to do so, we first need to make our file accessible to other users.

KOrganizer keeps its default calendar in a hidden directory within each user's /home. Multiple users on one machine and

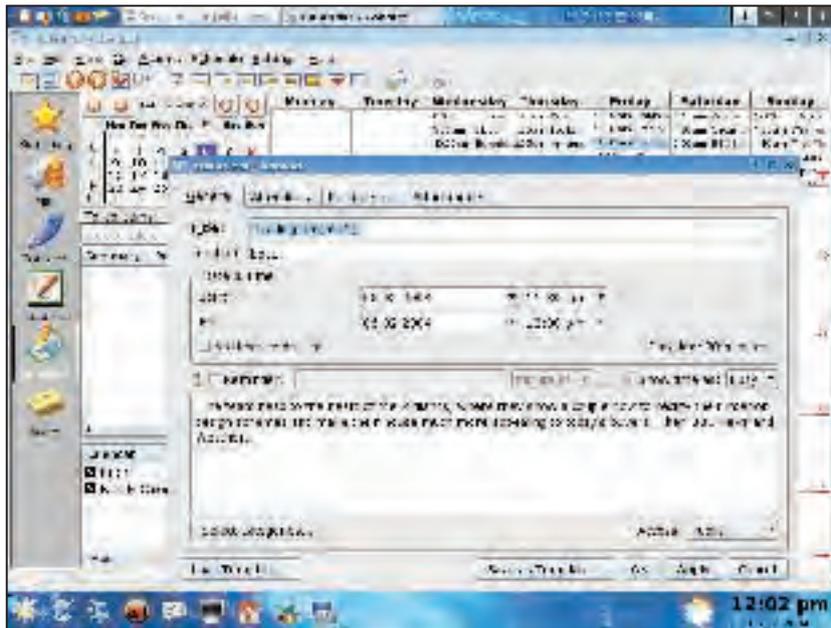
You can rename imported files.

TIP

It is possible to add other icons to any of these numbered sections by right-clicking on a toolbar and selecting 'Configure Toolbars..'. Available icons are listed on the left. Simply highlight what you'd like to add, remove or move and click the appropriate arrow key.



TUTORIAL Beginners' Linux: KOrganizer



www.project24.info provides a number of services including UK TV listings in .ics format.

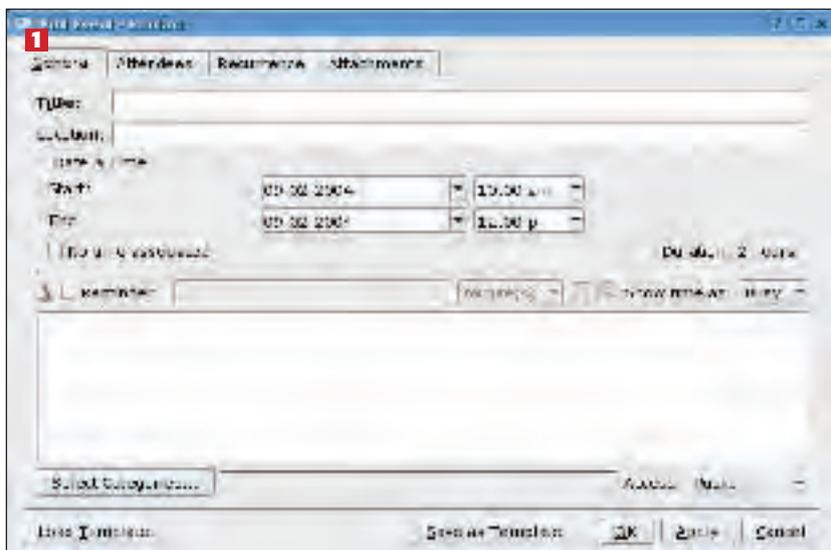
◀ across a small network will access our calendar (without the need to install a Groupware server!), so the first task is to move each user's calendar to a location accessible locally and on the network. Highlight the default entry and select 'Edit'. Changing the location is now simply a case of adding a new path under the 'Calendar in local file resource settings' section. Change it to something like '/home/user/filename.ics' which will ensure that others can read the file, but not write to it. Unless there's a good reason to do so, it's best to keep iCalendar as the format for the file, and you should manually add the extension as it doesn't get done automatically.

Now another user can select 'Add', input the path to your calendar and view it in the normal way. As the calendars are stored separately and loaded dynamically, the most up-to-date information is available at all times.

It is also possible to hook up to a remote calendar on the Internet which makes getting things like TV listings and other dynamic data a possibility. It also means you can subscribe – in Apple Mac OS X-style – to webcalendars on the Internet and, with the right service, even post your own public calendars, full of the events, occurrences and anniversaries that you have put together yourself. The procedure is almost the same as for a local file, but instead of an ordinary path, you must input a fully

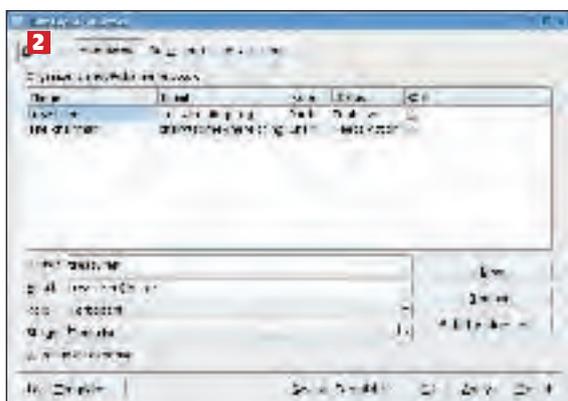
RUNNING YOUR LIFE

The more you put in, the more you get out...



The creation of To Do entries and Event entries is a similar process, so we'll go through that here. In the main toolbar select New Event (or New To Do) to launch the Event dialog.

1 In this tab, we add the basic details of the event such as location, date, start and finish times (if appropriate). We can also configure a reminder to go off, by either playing a sound or launching an application, a set period (minutes, hours or days) before the date of the event. We'll also use this dialog to assign a category to the entry such as Birthdays, Entertainment, Holiday or Business. Categories are vital for filtering events, but as you can set the calendar to display events using specific colours (do 'Settings>Configure KOrganizer>Colours' and choose the category and colour you want to associate with it) they also provide a very quick, visual overview of your diary. One thing to note – at least its the case we have found with degrading a file from 3.2 to an earlier version – is that importing a calendar doesn't import the colour values you may have initially set, so if you work on two machines with the same file, colour associations will need to be set up on both. You may also notice, depending on your choice of distribution, that the Month View doesn't pick up these colours. If this is the case, go into the Configure KOrganizer dialog again and select views. Among the options you will find 'Month view uses category colours'. Click on the option and hit 'OK'. You may have to refresh the view (simply select a different view and then go back) to see the effect of these changes.



2 The second tab is for – optionally – setting up meeting Attendees. This is vital when attempting to gather together a group of busy people. Attendees can be assigned roles within the meeting and, with a provided email address, can be sent notification quickly if, for instance, the time or venue of a meeting is altered. Once the

qualified URL for downloading and (if necessary) uploading files and select how often the remote site should be polled. If you want to be really smart, it's possible to store a calendar on your personal web space and link there, giving you access from any web-connected – and suitably configured – computer.

8 - The main window

This, the most prominent part of the application, is where you can get both a broad overview and the close-up look of your up-and-coming events. It is where all our colour-coded, prioritised reminders can gang up on us. The standard view displays the current month and adding events is simply a case of double-clicking inside a date box to open the dialog and type in the data (see above). However, while you may start and end the working day in the 'Month View', it's possible that some of the alternative methods of visualising your time will come to dominate your usage, depending on the way you work.

If work days are filled with meetings, lunches and deadlines, the day view with its shocking red time bar (updated every minute!) will allow you to stay on top events, or at least see how swiftly they're approaching.

By default the day view breaks up the day into 15-minute slots and double clicking (or right-clicking and selecting 'New

Event') any available space will bring up the Event dialog. In the day view, you should notice the 'working day' is highlighted in a nice shade of lilac, and by default this runs from 8.00am to 5.00pm. If this is at odds with the hours you keep, you can adjust the start and end of the work day by doing 'Settings>Configure KOrganizer>Time & Date' and changing the numbers. You can also use this dialog to set the default appointment time and duration, and opt to exclude the working day highlight on Saturdays and holidays.

At the top of the main window is a block view of appointments, providing more details on the day's events. Double-clicking an event, again, opens the Event dialog for editing. Both the time bar and appointment blocks are also features of the other (non-month) views.

Conclusion

This has been a basic guide to a very powerful application. Even if you've yet to upgrade – and we would recommend it for the 'separate but together' calendars alone – *KOrganizer* (and other PIM software) can make a difference to your quality of life. Less stress is always a bonus! However, when you integrate the calendar portion of *Kontact* into the rest of the suite it really comes into its own... [LXF](#)

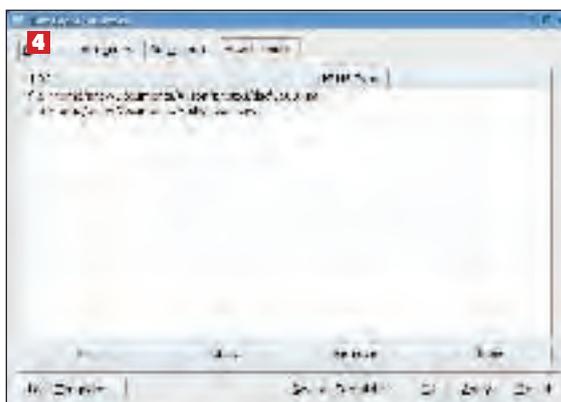
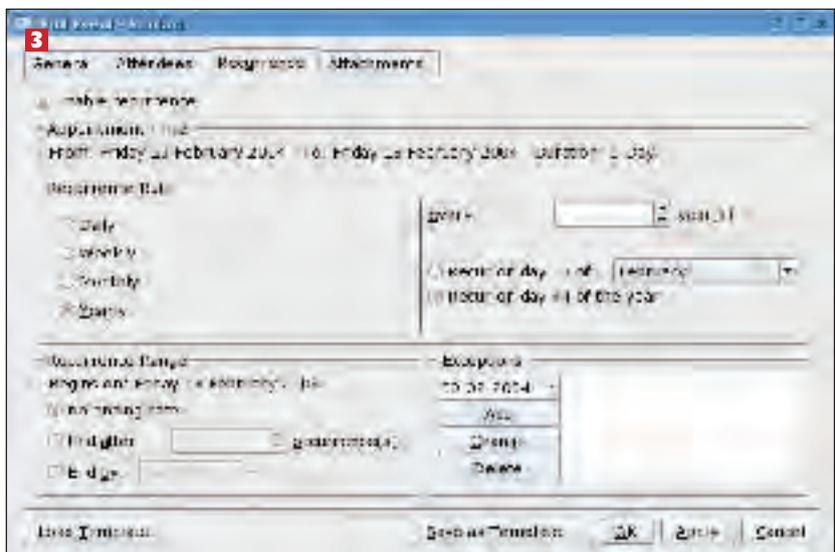
appropriate information has been entered, every member of the meeting can be contacted using the 'Publish' icon (see **4** on the annotated screenshot) to launch *Kmail* with addresses, subject, agenda (notes) and, if necessary, attached documents added and ready to send. The Publish tool also attaches an iCalendar file to the email which can then be merged into the recipients' calendars.

3 The third tab is used for setting recurring events such as regular budget meetings, anniversaries, birthdays, etc; with options including frequency of occurrence, exceptions to the general rule, and an optional 'end date'. If the event comes under the To Do heading, the Recurrence tab is absent.

4 The final section allows us to attach documents to a specific event. These can be anything, agenda, a presentation, photos, evidence, etc. The only difficulty here is that the *KOrganizer* developers have yet to implement a browse icon on the file dialog (send in a feature request!) so, in order to attach anything, we need to know its complete path and filename. Simply select 'Add..' and insert the information. You can use this dialog to edit (the location) or remove files, and selecting an attachment and clicking 'Show..' will open the file in its associated application. This option is not available in older versions of the software though.

NEXT MONTH

You'll have noticed that *Evolution* has been mentioned a couple of times in this tutorial. For those among you who have a preference for the GNOME desktop, we haven't forgotten that your life might need some organising too – we'll be giving *Evolution* a similar treatment next issue.





FLEX & BISON

Compiler writing

Like some sort of second-class citizen, SKYLang gets relegated to just two pages this month. But **Paul Hudson** thinks it'll be two pages you won't want to miss...

Last issue, we looked at how to add functions to SKYLang, and you'll be pleased to know that the majority of the hard work is already past. However, our initial implementation of functions is perfectly acceptable if all you want to do is call a function, but do nothing with its return value. In practice though, you want to be able to read values back from your function calls so that they can play a meaningful part in your script. Thus, the act of accepting return values from functions is the topic of this month's installment, and fortunately for you it's not all that difficult!

Functions as statements

In our current scheme a function call is a statement by itself, which means that **foo()**; is treated in the same way as **echo \$bar**;. This is a very simple way to treat them, and it works fine if you don't care about return values, but we're going to rewrite that so that our functions become fully fledged members of the script. By changing them from statements to expressions, we'll be able to write code like **echo foo() + bar()**, as they will be treated like any other value.

Some languages – such as Pascal, for instance – will split the return value of a function from actually exiting the function; which looks like this:

```
function somefunc() : integer;
var
  i: Integer;
begin
  result := 556;
  i := 1;
end;
```

The return value from that function (and indeed all functions) is stored in the result variable, but assigning to it doesn't make the function end. Instead, **i** will be set to 1, then the function ends and the return value is sent back. Other languages, well, in fact *most* languages, return a value and exit the function at the same time. In C, this is accomplished with the **"return"** keyword, which is what we'll be implementing here.

Thanks to the way **Bison** works, effecting this change is as simple as copying the **T_FUNCTION_NAME T_OBRACK T_CBRACK T_SEMICOLON** action to the expressions block in `skylang.y`. I put mine under the **T_STRING** action. Note that because our function call is no longer an entire statement, it must not end in a semicolon – snip off the **T_SEMICOLON** from the action once copied.

In order to handle return values, we need a new **"return"** keyword, so add **T_RETURN** to the end of **%token <var>** in `skylang.y`. We need to back this up with a lexer match, so add the following below the decrement match in `skylang.l`:

```
"return" {
    return T_RETURN;
}
```

The return operation itself will be done with the **SOP_RETURN** opcode, so you will need to add that to the end of the **SKYOpType** enum in `skylang.h`.

The details

That's the obvious stuff out of the way – now onto the actual details. To pass a return value back from a function call, we should make the **execute()** function of our `oparray` return a value. To make a function return a value inside SKYLang, we

DOWNLOAD THE CODE

You can grab the code for this tutorial from the *Linux Format* website: www.linuxformat.co.uk/compilers/lxf52.zip

need to use the operation result as the value of the function expression. Finally, to handle the **SOP_RETURN** code, we also need to have a case in our main `execute()` function.

To make the first change, you need to edit the function prototype for `execute()` in `skylang.h`. It's currently `void execute();` but it needs to be `SKYVar* execute();` so that the function returns a variable for use in other operations. This change also needs to be made in `skylang.y`, so change `void SKYOpArray::execute() {` to `SKYVar* SKYOpArray::execute() {`. As we now need to always return a value from this function, add this next line of code beneath the definition of "jump" in `execute()`:

```
SKYVar* return_value = NULL;
```

This needs to be backed up by `return return_value;` at the very end of the `execute()` function, so that it returns an empty value by default. In order to make the return value from our function equal to the parameter passed to `return` (eg `return 1;`) we need to add one simple line to the end of the **T_RETURN** action:

```
$$ = new_op->result;
```

With that in place, the return value from the function is used as the value for the expression, which means in the statement `echo foo();` the return value of `foo()` is used as `op1` in the call to `echo`.

At the bottom of the current list of statements, you need to add an action for **T_RETURN**, which looks like this:

```
T_RETURN expression T_SEMICOLON {
```

```
SKYOp *new_op = new SKYOp;
```

```
CurrentOpArray->opcodes.push_back(new_op);
```

```
new_op->opcode = SOP_RETURN;
```

```
new_op->op1 = $$;
```

```
};
```

In that code, the return statement is added to the current `oparray`, using the value of the expression as `op1`. Keep in mind that **expression** is calculated beforehand using our existing rules, so we can write `return 1 + 1 + 1 * 3;` and it will work fine.

Now, to add the `execute()` case for **SOP_RETURN**. As you've just seen, the actual evaluation of values is done before we execute the actual return, meaning that the value to return has already been pre-calculated and placed into `op1` of the **SOP_RETURN** statement. As the return keyword returns whatever is used with its only "parameter", this is very simple:

```
case SOP_RETURN:
```

```
return op->op1;
```

So far, the changes have all been very minor, but there is still one key component missing: actually handling the return value from `execute()` and assigning it for use. Now, if you're thinking, "but didn't that `$$ = new_op->result` line do that already?" then you're sort of on the right tracks, but you're getting derailed somewhere around the difference between the result of a function and the result of an opcode.

Put simply, when the `execute()` function returns a value, it returns control of the interpreter back to the **SOP_CALLFUNC** case block. However, the return value is still floating around in the C code – it has yet to actually be assigned to the return value in our opcode. To handle this, our **SOP_CALLFUNC** block needs to get a little more complicated, as we need to capture the return value from `execute()`. However, it's not just as easy as assigning the return value to `op->result` because that would change `op->result` but not anything else that also points to the same value – it would essentially make `op->result` point elsewhere. The easiest way around this is just to copy all the values across, which makes the new **SOP_CALLFUNC** case look like this:

```
case SOP_CALLFUNC:
```

```
if (Functions[op->op1->charval] != NULL) {
    return_value = Functions[op->op1->charval]->oparray-
>execute();
    op->result->intval = return_value->intval;
    op->result->charval = return_value->charval;
    op->result->floatval = return_value->floatval;
    op->result->type = return_value->type;
}
break;
```

Testing it out

Although that was a bit rushed, I hope you've managed to follow along. Making your functions return a value isn't a paradigm-shifting change by any means, but there are enough potential slip-up points in there that you might have problems. As always, I recommend you compare your source code to my own (see the box, *Download the code*) to make sure that *a)* you didn't make any mistakes or miss a point from previous issues; or that *b)* I haven't slipped in a little change and forgot to mention it. To verify you have your code working correct, try running this script:

```
function foo() {
    echo "a";
    echo "b";
    return 2 + 2 + 2;
    echo "c";
}
```

```
echo "1";
echo foo() + foo();
echo "3";
```

If your compiler is working properly, that should output **1, a, b, a, b, 12, 3**, showing that functions are operating as expressions as planned, and also that return values are being sent back properly. Note that **c** is never printed out, as calling `return` from inside `foo()` causes SKYLang to exit the function immediately.

All being well you should be able to return all sorts of values from your functions, even the values returned by other functions. This is demonstrated here:

```
function foo() {
    return bar();
}
function bar() {
    return 1;
}
echo foo() + foo();
```

There `foo()` is called, which in turn calls `bar()`, which returns **1**. As a result, the above script should output **2**.

Conclusion

Although we've yet to implement function parameters, handling return values is a great step forward that edges us all the closer to finishing SKYLang. Although functions are really only complete once they accept parameters, having them return values like above does mean that code re-use is easy to do. There are quite a few unanswered questions, such as what happens if you set a value in the function – is it also set in the global scope, or is it only set at the function scope?

That sort of question is best answered once functions can accept parameters, because that requires a local symbol table also. As a result, lots of things come together, which should make next month *much* harder – you have been warned! [LXF](#)

NEXT MONTH

The next step is to allow functions to handle parameters, which adds all sorts of complications – but also all sorts of new possibilities for our fledgling language.



TROUT WARS

RAIDERS OF THE LOST POND

CODING WITH SDL

Game programming

So you want to write your own game? Ever the expert at playing with himself, **Paul Hudson** introduces us to the subject in the first part of our new series...

An oft-quoted phrase is “everyone has a book inside them,” and that may well be true. However, in the world of geeks, I think it can safely be rewritten as “everyone has a game inside them,” and that includes you too! In this series we will be looking at creating a game from scratch using C++ and the Simple DirectMedia Layer (SDL). If you read “C++” there and considered putting flicking past, relax – this series is aimed at C++ novices, which means if you are proficient in the use of other programming languages you should be fine here.

Now, to whet your programming tastebuds: we are going to create a side-scrolling 2D game with sprites, animations, music and sound effects, and we are going to do it all in using simple step-by-step programming that hopefully everyone will be able to follow. Before you get *too* excited, I should let you know that this issue we’re just going to be discussing how it will work and doing some *very* basic ground work – things start getting playable by next issue!

When I was talking this idea over with some of the chaps on *PC Format* magazine, I was told “no one plays 2D games any more, teach them 3D!” If you’re thinking the same, I’ll say to you what I said to them: teaching 3D games programming involves more theory than programming. That said, we may get to 3D if there is any interest – please do write in to tell us what you want to know.

SDL history

Simple DirectMedia Layer is not a particularly accurate moniker, as it is by no means simple. Yes, it is quite easy to learn and use, but it can do a massive amount: render 2D and 3D graphics, play back music files and sound effects, read input from various devices, render fonts on screen, and even handle network connections for multiplayer games. Furthermore, it can do so whilst being entirely cross-platform, which means the game we are going to develop will work on Linux, Mac OS X, Windows, BSD, and others; so, *simple* is a description of its learning curve, as opposed to its capabilities! To give you an idea of what it’s capable of, SDL was used by Loki Games to make its port of the popular game *Civilization: Call to Power*, pictured opposite.

The core functionality of SDL allows you to do graphical rendering and handle input, but there are various modules you can add on to perform extra tasks. *SDL_Mixer*, for example, gives you access to audio hardware so you can play back MP3s and such, and *SDL_ttf* lets you render TrueType fonts to the screen.

SDL itself was written by Sam Lantinga, the lead programmer from Loki Software who worked on *Call to Power*, *Railroad Tycoon II*, *Tribes 2*, and more. Since then he’s moved on to Blizzard Entertainment where he has been working on *World of Warcraft* among other things, which puts him at the top of the games industry. Sam also wrote several of the most popular libraries for SDL, so we all have a lot to thank him for!

First steps

The basic SDL program we’re going to create to start with looks like this in pseudocode:

```
Init graphics and sound
While (true) {
  Render scene
  If close button has been pressed, exit
}
```

Looks easy, huh? Sadly, it’s not that easy in C++, but it certainly shouldn’t tax you. In order to make the code as friendly as we can, we will be using classes and objects to break up the game into logical parts. The master part will be the class **CTWGame**, which represents the game itself, and will handle creating the game window, creating levels and loading music, etc. Later on, we will also have classes for the player, enemies, and other game objects as this makes life much easier later on.

If you hadn’t guessed already, the game we’ll be working towards will be called *Trout Wars: Raiders of the Lost Pond* – but you’re welcome to pick your own title and customise everything printed here for your own needs. But we’re getting ahead of ourselves – the very first thing we’re going to do is create the most simple SDL application, which means we’ll create a window for our game then end the program immediately. To get this

simple task done, create the two files we'll be working with throughout this project: `TroutWars.h` and `TroutWars.cpp`.

Here's how the first version of `TroutWars.cpp` should look:

```
#include "TroutWars.h"
int main(int argc, char *argv[]) {
    SDL_Init(SDL_INIT_VIDEO);
    atexit(SDL_Quit);
    SDL_SetVideoMode(SCREEN_WIDTH, SCREEN_HEIGHT, 32,
        SDL_HWSURFACE|SDL_DOUBLEBUF);
    return 0;
}
```

And here's the code for `TroutWars.h`:

```
#include <SDL.h>
#include <stdlib.h>
#define SCREEN_WIDTH 800
#define SCREEN_HEIGHT 600
```

It's just a handful of code lines across both files, and yet if you compile and run it you should see a black window pop up temporarily, then close itself. So, how do so few lines of code get us a working SDL window? Well, line one calls the `SDL_Init()` function, which starts up the SDL library, as well as any extensions we want to use. `SDL_INIT_VIDEO` is passed in, which means we want to initialise the video – we could also have passed in other parameters by combining them with a bitwise **OR** - |. That is, to initialise video *and* sound, we would use `SDL_INIT_AUDIO|SDL_INIT_VIDEO`. Sound is not used at this point, so I'd recommend staying away from anything other than `SDL_INIT_VIDEO` for the time being.

The next line calls the `atexit()` function, which is a standard C call provided by `stdlib.h` to register shutdown functions. The `SDL_Quit()` function terminates the SDL program, cleans up the library, and frees up any internal SDL resources that were allocated, so we want to call it whenever the application shuts down. Rather than having to catch every line of code that exits the application and put `SDL_Quit()` in there, it's much easier to use `atexit()`, which does the same thing without the hassle!

The third function call is to `SDL_SetVideoMode()`, which creates our game window and does so by taking four parameters: width of screen to create, height of screen, number of bits per pixel, and special flags to use. Calling `SDL_SetVideoMode()` actually returns a handle to the screen for us to use for drawing, but this we don't need this just yet. The fourth parameter, special flags, has several options you can use, but the key ones are:

SDL_SWSURFACE Create the video surface in system memory
SDL_HWSURFACE Create the video surface in video memory
SDL_ANYFORMAT If a video surface of the requested bits-per-pixel is not available, use whatever is available
SDL_DOUBLEBUF Enable hardware double buffering
SDL_FULLSCREEN SDL will attempt to use a full screen mode
SDL_OPENGL Create an OpenGL rendering context.

You can combine these together in various ways to get your game to work as you want it to. Using `SDL_HWSURFACE` is generally best, as is `SDL_DOUBLEBUF`. Many people find that using `SDL_FULLSCREEN` with `SDL_DOUBLEBUF` causes a flickering that can only be fixed by dropping `SDL_DOUBLEBUF` – your mileage may vary.

Note that the first two parameters to `SDL_SetVideoMode()` use our macros from `TroutWars.h`. **SCREEN_WIDTH** and **SCREEN_HEIGHT** have been defined as **800** and **600**



respectively, so these will be passed into the function. These two will be used in various other places so that all game information and screen size can be changed from just one location.

All the code is wrapped up in the function `main()`, which is the default function run when a program starts up. It returns an int (short for integer – a whole number), which is passed back using the `return 0;` statement. Zero, confusingly enough, means “everything went OK”, so by returning **0** back, we're saying everything worked as planned. The two parameters that `main()` accepts are the standard `int argc`, which is the number of arguments passed in, and `char *argv[]`, which is a pointer to a character string containing the arguments themselves. We won't be using these, so you can safely ignore it all.

`TroutWars.h` also contains an include for `SDL.h`, which contains the core SDL header information. You may find this needs to be `SDL/SDL.h` and not `SDL.h`, depending on your system configuration – if you get many errors from your compilation, this is the first thing you should try.

Adding some class

You already have the most basic SDL game – it shows a blank window and exits immediately. This isn't much fun, as I'm sure you'll know as you've probably ‘played’ it a few times already to make sure it works, but before we go about starting to add things to the game we first need to implement the class architecture mentioned earlier. That is, we don't want to have the `main()` function knowing about the screen size or the bits-per-pixel used. Instead, this is all going to be managed by the `CTWGame` class, so what we're going to do is change our existing code so that it all runs through this class, and, at the same time, we're going to add a little more functionality:

- Both `SDL_Init()` and `SDL_SetVideoMode()` return data. These should be checked for a non-error response and stored somewhere if appropriate
- Our game should loop until it is exited. During this time, it should keep drawing a black screen
- The input required to make the game exit will be the user clicking the close button.

Adding these three pieces of functionality – as well as the class structure – will be the most difficult thing we'll have to do for a while, because it makes our simple SDL code from before look a ➤

Civilization: Call to Power was one of the finest games that Loki ported to Linux, and it was written entirely in SDL.

COMPILING YOUR GAME

Once you have SDL installed correctly, compiling is quite easy. First, make sure you have both `g++` and `sdl-config` installed on your system – try running them with `--version` and make sure you don't get any errors.

If you have both of these two, compiling is a case of running this:

```
g++ TroutWars.cpp -o
TroutWars `sdl-config
--libs`
```

That should create an executable called `TroutWars`, which you can run using `./TroutWars`.

Note that the quotes around the `sdl-config` part are actually backticks, which means that `bash` will execute `sdl-config --libs` and put its output into `g++`. If you don't have `sdl-config`, you should be able to use

```
-ISDL -ISDL_mixer -
ISDL_ttf
```

to get the same results.

TUTORIAL Game programming

◀◀ lot more intimidating. Hopefully you'll be able to see that it's really not so different once you get past the object-orientation! Here's the new TroutWars.h file:

```
#include <stdlib.h>
#include <time.h>
#include <SDL.h>
#define SCREEN_WIDTH 800
#define SCREEN_HEIGHT 600
class CTWGame {
public:
    SDL_Surface* sfcScreen;
    CTWGame();
    void ClearScreen();
    void DrawScene();
    void Play();
};
```

You'll notice there are two extra files included – they are used later. The main thing in TroutWars.h now is there **class CTWGame**, which is our game. Inside, there is one variable, **sfcScreen**, which you can ignore for now, and functions: **CTWGame()**, the constructor(); **ClearScreen()**, which draws a large black box over the screen every frame to clear it; **DrawScene()**, which is where scene drawing code will be placed; and also **Play()**, which will host the main game loop.

Here's the new code for TroutWars.cpp:

```
#include "TroutWars.h"
CTWGame::CTWGame() {
    srand((unsigned)time(NULL));
    if ( SDL_Init(SDL_INIT_VIDEO) < 0 ) {
        printf("Unable to initialise SDL: %s\n", SDL_GetError());
        exit(1);
    }
    atexit(SDL_Quit);
    sfcScreen = SDL_SetVideoMode(SCREEN_WIDTH,
    SCREEN_HEIGHT, 32, SDL_HWSURFACE | SDL_DOUBLEBUF);
    if ( sfcScreen == NULL ) {
        printf("Unable to set SCREEN_WIDTHxSCREEN_HEIGHT
    video: %s\n", SDL_GetError());
        exit(1);
    }
}
void CTWGame::ClearScreen() {
    SDL_FillRect(sfcScreen, NULL, 0);
```

```
}
void CTWGame::DrawScene() {
    ClearScreen();
    SDL_Flip(sfcScreen);
}
void CTWGame::Play() {
    int done=0;
    while(done == 0) {
        SDL_Event* event;
        while ( SDL_PollEvent(event) ) {
            if ( event->type == SDL_QUIT ) done = 1;
        }
        DrawScene();
    }
}
int main(int argc, char *argv[]) {
    CTWGame* game = new CTWGame;
    game->Play();
    return 0;
}
```

Although that's about nine times longer than the previous code, you might be surprised to see that compiling and running it produces almost exactly the same result. The only difference is that the black screen doesn't go away immediately – you need to click the close button to get rid of it.

How it works

So, what does the new code do to make it worth the extra length? I'll break it down starting at **main()**, as that's where the program starts. Notice now that to get the game started we just create a new **game** variable of type **CTWGame**, and call its **Play()** function. This is much better than **main()** having to know how big the screen should be, etc, and having this code encapsulation will prove to be a boon later on.

When we call **new CTWGame**, the class constructor is called. If you're new to C++, the constructor is a function that sets up initial values for the object, and is always named the same as the class itself. Thus, our CTWGame constructor is **CTWGame::CTWGame()**, and inside that function you should recognise most of the calls.

The first line is new, and it calls both **srand()** and **time()**. The **srand()** function seeds the random number generator in C++, and you need to pass a pseudo-random number into it to get started. The best way to do this is to use the **time()** function, which returns the current system time. The (unsigned) part before the **time()** call is there to avoid a compiler warning about type conversion, and can be safely removed. You should only ever call **srand()** once, as calling it more than once effectively resets the random number generator, giving you the same 'random' numbers all over again. Using **srand()** and **time()** requires inclusion of **time.h** and **stdlib.h** in TroutWars.h.

Moving on, **SDL_Init()** is called like last time, except now we check whether the return value is less than **0**. If **SDL_Init()** has a serious error, it will return less than **0** and also fill an internal string with the error reference. So, to check whether SDL was initialised successfully, we check whether the value was less than **0** and print out an error message if it was. Printing the error message is done using **printf()**, and passing in **SDL_GetError()** as the parameter – the **%s** is replaced by **SDL_GetError()** at runtime, making a meaningful error message. Finally, we call **exit()**, another function brought in with **stdlib.h**, which terminates the program immediately passing back its parameter as the return code. Using **exit(1)**

INSTALLING SDL

An integral part of many modern distros

If you're using a home distro such as Mandrake, SUSE, or Fedora, you're likely already have SDL installed. If you don't, but you are using a home distro, you should be able to install SDL and all its accompanying libraries using your installation manager. On Mandrake, for example, you would install **libSDL1.2**, **libSDL_mixer1.2**, **libSDL1.2-devel**, **libSDL_ttf2.0**, and **libSDL_image1.2**, although your distro may have different version numbers. **NOTE:** the **-devel** version of SDL is required to get the header and library files, otherwise all you will have is the ability to play SDL games rather than make them!

If your distro does not have SDL available for binary install, you have two options: go to

<http://libsdl.org> and grab the source code yourself (www.libsdl.org/download-1.2.php), or alternatively you can try installing the pre-packaged version. You will need to do the same for **SDL_ttf** (www.libsdl.org/projects/SDL_ttf), **SDL_net** (www.libsdl.org/projects/SDL_net), and **SDL_mixer** (http://www.libsdl.org/projects/SDL_mixer) too.

Once you have extracted the source tarballs, compiling consists of these commands:

```
cd <SDL source directory>
./autogen.sh
./configure
make
make install
```

means, “terminate immediately, and pass 1 back to the system”. In the same way that the **0** in **return 0**; means, “everything went OK”, returning **1** means there was a problem.

The next line calls **atexit()**, which is unchanged from before. The line after calls **SDL_SetVideoMode()**, which is also unchanged from before with the exception that we now catch its return value and place it into **sfcScreen**. If you recall, **sfcScreen** was defined in **CTWGame** as **SDL_Surface***, which means it is a graphical surface that can be drawn upon or used to draw somewhere else. All graphics in SDL are handled using surfaces, and the return value from **SDL_SetVideoMode()** is the primary surface of our window – anything drawn onto this surface is shown on screen.

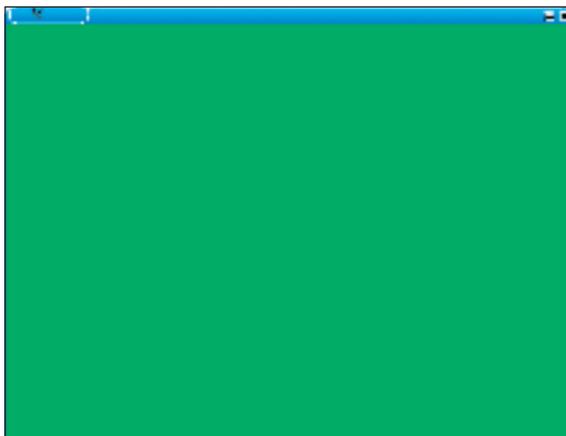
If the call to **SDL_SetVideoMode()** failed to set the requested video mode, it returns **NULL**. The next four lines of code are there to handle this eventuality, and equate to, “if **sfcScreen** is null there was a problem, so print out the error and exit immediately.” This error handling is the same as seen when **SDL_Init()** fails.

With the constructor call complete, **Play()** is called from **main()**. This function currently does just three things: loop around and around, check whether the close button has been pressed, and draws the scene.

The **done** integer variable is the key to the loop. If **done** is **0**, which it is set to by default on the first line, it means the loop should continue. Each time the loop iterates, we create an **SDL_Event** called **event**, which is used to handle events coming into the game window. These events include things like the mouse being moved, keys being pressed, joysticks being moved, etc, but right now we’re just interested in one event: the close button being pressed.

The **SDL_PollEvent()** function returns one event from the list of events that are currently waiting. There are usually quite a few events waiting, of which our ‘close button pressed’ event might be the very last one, so we use a loop to call **SDL_PollEvent()** repeatedly, passing in our **SDL_Event** object. The function returns true when an event was passed back, and also fills our event variable with information about the event that occurred. In order to handle each and every event, the **while** loop is: **while (SDL_PollEvent(event))**, which means “while there are still events to handle”.

Inside the loop, we look up the event’s type variable, and compare it against the SDL constant **SDL_Quit**. If they match, **done** is set to **1** and therefore the parent loop terminates. There are dozens of types of event types, from **SDL_KEYDOWN** for key presses to **SDL_JOYHATMOTION** for responding to a user pressing the hat switch on their joystick.



Using static function variables, you can get the colour to cycle; but it’s still quite a way from being fun!

The final part of the loop calls **DrawScene()**, so the program jumps to **CTWGame::DrawScene()**. This function, in turn, calls **CTWGame::ClearScreen()**, so the program jumps again. **ClearScreen()** uses a new SDL function, **SDL_FillRect()**, which takes a surface to draw upon, a rectangle to draw, and a colour. Normally you can use this function to draw coloured filled rectangles wherever you want on the screen, but if you pass **NULL** and **0** to the function it draws a rectangle over the entire screen and colours it black – perfect for our use.

With **ClearScreen()** finished, control jumps back to **DrawScene()**, where another new function is called: **SDL_Flip()**. If you remember our call to **SDL_SetVideoMode()**, we passed in the parameter **SDL_DOUBLEBUF** to make the screen double-buffered. Double buffering is a very simple drawing technique which uses two drawing surfaces rather than just one for the primary display. All drawing actions, such as our rectangle drawing, are done on what’s called the back buffer – an invisible screen stored in memory. At the same time as this, there is a front buffer being displayed on screen that contains the last frame. Once drawing is complete, the buffers are flipped over so that the back buffer becomes the new front buffer and *vice versa*, updating the screen. If you don’t use double buffering, all drawing is done live to the front buffer, which nearly always results in visible artefacts as the user’s eyes see part of the old picture and part of the new picture. To flip the front and back buffers for a surface, we use **SDL_Flip()** and pass in the **SDL_Surface**.

That’s **DrawScene()** complete, so control goes back to **Play()**, where the loop repeats itself again (and again, and again, and again...) so we’re all finished. When the close button is pressed, **done** is set to **1**, so the **Play()** loop terminates, **Play()** returns control back to **main()**, and **main()** exits.

Is the extra code worthwhile? Well, you’ve got to admit that it’s pretty cool to be able to say **Game>Play()**, and keeping **Play()** separate from the actual scene drawing is good too – this will become increasingly apparent as more code is added.

Last tweaks

I’d hate to end the first installment of this new series without actually showing anything other than a blank black screen, so replace your **ClearScreen()** function with this code:

```
void CTWGame::ClearScreen() {
    static int red = 0, blue = 0, green = 0;
    red++; blue += 2; green += 4;
    if (red > 255) red = 0;
    if (green > 255) green = 0;
    if (blue > 255) blue = 0;
    SDL_FillRect(sfcScreen, NULL, SDL_MapRGB(sfcScreen-
    >format, red, green, blue));
}
```

This uses a new function, **SDL_MapRGB()**, to create a colour based upon three integers: red, green, and blue. These are defined as static, which means their values are stored for the next time **ClearScreen()** is run. As a result, this function starts red, green, and blue off at **0** each, then increments them by a different amount of each time the function runs. This is then all passed into **SDL_MapRGB()** to create a new colour to use for our rectangle. **SDL_MapRGB()**’s first parameter is a pixel format, and you should always pass in the **format** variable of the surface you want the colour for.

This time, you should get the same empty window, but at least it changes colour smoothly! [LXF](#)

SUPPORT SDL

SDL is wholly supported by funding from the community, so why donate just a little money to help pay for development costs? The SDL homepage is at www.libsdl.org and there are instructions online for how to donate.

NEXT MONTH

Our game isn’t much good right now unless you make trance visualisations for your local dance club. However, we’ve put down a firm grounding for future issues, which means putting in the hard work now will pay off later. We have *big* plans for this tutorial series, so write in and tell us what you’d like to see included. Next issue we’re going to create a **CPlayer** class that will load a sprite from disc and move it around the screen based upon user input. If there’s space we’ll also be looking at SDL’s text-handling capabilities with the **SDL_ttf** library.

PHP VERSION 5

Practical PHP programming

Get your company organised with LDAP – Nick Veitch and Paul Hudson join forces to show you how...

Although databases allow you to do many weird and wonderful things with your information, their limits are well-defined and well-adhered to – there are several areas where a standard database just wouldn't work, and a specialised database system is used instead. The most common of these is LDAP, the *Lightweight Directory Access Protocol*, which is widely used to resolve name and contact information for people. When you type "Smi" into the address bar of an email using Mozilla and it fills in "th, John", that's an example of LDAP in action.

Directories are in common use, even in small organisations, as they are a quick and convenient way of centrally storing data on company assets. The obvious application is an employee directory, which can provide phone details, location information and even relationships. LDAP is also often used as an authentication and authorisation tool – providing a single repository allowing other services to validate requests – eg email servers, FTP, even printing.

There's nothing to stop you including hierarchical data for other objects too – printers are again a good example, as are servers of different types.

What makes a directory system like LDAP more specific in application than a database is the way that the data is organised. LDAP isn't really a database at all (in fact many implementations use external database applications to actually store the data), but

a protocol, and one that allows you to access the data in a particular way. LDAP data is hierarchical. The root DN (or Distinguished Name) of a directory establishes a scope, and additional objects and containers are placed on a tree that very much resembles a filesystem. The DN of an object, eg **cn=Nick Veitch,ou=people,dc=ixf,dc=org** establishes its place in that tree. This has some interesting effects – it means for example, there can be two entries for **Nick Veitch** (shudder), as long as they don't conflict in the hierarchy.

Each object has attributes, which loosely follow the idea of a key/value pair, which you are probably familiar with. LDAP actually allows for multiple values for a particular attribute key. This is very useful – for instance, an individual may have more than one phone number or email address – but it does mean that you need to deal with directories differently from a standard database.

PHP and LDAP

Fortunately for everyone, PHP has a special set of functions for dealing with LDAP. They are included with the main code, but aren't compiled in by default. If you have compiled your own version of PHP from source, you may need to reconfigure it with the option **--with-ldap** added to your usual configure options. For those with a package-based distribution (Mandrake/Red Hat/Debian or whatever) you'll probably find your vendor includes a *php-ldap* package to allow these functions to work.

Either way, you will need a local version of an implementation of LDAP available to provide access to the libraries. Most distributions come with *OpenLDAP*, which is perfect for this purpose. If not you can grab the latest version from <ftp://ftp.openldap.org/pub/openldap/openldap-stable.tgz>.

It would be advantageous for the purposes of this tutorial to have your own accessible LDAP server running on a local server under your control, particularly for the later sections covering authorisation. But for basic searching, you don't need your own server – there are plenty of servers open to the public! Since the data in the LDAP server is usually something that a company or organisation wants to publish, you'll find many LDAP servers open to public access for querying – see the box *LDAP resources* over the page for more details.

Starting out with LDAP

Before we write any code, we need to know at least two things about the LDAP server that we intend to develop for – its address and the root DN. As we mentioned before, the DN establishes the context of the server. Confusingly, there are currently two different ways this is normally done.

As LDAP grew out of X.500, it used to be common to define the root DN in terms of an organisation and a geographical location, eg **o=Future Publishing,c=UK**. This sort of made sense, because different directories could be maintained in different geographical locations for the same company. The downside is that there is a possibility of conflicts, because the actual organisation name is fairly arbitrary.

The more modern method is to tie the root DN more obviously with something unique, and taking a cue from the Internet, LDAP has borrowed the idea of using domain names. So you will often find organisations that use **dc=futurenet, dc=co.uk** instead. Some even combine the two approaches, so you may even come across **o=bath.ac.uk**, for example. It doesn't make a great deal of difference, as long as you know how the server defines itself in terms of a root DN.

Other information that can be useful is an idea of the structure of the data (*ie* an outline of the tree) and the schemas used. These latter describe the attributes given to individual objects, which are strictly controlled. You can't assign any old attributes (well, not easily), but the plus side is it gives you an idea of what information you can search for. Both of these pieces of information can usually be found out if you know the first two, just by exploring. For reference, here is a typical set of data from an LDAP server:

```
dn: dc=lx,dc=com
objectClass: domain
dc: lx
structuralObjectClass: domain
entryUUID: 52dd1f14-fbce-1027-8f27-eba47dc0e1f6
creatorsName: cn=admin,ou=admin,dc=lx,dc=com
modifiersName: cn=admin,ou=admin,dc=lx,dc=com
createTimestamp: 20040225110553Z
modifyTimestamp: 20040225110553Z
entryCSN: 2004022511:05:53Z#0x0001#0#0000

dn: ou=people,dc=lx,dc=com
objectClass: top
objectClass: organizationalUnit
ou: people
description: Container for people
structuralObjectClass: organizationalUnit
```

```
entryUUID: 52e1d036-fbce-1027-8f28-eba47dc0e1f6
creatorsName: cn=admin,ou=admin,dc=lx,dc=com
modifiersName: cn=admin,ou=admin,dc=lx,dc=com
createTimestamp: 20040225110553Z
modifyTimestamp: 20040225110553Z
entryCSN: 2004022511:05:53Z#0x0002#0#0000

dn: ou=admin,dc=lx,dc=com
objectClass: top
objectClass: organizationalUnit
ou: admin
description: Container for Administrators
structuralObjectClass: organizationalUnit
entryUUID: 52e290de-fbce-1027-8f2a-eba47dc0e1f6
creatorsName: cn=admin,ou=admin,dc=lx,dc=com
modifiersName: cn=admin,ou=admin,dc=lx,dc=com
createTimestamp: 20040225110553Z
modifyTimestamp: 20040225110553Z
entryCSN: 2004022511:05:53Z#0x0004#0#0000

dn: cn=admin,ou=admin,dc=lx,dc=com
cn: admin
objectClass: top
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
userPassword:: e1NTSEF9bnpGYIZMUlhpUWxbDJGamZzNXy=
uid: admin
mail: admin@lx.com
sn: admin
structuralObjectClass: inetOrgPerson
entryUUID: 52e31130-fbce-1027-8f2b-eba47dc0e1f6
creatorsName: cn=admin,ou=admin,dc=lx,dc=com
modifiersName: cn=admin,ou=admin,dc=lx,dc=com
createTimestamp: 20040225110553Z
modifyTimestamp: 20040225110553Z
entryCSN: 2004022511:05:53Z#0x0005#0#0000

dn: cn=Nick Veitch,ou=people,dc=lx,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
uid: nveitch
cn: Nick Veitch
givenName: Nick
sn: Veitch
mail: nveitch@lx.com
telephoneNumber: 2335
l: Bath
structuralObjectClass: inetOrgPerson
entryUUID: 52e58ffa-fbce-1027-8f2c-eba47dc0e1f6
creatorsName: cn=admin,ou=admin,dc=lx,dc=com
modifiersName: cn=admin,ou=admin,dc=lx,dc=com
createTimestamp: 20040225110553Z
modifyTimestamp: 20040225110553Z
entryCSN: 2004022511:05:53Z#0x0006#0#0000

dn: cn=Paul Hudson,ou=people,dc=lx,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
```

STILL ON PHP 4?

If you have yet to migrate to PHP 5, you should at least consider upgrading to PHP 4.3.5. This recently released update to PHP 4 includes dozens of bug fixes, as well as newer builds of both PCRE and GD. PHP 4.3.5 is a recommended upgrade for all PHP 4 users. Get it now!



```

uid: phudson
cn: Paul Hudson
givenName: Paul
sn: Hudson
mail: hudzilla@lxf.com
telephoneNumber: 2335
l: Bath
structuralObjectClass: inetOrgPerson
entryUUID: 52e76726-fbce-1027-8f2d-eba47dc0e1f6
creatorsName: cn=admin,ou=admin,dc=lxf,dc=com
modifiersName: cn=admin,ou=admin,dc=lxf,dc=com
createTimestamp: 20040225110553Z
modifyTimestamp: 20040225110553Z
entryCSN: 2004022511:05:53Z#0x0007#0#0000

```

This is a typical LDIF representation of the data. At the start it defines the domain that the whole directory refers to – the start of the tree, if you like. The DN that begins **ou=people** is a container that will be used to include all the people. After the DN at the beginning of each entry is a list of attributes. The attributes are largely determined by the schemas, which relate to the different classes, also specified as attributes in each entry. Because of the schema system, there is a convention for these attributes, which makes it a lot easier to know what you will be dealing with!

You'll also see that every entry includes timestamps indicating when it was created and last modified, as well as a record of who made these additions or changes. These are added automatically to every entry by the LDAP server, and while you may want to search on them, you won't need to worry about changing them with your own code.

Of course, if you have set up your own LDAP server, you'll know all of this information already. Unfortunately, this isn't really the place to go into great detail on the subject, but check out the LDAP resources box for more help on this.

Connecting and searching

Time, at last, for some more code. Here we'll implement a simple search, returning pertinent details from the LDAP directory:

```

<?php
// variables relating to the LDAP search
$daphost='192.168.130.91';
$ldapDN='dc=lxf,dc=com';
$filter='sn=';
echo "<h1>LDAP Query</h1>";
echo "Attempting to connect ...<br />";
$connection=ldap_connect($daphost);
echo "Connect result is $connection<br />";
if ($connection) {
    echo "Binding ...";
    ldap_set_option($connection,
LDAP_OPT_PROTOCOL_VERSION, 3);
$result=ldap_bind($connection);
echo "Bind result is $result<br />";
echo "Searching for $filter..";
$sr=ldap_search($connection,$ldapDN , $filter);
echo "Search result is $sr<br />";
$count=ldap_count_entries($connection, $sr);
$info = ldap_get_entries($connection, $sr);
echo "Data for $count items returned:<p /><hr />";
for ($i=0; $i<$count; ++$i) {
    echo "Distinguished Name is: {$info[$i][\"dn\"]}<br />";
    echo "Common Name is: {$info[$i][\"cn\"]}<br />";
}
}

```

LDAP RESOURCES

www.openldap.org/ – the home of the *OpenLDAP* implementation

www.php.net/manual/en/ref.ldap.php – the PHP site has plenty of function reference material for LDAP

<http://docs.sun.com/db/prod/5320#hic> – Sun's development of the original *iPlanet* code. Includes plenty of documentation

www.emailman.com/ldap/public.html – A list of public access LDAP servers, mostly US universities

```

echo "First Email address is: {$info[$i][\"mail\"]}<br /><br />";
}
echo "Closing connection";
ldap_close($connection);
} else {
echo "<h3>Failed to connect to LDAP server</h3>";
}
?>

```

The usual sequence with a search is to connect to the server, **bind** to it, perform the search, retrieve the results, process them, then close the connection.

ldap_connect() takes one argument, the address or hostname of the server, returning a boolean result.

We have then used the **ldap_set_option()** function to force the connection to use v3 of the LDAP protocol. This isn't always necessary, but if you have problems authenticating to the server, this can prevent some of the common causes.

Binding is effectively the 'login' phase. In most cases an anonymous **bind** is okay, because LDAP is usually configured to allow read access to anonymous users. **ldap_bind()** optionally takes arguments specifying a username and password.

While the connection is live, you can perform a number of functions. In this case we want to do a simple search. The **ldap_search()** function requires a handle to the connection (returned as a resource to **ldap_connect()**), the DN or scope to search, and the filter to use. Normally the DN is the root DN of the system, but depending on the structure of the LDAP data, you may want to only search a particular group (see also **ldap_list()**).

The final option is the filter. For this, you really need to know what attributes each entry in the LDAP is likely to contain. Fortunately, because of the schema, there are certain things you can have a fair idea of. People will usually have a **cn** (common name) and/or **sn** (Surname) attribute. It really depends what the purpose of your script is.

You should be aware that the results array returned also contains plenty of arrays. Each DN object returned is an array of attributes, and each attribute is actually an array too. This is true even if there is only a single entry for an attribute, so if your script is printing out people's phone numbers as 'Array', you know what has gone wrong!

Advanced searching

First, some more on filters. These are the common search descriptors for LDAP, and can be combined in many ways: **((cn=Edward*)(cn=Edwin*))** will match all names beginning Edward or Edwin, whereas this second example

(&(cn=Edward*)(ou=Accounts)) matches everyone called Edward in the Accounts group. It is possible to build even more complex queries, but usually a simple combination will be all that is required.

Another problem you may run across is partial data. This occurs because, to avoid overloads, LDAP servers usually specify a maximum number of records they will return, and often a timeout value too. To restrict your search (and this minimise the number of returns) it could be useful to employ **ldap_list()** as a replacement for the **ldap_search()** function. This takes the same parameters as the latter, but limits the search to one level below the base DN.

A further technique to minimise the load is to specify exactly which attributes and values are required. Complex directories include all sorts of information, little of which will actually be of use depending on the purpose of the search. **ldap_search()** and **ldap_list()** take an optional fourth parameter, an array of the attributes required. The result array will then contain just those attributes. Our modified code would look like this:

```
....
$filter="((cn=*nic*)(cn=paul*))"
$required=("cn";mail")
$sr=ldap_list($connection,$ldapDN, $filter, $required);
echo "Search result is $sr <br />";
$count=ldap_count_entries($connection, $sr);
$result=ldap_sort($connection,$sr,"cn");
$info = ldap_get_entries($connection, $sr);
....
```

Results from LDAP are unsorted – they are returned as they appear in the directory. Here we have used the **ldap_sort()** function to sort the results by 'cn' before returning them. The sorting must be done before issuing the **ldap_get_entries()** function. Although the sort should theoretically support sorting on multiple attributes, in practice this isn't supported by all LDAP servers, so your mileage may vary.

Comparing attributes

Occasionally you may simply want to compare some attribute. For example, a common use of LDAP is as a repository for user account and password data. If you wanted to unify your PHP script with the data, you could set up a simple script to compare user-entered password data with that existing in the directory.

Assuming we have already established a connection and a **bind** to the directory:

```
$dn = "cn=John Smith, ou=people, dc=ixf, dc=com";
$attr = "userPassword";
$value = "oscarthedog";
// compare value
$sr=ldap_compare($connection, $dn, $attr, $value);

if ($r == -1) {
    echo "Error: " . ldap_error($connection);
} elseif ($r == true) {
    echo "Hurrah!";
} elseif ($r == false) {
    echo "Boo!";
}
```

Modifying data

In order to modify data held on an LDAP directory, two additional factors have to be considered. Firstly, the permissions to do so

must exist on the LDAP server itself. In practice, users are often allowed to change their own data, and administrators can change anyone's data. In order to work with such a server, you need to authenticate yourself at the **bind** stage, by supplying a user identity and a password. The user identity should be a proper DN for an entry in the directory.

```
<?php
$ldapuser= 'cn=admin,ou=admin,dc=ixf,dc=com';
$dappasswd='secret';
$daphost='192.168.130.91';
$ldapDN='dc=ixf,dc=com';

echo "<h1>LDAP Adding data</h1>";
echo "Attempting to connect ...<br />";
$connection=ldap_connect($daphost);
echo "Connect result is $connection <br />";
if ($connection) {
    echo "Binding ...";
    ldap_set_option($connection,
LDAP_OPT_PROTOCOL_VERSION, 3);
$result=ldap_bind($connection,$ldapuser,$dappasswd);
echo "Bind result is $result <br />";
$entry["cn"]="Jill Zemanova";
$entry["sn"]="Zemanova";
$entry["mail"]="jzem@ixf.com";
$entry["objectclass"][0]="person";
$entry["objectclass"][1]="organizationalPerson";
$entry["objectclass"][2]="inetOrgPerson";
$entry["uid"]="jzem";
$entry["givenName"]="Jill";
$entry["telephoneNumber"]="1234";
$result=(int)ldap_add($connection, "cn=Jill Zemanova,
ou=people, dc=ixf,dc=com", $entry);
echo "Adding record result = $r;
ldap_close($connection);

} else {
echo "<h3>Failed to connect to LDAP server</h3>";
}
?>
```

The **ldap_bind()** function takes care of authentication. Assuming you have the permissions, you can then alter data. Here we have supplied the password as cleartext in a variable – in practice you would probably want this entered differently. The LDAP server takes care of checking the password itself, so the server does the work of checking against whatever encryption/hash mechanisms are used.

Adding data is one of the simplest processes to achieve. The **ldap_add()** function takes parameters for connection resource, the DN of the object and an array containing the attribute data to add. The LDAP server itself will sanity check it, and may refuse to add the data if it doesn't follow the schema (eg an object with the objectclass **person** must have an entry for **cn**).

Coming to a close

With over 40 functions dedicated to the task, it's hard to describe using LDAP with PHP as straightforward. Having said that, your average LDAP script uses only a handful of the 40 functions, which means that the hardest thing is often just remembering the contents of the array return values! 

NEXT MONTH

We'll be looking at the Simple Network Management Protocol (SNMP) and how to use it with PHP.

TUTORIAL GIMP



IMAGE CONSTRUCTION

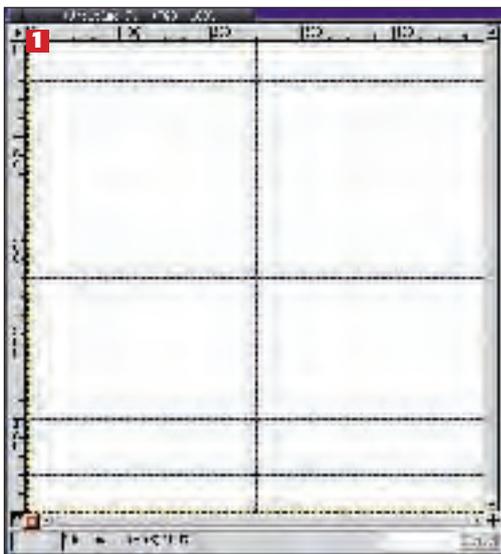
3D Design with The GIMP



Perspective views can turn any image into 3D artwork. All it takes is a little organisation, a little crafty lighting and **Michael J Hammel's** expert tutelage...

The *GIMP* isn't designed to do 3D work – such tasks are generally left to tools like the recently open-sourced *Blender* (Jono Bacon's tutorials are in *Linux Format* issues 39–50), SoftImage's *Maya* and SideFX Software's *Houdini*.

But not all 3D tasks require some complex tools. In this issue's set of tutorials we'll look at the processes for creating some basic 3D shapes using *The GIMP* and how such simple shapes are processes can be used for meaningful work.



SIMPLE SHAPES: 3D CONE WITH TEXTURE

Aside from a sphere, the cone is the most simplistic 3D shape we can make with *The GIMP*. It requires only a little optical illusion through the use of a merged selection and properly applied gradient.

Guides

1 Start with a new window with a single white background layer. Make the image square, 500x500 pixels for this example. Now add some guides to make it easy to outline the cone. Use **Guides> Center Guide** if you have it, or simply place a vertical guide offset 250 pixels from the left edge. Add two more guides offset 40 pixels from the top and bottom. Add two more guides 60 pixels offset from each side of the window. Finally, add another guide 350 pixels from the top. Your canvas should look like this example (the horizontally centered guide will be there if you used **Guides> Center Guide**, otherwise it won't).

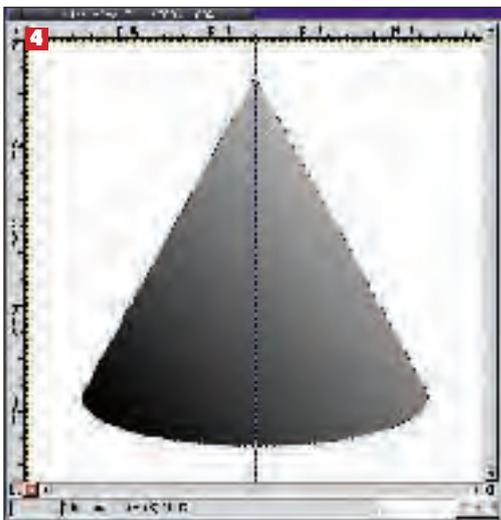


Triangular Selection

2 Choose the Bezier Selection Tool from the Toolbox. Click on the guides where they intersect at X/Y locations 250/40, 60/400, and 440/400. Be sure to close the triangle by clicking on 250/40 again, then click inside the outline to create the selection.

Cone Selection

3 Select the Elliptical Selection Tool from the Toolbox. Starting at the guide intersection at X/Y location 60/350, hold the Shift key down, and then press and hold the left mouse button down. While holding the mouse button down, release the Shift key and drag the mouse to the guide intersection at 440/450. When you release the mouse button, the elliptical selection will be merged with the triangular one.

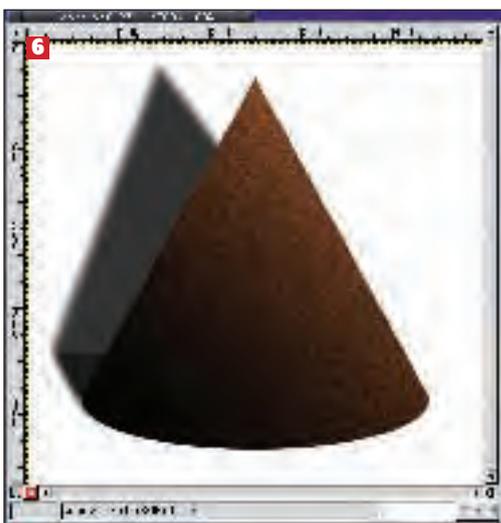
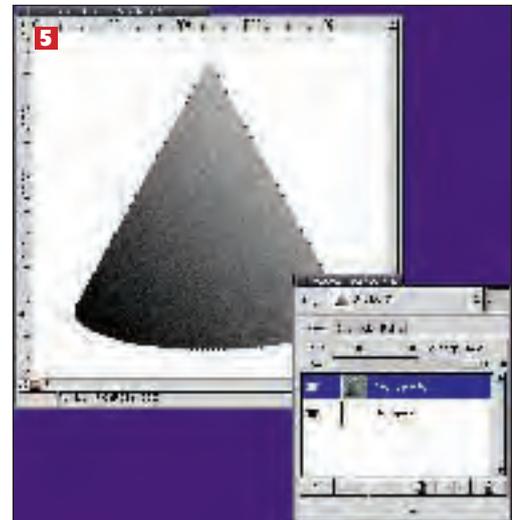


Gradient

4 Add a new, transparent layer to the image. Set the foreground colour to white and background to black (which is opposite of the default setting). Select the Gradient Tool in the Toolbox. Click at the guide intersection at 440/40 in the upper right of the image, then drag to the guide intersection at 60/400 in the lower left. This causes a lighter shade of grey on the right side of the cone and black on the lower left of the cone.

Noisify and Multiply

5 Turn off the guide visibility (View > Toggle Guides). Turn on the Keep Transparency For This Layer option (see the Layers and Channels Dialog). Add noise to the image (Filters > Noise > Noisify) set to a low level: about 0.10 for all channels. Duplicate this layer and set the new layer to Multiply. Reduce the new layers Opacity to about 70%.



Colorize

6 Duplicate this new layer. Set the foreground color to #cc4a00. Fill this new duplicate layer with this color. Remove the selection (Ctrl-Shift-A). Add a drop shadow (Script-FU > Shadows > Drop Shadow).

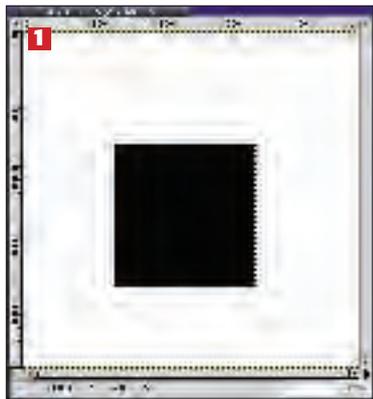


TUTORIAL GIMP

SIMPLE SHAPES: 3D BOX

While the cone was simple, it wasn't exactly perfect. The top of the cone was a perfect point and made that part of the effect appear a bit 2D. The drop shadow (which we angled at the bottom to make the cone appear to be sitting on a table) helped,

but we can do better. In this tutorial, we'll create a box using multiple square pieces that we'll align and shade with gradients to get an even better effect. This effect is just as easy as the last, but does require a steady hand to align the layers manually.

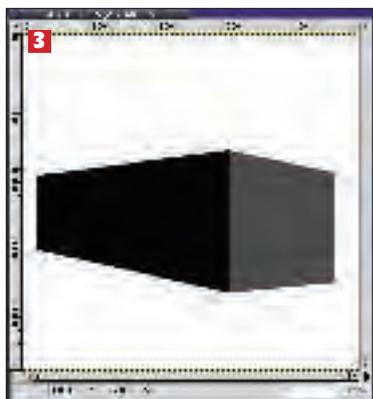


Black square

1 Start with a new canvas, preferably square – but it doesn't have to be for this tutorial. Add a transparent layer. Select the Rectangular Selection tool from the Toolbox. Create a square selection by holding down the Shift key while you drag the mouse. The size doesn't matter – we'll be modifying this in a moment anyway. Fill the selection with black.

Side Perspective

2 Copy and paste (Ctrl-C, Ctrl-V) the selection to create a new layer with the black square. Turn off the original layers' visibility for a moment while we work on the new layer. Double-click on the Transforms Tool from the Toolbox and select its Perspective Transform. Click inside the black square to show the drag boxes on the corner of the square. Drag the left side squares toward each other and the left side. Leave the right side boxes where they are. Click on the Transform button. This is the left side of the box.

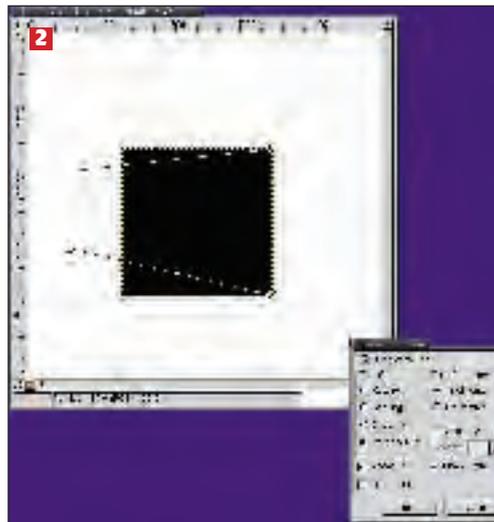


Side Perspective

3 Paste another black square into the image as a new layer. Open the Transform Tools>Tool Options window again and click on the black square. Move the right side drag boxes toward each other a small amount, with top box moved down more than bottom is moved up, and in toward the left side of the black square. Click on the Transform button. Reduce the opacity of this layer to 70%. Move the layer so its left edge butts up against the right side of the previous layer. You may have to move both layers a bit to make them both fit in the canvas properly. This is the front of the box.

Side Perspective

4 Make sure the Foreground color is set to black, and the Background is set to white. Make the left side of the box layer active. Drag from the left of this layer to the right using the Gradient tool. Make the front size of the box layer active, and drag the gradient from outside the right side of this layer to outside the left side of the layer. The left side will go from black to white, but the front will go from dark grey to light grey. This makes the edge where the side and front meet stand out a bit more while still giving the appearance of directed lighting.



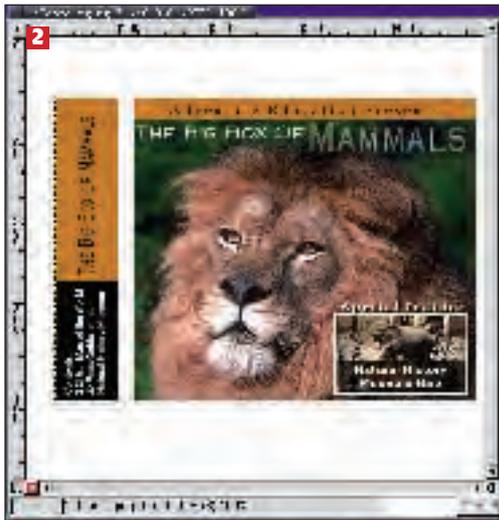
REAL WORLD: 3D PACKAGING

Our first design used shadows (via gradients) to simulate the effect of three dimensions. Our second design used the Transform Tool to change the perspective of an object (a square) to give the appearance of depth. Lighting (via gradients) enhanced the result. Transforms are powerful tools in simulating depth in *The GIMP*, as these next two real-world tutorials will show. This Packaging method described here is sometimes used in *Linux Format* to create images of books or boxed products for reviews.

Box cover

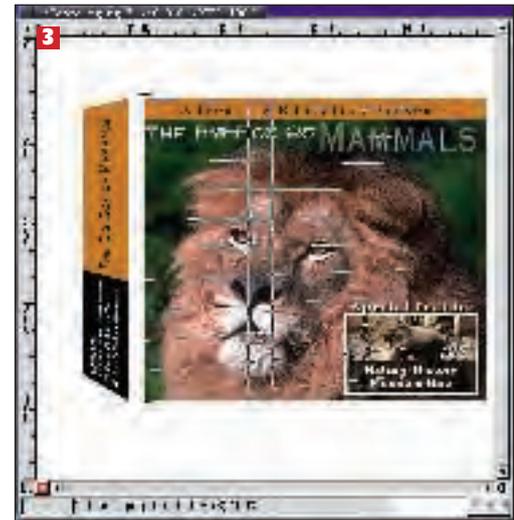
1 Start with a blank canvas at 500x500 pixels. Open an existing image to use as the front cover of a box. Make it about 375 pixels wide. Here we've created a separate image for our cover and pasted it into our blank canvas. Move the image to the right side a bit to leave room for the side of the box we're about to create.





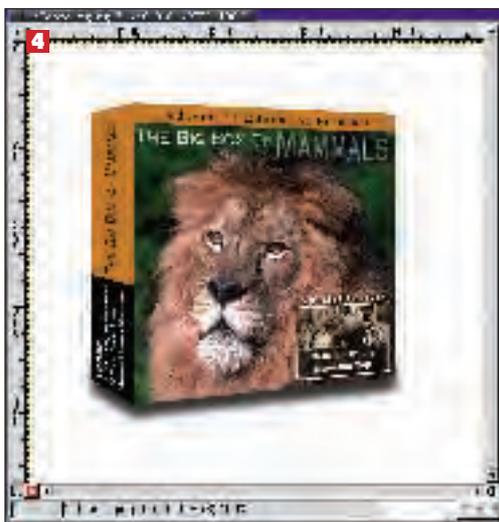
Box side

2 Duplicate the cover layer and resize its width only to about 70 pixels. Select the entire layer (Ctrl-A) and fill it with a solid colour. Adding multiple sections of colour will add to the final effect. Add a text layer and rotate it 270 degrees. Merge the side and text layers into a single layer. You should now have a side layer and a cover layer along with the white background layer. Make sure the side and front layers are aligned at their tops.



Perspective both

3 Add two horizontal guides, one 30 pixels above the bottom of the side layer and one 15 pixels below the top. Add a vertical guide 20 pixels to the right of the left side of the side layer. Select the Transform Tool from the Toolbox and choose the Perspective transform. With the side layer active, click on it once to display drag boxes. Grab the upper left drag box and move it down and in to where the upper horizontal and vertical guides intersect. Grab the lower left drag box and move it to where the lower horizontal and vertical guides intersect. Click on the Transform button to perform the transformation. Add a vertical guide about 90 pixels in from the right side of the cover layer. With the cover layer active, click on it once and drag the upper right and lower right drag boxes to the intersections of the upper and lower horizontal guides with the new vertical guide. Click the Transform button to perform the transformation.



Drop shadowed

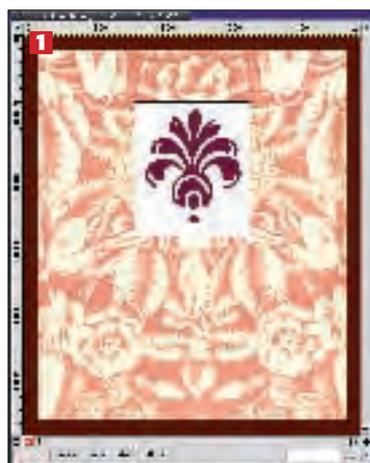
4 Now align the left edge of the cover layer with the right edge of the side layer. Make sure the tops of both layers are aligned as well – if all went well, the sides you are aligning should be the same size so the corners of the front and side will align perfectly. Merge the side and top layers. Make a rectangular selection of the bottom 1/5th of this merged layer, copy and paste it back in. Make a drop shadow this is new layer, then delete that layer (keeping the shadow). Move the shadow layer below the merged side/top layer. Adjust the position of the shadow layer as needed.

REAL WORLD: 3D LAYERING

The last real world example uses the exact same methods as the previous tutorial. In this case, a layered card design for an invitation is shown in an exploded view so that the multiple layers become obvious.

Head-on view

1 Starting with an image with multiple layers, we show a head on view of a layered card design. Viewers see only the two dimensions of this image and miss the 3D aspect the actual card will have. In order to fix this, we show an exploded view that uses perspective transforms on each layer. The key here is that each of the multiple layers of the card have their own layers in *The GIMP*.



Exploded view

2 Each layer is transformed manually. The trick here is one of trial and error – experimenting with the transformed edges so they look real and in proportion to each other. This is very much an artists playground – while we could explain how to do this mathematically and by using the measure tool and lots of guides, its much quicker to experiment, undo, and experiment again to find the right transforms. Adding shadows enhances the 3D view.

TEMPLATE PROCESSING

Using the Template Toolkit

PART 1 – BASICS Dave Cross looks at a good way to present data from various sources.

Template processing is a method of producing output that takes some fixed (or boilerplate) text and puts variable data values inside it. The most obvious example is that of a form letter. When you get that prize draw win notification from an organisation like *Readers Digest*, it has been made to look as though it is personal to you when, in fact, a few pieces of information about you have been inserted into a letter template.

The Template Toolkit is a piece of software that allows you to carry out powerful template processing operations. It is written in Perl, so it runs on pretty much any computing platform that you can name, but you don't actually need to know any Perl in order to use it. If, however, you do know Perl, then a whole extra level of power is opened up to you.

Using the Template Toolkit

Installing *TT* gives you two new command-line utilities – *tpage* and *ttree*. You use *tpage* to process a single template file and write the output to STDOUT. If your needs are more sophisticated, *ttree* can process a whole directory tree of templates and write the output to another directory. We'll start by using *tpage*.

The **tpage** command takes only one argument, **--define**. This allows you to define data variables to be inserted into your template. You can use many variable definitions at once. You would therefore use it from the command line something like this

```
tpage --define name='Mr Cross' --define amount='100'
--define due='1st April' letter.tt
```

This defines three variables called **name**, **amount** and **due**; which can now be referenced within the template. The template itself is in the file *letter.tt* and it looks like this:

```
Dear [% name %],

According to our records you owe us £[% amount %].

Please pay before [% due %] or we will send the boys round.

Regards.
```

You can see from this simple example that the places where you want your variables to be inserted are marked with **[% ... %]**

tags. These are known as template *directives*. We'll look at many different types of directives over the next few months, but currently we are looking at the simplest of directives which just contains the name of a variable. When the template is processed, this directive will be replaced by the current value of the variable.

So, if we process this template using the command-line we saw above, we'll get this output:

Dear Mr Cross,

According to our records you owe us £100.

Please pay before 1st April or we will send the boys round.

Regards.

Using Plugins

Obviously, if we had to type a command-line like that each and every time we had to send a form letter, we really haven't gained much advantage or efficiency over typing them all in manually. In a while, we'll see a way to make this easier, but first let's take a moment to tidy up the output a little.

Currently, our template just takes the values we give it and reproduces them. For the amount due, it might look better if we forced the value to always have two decimal places. We can do this by using the *Format* plugin.

In *TT*, plugins are a way for the template processor to provide extra functionality by interacting with external resources. In this case the *Format* plugin provides an interface to a C-style **printf** feature. We change the template so that it looks like this:

```
[% USE money=format("%.2f") -%]

Dear [% name %],

According to our records you owe us £[% money(amount) %].

Please pay before [% due %] or we will send the boys round.

Regards.
```

We've added a **USE** directive which loads the *Format* plugin and

we've used the plugin to create a formatting function called **money**. Any expression that is passed to this function is formatted according to the format string that was used when **money** was created. In this case, the **"%.2f"** ensures that the amount has two decimal places. If we process our template using the same call to *tpage* as before, we now get slightly changed output.

Dear Mr Cross,

According to our records you owe us £100.00.

Please pay before 1st April or we will send the boys round.

Regards.

Another small change that we've made to our template is the addition of a **'-'** character at the end of the **USE** directive. Generally, all whitespace outside of a directive is passed straight through to the output, but the **'-'** character tells *TT* to ignore the whitespace (including newlines) that follow the directive. The net effect of this is that adding the **USE** directive doesn't add a blank line to the output. *TT* comes with a useful set of standard plugins. We'll see some more of them in the next few sections.

Reading data from a file

As I mentioned before, we haven't really gained much if we have to type in a long command each time we want to process a form letter. It would be far easier if we could read the data in from some external source. And, of course, we can. We can read data from all sorts of external sources. We'll start by reading the data from a text file.

We'll assume that our data file has the following format:

```
name : amount : due
```

```
Mr Cross : 10 : 1st April
```

```
Mr Smith : 20 : 1st March
```

```
Mr Jones : 50 : 1st February
```

The first line of the file contains the names of the fields and the other lines contain the actual data. We can now change our template to look like this:

```
[% USE money = format("%.2f") -%]
```

```
[% USE debtors = datafile(file) -%]
```

```
[% FOREACH debtor = debtors %]
```

```
Dear [% debtor.name %],
```

GETTING THE TEMPLATE TOOLKIT

Download resources

If you're happy installing Perl modules from the Comprehensive Perl Archive Network then you can go to <http://search.cpan.org/dist/Template-Toolkit/> and download it from there. It is installed in exactly the same way as most other Perl modules.

If you'd rather not get involved with CPAN, then there are also Debian packages and RPMs available. The Debian packages can be downloaded from the Debian packages repository at <http://packages.debian.org/unstable/interpreters/libtemplate-perl> and <http://packages.debian.org/unstable/doc/libtemplate-perl-doc>. The RPMs can be found by searching at <http://rpmfind.net>. The package name is *perl-Template-Toolkit*.

If you need any more information about installing *Template Toolkit* or you want to use the bleeding-edge CVS versions, then you can access those at the official *Template Toolkit* website at www.template-toolkit.org/.

```
According to our records you owe us
```

```
£[% money(debtor.amount) %].
```

```
Please pay before [% debtor.date %] or we will send the boys round.
```

```
Regards.
```

```
[%- END %]
```

As you'll see, we've added another directive which loads the *Datafile* plugin. This plugin opens the given data file and returns an iterator object which can be used to access the data in the file. We assign this iterator object to the variable **debtors**. Note that the name of the file to be used is in the variable **file**. This is now the only variable that we need to pass to the template processor.

Having opened the data file we can process it a row at a time using the **FOREACH** directive. **FOREACH** iterates across a list, setting the loop variable to each element of the list in turn. In this example the loop variable is called **debtor** and each time round the loop it gets one of the values returned by the **debtors** iterator. These values are objects that contain the individual data items from each row in the data file. You can access these items using a dot notation, for example the name from the current row of data is in **debtor.name**.

We have changed the names of the variables used in the rest of the template to reflect this.

The end of the **FOREACH** loop is marked with an **END** directive and notice the **'-'** and the start of that directive which strips out any whitespace preceding it (*ie* the newline before it.)

Of course, the random format that I chose for the data file just happened to be the default format for the *Datafile* plugin (using colons as the delimiter) but it's easy to use an alternative delimiter. For example, if our debtors file had been delimited with pipe characters, we would have used code like this:

```
[% USE debtors = datafile(file, delim => '|') %]
```

The delimiter can be surrounded by optional whitespace which is removed from the values. The first row must contain the names of the data items and any blank lines or comment lines (which start with a **#** character) are ignored.

Splitting the output

The remaining problem is that this prints out all of the letters in a continuous piece of text, but actually we want each letter on a separate page. We can achieve this by inserting a form-feed character (**0x0C**) at the end of each page like this:

```
[% USE money = format("%.2f") -%]
```

```
[% USE debtors = datafile(file) -%]
```

```
[% FOREACH debtor = debtors -%]
```

```
Dear [% debtor.name %],
```

```
According to our records you owe us
```

```
£[% money(debtor.amount) %].
```

```
Please pay before [% debtor.date %] or we will send the boys round.
```

```
Regards.
```

```
[% UNLESS loop.last -%]
```

```
^L
```

```
[%- END %]
```

```
[%- END %]
```



TEMPLATE TOOLKIT DOCUMENTATION

Where to go for more information

This article has just scratched the surface of the *Template Toolkit*. We'll be going into more detail later in this series, but in the meantime you can get more information from a number of sources.

Distribution Documentation

The *Template Toolkit* distribution comes with a large number of manual pages. The best place for a beginner to start is probably with **man Template::Manual** which is a guide to all the other *Template Toolkit* manual pages and **man Template::Tutorial** which introduces a couple of tutorials that are part of the *Template Toolkit* documentation set.

Web Page

The official *Template Toolkit* web page is at <http://template-toolkit.org/> (or <http://tt2.org/> if

you don't like typing). Here you'll find all the man pages online together with a number of talks about *TT*, the latest version of the software and a number of other interesting things.

Mailing list

There is a mailing list for the discussion of things relating to *TT*. You can subscribe to it at www.template-toolkit.org/mailman/listinfo/templates. All of the core *TT* developers are regular posters to this list.

Book

The book *Perl Template Toolkit* by Darren Chamberlain, David Cross and Andy Wardley (ISBN 0-596-00476-1) has recently published by O'Reilly. You can get more details about it from www.oreilly.com/catalog/peritt.

◀◀ The `^L` is the representation of a form-feed character when displayed in many text editors.

There are a couple of other changes that I've made to the template. These prevent an extra form-feed being output after the last letter. A printer will automatically perform a form feed at the end a print job, so if our output contains one as well there will be two form-feeds in the job and an extra (blank) sheet of paper will be used. We can prevent that using the code shown.

The code uses an **UNLESS** directive to optionally display part of the template. **UNLESS** works like **IF** but the logic is reversed, the code within the block is executed only if the **UNLESS** condition is false.

In this condition we check the **loop** variable. This is a special *TT* internal variable which contains various interesting pieces of information about the current **FOREACH** loop. There are boolean flags **first** and **last** which return true only if you are in the first or last iteration respectively. The **size** data item contains the number of elements in the list and the **count** and **index** items contain slightly different views of the current iteration. **count** gives you the number of the iteration (from one to **size**) and **index** gives you one less than that number (useful if you're a Perl programmer and used to array indexes that start from zero).

In our current example we use **loop.last** to avoid printing the extra form-feed. That's all fine if you have your data in a suitable data file. But maybe it's stored in a database instead.

Accessing a database

The great thing about *TT*'s plugin system is that it's very easy for your template to get data from all sorts of interesting and useful places. For example, there's a plugin to Perl's database interface system, *DBI*. This allows your template to access data stored in most kinds of database.

Assuming that our data is stored in a table called **debtors** in a *MySQL* database called **accounts**, we can change our template to look like this;

```
[% USE money = format('%2f') -%]
[% USE DBI(database = 'dbi:mysql:accounts'
username = 'acc_user'
password = 'sekrit') -%]
[% FOREACH debtor = DBI.query('select name, amount, due
from debtors') -%]
```

```
Dear [% debtor.name %],
```

```
According to our records you owe us
£[% money(debtor.amount) %].
```

```
Please pay before [% debtor.due %] or we will send the boys
round.
```

```
Regards.
```

```
[% UNLESS loop.last -%]
```

```
[%- END %]
```

```
[%- END %]
```

It's interesting to note just how few changes we have had to make here. All of the loop code is identical, it's just the code that sets up the loop that is different.

The new code is simple enough to follow. We load the *DBI* plugin, passing it the various parameters required to connect to the database. We then run an SQL query against the database. This returns an iterator object that we can use in a **FOREACH** directive in exactly the same way that we used the iterator that the *Datafile* plugin created. Each value returned by the iterator is an object which contains data items for each of the columns selected from the database. The names of these data items are given by the column names in the SQL statement.

Formatting dates

There's another formatting improvement that we can make at this time – we can reformat the date. If we assume that the due date is stored in a database date column, then most databases will return it in the format `YYYY-MM-DD` which isn't very user-friendly. We can use the *Date* plugin to fix that. We can load the *Date* plugin with a directive like this:

```
[% USE date(format = '%d %B') -%]
```

And then use it with a directive like this:

```
[% date.format(debtor.due) %]
```

Here we have given the *Date* plugin a format string. This is in the same format as the format strings used by the Unix **date** command. Our `'%d %B'` format gives us the day of the month followed by the full month name. Having seeded the plugin with that format, any dates that are passed to the **date.format** function are converted to that format.

If you want to override the default format at any time, you can pass a new format definition as the second argument to the **date.format** function like this:

```
[% date.format(debtor.due, "%A, %d %B %Y") %]
```

As **%A** is the full weekday name and **%Y** gives the year, this example will display the due date in the format **Sunday, 01 February 2004**.

One small problem with the *Date* plugin is that it is a little fussy about the format of the date. It only accepts dates as either the number of seconds since 1st Jan 1970 (known as the *Unix Epoch*) or in the format `h:m:s d/m/y` (hours:minutes:seconds day/month/year). That's not the format that we're currently getting from the database so we'll need to use *MySQL*'s **date_format** function to correct that.

```
[% FOREACH debtor = DBI.query('select name, amount,
date_format(due, "%h:%i:%s
%d/%m/%Y")
as due
from debtors') -%]
```

Two things to notice here. Firstly, whilst the *MySQL* date format strings look a lot like the standard Unix ones used by *TT*, they are actually different (**%i** for minute instead of **%M**, for example). Secondly, we've added a column alias (**as due**) to the date column so the *DBI* plugin continues to use that name for the column.

Dealing with XML

So we've managed to extract our date from data files and databases. What if our data is stored as an XML document? As you'd expect, *TT* has plugins to handle that too. In this example I'll use the *XML.XPath* plugin to access data within a document. Let's assume that our debtor document looks like this:

```
<debtors>
<debtor>
<name>Mr Cross</name>
<amount>10</amount>
<date>1st March</date>
</debtor>
<debtor>
<name>Mr Smith</name>
<amount>20</amount>
<date>1st February</date>
</debtor>
<debtor>
<name>Mr Jones</name>
<amount>50</amount>
<date>1st February</date>
</debtor>
</debtors>
```

Here's the template that we'll use to process it.

```
[% USE money = format("%.2f") -%]
[% USE debtors = XML.XPath(file) -%]
[% FOREACH debtor = debtors.findnodes('/debtors/debtor') -%]
Dear [% debtor.findvalue('name') %],

According to our records you owe us
£[% money(debtor.findvalue('amount')) %].

Please pay before [% debtor.findvalue('date') %] or we will
send the boys round.

Regards.
[% UNLESS loop.last -%]
^L
[%- END %]
[%- END %]
```

The template first creates a **XML.XPath** object by passing the name of the XML file to the *XML.XPath* plugin in the **USE** directive. This **XML.XPath** object can then respond to **XPath** queries in a number of ways. The first way that we use is to use the **findnodes** function to get a list of all of the nodes matching the query **/debtors/debtor**. This returns an iterator that we can use in a **FOREACH** loop like all of the iterators that we have seen previously. Each object returned by this iterator is another **XML.XPath** object representing one of the nodes in the set. In this case each element is a **debtor** node that is contained within the main **debtors** node. We can then use the **findvalue** function to get the text value contained in the various nodes that we are interested in. [LXF](#)

NEXT MONTH

We've only done very simple things with *TT* this month. Next month we'll get a bit more complicated and in the process we'll start doing things that are a bit more useful than sending begging letters.

At Hostway we understand that every hosting customer has unique demands. But there are some things you all agree on.

You want rock solid systems reliability so you can be confident your website is up first time, every time someone visits. You want the highest quality support and customer service so whenever you need help or have a problem, there's someone with a solution on the other end of the line when you call. You also want innovative, easy-to-use administration tools because your time is better spent creatively, or on managing your business. That's not all though. You want support for the latest technologies, and the flexibility that will allow you to build your website and sustain it as it grows, with minimum hassle. And you want all of this at a competitive price.

Too much to ask for? No, we don't think so either...

unique customers - unique hosting

A FREE MONTH
of web hosting.
When ordering online
use "FUTLIN052".

	KEY FEATURES	DISK SPACE	MONTHLY TRAFFIC (Gb)	SHELL ACCESS	PHP4 SCRIPTING	MYSQL 3.23 DATA AREA	ADDITIONAL FEATURES	ACTIVE CGI LIBRARY	CRONJOBS	XML SUPPORT	PYTHON	FRONTPAGE EXTENSIONS	EMAIL FEATURES	POP3/SMTP EMAIL ACCOUNTS	UNLIMITED ALIASES	EMAIL FORWARDING	MEDIA FEATURES	HTTP STREAMING	REAL SERVER PROFESSIONAL	MONTHLY FEE*	ANNUAL FEE*	SETUP FEE*
LINUX SILVER	100	1	Yes	Yes	0	Yes	Yes	Yes	Yes	Yes	Yes	5	Yes	Yes	Yes	No	£7.95	£79.95	FREE			
LINUX GOLD	200	4	Yes	Yes	1	Yes	Yes	Yes	Yes	Yes	Yes	10	Yes	Yes	Yes	No	£14.95	£134.95	FREE			
LINUX GOLD PLUS	200	6	Yes	Yes	1	Yes	Yes	Yes	Yes	Yes	Yes	30	Yes	Yes	Yes	No	£19.95	£179.95	FREE			
LINUX PLATINUM	500	10	Yes	Yes	2	Yes	Yes	Yes	Yes	Yes	Yes	50	Yes	Yes	Yes	No	£49.95	£449.95	FREE			
LINUX PLATINUM PLUS	500	15	Yes	Yes	2	Yes	Yes	Yes	Yes	Yes	Yes	100	Yes	Yes	Yes	Yes	£89.95	£809.95	FREE			

All plans come with a 30 day money-back guarantee | * excluding VAT & subject to change.

call us now on 0808 180 1880 • sales@hostway.co.uk • www.hostway.co.uk

domain registration • web hosting • resellers • e-commerce • dedicated servers | chicago • london • amsterdam • seoul • sydney • vancouver • tampa

LXF52 APRIL 2004 **85**
HOSTWAY
THE HOSTING COMPANY

NETWORK MANAGEMENT

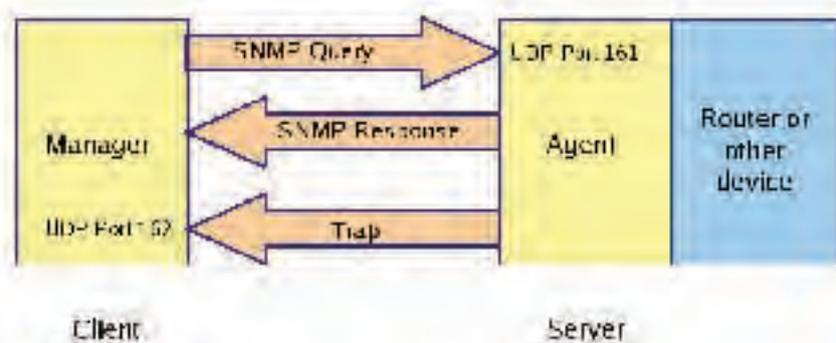
SNMP: Simple Network Management Protocol

PART 1 Following on from his series on IP networking, **Dr Chris Brown** continues with a further two-part tutorial to make sure that **SNMP** is in fact “simple.”

The purpose of SNMP is to support the remote management of bridges, routers, servers, and other components of a company's network infrastructure. As internet protocols go, SNMP has been around a good while. The earliest official document describing SNMP is RFC1065 which dates from 1988. SNMP has remained under active development and there have been two new major releases, version 2 (1993) and version 3 which emerged around the turn of the century.

RFC (Request for Comments) documents are the ‘official’ way to propose new Internet protocols. Their issue is regulated by the IETF (Internet Engineering Task Force) and over time some of them become adopted as formal Internet standards. There are many RFCs describing various aspects of SNMP. Not all RFCs aspire to become full standards – some are of a more informational nature. One of my favourites is RFC 1118, entitled *Hitch-Hiker's Guide to the Internet*. RFCs are available from many internet archives but a good place to start is www.ietf.org/rfc.html.

Fig1 **SNMP architecture: managers and agents.**



This month, we'll begin by taking a look at the architecture of SNMP, and in next month's issue we'll examine the NET-SNMP package, a popular Open Source SNMP toolkit for Linux.

As shown in **Fig1**, SNMP is designed around a simple client/server architecture. The server – which is called an SNMP agent – is a piece of software that ‘instruments’ (captures information about) the device being managed. In the case of a router, for example, the agent is part of the software built into the device. The client side of the architecture is an application that is usually called a network management system (NMS) or simply a manager. It usually runs on the network administrator's desktop. The NMS gathers data from the various agents it's responsible for, and usually displays them in a nice graphical format. The type of data available from a device depends on the device, and there's generally a lot of it, but to give you a quick idea, for a router you can typically pull back data on the status of a device's network interfaces (whether they're up or down), the amount of traffic they're handling, the number of errors they've seen, the entries in the routing tables, and much else besides. Each of these pieces of information has a well-defined name (we'll get to those later) and perhaps the most fundamental operation of SNMP is to issue **SNMP-GET** requests to retrieve named items of data from the agent. As shown in **Fig1**, the agent sends an **SNMP RESPONSE** back, with the requested data.

This **GET/RESPONSE** mechanism essentially provides a passive way of monitoring a device. If you want to find out if a router is in trouble, (for example, if it is having to discard lots of packets) you'd have to repeatedly poll it to see if the discarded packet count was getting too high. This is wasteful on network bandwidth, and of course constant polling is not the best thing to do to a router which may be already stretched to the limit. So

there's another mechanism in SNMP, called a *trap*, which lets an agent actively notify a manager if something significant has happened. Whereas an SNMP agent only sends **SNMP RESPONSE** messages in response to a request from the manager, *traps* are unsolicited and asynchronous – they can occur at any time.

Whether SNMP is really deserving of being called 'simple' is open to debate. There's certainly evidence that the designers of the early versions of the protocol had simplicity as a design goal, but later versions have (perhaps inevitably) made some sacrifices of simplicity for the sake of flexibility. In particular, the encrypted authentication mechanisms introduced in version 3 of the protocol, while perhaps necessary, stretch the notion of simplicity beyond reasonable bounds. Still, the name stays. After all, who would want to use something called the "Complicated Network Management Protocol"? Marketing is everything!

The Management Information Base

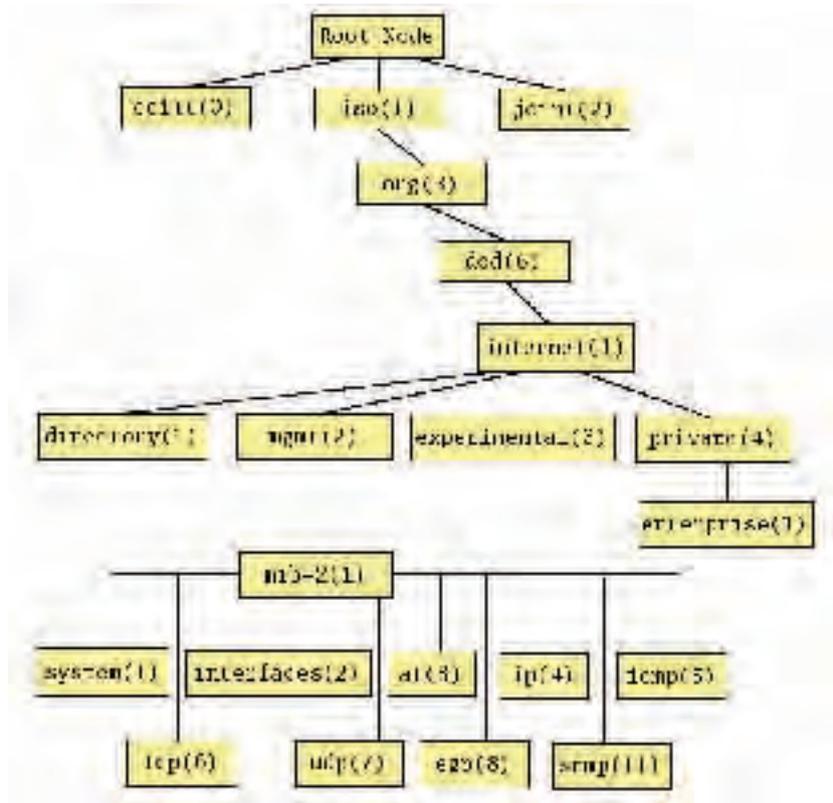
The protocol SNMP protocol itself (the format of the messages used to **GET** and **SET** variables in an agent) is fairly straightforward. The real complexity lies in the structure of the information that can be managed, and the way in which variables are defined and named. The variables are organised in a tree structure and are named using a scheme akin to pathnames in the Unix filesystem. For example, whereas a file might be named `/usr/local/bin/snmpset`, an SNMP variable might be named **.iso.org.dod.internet.mgmt.mib-2.ip.ipDefaultTTL**. You'll notice that we use '.' instead of '/' to separate components of the names, and that they are usually rather long. Each of the components of the name can also be represented by a number, so for example the variable named above can also be written as **.1.3.6.1.2.1.4.2**. The numeric form of the name is obviously much shorter but it's darn hard to remember and copy down accurately. These names are more formally known as 'Object Identifiers' (OIDs).

Time for a little alphabet soup. The structure and naming of all these variables is defined in documents called MIBs, which stands for Management Information Base. The structure within the MIBs is known as the SMI, which stands for Structure of Management Information. The language in which the MIBs are written is called ASN.1, which stands for Abstract Syntax Notation (a worryingly meaningless name, if you ask us.). Finally, the way in which information is actually encoded into messages and passed between agent and manager is known as BER, which (*phew!*) stands for Basic Encoding Rules. You are paying attention, aren't you? There'll be questions on this in the exam! ASN.1 and BER are formal notations used for describing data transmitted by telecommunications protocols. They date from 1984 and were not invented as part of SNMP, but were simply adopted by it. If we were designing a notation for MIBS today, my guess is we'd be using XML instead.

To give you an idea of how variables are defined in MIBs and to see ASN.1 in action, let's take a look in a MIB called IP-MIB which defines variables associated with the IP protocol. We will home in on the definition of the **ipDefaultTTL** variable we mentioned earlier. It looks like this:

ipDefaultTTL OBJECT-TYPE

```
SYNTAX      INTEGER (1..255)
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
```



"The default value inserted into the Time-To-Live field of the IP header of datagrams originated at this entity, whenever a TTL value is not supplied by the transport layer protocol."
 ::= { ip 2 }

Fig2 Top levels of the object tree.

As this example above shows, every variable has the following properties:

- A text name
- A syntax (data type)
- An access level (none, read-only, or read-write)
- A status (for example: mandatory, optional or obsolete)
- A text description

SNMP defines a number of data types for its variables but many of them are really just integers. For example, some variables are counters – integer values that can only increase, such as the number of packets a network interface has received. Other variables are gauges – integers whose value can go down as well as up, such as the speed of an interface, or the air temperature inside a cabinet. Sometimes integers are used simply as 'enumerations' – variables that can take on one of a small set of defined values that indicate the status of something, such as whether an interface is up, down, dormant, or in a test mode. Other data types you'll see include 'Timeticks' (integers again, representing time in hundredths of a second), IP addresses, and 'OCTET STRINGS' – typically text strings. For example, the personal contact information for a managed device is represented as an octet string.

The MIB entry shown above also tells us where this variable sits in the overall tree. The line **::= { ip 2 }** says that this is node number 2 under the **ip** node in the tree. Higher up in the same MIB you'll find a line that says:

```
ip OBJECT IDENTIFIER ::= { mib-2 4 }
```

which defines **ip** as node #4 under the **mib-2** node in the tree. ➤

TUTORIAL SNMP

◀◀ If we were to dig around in a few other MIB files, we'd find that **mib-2** was defined, in turn, as node **1** under the **mgmt** node, that **mgmt** was defined as node **2** under the **internet** node, and so on. Eventually, we'd work our way to the top of the tree and discover that the full path name of this variable is **.iso.org.dod.internet.mgmt.mib-2.ip.ipDefaultTTL**, or its numerical equivalent: **.1.3.6.1.2.1.4.2**.

As you explore the MIBs, you begin to get a feel for the true scope of SNMP. **Fig2** shows just a fraction of the top levels of the object tree. Actually, the top four levels of this tree (down to the box labelled **internet(1)**) are not very interesting – there are no other branches in these levels with anything useful in them. Part of the reason why these upper levels exist is that SNMP was originally conceived at a time when networking models and protocols other than TCP/IP were being promoted, and space was left in the object tree to accommodate them. As it turned out, TCP/IP remained dominant, and SNMP development focussed on that protocol suite. Still, the upper levels of the object tree remain as a sort of historical wasteland. At any rate, all the object ID's you're likely to meet in SNMP start **.1.3.6.1**, or in textual form, **.iso.org.dod.internet**.

Below that level, things start to get interesting. There's a very important node below the **internet** node called **mib-2**. Under there, you'll find subtrees that contain information on device interfaces, together with configuration data and traffic statistics for the major internet protocols – **IP**, **ICMP**, **TCP**, **UDP** and so on. This stuff is all defined in the RFCs and makes up what we might call the 'official' part of the MIB tree.

Another important branch is the one under the **enterprise** node (**.1.3.6.1.4.1**) shown in **Fig2**. This is where vendors of networking hardware get to register their own **enterprise-specific** pieces of the SNMP namespace, and define their own variables. Cisco Systems (as you might expect for a leading vendor of networking products) has a huge number of MIBs defined under its registered subtree **.1.3.6.1.4.1.9**. Any equipment vendor can register MIBs under the enterprise subtree and a great many have. If you look at the register of private enterprise numbers maintained by the IANA (Internet Assigned Numbers Authority) you'll find over 19,000 entries. (You can find it at www.iana.org/assignments/enterprise-numbers). Let's be clear here – we're not talking about a total of 19,000 nodes in the whole of the subtree under the **enterprise** node, we're talking about 19,000 separate organisations, each with MIBs of their own under the enterprises subtree. The more you learn, the less SNMP seems to deserve to be called 'simple'.

Comparing the simplicity of the SNMP protocol with the complexity of the MIBs somehow reminds me of a case once put forward by a colleague: that binary is the easiest language to program computers in. The argument hinges on the observation

that if you're writing in binary, there are only two mistakes you can make: You can put a 1 where you should have put a 0, and you can put a 0 where you should have put a 1. I feel sure there is a weakness somewhere, but I find it hard to put my finger on it...

Tables

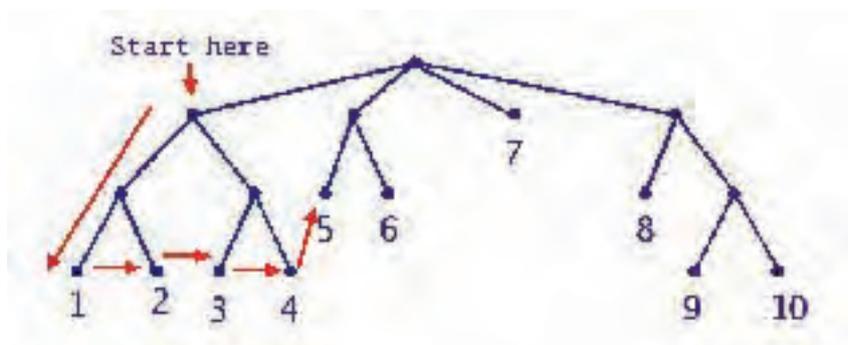
One of the hardest things to get your brain around in SNMP is tables. You'll find a good example of a table in an important MIB called **IF-MIB** (it's rooted at **.1.3.6.1.2.1.2** in the tree – see if you can find it in **Fig2**). This MIB contains variables that describe the operating status of a device's network interfaces. Now, a device such as a router has, of course, several network interfaces, and they are represented by the rows in a table. The columns of the table represent the various parameters monitored for each interface – its speed, its physical address, the number of bytes it has transmitted, and so on. This leads us to the structure shown (in part) in **Fig5**. The table as a whole is called **ifTable** and its OID is **.1.3.6.1.2.1.2.2**. The table is defined as a 'SEQUENCE OF IfEntry' ('SEQUENCE OF' is a piece of ASN.1 notation meaning 'array of'). An **ifEntry** defines one row of the table (it's defined as a 'SEQUENCE', which is a bit like a structure in C or C++; that is, it consists of a list of named objects which can be of a variety of types). To access individual cells of the table, such as the one shown shaded in the figure, we append the column identifier followed by the row identifier (the interface number) to the OID, giving a complete OID such as **.1.3.6.1.2.1.2.2.1.5.3**. Snappy, huh?

Walking the MIB

The simplest query from an SNMP manager to an agent is an **SNMP GET** request for a single named variable within the MIB. That's fine if you know exactly what OID you're looking for. However, there's another form of query called **GET-NEXT** that asks the agent to return the object in the MIB tree that follows the one you asked for. At first sight this feature of SNMP looks rather bizarre and not especially useful, but by issuing repeated **GET-NEXT** commands you can easily retrieve a series of related entries in a MIB. This is known as 'walking' a part of the MIB tree. **Fig3** shows the implied ordering of the objects within the MIB and tries to illustrate the notion of the 'following' entry. Walking this tree from the node labelled 'Start here' would traverse the tree in the order shown by the arrows, stopping when we reached a node outside of the subtree we started at.

A comparison with the ordering of words in a dictionary might help us understand the way **GET-NEXT** works. The alphabetic, or 'lexicographic' order of the words in a dictionary, which we're all familiar with, can be described in terms of their positions within a tree. The top level of the tree branches out 26 ways, labelled **a** to **z**, representing the first letter of the word. Then below each of those branches, there are potentially another 26 branches labelled **a** to **z** that represent the second letter of the word. Some of those branches are absent, of course. For example, there are no words beginning with **bb** or **bc**. The bottom levels of the tree (the 'leaf nodes') represent actual, complete words. We could point to the node labelled **ba**, and there would be plenty of nodes beneath it (all the words beginning **ba**) but **ba** itself is not a word and does not appear in the dictionary. In our comparison, a **GET-NEXT** operation on the node **ba** would return the first word in the dictionary beginning with those two letters. In my dictionary, it's 'baal'. A **GET-NEXT** of 'baal' would return 'baba'; a **GET-NEXT** on 'baba' would return 'babble', and so on until we reached 'bazooka'. A **GET-NEXT** on 'bazooka' yields the unlikely

Fig3 Walking the MIB.



looking word 'bdellium' (no, I have not made that up) but of course this word does not lie under the **ba** subtree and if we were using **GET-NEXT** to walk the **ba** tree, we would normally stop at this point.

A common use of **GET-NEXT** is to read back the contents of an entire table. If you make a **GET-NEXT** request for the **ifTable** using the OID **1.3.6.1.2.1.2.2.1**, you'll get back the entry for the top left corner of the table, which if you refer back to **Fig5** will (probably) be **1.3.6.1.2.1.2.2.1.1.1**. Repeated use of **GET-NEXT** will then traverse the entire table, starting by scanning down the first column, then the second, and so on. This is an important use of **GET-NEXT**, because otherwise you'd have no easy way of knowing how many rows the table had. Even if you did, you would not necessarily know the values of the indexes which identify the rows, because they do not have to be consecutive integers.

As an extreme case, if you walk a MIB starting at the OID **.1**, you will eventually traverse the whole of the agent's MIB.

Getting the big picture

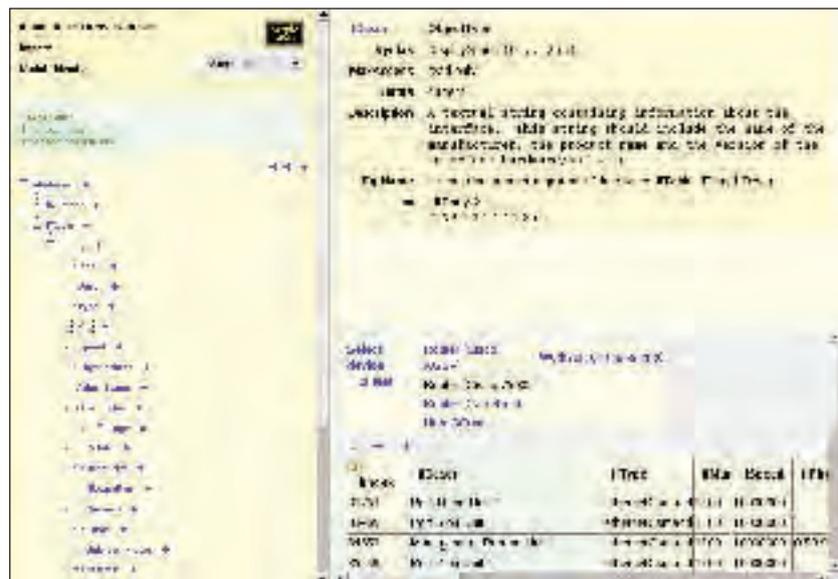
There is a saying that "you can't see the wood for the trees", meaning that sometimes the details get in the way of seeing the whole picture. In the case of SNMP, it is quite hard to get an appreciation of the potential size of the object tree by looking at the individual MIBs. However, you can get a fair idea of their scope by visiting www.simpleweb.org/ietf/mibs. Here you'll find a comprehensive database of MIBs, including all the standard internet ones and many vendor-specific ones. You can select a MIB, drill down into it, view the detailed descriptions of all the variables in it, and even query the values of those variables on a few live test devices. The site also carries copies of all the SNMP-related RFCs.

Take a look at the browser screenshot in **Fig4**, taken from the simpleweb.org site. Here we've selected the MIB **IF-MIB** for inspection. The frame on the left shows a tree view of this MIB. We've expanded the interfaces, **ifTable** and **ifEntry** nodes and have selected the entry called **ifDescr** for detailed examination. To fit this into the big picture, imagine pinning the **interfaces** node at the top of this subtree onto the **interfaces** node at the bottom of **Fig2**. In the top-right corner of the screenshot in **Fig4**, you'll see the detailed description of this node: its syntax, access level, status, description, and so on. This is taken directly from the definition in the MIB. In the bottom-right corner you'll see a short list of live test devices available on the site. In the screenshot, I queried the **3com** hub (www.3com.com/), and you can see just a small corner of the **ifTable** returned by this device. You really should visit this superb website and explore its range of features for yourself.

Monitoring and management

A lot of the emphasis in SNMP is on gathering data, and you might be forgiven for thinking that the 'M' stands for 'monitoring' rather than 'management'. In fact, most of the variables defined in the MIBs are declared as 'read-only'. However, you can use SNMP for management, too, because in addition to **SNMP-GET**, there's an **SNMP-SET** operation, which allows you to set values for those variables which have read-write status.

There are no explicit imperative commands in SNMP – that is, there are no commands which say "turn this interface on" or "turn that interface off". Instead, you control devices by changing the values of the appropriate variables. For example, one of the columns in the **ifTable** we met earlier contains a variable called



ifAdminStatus. You won't find it in **Fig5** because it's in a column off the right-hand end of the figure, but you can see it in the left-hand pane of **Fig4**. This variable is read-write. A value of **1** indicates that the interface should be up (active) and a value of **2** indicates that it should be down (inactive). So, simply setting the variable to one of these values turns the interface on or off.

Fig4 Screenshot from www.simpleweb.org.

Access Control

It may have crossed your mind by now to wonder if SNMP provides any access controls. Can anyone able to reach your network devices read all their operating parameters, or worse, turn them all off? Well, no. There is a primitive form of access control in versions 1 and 2 of SNMP, and it's built around the notion of 'communities'. A community, despite its rather grand name, is really just a password that must be supplied with each **GET** or **SET** command. At minimum, an SNMP agent usually defines two community strings: one allowing read-only access and the other allowing read-write access. Setting these to some difficult-to-guess values provides a measure of access control. Unfortunately, the community strings are passed as clear text in the SNMP messages, which makes them pretty easy to snoop off the network with a basic packet monitoring application. Version 3 of the SNMP protocol addresses this deficiency, at the expense of considerable extra complexity, by including an encryption scheme for communities for extra security. [LXF](#)

NEXT MONTH

Well, this month's piece has turned out quite heavy on the theory. Congratulations if you made it this far! Next issue, we'll get our hands dirty with a real implementation by installing the **NET-SNMP** package. This package includes an SNMP agent that runs on Unix (including Linux of course), a set of command-line tools for getting and setting SNMP data, and a software toolkit that helps you write your own agents and managers.

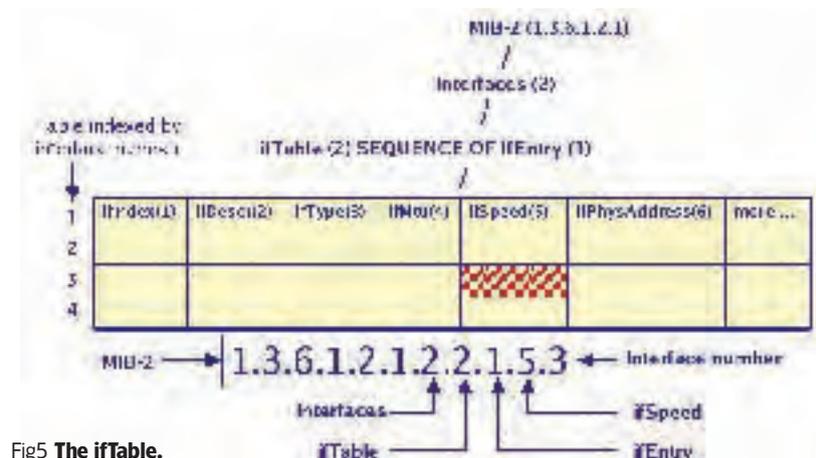


Fig5 The ifTable.

Answers

If you are really stuck and the HOWTOs yield no good result, why not write in? Our resident experts will answer even your most complicated problems!

Our experts

Whatever your question is, we can find an expert to answer it – from installation and modem woes to network administrations, we can find the answer for you – just fire off a letter or email and it'll all be taken care of.

LXF answers guy

David Coulson is a networking and security guru with plenty of sysadmin experience to boot.



Nick Veitch is the editor of the magazine, and answers your easy questions! Or indeed anything to do with *Grub*, *LILO*, *netatalk*, vi...



Hans Huberland is Rackspace Managed Hosting's Linux expert. Send any Linux system admin questions to sysadminqa@rackspace.co.uk



Kernel calamity

Q I downloaded kernel 2.6.1 from www.kernel.org and tried to compile kernel 2.6.1 on Red Hat 9 (existing kernel 2.4.8-20) using *grub* bootloader. The error message came out during I booted with kernel 2.6.1 was:

```
mount:error 19 mounting ext3
pivotroot : pivot root (/sysroot,
/sysroot/initrd) failed :2
umount /initrd/proc failed :2
Freeing unused kernel memory :
216 k freed
kernel panic : no init found. Try
passing init = option to kernel.
```

Is there any missing step that has caused this problem? I have done step by step below here :

```
# make config
# make menuconfig
I edited Makefile (EXTRAVERSION
= via)
# make dep
# make clean
# make bzImage
# make modules
# make modules_install
# cp /usr/src/linux-2.6.1/arch/i386/
boot/bzImage /boot/vmlinuz-2.6.1 via
# mkinitrd /boot/initrd-2.6.1 via
2.6.1 via
(/lib/modules/2.6.1 via)
- Copy /usr/src/linux-2.6.1/
System.map to System.map-2.6.1 via
and .config to config-2.6.1 via into
/boot
```

```
Added /etc/grub.conf
Title New RH linux (2.6.1 via)
root (hd0,0)
kernel /vmlinuz-2.6.1 via ro
root=LABEL=/
initrd /initrd-2.6.1 via.img
```

Any suggestions to solve this?

Chen Tu Yi/Supriady, via email

A It would appear that your kernel can't mount the *ext3* root filesystem, so possibly you are missing the *ext3* extensions

from your kernel. You'll want to ensure you have the latest packages required by 2.6, such as *mod-utils* and *procps*, as you may experience unwanted side-effects of the kernel change.

If you have a recovery kernel on your *grub* installation, then you can boot from that and recompile your kernel. Otherwise, you can boot off the Red Hat 9 installation disk, or a recovery disk, point *grub* to your old kernel and reboot the system.

Pande-modem-ium

Q I really hope that you can help. I purchased *LXF47* in order to be able to install Linux Mandrake 9.2. I also bought a copy of the *Complete Linux Handbook*. This is the first time that I have ever tried Linux. I have been using Windows since 1997 and have had very few problems but the reports on Linux were simply too good to ignore. So I installed Mandrake using a floppy and then from the DVD. At first all appeared to go well.

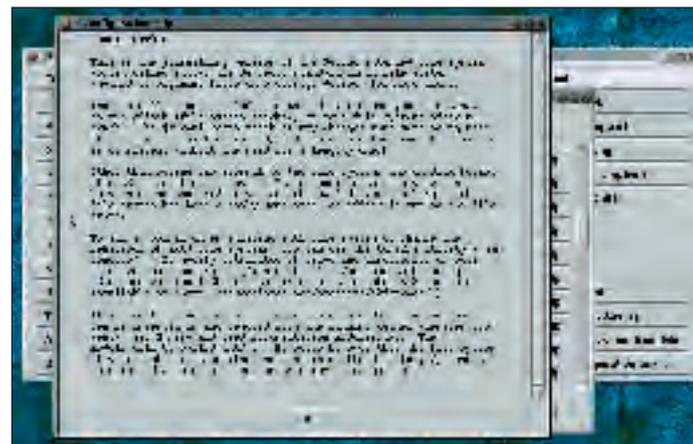
Then I hit problems. No matter what I try, I simply cannot connect to the Internet. I have searched high and low for help on the Internet but checking out all the

FAQs the answers seem far, far too technical or just not relevant; sometimes I simply don't understand them. I'm neither a beginner, nor am I some kind of computer genius, but many of the answers just blind me with science.

I am very disappointed, as I was lead to believe that Mandrake was the easiest distro for beginners. I have tried everything and have come up with nothing! My modem dials out, and then after a few seconds tells me I'm not connected. It gives me no clues whatsoever as to what the problem might be.

I've tried reinstalling Linux, but this has not worked either. The help files just don't help. I really am getting desperate, as I was so looking forward to giving Linux a try. I've tried the install wizard but am now confused by which choices I should select, there seems to be many things to select, but no explanation as to what each does. My system runs perfectly under Windows XP.

My system is as follows: Athlon 1GB, 640MB RAM, 16MB Voodoo 3 video, Hercules Fortissimo 11 sound, 400watt PSU, 120GB Samsung HDD, 20GB Western Digital HDD,



The *ext3* filesystem is a block driver that sits on top of *ext2*, so to mount a filesystem *ext3* you need kernel support for it.

Plextor CD-RW , Hitachi DVD, Zoom 56k external modem, Epson 830 printer and Epson 1260 scanner. Linux recognises all of these.

Please, please help. I really would like to get in to Linux.

Mark W Platt, via email

A As you've got an external serial modem, it should be very straightforward to connect to the Internet. You can use either the KDE PPP dialler, or *gnome-ppp*, which will connect to your modem and dial out to the Internet.

You will need to select the correct authentication method for your ISP: either PAP or CHAP. Both of the above tools can give debugging output, so you can follow the PPP connection process through the steps until you get a successful connection. You may also want to make notes of the configuration on your Windows installation to ensure that you're formatting the username correctly.

More MDK muddle

Q I purchased your December issue – LXF47 – due to the tasty Mandrake Linux 9.2 DVD. However, I was disappointed

to find out that during install, my Seagate ST380013 80GB SATA hard drive was not recognised (my mobo is a Gigabyte 8S655FX with SiS ATA drivers). I have scoured the Internet and only found minor hints on how to get Mandrake up and running on an SATA drive, and was wondering if you could please help me out? So far, the closest advice I have heard is to disable the DMA at the install prompt, but being a Linux newbie, I have no idea how to do that or what that does. Any guidance would be much appreciated.

Daryl, Australia

A Disabling DMA probably isn't going to help you too much, as the kernel should at least be able to recognise the drive with DMA enabled – after all, you can't have DMA enabled on something that doesn't exist.

You require support for SATA in the kernel, specifically the new controllers, so it may be likely that Mandrake 9.2 does not support the specific SATA controller on your board. Formulating a workaround yourself would be quite complex, but if you haven't updated the kernel so far, this



For those of us who can't manage with 11Mbit, the 802.11g specification is supported under Linux thanks to the Prism54.org project.

would be the first thing to try. Mandrake have released several kernel updates which you will be able to install using the package manager if you have set an external source.

PC card invisible?

Q As a new reader of your magazine and to Linux, I was hopping that you may

be able to help with a problem I am experiencing. I have just installed a copy of Red Hat Linux 9 on my laptop, and I am having problems getting my 3Com Wireless PC Card (3CRWE154G72) working. The OS (Red Hat Linux 9) appears not to recognise the PC Card, although I know it to be working correctly (checked the card within an XP machine). The laptop beeps when the card is inserted and beeps when it's removed.

Can you please explain (plain English) how do I get the OS (Red Hat 9) to see the PC card? How do I configure it within 'Networks'? Do I create a profile etc?

David Byfield, via email

A Red Hat 9 does not support 802.11g wireless adaptors 'out of the box', as they didn't exist when it was created. What you'll find useful is the prism54 project over at <http://prism54.org>, which offers some drivers for the chipset that your card has. It is development code, so it's not going to be as simple as clicking on a few things to make it all work smoothly. The project has fairly extensive documentation explaining how to get everything running, however it does require some small degree of knowledge of kernel compilation and patching.

Even with a more recent distribution, such as Mandrake 9.2 or Fedora, it takes some effort to get it all working. However, it doesn't seem particularly invasive and if you have some idea of how to compile everything, it shouldn't be too difficult to get the card working.

A QUICK REFERENCE TO: DVDs

Watching movies on a computer can be great fun: whether it's having your favourite movie on your laptop to take on a trip, or using the company's LCD projector to pretend that you are at the cinema watching the latest release. Of course, playing DVDs with Linux is not without its problems, particularly down to the use of CSS, FIXME, an encryption method used by studios to stop people copying DVDs digitally. However, this encryption is rather poor by most standards, and it didn't take 'DVD Jon', a young Norwegian, very long to create a crack for it. While this resulted in him being prosecuted, we now have an Open Source way of decrypting DVDs on Linux systems, which is usually distributed in the form of *libdecss*.

There are a number of players for Linux which can handle DVDs and DeCSS, the most popular of which are *Xine* <http://xine.sf.net/> and *MPlayer* www.mplayerhq.hu/.

Both are on this month's coverdiscs, and have their pros and cons, but generally, they work very well and each can take advantage of various DVD capabilities that the other lacks. A particularly problematic DVD feature is interactive menus; *Xine's dvdnav* plugin <http://dvd.sf.net/> is a good choice, although it does have a number of minor issues when dealing with some DVDs. Sadly, by getting DVD menus to work, we lose out on some speed improvements which the standard d4d player provides. *dvdnav* has been included into the main *Xine* distribution over the last few months, although it does not natively support DeCSS due to legal issues surrounding the distribution of the code within the US. However, individual users are able to install the DeCSS libraries on their own, which *Xine* will use.

Many people find *MPlayer* to be much faster than *Xine*, not to mention simpler to install, as it does not have a whole array of

dependencies. Both *MPlayer* and *Xine* can play other formats, such as AVIs and DivX (!) files, so we don't need to have half-a-dozen different applications to play various formats.

It is also possible to copy DVDs and create DivX files, which can be stored locally on a system and played without the DVD being present. A DVD in DivX format may take up around 1Gb, although this depends upon compression levels and the resolution used, so we could copy a selection of DVDs and store them all on a large hard disk for quick access. The most popular package to do this is *transcode* www.theorie.physik.uni-goettingen.de/~ostreich/transcode/ which can process a large number of video and audio formats into other formats. A DVD can be converted to a format that can be played on almost anything, even a low-end system without fancy software decoders.

FREQUENTLY ASKED QUESTIONS MOZILLA

FAQ MOZILLA IS SORT OF LIKE NETSCAPE, ISN'T IT?

Mozilla was developed using the source code originally used to build *Netscape 4*. However, it has practically all been rewritten, so apart from the general look-and-feel, it's actually 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 distributions 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 generates optimised binaries, it's often a good idea to just grab the binary tarballs. Not to mention the fact that 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 its 1.0 build quite a while ago, and the current stable release is 1.6. 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.

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 websites. 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* is completely compliant with the HTML, XHTML and CSS standards, as well as being much faster than the HTML rendering system that 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 that it can do with CSS to make your sites even more exciting.

FAQ ARE THERE ANY THIRD-PARTY ADDITIONS FOR MOZILLA?

There are a multitude of third-party projects for *Mozilla*, many of which can be found in the Internet at www.mozdev.org/. Many of these add extra capabilities to *Mozilla*, or extend upon those which are currently available. Indeed, a number of

www.mozdev.org/ projects have been merged into the main *Mozilla* tree, exposing these projects to a much wider audience.

Visit www.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, but others are extras which you wonder how you did without in the past.

FAQ 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 <http://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 by those who can fix it.

FAQ SO MOZILLA IS THE SAME AS EVOLUTION?

Not at all. *Evolution* is a PIM, where as *Mozilla* is a web browser. *Evolution* is more comparable to something like *Microsoft Outlook*, rather than *Mozilla*. Of course, one can use both *Mozilla* and *Evolution*, effectively replacing the need to have *Internet Explorer* and *Outlook* available. See the next issue of LXF for an *Evolution* tutorial.



Mozilla is a great Open Source web browser and mail client, but it also makes an alternative to Internet Explorer on Windows boxes.

« Laptop install

Q I have a Sony VAIO PCG-FR315S laptop that came with Windows XP preinstalled on it. I am a Linux fan (although still rather a newbie) and I have Linux Red Hat 9.0 installed on my desktop PC.

I didn't think that installing Linux on a VAIO was an easy job because of the rather unusual components that it had, until I read Jono Bacon's article about Red Hat

AS in LXF49 and saw that he was actually using a Sony VAIO, albeit not the same model that I have.

So before I commit myself to partitioning the hard disk and installing Linux, the question is: will it recognise all the components of my laptop? I searched the Internet and couldn't find the drivers for the specific model I have.

A *Ayman Sayed Darwish*, via email www.linux-on-laptops.com/ is a great resource for those who are attempting to install

Linux on laptops, as many people have documented their experiences, both good and bad. However, you are correct in that there is no listing for the specific Vaio you have.

However, a quick search with Google found a great resource at <http://s.robby.free.fr/sony.html>, where someone has installed Debian onto one of the laptops. While this is slightly different to Red Hat 9, they have listed all of the components on the laptop, which mostly appear to be supported by the standard Red Hat

9.0 kernel. Once you get the basics installed, you can then rebuild the kernel and add in support for anything that is missing.

With laptops, it's often a case of getting the base system installed, then working through each component which doesn't quite work right, until everything is happy. For laptops with very new components and chipsets, it is often a case of following kernel development or running some code that hasn't quite made it into standard releases at this stage.

Posting to the forum

The LXF online community

Not only do our popular forums at www.linuxformat.co.uk have sections dedicated to your technical queries, hardware, programming languages and general help; but also there's always a lively discussion going on!

Distro-saster

Q My problem is: after deciding to update my version of Linux, I thought I'd try some of the other distros to see what they had to offer; such as SUSE 9.0, Fedora Core 1 and Mandrake 9.2. I was currently using Mandrake 9.0 on an AMD XP 2000+ with 256DDR and 60GB HD.

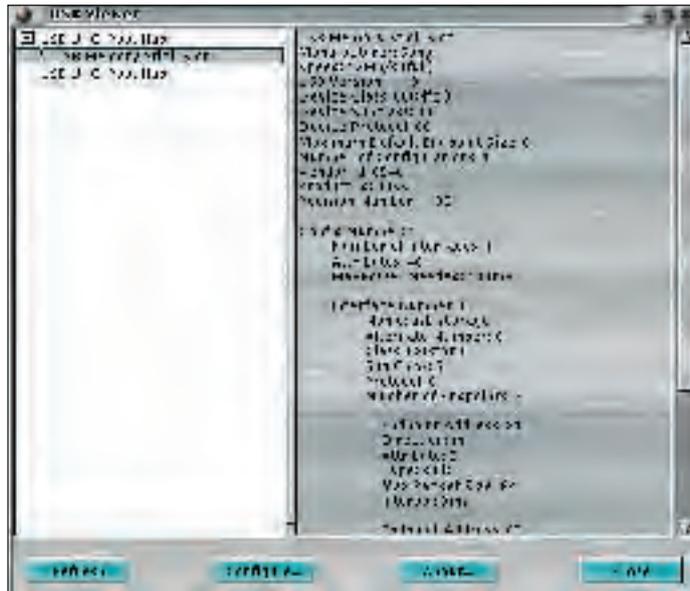
This was all well and good until I tried to setup file sharing on my home network of four computers, the other three running Windows XP Pro. I found in Fedora that this was easy, and files could be able to be read/write if I wanted, but I didn't like the way the system was set out, so I tried SUSE 9.0 (after a three-hour FTP install). This time, it didn't work properly with my 128MB GeForce FX, plus the networking wasn't able to do read/write. The same network problem was present in Mandrake 9.2. Have you any ideas what packages have been used in Fedora with networking that differ from Mandrake? This problem has driven me up the wall for months!

Garf, via email

A Windows networking in Linux is provided through the Samba system, which offers both clients to connect to Windows systems, and a server that allows Windows boxes to connect to our Linux system.

It should work with Mandrake 9.2 'out of the box', as it handles Windows filesystems through Samba as part of the standard installation. You may want to compare the configuration files in Fedora to those in Mandrake, specifically /etc/fstab and /etc/samba.conf.

If you can post more information regarding your problem mounting Windows filesystems, such as error outputs or specific information regarding what you are mounting, on the LXF forums, we can catch up with the information.



Supporting laptops under Linux can become quite frustrating, as they often contain custom chipsets or unusual devices.

Beginner's blues

Q I am a complete novice to Linux. I have been trying to learn and am getting quite far, but I find it frustrating when I try to install software. I have the Fedora Core distros.

When you install something on Windows it is all done via .exe files that just work via click-and-run?

Why is there not something like this for Linux? I know that RPMs are similar, but most of the time when I try to install a product such as MPlayer, there are lots of things I need to install and you need to know which order to install them!

I am so sure that if there was an install program for Linux distros (that was compatible with all distros) that everyone would start to transfer over to Linux from Windows. I have been trying for a few years now but won't until I am fully competent with installing apps.

Peter Sarson, via email

A Windows .exe files are simply binaries which you can execute. Some of these are self-extracting installation processes, which help you to install software.

RPMs, on the other hand, are more like zip files containing programs, their configuration files and anything else that is bundled with them. You install them through the 'rpm' tool, which installs all of the software and runs any post-installation configuration scripts that are needed to get the software working. Most distros come with a graphical

method to install and maintain packages, as well as downloading all the required packages you need to get something working. This requires less effort than under Windows: all you do is tell it what you want to install and let it download it. After that, everything is installed and ready to go.

Submission advice

We are happy to answer all sorts of Linux related questions. If we don't know the answer, we'll find out for you! But in order to give you the best service, it helps a lot if you read the following submission advice.

- Please be sure to include any relevant details of your system. 'I can't get X to work' doesn't really mean anything to us if we don't know things like what version of X you are trying to run, what hardware you are running on.
- Be specific about your problem. Things like 'it doesn't work' or 'I get an error' aren't all that helpful. In what way does something not work? What were you expecting to happen? What does the error message actually say?
- Please remember that the people who write this magazine are NOT the authors or developers of Linux, any particular package or distro. Sometimes the people responsible for software have more information available on websites etc. Try reading the documentation!

We will try and answer all questions. If we don't answer yours specifically, you'll probably find we've answered one just like it. We can't really give personal replies to all your questions.

WRITE TO US AT:
Linux Format, Future Publishing, 30
Monmouth Street, Bath BA1 2BW or
email: lxf.answers@futurenet.co.uk

New Linux Desktop



Linuxcom 518 Series
Smartest Media PC

"No more worm's or blue screen of death it's Linux"

Features 2 Ghz AMD Duron
266 Mhz Front-Side Bus
3D/2D Graphical Accelerator,
Motion Compensation
MPEG1/MPR11 Video
Ethernet/Broadband Ready

17" LCD Flat Screen
DVD/CD ReWriter+/-
80Gbyte Ultra DMA ATA 133 IDE
DDR 266Mhz memory
Wireless Keyboard and Mouse
Mandrake10.0

Apache 2.0 Server
OpenOffice 1.1
K3b DVD/CD Copy
GnuCash, Gimp
Thunderbird Anti-spam

Package 517 **£549+vat** (£645)
with 19" LCD Flat screen

Package 519 **£620+vat** (£728)
Desktop only

Package 501 **£310+vat** (£365)

Extras



Wireless usb £27 +vat
ISDN pci card @22 +vat
ADSL pci Card £21 +vat



www.lioncom.co.uk
sales@lioncom.co.uk

tel: 0044 1689 889003



ANSWERS



Dual-boot duel

Q Some of the workstations in my business are dual-boot, due to the fact I'm working on migrating from Microsoft to Linux systems. This means, when trying to work on other areas (or when playing games) we have to reboot to access Microsoft Windows XP Professional. I understand that Linux was never designed to require rebooting... not frequently at least. It has been recommended that I run the File System Check whenever I enter Linux again to ensure its performance doesn't degrade (it appears to be running rather slow (SUSE Linux 8.2 Professional on AMD 2100+, 1028MB RAM, 2x 60GB workstations), especially when booting (I was led to believe that ten seconds wasn't far-fetched, plus it feels much slower than it used to). I've attempted to run the **fsck** command and the variations such as **reiserfsck**, but I'm continually told (understandably) that I cannot perform the command as the hard disk/partition is mounted. I've received hints about using a single user mode to perform this task, but I know nothing about this – all my systems have a minimum of 2 users – root and the user of that particular machine. How can I run the **fsck** commands when rebooting?

Phil, via email

A You're quite right Phil, Linux was not designed to need frequent reboots; however, this does not mean that it does not work just as well as a desktop that is switched off each night.

I don't think it is necessary to run a filesystem check after every reboot if you're rebooting regularly. However, if you only reboot the server once every few months, then it's a good idea. I would agree that entering into single-user mode first is a safe way to do this. How to go about doing this will depend on which bootloader you have running on your system. Chances are you are using either *grub* or *Lilo*.

Lilo usually presents you with a **boot:** prompt. Hitting **Tab** twice here will show you which modes you can boot into (for example: Windows, Linux). To boot into a single user mode linux you should type **linux -s** at this prompt. In most default configurations, you should not be prompted for a password and you should be taken directly to a shell. Here you will be able to do all the **fscking** your heart desires.

If your system runs *grub*, things are a little different. You will probably be presented with a text-based or graphical menus from which you can select an operating system to boot into. Using the cursor keys, you should highlight the entry for your Linux system and press **E** (for edit). *grub* should then show several lines, telling you which disk drive and which kernel to use amongst other things. Here is an example showing my default configuration:

```
root (hd0,0)
kernel /vmlinuz-2.4.21-9.ELsmp ro
root=/dev/rd/c0d0p2
initrd /initrd-2.4.21-9.ELsmp.img
```

Find the kernel line and press the **A** key to append data onto this line. Add a **-s** onto the end of this list and press **Enter** to accept the change. Pressing **B** now should let you boot the system with the temporary change you have made to the configuration.

PPPProblems

Q I am having some problems with *pppd*. I have set up a private class A LAN, and the Linux box runs kernel 2.4.22, *Samba*, *DHCP*, *Apache* and *PostgreSQL*. I wish to get this server to automatically dial out to our Irish ISP on demand, or even a few machines that are hooked up to the private LAN to connect to the Internet (Microsoft-based OS) via the Linux server. Mind you, I have been reading countless HOWTOs and rolled up my sleeves to get to the nitty-gritty.

1 The funny thing with my current setup is that from a Windows PC, using *putty* to *ssh* to the Linux box it automatically dials out (I don't want this to happen!) which has been bugging me and don't know how to stop it!? I tried using *active-filter* keyword in the *pppd* script but that prevented anyone dialling out!

2 There are some sites that I can access from the Windows PC, some others are inaccessible, yet on the Linux box I can access all sites, I suspect DNS is the problem here? What do I do?

Samba 3.0.0 acts as PDC, *pppd's* version is 2.4.1 and the distribution is SUSE 8.2 Pro. I have enclosed the following files to aid in your understanding of my problem.

```
/etc/resolv.conf:
domain eircom.net
nameserver 159.134.237.6
nameserver 159.134.248.17
```

```
/etc/ppp/options:
demand
idle 15
noauth
usepeerdps
passive
defaultroute
lock
```

```
/etc/hosts
127.0.0.1 localhost
159.134.238.175 eircom.net
xx.xx.xx.xx linuxserver linsrvr
```

Hope you can help me out! Thanks!
Tom, via email

A Thanks for your detailed question. I believe that both problems are DNS-related. In the first part of the question when your Linux server receives an *ssh* session request, it is probably trying to do a reverse DNS lookup request for the IP address that you are connecting from. If you do not have a local DNS server, this action will trigger your PPP daemon to dial out to resolve the DNS. The easiest solution here is to add entries for all the IPs on your network

ABOUT RACKSPACE MANAGED HOSTING

Rackspace Managed Hosting[®] provides complete managed hosting solutions with Fanatical Support to small to medium-sized customers as well as large enterprises. All customer platforms include state-of-the-art data centres,

customised servers, burstable connectivity, 99.999% uptime SLA, a dedicated account manager, instant emergency response and access to live expert technicians 24x7 for support of all hardware and core software. Founded

in 1998 with locations in San Antonio, Texas and London, the company manages over 7,000 servers for customers in more than 80 countries. For more information, visit www.rackspace.co.uk and see this month's *Linux Pro*.



WIN an Archos AV140 Video Recorder WITH RACKSPACE MANAGED HOSTING

www.rackspace.co.uk

Every month, the best question

related to Systems Administration that a LXF reader sends in wins a prize: this month, you have the chance to win the awesome Archos AV140 Video Recorder.

The AV140 combines a modular MP4 video player and recorder, MP3 music player and recorder, digital camera and camcorder, digital photo wallet and data storage in a compact device that fits in the palm of your hand and is compatible with Linux. Based on a 40GB hard disk that is completely compatible with Linux, Windows and Mac OS, all you have to do is connect it to your computer and it will be mounted as an additional hard disk – on which you can store any type of data. USB 2.0 auto-connection ensures fast data transfer to and from your PC.

You can now enjoy 160 hours of near-DVD quality MP4 video with

MP3 sound wherever you are. You can record and playback to your television, VCR or camcorder directly in MP4 format. Playback is an AVI format file with MPEG4-SP VHS quality video and MP3 stereo sound, with the ability to record up to 2,000 hours of voice audio. See the webpage at www.archos.com/products/prw_500432_specs.html for more information on this prize.

Email all your sysadmin questions and problems to our expert system administrator Hans Huberland: sysadminqa@rackspace.co.uk
Every published query on these system admin pages wins a Rackspace T-shirt.

WIN!



★ Star Question – AV140 winner!

This issue's lucky winner is **Scott** – your new portable multimedia player/recorder will be with you shortly!

Q I have just hired a new Sysadmin who will be working for us from a remote location. I want to get him up-to-speed by walking him through how our systems are configured, how we handle certain functions, and so on. Ideally, I'd like to have him see my console session as I perform tasks and explain to him on the phone what I'm doing, I guess similar to VNC, but not in a GUI. Is there any way this can be done?
Scott, via email

A If you look at the man page for the script command, you'll see a neat little trick documented involving the **mkfifo** command. To sum it up, you should be able to log into a system where the demonstration will be done and type the following:

```
$ cd /tmp
$ mkfifo sharesession; script -f sharesession
```

What you're doing here is creating a named pipe called **sharesession** and then running script, which will save

all output on your terminal to a file called **sharesession**. The **-f** option causes script to flush output after each write, otherwise script will only write to **sharesession** when the forked shell has exited.

Your new System Administrator should now be able to log into the machine and do the following:

```
$ cat /tmp/sharesession
```

This will allow your new System Administrator to view what it is you are doing in real-time. Alternatively, if you would like your new administrator

to be able to interact with you and also use his keyboard to issue commands, there is a script called **kibitz** which comes with the *Expect* package (<http://expect.nist.gov>) which ships as standard with many distributions. To use it both you and your new administrator need to be logged into the server, then just type:

```
$ kibitz username
```

He will then see a prompt telling him what he needs to type; and once he has done that, the pair of you will be sharing a terminal.

to /etc/hosts. However, it would be more technically correct (and more efficient if you were doing this on a larger scale) to set up a local caching DNS server. This is pretty straightforward and *BIND*'s configuration does this by default. If you are going to use a local caching DNS server, you will just need to add entries for the IPs on your LAN onto this DNS server. Ideally, all the workstations should be set to use this DNS server using DHCP. I'm not certain why some sites are responding from Windows and

others are not, but if you would like to verify that it is in fact DNS-related, you could try pinging the domain name from a Linux box and taking note of the IP address which is resolved.

Then try pinging the same name and IP from the Windows machine. If the domain name ping does not return any packets, but the IP address does, then the problem is most definitely DNS-related and my suggestions above should sort your problem out with a minimum of fuss.



Captain's login

Q I use Red Hat Linux 9, with a 3Com 3C90X NIC card ADSL modem and SMC router. I used to be able to log into the Net and it does not give me any problems. But recently I used the LXF49 coverdisc to install Fedora Core 1 and I am facing problems.

The problem is that they do not assign me a IP address to my NIC card. How can I assign a permanent IP address when I boot up? When I use `ifconfig` to assign a IP address to the NIC card, after I rebooted, the IP address was gone. (when using RH 9, when I boot up, my router light blinks, so I know that I could connect to the Internet.) Following creating a folder with space, using `Xterminal` I could not get into the dir using `cd`.

After installing a Tarball program, the new programs files is not in the launcher menu. How can I create a launcher in the Launchers menu and modify a launcher's properties...?

Lai Ying Wah, via email

A When you use `ifconfig`, it simply sets the IP address on the box and doesn't save it once the box has been rebooted. You can set the IP address which is to be set a boot time by modifying `/etc/sysconfig/network-scripts/ifcfg-eth0`. You should also be able to set it up to use DHCP using the `dhclient` script, which will perform a DHCP request on the network for you if that is supported by your ISP.

You can access files and directories containing a space by putting the filename "in quotes". You can also precede the space with a `\` character, which will escape it for you.

You will be able to add a menu item using the 'Desktop Preferences', although you will manually point it to the binary. Most distributions handle the menus properly, so if you can install the packages for your distribution rather than compiling things: then, you won't have the extra hassle. Of course, some source distributions come with the launcher script already, so if you point them to the right place for your GNOME or

KDE installation, it should work fine 'out of the box'.

Deeply dippy DHCP

Q I've been reading everything on networking and TCP/IP, but there is a particular problem that I do not understand. I give an example:

I have a small network with an E-Tech ADSL router that I also configured as DHCP server, and this works OK. In the network I have two PCs: PC A is a file and print server and runs `Samba` and exports two NFS points. Up to now it has a fixed IP address.

PC B is workstation and mounts the two NFS points in its file system and also has a fixed IP address. Both PCs resolve the names via `hosts` file.

Now when I change PC A to DHCP, it gets an IP address alright from the router, but how should PC B now resolve the name? When I do this now, the whole NFS story goes out of the window, which I understand why it happens, but that's all.

From my former capacity in a professional environment (database applications), I seem to remember that a machine which has to be a server of some sort needs to have a fixed IP address. From what I see, I could believe that, but I am still wondering. Can you enlighten me about this?

Herman Viaene, via email

A If you are using NFS, you really want to avoid using DHCP. It's not improbable that your IP address may change while the box is up, which will disrupt the existing NFS mount on the client. As you have two systems, there isn't a whole lot to be gained by making them work using DHCP, as it's not a complete nightmare to configure both systems for static addressing.

You should be able to modify your router to leave a block of addresses at the front of the range for static assignment, and use the rest for DHCP in the event that you have a temporary system or a workstation that doesn't require any static addressing. [LXF](#)

A HELPING HAND WITH FEDORA

BIOS change?!

The Red Hat Fedora distro from the LXF49 cover discs all works well except that I cannot find a way to mount my Windows partition which is `hdb1 vfat`. Is this possible as I want to copy some jpeg files over to experiment with `The GIMP`? I could do it via floppies, but surely there is a way without this bother? I'm running a 750MHz Duron with 128 MB of RAM; the 20GB hard disk is partitioned as follows:

- `hdb1 vfat 7319 MB running Windows 98SE`
- `hdb2 ext3 102 MB /boot`
- `hdb3 ext3 11868 MB /`
- `hdb4 Extended 251 MB`
- `hdb5 swap 251 MB`

`hdb4` and `5` seem to be on top of each other! Before installing Fedora, my HD was primary master but has now become secondary master – a change to the BIOS?! Strange, but it all still works OK! When starting Fedora, it reports: Starting `smartd`...[failed]

Peter Lonsdale, East Devon.

An extended partition is not a partition that stores data. It simply stores lots of smaller partitions, known as logical partitions; this is to get around the ancient four-partition limitation from the early days of DOS. You should be able to mount `/dev/hdb1` with:

```
# mount -t vfat /dev/hdb1 /mnt/dos
```

If that does not work, providing us with the error output – or posting it to the LXF forum – should provide some answers. The hard disk can't change from primary to secondary, or master to slave without something being changed in the BIOS or in the physical configuration of the IDE system.

Install ire

I have been trying to install some apps like `Rekall` and `kSQLanalyzer` from LXF49's discs, and while the `./configure` stage seems to complete without any errors, the make stage bombs out with:

```
/usr/bin/ld: cannot find -lpng
collect2: ld returned 1 exit status
make[3]: *** [libkbase.la] Error 1
make[3]: Leaving directory
```



Fedora is a freely distributed fork of the Red Hat distribution, and should be considered for new installations rather than Red Hat 9.0

```
~/etc/rekall/rekall-2.2.0-beta0/libs/kbase'
make[2]: *** [all-recursive] Error 1
make[2]: Leaving directory
~/etc/rekall/rekall-2.2.0-beta0/libs'
make[1]: *** [all-recursive] Error 1
make[1]: Leaving directory
~/etc/rekall/rekall-2.2.0-beta0'
make: *** [all] Error 2
```

I'm running a standard install of RH 8, upgraded to 9, and now Fedora. Chor Hin Ong, via email

A You are missing `libpng`, which is required when compiling this code. If you install the `libpng` package, along with the `libpng-devel` package containing the headers, then you should be able to build the code correctly. You may be missing other libraries and headers required by code you are compiling, however you will have to examine the error output from each build to establish what you need.

User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally natter about all things Linux. You can find lots more information online at: www.lug.org.uk

1 HAMPSHIRE

URL www.hants.lug.org.uk
Contact Hugo Mills

2 BRISTOL & BATH

URL www.bristol.lug.org.uk

3 SCOTTISH

URL www.scottish.lug.org.uk

4 OXFORD

URL www.oxford.lug.org.uk
Contact Alasdair G Kergon

5 KENT

URL www.kent.lug.org.uk
Contact Kevin Groves

6 BRIGHTON

URL www.brighton.lug.org.uk
Contact Johnathan Swan

7 WORCESTERSHIRE

URL www.worcs.lug.org.uk

8 NORTHANTS

URL www.northants.lug.org.uk
Contact Kevin Taylor

9 ANGLIAN

URL www.anglian.lug.org.uk
Contact Martyn Drake

10 MILTON KEYNES

URL www.mk.lug.org.uk
Contact Denny De La Haye

11 SCUNTHORPE & DONCASTER

URL www.scundog.org
Contact Shaun Holt – shaun@scundog.org

12 MORAY

URL www.moray.lug.org.uk
Contact Stewart Watson

13 WEST WALES

URL www.westwales.lug.org.uk
Contact Dan Field

14 WOLVES

URL www.wolveslug.org.uk
Contact Jono Bacon

15 PETERBOROUGH

URL www.peterboro.lug.org.uk
Contact Steve Gallagher

16 EDINBURGH

URL www.edinburgh.lug.org.uk
Contact Alistair Murray

17 TYNESIDE

URL www.tyneside.lug.org.uk
Contact Brian Ronald

18 LEICESTER

URL www.leicester.lug.org.uk
Contact Clive Jones

19 GREATER LONDON

URL <http://glug.linux.co.uk/>
Contact John Southern

20 SURREY

URL www.surrey.lug.org.uk
Contact Jay Bennie

21 CAMBRIDGE

URL www.cam-lug.org.uk

22 DEVON & CORNWALL

URL www.dclug.org.uk
Contact Simon Waters

23 FALKIRK

URL www.falkirk.lug.org.uk

24 MANCHESTER

URL www.manlug.mcc.ac.uk
Contact John Heaton, Owen Le Blanc

25 HERTFORDSHIRE

URL www.herts.lug.org.uk
Contact Nicolas Pike

26 WEST YORKSHIRE

URL www.wylug.lug.org.uk
Contact Jim Jackson

27 SHEFFIELD

URL www.shefflug.co.uk
Contact Richard Ibbotson

28 STAFFORDSHIRE

URL www.staffslug.org.uk

29 NORTH EAST

URL www.shofar.uklinux.net/NELUG

30 LONDON

URL www.lonix.org.uk

31 BERKSHIRE & THAMES VALLEY

URL www.sclug.org.uk

32 LIVERPOOL OPENSOURCE

URL http://linux.liv.ac.uk/_liv_linux_ug/
Contact Simon Hood

33 DEAL AMIGA CLUB

Email superhighwayman@hotmail.com
Contact John Worthington

34 CHESTERFIELD

Email spirelug@yahoo.co.uk
Contact Robin Needham

35 SOUTH DERBYSHIRE

URL www.sderby.lug.org.uk
Contact Dominic Knight

36 BELFAST (BLUG)

URL www.belfastlinux.cx
Email russell@belfastlinux.org

37 WILTSHIRE

URL www.wiltshire.lug.org.uk
Contact Jason Rudgard

38 SOUTH LONDON

URL www.sl.lug.org.uk
Email edo@perceptiondm.com

39 CHESHIRE

URL www.sc.lug.org.uk
Contact Anthony Prime – enquiry@sc.lug.org.uk

40 NORTH WALES

URL www.northwales.lug.org.uk
Contact Andy Hutchings A-Wing deltaone@virgin.net

41 MIDLANDS

URL <http://midlandslug.port5.com/>
Contact Pete Thompson

42 CUMBRIA

URL www.cumbria.lug.org.uk
Contact Jamie Dainton

43 DORSET

URL www.dorset.lug.org.uk
Contact John Robinson

44 SHROPSHIRE

URL www.shropshire.lug.org.uk
Email shropshire@lug.org.uk

45 SOUTH WEST

URL www.southwest.lug.org.uk
Email southwest@lug.org.uk

46 SOUTH WALES

URL www.swlug.org.uk

47 NORTH LONDON

URL www.kemputing.net/lug/anlug-aims.html
Email jason@voyagercomputers.co.uk

48 MALVERN

URL www.malvernlug.cstf.co.uk
Contact Phil Ironside – phil@cstf.co.uk

49 HUDDERSFIELD

URL www.hud.lug.org.uk
Contact Dave Naylor – knocker@caramboo.com

50 NOTTINGHAM

URL www.nottingham.lug.org.uk

51 ST ALBANS & LUTON

URL www.lust.lug.org.uk
Contact Michael Culverhouse – mike@easily.co.uk

52 WREXHAM

Contact Paul Kersey-Smith
Email paul@pkls.fsnet.co.uk

53 PRESTON & LANCS

URL www.preston.lug.org.uk
Contact Phil Robinson

54 DERRY

URL www.derry.lug.org.uk

55 ISLE OF WIGHT

URL www.iow.lug.org.uk
Contact David Groom – info@iow.lug.org.uk

56 SCARBOROUGH

URL www.scarborough.lug.org.uk

57 BLACKBURN

Email matt@consultmatt.co.uk

58 YORK

URL www.york.lug.org.uk

59 LINCS

URL www.lincs.lug.org.uk

**60 HULL**URL www.hull.lug.org.uk**61 WALTON-ON-THAMES**Contact William Mutch
Email rael@freeuk.com**62 GLOUCS & COTSWOLDS**URL www.gloucs.lug.org.uk**63 WEST OF SCOTLAND**URL www.wos.lug.org.uk**64 SOUTH STAFFORDSHIRE**URL www.staffs.lug.org.uk**65 MANSFIELD**URL www.mansfield.lug.org.uk**66 BORDERS**URL www.linux.bordernet.co.uk**67 BIRMINGHAM**URL www.sb.lug.org.uk**68 COVENTRY**Email info@coventry.lug.org.uk**69 NEWARK & LINCOLN**URL www.newlinc.lug.org.uk**70 BEDFORDSHIRE**URL www.beds.lug.org.uk**71 LINCOLN**URL www.lincoln.lug.org.uk**72 LOUGHBOROUGH**URL www.loughborough.lug.org.uk**73 EXETER UNIVERSITY**URL www.euslug.lug.org.ukEmail N.J.Murison@exeter.ac.uk**74 SUNDERLAND**Email thomas.croucher@sunderland.ac.uk**75 EAST YORKSHIRE**Email sharkonline@whataemail.com**76 CLEVELAND OPEN SOURCE GROUP**Email openlug@digitalmedia.co.uk**77 BEVERLEY**URL www.beverley.lug.org.uk**REVISED
DETAILS****78 DUNDEE & TAYSIDE**URL www.dundee.lug.org.uk**79 SUSSEX**URL <http://sussex.lug.org.uk/>**80 WIGAN & ST HELENS**Email paulf.johnson@ukonline.co.uk**81 BRIXTON**URL www.communitytechnology.org.uk/~linuxhome**82 ST.ANDREWS, FIFE**URL www.standrews.lug.org.ukEmail stuart@nx14.com**83 NUNEATON**URL www.nuneaton.lug.org.uk**84 ISLE OF MAN**URL www.iom.lug.org.ukEmail helix@manx.net**85 AYLESBURY**URL www.aylesbury.lug.org.ukEmail drbond@educational-computing.co.uk**86 LANCASHIRE**URL www.lancasterlug.org.uk**87 EAST LONDON**URL www.eastlondon.lug.org.uk

Contact Jonathan Spriggs

88 ORMSKIRKEmail rob@northwestlinux.co.uk**89 HEREFORD**URL www.hereford.lug.org.uk/Email rbjh@good-news.fsnet.co.uk**90 EAST HERTS**Email madtom1999@yahoo.com**91 SWINDON**Email nick.trueman@ntlworld.com**92 MENAI**URL www.menai.lug.org.uk**93 ABERDEEN**URL www.aberdeen.lug.org.uk**94 SHETLAND**URL www.shetland.lug.org.ukEmail c_s_s_butler@yahoo.com**95 GLASTONBURY**URL www.glastonbury.lug.org.uk

Contact Steve Leonard-Clarke

96 SOUTHEND-ON-SEAURL www.sos.lug.org.uk

Contact Derek Shaw

97 ORPINGTONURL www.orpington.lug.org.uk

Contact Barry Schofield

YOUNG LINUXURL www.young.lug.org.uk**SCHOOLS**URL www.schools.lug.org.uk

LINUX USER GROUPS

LUGS OF THE MONTH

THE GREAT 2004 LUG EXPLOSION!

It's spring, and all over the UK there are unprecedented numbers of new LUGs springing up! These groups will be included on our map next month, but if anyone from these groups would like us to publicise their activities in this *LUG of the Month* panel – get in touch with us at linuxformat@futurenet.co.uk.

CHICHESTER HARBOUR MARITIME GEEKS

www.mgeeks.lug.org.uk

Hugh Saunders hugh@mjr.org

BLACKPOOL

www.blackpool.lug.org.uk

Chris Hopley chop@blackpool.ac.uk

CHANNEL ISLES www.ci.lug.org.uk

Myles Hartley myles.hartley@db.com

LEAMINGTON SPA

www.leamington.lug.org.uk/

Ian Sixsmith ian@siim.info

NORTHWEST LONDON

www.northwestlondon.lug.org.uk

D Braganza diamond@fastmail.fm

RUGBY www.rugby.lug.org.uk/

Matt Grice

rugbylinux@traveller-software.co.uk

STUDENTS www.students.lug.org.uk

Jamie Davies

jamie.davies@opensourceware.net

THANET (MARGATE)

www.thanet.lug.org.uk

Patrick Regnoul patrick@patricknet.net

TRURO SCHOOL

www.tslug.lug.org.uk

Tim Sheridan tim@rapidgn.com

WEST LOTHIAN

www.westlothian.lug.org.uk

Ben Drysdale ben@150bpm.co.uk

We also need suggestions from every reader on how to better represent both UK and world LUGs in *LXF* – send 'em in to the same email address, and we promise to use all the best ideas!

Worldwide Linux User Groups

Free Software users across the globe

Africa

EGYPT

URL www.linux-egypt.org

GAUTENG, SOUTH AFRICA

URL www.glug.org.za

Email glugmin@revolution.org.za

THE LORD'S ABODE, JO'BURG, SA

Email Andrew Gargan avrin17@iname.com

Australia

ADELAIDE

URL www.linuxsa.org.au

Email mtippet@anu.edu.au

ALICE SPRINGS

URL www.aslug.org.au

MELBOURNE, VICTORIA

URL www.luvasn.au

Contact luv-committee@luvasn.au

PERTH

URL <http://plug.linux.org.au/>

SYDNEY

URL www.slug.org.au

Europe

COSTA DEL SOL (English speaking)

URL www.fuengirola.lug.org.uk

DENMARK

AIssund www.alslug.dk

Esbjerg www.eslug.dk

Fyns www.flug.dk

Midt-og Vestjylland www.mvjlug.dk

Nordjylland www.njlug.dk

Skåne Sjælland www.sslug.dk

Trekantsområdet www.tlug.dk

Vest-fyn www.haarby-net.dk/vflug

Århus www.aalug.dk

EIRE

URL www.linux.ie

Email root@linux.ie

URL www.dilu.org

Contact glossary@dilu.org

MILUG (Longford)

URL <http://midlands.linux.ie>

Contact midlands@linux.ie

Middle East

ISRAEL

URL www.iglu.org.il/IGLU/

Contact webmaster@iglu.org.il

PALESTINE

URL www.lugps.org

Email isam@planet.edu

Asia

HONG KONG (multilingual)

URL www.linux.org.hk

SINGAPORE – SLUG

URL www.lugs.org.sg

SRI LANKA

URL www.lklug.pdn.ac.lk

MYANMAR (formerly BURMA)

URL www.myanmarlug.org

Email afyde@balug.org

PAKISTAN

URL www.linuxpakistan.net

Email tux@clug.org

HYDERABAD, SINDH, INDUS VALLEY

URL www.geocities.com/slug_pk/

KASHMIR

Coming soon!

China

BEIJING (GB encoding, but mostly written in Chinese)

URL <http://mud.263.net.cn/~linux>

CHINESE LINUX USER GROUP

URL www.linux.org.cn

NANJING

URL <http://jllib.jlonline.com/njlug>

India

LINUX INDIA

URL <http://linux-india.org>

ALIGARH LUG

URL <http://linux.amupost.com>

BOMBAY

URL www.ilug-bom.org.in

CHANDIGARH

URL www.geocities.com/vipinb

CHENNAI AND MADRAS

URL www.chennaiug.org/

CYBERABAD (CLUG)

URL <http://seeknew.freesevers.com/clug/>

DELHI

URL www.linux-delhi.org

KOLKATA

URL www.ilug-cal.org

MADURI

URL <http://linuxmadurai.tripod.com>

NORTHERN INDIA LINUX

URL <http://groups.yahoo.com/group/lug-northindia>

Spreading the word

Jono Bacon identifies the importance of using relevant case studies in advocating Linux use to businesses of all sizes.

As has been ranted here on many an occasion, case studies are the lifeblood of a Linux proposal. A business wants to see that Linux has worked for someone else before they will even consider the idea. This is all very well, but where do we find out case study details?

The Internet should be your first port of call. There are many successful case studies documented for different types of organisation. A bit of Googling will no doubt bring up something relevant. It is also recommended that you regularly check the big news sites such as *Linux Today*, *The Register* and *Slashdot*.

Another method is to speak to people that you know who may use Linux in a similar environment. The key word here is 'similar'; a case study of how a particle acceleration unit uses Linux is not exactly comparable to that of a small business. Remember to keep the case studies similar as this will sow the seeds of Linux viability easier.

The information contained within the case study is quite important. Costings, support, training and stability are all more

important than software used and the ethical factors of Linux. Most businesses probably don't care about the ethical stance of Linux – they want security, stability and demonstrable improvement in cost efficiency. If all of this can be achieved at a low price, it is likely that Linux will be of interest to a business.

There's one more important source of case studies: our *Linux Pro* supplement. Over the last couple of years, we have included case studies of Linux migrations and use of Open Source software from organisations as diverse as Insurance companies and Premiership football clubs.

Always remember what the business wants of Linux, and try to tune a case study around those needs. Certainly do not over-embellish a case study, just pick a number of relevant cases that are pertinent to the business in question.

Don't forget to check out the new Business LUG site featured in *LXF51*'s User Groups (www.businesslug.org); and as ever, please send your comments, advocacy stories and suggestions to spreadingtheword@jonobacon.org **LXF**

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

Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format DVD*. *PwManager* information here is pertinent to CD readers too.

DESKTOP KDE3.2

The CDs this issue contain the source code for the newly released KDE 3.2. Installing this way can be a lengthy task, easily taking a full day to compile and install all of the components on your system. Fortunately for those of you with the DVD edition, you can install KDE from binary packages in a fraction of the time.

If you are using an RPM-based distro, **cd** to the appropriate directory of the DVD and type one of the following commands:

```
rpm -Uhv *.rpm
```

```
rpm -Fhv *.rpm
```

What is the difference? The first command installs all of the packages in

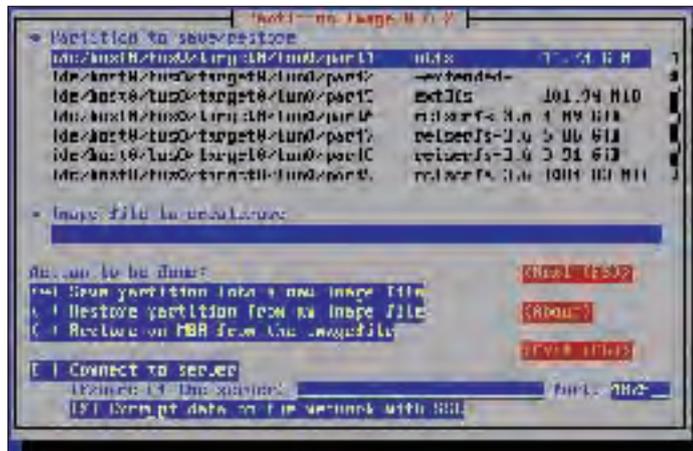


KDE 3.2's new *Kontakt* 'meta-app' integrates *KMail*, *KOrganizer* and *KAddressbook* in one groupware package, and you can still use the components separately.

the directory, as updates to any existing packages, new installations otherwise, hence the **U** for update. The second command works similarly, but only installs packages that are an update of a currently installed package. The **F** is for freshen, it will only upgrade, not install anything new. If you have a full KDE installation already, it won't make any difference which you use.

The Debian, Fedora Core and SUSE packages came from the official section of the KDE FTP site. The Slackware packages were provided in the contrib directory, so should be considered unofficial. Mandrake hasn't released any KDE 3.2 packages for 9.2, and is unlikely to as developers are concentrating their efforts on finishing Mandrake 10.0. The Mandrake packages on the DVD are from a third-party source (ftp://mandrake.contactel.cz/people/bluehawk/kde32-92) and are backports of the Cooker packages (effectively 10.0 beta) to 9.2. Some reports indicate that you may have to remove your existing KDE installation before installing from these.

Whichever installation method you use – source, official packages or unofficial packages – this is a major system upgrade. It would be a wise move to make a full backup before proceeding. The easiest way to do this, provided you have another partition that can hold the backup file, is to boot from



It may be wise to backup your root partition before installing KDE 3.2, and Partition Image is just the program for it.

the DVD into the *System Rescue CD* environment and use partition image to make a cloned copy of your root partition. Type **partimage --help** to see the options, it's fairly straightforward. Running *partimage* with no arguments gives a *curses* GUI from which you can select the various options.

DESKTOP MYPASSWORDS SAFE

DESKTOP PWMANAGER

Do you feel like you're getting snowed under with passwords? Every website you visit seems to need a login, every mobile phone has passwords now. Keeping track of all of these can be a nightmare. You have several choices. You could write them all on Post-it notes and stick them on the side of your computer or monitor (don't laugh, a lot of people do this). You could use the same password for everything, but this is very insecure (as is the previous option) and each organisation puts different restrictions of password lengths and contents, so you still end up with several passwords. Alternatively, you could store them all

in one secure location. KDE 3.2 has a partial solution to this problem with its *KDEWallet* system, which keeps web site login and password details in an encrypted file, so you only need to remember one password. But it doesn't help with those sites that use select gadgets instead of text fields for data entry, and it won't keep the PIN for your mobile phone.

What you really need is a way of keeping all this private information in an encrypted form, protected by a master password. Because you only



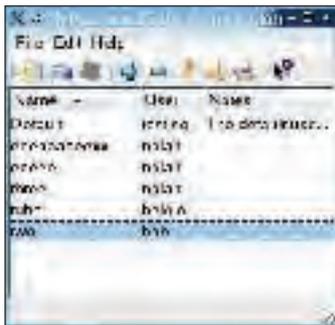
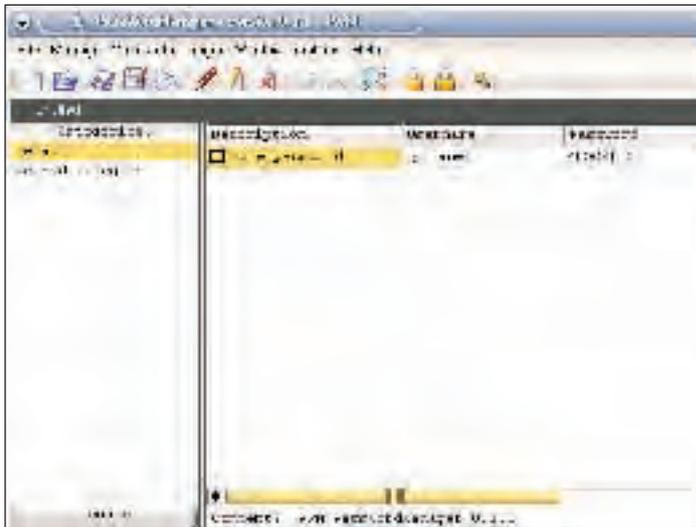
Wherever you see this logo it means there's related stuff on the DVD

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.

WARNING: KMAIL CAN LOSE MAIL!

A bug in *KMail* surfaced shortly after release: this can result in the loss of mail under certain circumstances. A patch has been released and is included on the CD. To apply it, copy **kdepim-3.2.0-HOT_FIX_FOR_KMAIL.patch** to the same directory as **kdepim-3.2.0.tar.bz2**, unpack the original archive as usual, with **tar xjf kdepim-3.2.0.tar.bz2** and then apply the patch by typing **patch -b -p0 <kdepim-3.2.0-HOT_FIX_FOR_KMAIL.patch**. Now just continue with the compilation as before.



Keep all your passwords, PIN codes and other private information in one secure place, with either *MyPasswordSafe* or *PwManager*.

Both programs come with source code to be installed in the usual way. *PwManager* also comes with a range of binary packages for the main distros. *MyPasswordSafe* doesn't have RPMs but it does have a tarball of a compiled binary. You can unpack and run this without installing, which is an easy way of trying the program to see if you like it. If you do, you can either compile the source or copy the binary to somewhere in your path, such as `/usr/local/bin`. Both *MyPasswordSafe* archives contain the documentation as an HTML page. You will have to install *PwManager* to read its full documentation, though.

SERVER SPAMASSASSIN

Unsolicited Bulk Email (UBE) – or spam as it is more commonly known – is a real problem. A year ago, many pundits were forecasting that spam would account for more than 50% of the total email traffic on the Internet by the end of 2003. Unfortunately, they were right.

There are many different techniques that can be used to filter spam, and this program is able to use most of them. It is able to identify key phrases, block or allow mail on address, use black hole lists and also do Bayesian analysis of the mail's content. Most importantly, it is able to learn from the mails you receive to improve its accuracy. This is important because spammers are also interesting in anti-spam software, they are trying to find ways around it.

SpamAssassin works well with its 'out-of-the-box' configuration, but

CREATING A SYSTEM RESCUE CD ISO IMAGE FROM THE COVER DISC



***System Rescue CD* contains partitioning and rescue tools, and boots straight from the DVD.**

If you have the DVD issue you may want to create a CD version of the *System Rescue CD*. Or maybe you have the CD but want only the System Rescue files on a CD, so you can add others or burn it to a mini CD. In either case, creating the ISO image is quick and easy. As usual, we have provided scripts to do this from Linux or Windows, although the Windows script is redundant, since you can boot from the cover disc and run the Linux version.

To build the ISO image in Linux, type the following command in a terminal

```
sh /mnt/cdrom/Distros/SystemRescueCD/mkiso
```

This will create the ISO image in the current directory. If you want to create it somewhere else, give the path as an argument, for example:

```
sh /mnt/cdrom/Distros/SystemRescueCD/mkiso /tmp/iso
```

If you want to do this from within the System rescue environment, all you need to do differently is mount one of your hard disk partitions first. For instance, if your `/home` partition is on `/dev/hda6` and your home directory is in `/home/me`, type

```
mount /dev/hda6 /mnt/temp1
```

```
sh /mnt/cdrom/Distros/SystemRescueCD/mkiso /mnt/temp1/me
```

Windows users can create the CD image by double-clicking the `winmkiso` icon or running the script from a MSDOS prompt. In the latter case, you should change to the `SystemRescueCD` directory of the CD before running the script. You can specify a destination directory when running it from a DOS prompt. With no argument, or when run from the icon, it puts the ISO image in `C:`.

For example, if your DVD drive under Windows/DOS is `E:` and you want to save the ISO image to `D:\ISO`

```
E:
cd Distros\SystemRescueCD
winmkiso D:\ISO
```

INSTALLING KDE 3.2 ON GENTOO

If you installed Gentoo from our recent coverdiscs and want to upgrade to KDE 3.2, the process is quite painless. First you will need to run `emerge sync` while online to make sure you have the latest portage tree. Then copy the source tarball to your distfiles directory, this is `/usr/portage/distfiles` by default. Now, while you are online, run the following:

```
emerge -f kde
```

This will download any updated files, such as the *KMail* patch, and should only take a short while, even on dialup.

```
emerge -u kde
```

to compile and install the latest version with your Gentoo setup.

really shows its strength when you use its learning capabilities. It works by feeding it two sets of mails, in either `mbox` or `maildir` format, one is your genuine mails the other contains the spam mails it failed to identify. The `sa-learn` program will use this information to filter more effectively in future.

SpamAssassin is best used on a mail server, blocking mail as it arrives, but it can also be used just as effectively between a POP3 retrieval program and your mailer.

GAMES GCOMPRIS

One of the reasons for the ubiquity of Windows is familiarity. It is usually the first – and often the only – operating system that most people use. Increasing the usage of Linux, particularly among home users, must entail increasing exposure, and at an early stage of the computing experience. You don't get much earlier than two-years-old, which is the beginning of the target age group of *GCompris*. This is a complete

COVERDISC DVD

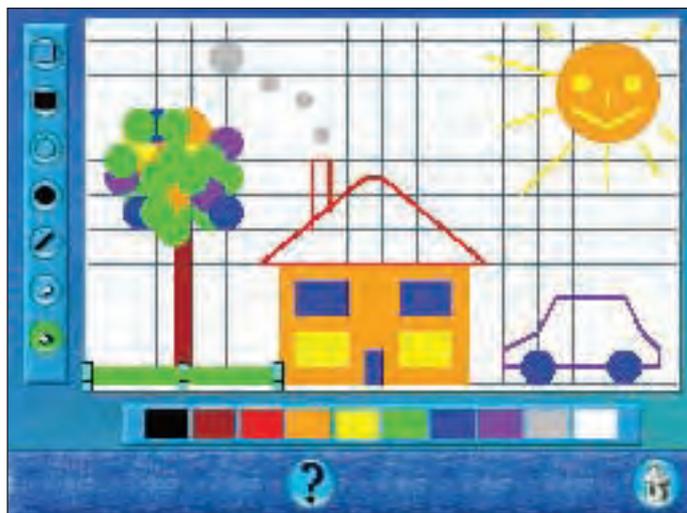
« educational suite for ages from 2 to 10. In contrast to the excellent *TuxPaint* *GCompris* provides a collection of various educational activities, accessed through a common interface. The activities are grouped into six categories; reading, amusement, experiences, board based, algebra and discover the computer.

The DVD contains a generic RPM package and a set of Debian packages, as well as the source code. It is a big package, around 30MB, but installation doesn't take long, even when compiling from source, because much of this is data rather than code. *GCompris* provides a great deal of audio, as well as visual, feedback for the youngsters using it. A large part of each package is taken

up with Ogg Vorbis files. *GCompris* has a long list of dependencies, but these are general packages that are included with your distro. You shouldn't need anything else to install it, but you may need the distribution discs handy.

While this may be a good way of exposing people to Linux at an early age (both for the children and their parents) the most important thing about *GCompris* is that it is fun and educational for children to play with. If you have children in this age range, do them a favour and install it. [LXF](#)

Keep the kids happy and gain impressionable converts for the penguin with *GCompris*.



DVD CONTENTS AT A GLANCE

Desktop

ChruppyInfoCenter	A system monitor that uses root-tail to display data
CLAWFileBrowser	An icon view browser
Cuetools	A set of utilities for working with cue files and TOC files
ESEKeyDaemon	A multimedia keyboard (userspace) driver for Linux
EvRouter	A Linux input event router
Filer	A small file manager
GKrellMCountdownPlugin	A tool to count down to a predetermined time in GKrellM 2.x
gkrellmSunClock	A gkrellm solar plugin
Gtk-ChangeTheme	A slick GTK+ 2.0 theme changer
Imwheel	Support for wheel and 4+ button mice in X11
jspFileBrowser	A Web-based file access and manipulation system
Killer-Applet	A KDE applet process killer
LinuxEasyAccessKeyboard	Linux support for Easy Access and Internet Keyboards
Lire	A pluggable log analyzer which supports over 30 log types
Mathomatic	A symbolic math program
MMKeyboardController	A project to enable mouse gestures and multimedia buttons
MyNetMonitor	A GTK 2 net monitor with system tray support
MyPasswordSafe	A password manager
PyBabelPhish	Language translation software
Roffit	A program that reads an nroff file and outputs an HTML file
Smb4k	A SMB share browser for KDE 3.1x
SMBWebClient	A Web interface to smbclient
TEAForLinux	Modest, simple editor with useful features for HTML editing
WMTimer	Alarm/chronograph/countdown timer dockapp for Windowmaker

Development

Argtable	ANSI C library for parsing GNU style command line arguments
GuardedMemoryMove	A buffer overflow exploit catcher
Gwenhywfar	A multi-purpose system abstraction library
IdeaJad	A free decompiler plugin for IntelliJ IDEA
Jaffa	Enterprise-focused Java framework for rapid development
JMP	A runtime profiler of the JVM
MegaUpload	A progress bar for PHP file uploads
PHPDoctor	A Javadoc-style comment parser for PHP code
PMD	A Java source code analyzer
PyInstallShield	A cross-platform installation program
Python-LDAP	LDAP modules for Python development/deployment
Shipper	Automates shipping of open-source project releases
SwingWT	An implementation of the Swing/AWT APIs over SWT
TOADC++GUILibrary	A library for creating graphical user interfaces
VisualEffectsEngine	A particle system engine

Distros

BG-RescueLinux	A 2-Disk/EI Torito Linux 2.4 rescue system
INSERT	A credit card-sized Linux rescue system
PLDRescueCD	50MB live CD distribution, based on PLD Linux

Games

GCompris	A complete educational suite for children from 2 to 10
GNOMEWarPad	A 'VGA Planets' client for the GNOME platform

Graphics

AcidRip	A GTK Perl DVD encoding frontend
Convex3D	A converter, viewer, and extractor for 3D file formats
FBgrab	A framebuffer screenshot program
Gecko*MultimediaSystem	A replacement for your VCR, DVD, MP3, and CD player
GOPchop	A GOP-accurate MPEG2-PS editor
IvyTV	Driver for iTVC15/CX23415/CX23416 MPEG-2 compression cards
PhotoOrganizer	PostgreSQL Backed Web Based Photo Manager
PHPPersonalMovieDatabase	Organise and display your movies list

Potrace	Transforms bitmaps into vector graphics
Q-DVD-Author	GUI for dvdauthor
Shalvideo	A TV recording application
SortimagesGUI	Sorts photos by capture date from the EXIF header
TileDriller	A simple tile paint program

Help

LDP	A complete mirror of the Linux Documentation Project
-----	--

Internet

BTManager	A Web-based BitTorrent client
ELOGElectronicWebLogbook	An electronic logbook with a Web interface
Galeon	A GNOME Web browser
Gossip	A GNOME instant messaging client
Lftp	Sophisticated command line based file transfer program
MydoomDeleter	Delete the Mydoom.(B) worm from a POP3 mailbox
Nicotine	A SoulSeek client written in Python
ShrikeIRCservices	A variant on Undernet-style services for TSora networks
SpamProbe	A content-based spam detection program

Office

CyberTester	Web application to create and conduct online tests and exams
eGroupWare	GroupWare suite with a calendar, mail, todo notes and more
Pointless	A presentation tool for the Unix world

Server

Bind	Berkeley Internet Name Domain
Davenport	WebDAV-based access to a Windows/Samba network
Dnsmasq	A small caching DNS proxy and DHCP server
HTML-Embperl	Embed Perl into HTML and XML
Lighttpd	A light httpd
MonkeyHTTPDaemon	An small and powerful Web server for Linux
PHPlist	A mailing list system implemented in PHP
phpMyAdmin	Handles the basic administration of MySQL over the WWW
Postfix	The Postfix MTA
SpamHammerd	A dynamic blacklist daemon for Postfix
UserManagerForPureFTPD	A user-friendly Web application to manage PureFTPd users
WebKnowledgeBase	Allows users to search and organize uploaded text and files

Sound

JavalImplementationOfSpeex	A Java port of the Speex speech codec
Mp3cdBrowser	A tool to organize large music collections
Pscdcover	A simple Perl/PostScript CD cover generator
RealRekord	A tool to record RealAudio streams from RealPlayer
Sonic-Rainbow	A multimedia player
SoundFontCombi	A synthesizer and SoundFont tool
ThundaauralJukebox	A jukebox server for use with a touchscreen
Xfce4XMMSControllerPlugin	An XMMS Controller for Xfce4
XMMS	The X MultiMedia System

System

ATAIdle	A utility to set power features of ATA drives
FWSnort	A tool that translates snort rules into an iptables ruleset
JavaThrottle	A pipe bandwidth limiter
Netdude	A framework for manipulating tcpdump trace files
PHPServiceMonitor	A basic service monitor
SoftwareSuspend	Software Suspend support for Linux 2.4 and later
Sysfence	A system resources guard for Linux
Umv	A tool that renames users and changes configuration files
Webmin	A Web-based interface for Unix system administration
Xinetd	Powerful inetd replacement

Essential disc info

Read this important information before you use your *Linux Format* coverdisc – CD or DVD. We've collated some helpful info to help you get the most from these jewels of data!

FINDING THE ESSENTIALS

Missing something?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases, the latest libraries and

other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

Package formats

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

Documentation

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

WHAT ARE ALL THESE FILES?

If you are new to Linux, you may find the profusion of different files and extensions confusing. As we try to give as many packages as possible for compatibility, there will often be two or three files in a directory covering different types of Linux, different architectures and usually source and binary versions – so which do you install? They can be identified by their filenames, and usually just by the file extensions.

Someap-1.0.1.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.0.1.i386.deb – The same, but a debian package.

Someap-1.0.1.tar.gz – This is usually source code.

Someap-1.0.1.tgz – Same as the above, tgz is abbreviated form of tar.gz

Someap-1.0.1.tar.bz2 – Same, but uses bzip2 compression instead of zip

Someap-1.0.1.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.0.1.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.0.1.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7.x PPC Linux.

Someap-devel-1.0.1.i386.rpm – A development version.

INSTALLING FROM TARBALLS

A tar ball is a two stage archive. First the files are archived into a single file with tar and then compressed with Gzip or Bzip2. To unpack, cd to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
```

```
tar xvf --bzip2 /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in .tar.gz or .tgz, and the second for Bzipped files, ending in .tar.bz2 or .tbz2. Naturally, you change the paths to suit the location and name of the archive. and replace /mnt/cdrom with whatever is applicable to your system (eg /cdrom). This normally unpacks the archive into a directory of the same name, enter that directory with:

```
cd progname-2.1.0
```

To compile and install the software, type the following three commands:

```
./configure
```

```
make
```

```
su -c "make install"
```

The last line will prompt you for the root password, as this stage must be run as root. If you are already logged in as root, just type **make install**. This will give you a default installation. If you want to change any aspect of the install, type

./configure --help to see the options available. For example, you are usually able to change the default location with the PREFIX argument. When you have finished installing, you may remove the source files with:

```
cd ..
```

```
rm -fr progname-2.1.0
```

You should also log out as root, before you do anything you may later regret.

DEFECTIVE CDs

In the unlikely event of your disc being defective please email our support team (support@futurenet.co.uk) for further assistance. If you would prefer to talk to a member of our reader support team please call **01225 822 743**.

CREATING INSTALL CDS WITH CDRECORD

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord --scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v
```

```
/path/to/image.iso
```

You can simplify the command by saving some default settings in `/etc/default/cdrecord`. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the "Burn" tab and the "ISO 9660 Image" gadget near the top of the window. Put the path to the image file in the gadget and press "Combust!".

Now put on the kettle while the CD is created for you.

Other OS?

You don't have to use Linux to burn the ISO to a disc. All 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 has one, and a DVD drive, and use the CD burning software on their computer. It can be Windows, MacOS, AmigaOS whatever.

No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions it is also possible to mount the images and do a network install, or even a local install from another disk partition. The methods often vary between distributions, so check on the distro vendors website for more information. [LXF](#)

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 program 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).



March 2004

Product code:
LXFB0051 (cd)
LXFD0051 (dvd)

DVD HIGHLIGHTS:
Java Desktop System,
8 mini distros, Scrubber,
Python 2.3.3, Bacula

CDs HIGHLIGHTS:
MandrakeMove,
Seapine Surround SCM
2.1, Simple CDR-X,
MythTV, Six mini distros,
Konserve, Tuxpaint,
Kernels 2.4.24 & 2.6.1,
Hardware applications

MAGAZINE FEATURING:
Get your hardware working,
Tadpole laptop, Xandros OS2,
Firewalls Roundup, PixiePlus,
Armari 4-way Opteron, Xinit
SPS440, Wine Rack, Opteron
vs Itanium2, MandrakeMove



February 2004

Product code:
LXFB0050 (cd)
LXFD0050 (dvd)

DVD HIGHLIGHTS:
EVERY SINGLE
HOTPICK EVER!
All the programs that
have featured in LXF's
Hot Picks selection
from issues 1 to 50

CDs HIGHLIGHTS:
50 HOTTEST PICKS!
Including: Anjuta, ZNES,
Planner, Scribus, Mozilla,
Anjuta, Exim, ProFTPD...

MAGAZINE FEATURING:
Big Hitters biog collection,
Extreme Programming Pros &
Cons, FORTRAN, CrossOver
Office, Systemax 6507 &
Evesham Reliance servers,
Astronomy apps, Blender



January 2004

Product code:
LXFB0049 (cd)
LXFD0049 (dvd)

DVD HIGHLIGHTS:
Pathatizer, Kopete, LDP,
K3b, AbiWord, Scribus,
CyberCafeOrganizer,
Warewulf, Antteater

CDs HIGHLIGHTS:
Fedora, Vega Strike,
Omnis Studio 3.1,
Mandrake 9.2 updates,
OpenOffice.org Quick-
starter, Recall, The GIMP,
NetPanzer

MAGAZINE FEATURING:
Java Desktop Environment,
Linux DVD Players, CD copy
protection feature, Maya 5,
Smoothwall Corp. Guardian,
What On Earth is the AFFS?
RH Enterprise Linux 3



Christmas 2003

Product code:
LXFB0048 (cd)
LXFD0048 (dvd)

DVD HIGHLIGHTS:
KDE3.2 alpha2,
DamnSmallLinux,
Audacity, Samba 3.0

CDs HIGHLIGHTS:
SUSE Live Eval 9.0
(boots from CD/ISO),
Slackware 9.1, gentoo
File Manager, Karchiver,
BitTorrent, Mozilla,
Opera, Firebird, Mambo
Open Source Site Server

MAGAZINE FEATURING:
OpenOffice.org & StarOffice,
HP Itanium server, KDE3.2,
MIDI sequencers, BitTorrent,
Majesty review, Scribus DTP
tutorial, Slackware 9.1 review,
Gigabyte SR125ED review



December 2003

Product code:
LXFB0047 (cd)
LXFD0047 (dvd)

DVD HIGHLIGHTS:
OpenOffice.org (PPC),
MailScanner, XFCE,
LDP, PuppyLinux,
Slurm, fiphoto,
Battlemech, XMMS,
Fandango, CVSToys,
Scrubber, Sylpheed

CDs HIGHLIGHTS:
Mandrake 9.2,
OpenOffice.org,
GNOME 2.4

MAGAZINE FEATURING:
Mandrake 9.2, Linux Expo UK
report, Lazarus, CVS, Xinit
SharqStation, MTStudio, SUSE
9.0 review, Scanners tutorial,
Newsreaders roundup, IBM's
Linux strategy, Nessus



November 2003

Product code:
LXFB0046 (cd)
LXFD0046 (dvd)

DVD HIGHLIGHTS:
RUNT, RPM Analyser,
Linux Easy Access
Keyboard, VCOol

CDs HIGHLIGHTS:
Gentoo Linux, Seapine
Surround SCM FULL
VERSION, Nano, Mass
Rotate, BZFlag,
Jahshaka, Kroupware,
ZoneMinder, BatchGIMP,
ConfluxGroupware

MAGAZINE FEATURING:
The new 2.6 kernel, Java IDEs
Roundup, PHP 5 preview,
Linux keyboards, Aethera, The
Grid, OOo 1.1, Armari Pro3D,
Regular Expressions tutorial,
New-look SUSE, ZoneMinder

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 & Europe £6.99, Rest of World £7.99

DVD – UK £6.99, US & Europe £7.99, Rest of World £8.99.

LINUX FORMAT

www.linuxformat.co.uk

EDITORIAL

Editor Nick Veitch nick.veitch@futurenet.co.uk
Reviews Editor Paul Hudson paul.hudson@futurenet.co.uk
Art Editor Julian Jefferson julian.jefferson@futurenet.co.uk
Production Editor Matt Nailon matthew.nailon@futurenet.co.uk

Editorial Contributors Mike Saunders, Chris Brown, Hoyt Duff, Jono Bacon, David Cartwright, David Coulson, Andy Chappelle, Michael J Hammel, Hans Huberland, Andy Hudson, Biagio Lucini, Marco Fioretti, Alan Smithie, Dave Cross

COVER CD PRODUCTION

CD Editor Neil Bothwick

ART CONTRIBUTORS

Art assistance Martin Mulchinock
Photography Amanda Thomas, Phil Sowels, Corbis, Digital Vision, Photodisc
Illustration Paul Bateman, Chris Winn, Shane Collinge

ADVERTISING SALES

Senior Sales Executive
 Andrew Hudson andrew.hudson@futurenet.co.uk
 (01225) 442244 Ext 2441

MARKETING AND PROMOTIONS

Marketing Manager Fiona Tully
Subscriptions Executive Marie Lock

PRODUCTION

Production Co-ordinator Diane Ross
Group Production Manager Clare Tovey
Commercial Buyer Mike Thorne
Advertising Production Co-ordinator Jo Crosby

MANAGEMENT

Publisher Kelley Corten
Group Publisher Dave Taylor
Publishing Director John Weir

DISTRIBUTION AND CIRCULATION

Circulation Manager Regina Erak
regina.erak@futurenet.co.uk

Distributed by Seymour Distribution, 86 Newman Street, London W1T 3EX **Tel** +44 (0)207 3968000

Overseas Distribution by Future Publishing Ltd. **Tel** 01225 442244.

Overseas Licences

International Licensing Director Simon Wear
simon.wear@futurenet.co.uk **Tel** +44(0)1225 822798
Fax +44 (0)1225 788105 **Cell phone** +44(0)777 5641493

Contact Details

Linux Format, 30 Monmouth Street, Bath BA1 2BW
Tel +44 (0)1225 442244 **Email** linuxformat@futurenet.co.uk

Subscriptions and Mail Order

Phone +44 (0)1458 271178. See page 102
Email linuxformat.subs@futurenet.co.uk

Copyright No part of this publication may be reproduced without written permission from our publisher. We assume all letters sent – by email, fax or post – are for publication unless otherwise stated, and reserve the right to edit contributions. All contributions to Linux Format are submitted and accepted on the basis of non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing. Linux Format recognises all copyrights in this issue. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight. We cannot be held responsible for mistakes or misprints.

All CD-ROM demos and reader submissions are supplied to us on the assumption they can be incorporated into a future covermounted CD-ROM, unless expressly stated to the contrary.

Disclaimer All tips in this magazine are used at your own risk. We accept no liability for any loss of data or damage to your computer, peripherals or software through the use of any tips or advice.

Printed in the UK by Midway Clark (Holt) and Southern Print

© 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 for people who share a passion. We aim to satisfy that passion by creating titles offering value for money, reliable information, smart buying advice and which are a pleasure to read. Today we publish more than 90 magazines in the UK, US, France and Italy. Over 80 international editions of our magazines are also published in 28 other countries across the world.

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



NEXT MONTH

ISSUE 53 ON SALE THURSDAY 15 APRIL 2004

BUILD THE ULTIMATE LINUX DESKTOP

We handpick the very best components supported by Linux to build the ultimate in performance desktops. Be here next time for the ultimate construction guide!

GET TO GRIPS WITH DATABASES

Our server school concentrates on one of the most common server tasks – setting up and using a database

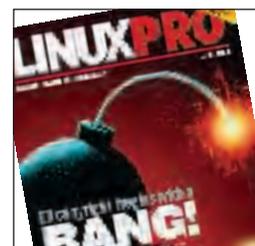
GIMP2

A radical overhaul of quite possibly the best-known piece of Linux software. Find out what's new and get hold of the latest version from our coverdiscs

PLUS

DON'T MISS NEXT MONTH'S LINUX PRO!

Dedicated supplement with pages of real-world Linux advice and case studies for IT professionals. Subjects covered include: Storage, Hosting, Security, Legal matters, Case studies and more!



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 102

ONLY
AVAILABLE WITH
LINUX FORMAT
MAGAZINE

LINUX PRO

FROM THE MAKERS OF LINUX FORMAT

APRIL 2004



EU copyright begins with a **BANG!**

Changes in copyright law heralded in the EU CD could leave individuals and companies vulnerable – equip yourself with the facts!

PLUS

Stuart Cohen talks OSDL

The Open Source Development Labs and the future of Linux

MontaVista: big on small things

Pioneer of embedded Linux, CEO Jim Ready speaks exclusively to us

The Oracle vision of a grid future

With 10g and a host of new applications, Oracle loves grids...

Using virtual machines to save cash

Maximising utilisation with VMware servers

PRACTICAL LINUX SOLUTIONS FOR I.T. PROFESSIONALS

Welcome

TWENTY PAGES OF REAL-WORLD LINUX FOR IT PROFESSIONALS

First off, I think we owe everyone an apology. For several years now we have been extolling the virtues of Linux and making statements to the effect that, aside from being more open, flexible and reliable than some OSes, it is also cheaper. A recent study by analysts IDC suggests that Linux is actually more expensive. The study (which I'm sure is quite coincidentally sponsored by Microsoft) looked at some specific scenarios and suggested that the Total Cost of Ownership, calculated over a carefully chosen period, saw Linux as costing a bit more. Well, what can I say but sorry? I'm sure our apologies will be echoed by companies such as IBM, Oracle, Fujitsu, HP and Novell in due course, who also seem to have been suffering from the same misapprehension, as indeed have the many tens of thousands of businesses of all sizes around the world who have already implemented Linux and found it to be more cost-effective. Sorry everyone. No, I don't care what your accountants say, they must be wrong – this study proves it. We look forward to the report that proves beyond question that running Microsoft Windows on as many servers as possible will solve world hunger and sort out the Middle-East peace process.

Back in the real world, there is plenty of thought provoking stuff in this issue, from Stuart Cohen's (CEO of the OSDL) perspective on future developments, to a look at how Oracle (who now use Linux as a reference platform) intends to offer a greater scale of products with an emphasis on grid computing. Plus there's a look at virtualising servers to maximise processing power, and an update on the rapidly growing, no-longer-niche market for embedded Linux technologies.

Nick Veitch Editor
nick.veitch@futurenet.co.uk



Personally, I look forward to the report which proves beyond question that running Microsoft Windows will solve world hunger and sort out the Middle-East peace process.

CONTACTS

EDITORIAL

Editor

Nick Veitch nick.veitch@futurenet.co.uk

Reviews editor

Paul Hudson paul.hudson@futurenet.co.uk

Art editor

Julian Jefferson julian.jefferson@futurenet.co.uk

Production editor

Matt Nailon matthew.nailon@futurenet.co.uk

Publisher

Kelley Corten

Group Publisher

Dave Taylor

Photography

Corbis, Digital Vion, Eyewire, Powerstock

ADVERTISING & MARKETING

Senior Sales Executive

Andrew Hudson

andrew.hudson@futurenet.co.uk

(01225) 442244 Ext 2441

Copyright No part of this publication may be reproduced without written permission from our publisher. All letters & contributions to *Linux Format* are submitted and accepted on the basis of non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing. *Linux Format* recognises all copyrights in this issue. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight. We cannot be held responsible for mistakes or misprints.

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.

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

CONTENTS

OSDL

Stuart Cohen talks about the future of Linux and the OSDL **p4**

COPYRIGHT

UK law surrounding copyright has changed significantly – be safe! **p8**

EMBEDDED FUTURES

MontaVista's Jim Ready discusses a rapidly growing field **p14**

CONSULTING THE ORACLE

Grid computing made reality **p16**

VMWARE

Virtualising for fun and profit **p18**



08



16

CASE STUDY: OSDL

DEVELOPING RELATIONS

NICK VEITCH investigates the Open Source Development Labs.



The Open Source Development Labs, www.osdl.org/, was set up in 2000 to promote the use of Linux in the Enterprise. With founding members including IBM, Intel and NEC, the idea was to create an organisation that could act as a focus for those wishing to use and develop Linux solutions. In its short history, it has attracted many more members and provided resources for many, many projects to be tested. *Linux Pro* recently met with Stuart Cohen, OSDL's CEO, for a chat about the future of Linux.

STUART COHEN: Our mission is really twofold. We talk about being the centre of gravity for the Linux community: that means being the one place where the users, developers and vendors can all come together.

It starts with being a participant and a member of the community – we've got to be sure that we're always doing our fair share of code submission – I would say we have five of the top twenty-five kernel developers in the world working for us right now. We're not delivering more code than any other organisation in the world, we don't even come close. But I would say we did our share, and we're doing a lot to help and support what's needed in the kernel areas. Obviously, with Linus and Andrew Morton, we are providing leadership and advice in terms of code submission, and we're very fortunate to have them working for us.

From a vendor standpoint, we have a very strong list of members. We had 24 companies at the beginning of 2003, and 35 at the beginning of 2004. Over the next year, we hope to get to between another 75 and 100 members; so we're significantly growing.

LXP: What's the motivation for vendors? TrollTech joined recently: what's in it for companies like that? They have obviously been part of the community for a long time anyway – why should they join now?

SC: The motivation a lot of them have is that first of all, the server market is growing and expanding rapidly. OSDL is becoming the one place where you can go – without getting a specific vendor bias – where you can find out what's going on in Linux. If you participate in our workgroups, whether that's datacentre or carrier-grade. Now our desktop announcement has created a lot of excitement: people want to talk about what the requirements are from a technical standpoint and a marketing standpoint. They want to network with their peers, hear what users have to say, share ideas. They want to make sure that when activities are going on, they get their ideas heard. I think that's why you're seeing a lot of people jump in.

We have these customer advisory councils that we have been running for a while now; we run one in the US that has met three times already, and there are about 20 Fortune100 companies involved in that. We are running our first one in Europe shortly and we're excited by that, and we have one planned for Japan. With three concurrent advisory councils running around the world, that really gives us a unique business perspective on what's going on and what's needed in the Open Source world, and also allows us to both validate and tie together a lot of the technical and marketing activities that are happening.

So that's really being the centre of gravity. The other piece of our mission is the acceleration of Linux. Things like the desktop initiative, like the legal defence fund. The papers that Eben Moglen (Columbia University law professor and EFF Counsel – <http://emoglen.law.columbia.edu/>) and Larry Rosen (General Counsel, Open Source Initiative – www.rosenlaw.com/) have put out for us have been very helpful. People take a lot of interest in that.

The value proposition of Linux is very strong – with customers, with vendors, with governments, with education and with small and large businesses. And they all have a different view of it: some look at it because they want



hardware flexibility, some want to drive costs down. With governments, there's a whole different aspect. Some like it just because it's Open Source – you can modify it and use it as you want. A lot of them look at it, especially in Asia, as an import/export issue. They want to create jobs, they want people to develop software to deploy in their own country or throughout the world. In countries like China, where the PC penetration is so low, the government is getting involved in developing on Linux, because they want to create a Linux box that is developed in China.

LXP: Do you see that there is any difference in attitudes around the world?

SC: Yes. Very different. In the US, it tends to be about flexibility. People are looking at Linux as giving them choice in terms of hardware or application stacks. Europe tends to be more focussed on the fact it is Open Source. In Japan and parts of Asia they want to create jobs. They want to develop software – they don't want to rely on the rest of the world for software and they don't want to pay the rest of the world for that software. If you look at usage models it's also different.

If you look at Linux today, a company of 100 employees, probably about 30 of them could be using Linux right now – the help centres, the call centres, the IT department, the receptionists, the engineering department, the shipping department. They typically need single applications and maybe a browser and email.

The next third, they're looking at Oracle, at SAP, at DB2. They're doing database calls and looking at documents and spreadsheets, maybe presentation tools. They're more like individual users. They are probably a year or two away from implementing Linux in a big way.

The third group are like local professionals. They're the spreadsheet jockeys – they're the ones worried about file management, document tracking; they're giving

presentations all the time, they're travelling all the time: they'll probably be the last to look at Linux. Replication, synchronisation, travel – all the issues that go with it – they probably have phones and PDAs, and that whole ecosystem is very strong. That's a generic statement though. Look at Europe and they want to push it a lot faster. In China, they want to move faster and move PC adoption along. In the US the market may be slower, even though it's bigger. There are a lot of units in that first third, so it may be a big number, but not ripple through business as fast as other places.

LXP: In the last two or three years, the desktop has become a real challenge in terms of... well, even a

BEING 'THE CENTRE OF GRAVITY' FOR THE LINUX COMMUNITY MEANS BEING THE PLACE WHERE USERS, DEVELOPERS AND VENDORS CAN COME TOGETHER

few years ago people dismissed the notion of Linux on the desktop completely, because they felt the application support wasn't there, and the hardware support wasn't there for a lot of stuff. Now, the underlying environment has changed substantially, but there still seems to be some kind of psychological block to implementing it. Do you think there are still valid technology issues to address?

SC: There are issues. The difference is that there are so many servers running Linux today that now there are so many clients connected to Linux, and people start thinking "Maybe I should run the clients on Linux." If you look at Linux as a desktop, it's behind where it needs to be in some ways; but even if it was an MS Office look- 

CASE STUDY: OSDL

alike, it probably wouldn't happen anyway because the cost of supporting both MS and Linux environments probably wouldn't be worth it. The first third are a great opportunity because they hang off the network – that fits with Linux very well.

The acceleration of Linux in the last year or so has been significant, probably due to the experience people had with Linux in the previous 12 months being so good – pilots were going better and easier, the reliability and performance were better than people thought.

LXP: There are currently a number of trials running with the UK government relating to Linux, and many more around Europe. Do you think that to some extent these are being used as a bargaining tool with suppliers rather than a serious evaluation?

SC: Well, it's the fastest-growing operating system in the server market (projected out for the next five years), so it isn't all hype. The other thing is that the numbers that tend to get published are analyst's figures, who tend to base them on boxes preloaded with Linux – they will tell you that for every box that ships there's probably another one downloaded from the Internet. That being said, are there still people out there using it to leverage Microsoft and Oracle and SAP and everyone else to get a better deal – well, that's business. Right now, Linux is a lower-cost solution, but I don't think that will carry the day for Linux. At some point, the platform flexibility, vendor flexibility, the upgradeability, the ability to roll out across enterprises are all things that will carry the day. The numbers speak for themselves, both now and in the future.

LXP: Do you think Linux is better placed to take advantage of the renewed interest in 64-bit computing. In the last few years, we've had the Itanium and the Opteron, and given the lack of much else that is able to run on them, is 64-bit going to be big for Linux?

SC: I think it will mirror the marketplace. I don't think Linux will be a single overwhelming leader in 64-bit. It doesn't hurt, and it does perform really well on 64-bit, but I don't think there'll be a situation where 90 per cent of the platform will be running on Linux.

LXP: Do you think that the way that Linux is so different in terms of the way it is used and distributed – that the community that surrounds it is so different from the model we have been used to – that it causes problems, particularly in perception?

SC: I think it's a lot more similar than it is different. At the end of the day, it's an OS that has applications running on top of it: if you're a business, you have wants and needs, and you have to figure out what they are and what solutions you want to have. Do you want to buy hardware or software off-the-shelf, or will you use custom stuff and develop yourself? People are actually asking, "How do I provide the best solution for my users?" There's no magic about Linux. It comes from the Open Source community and there are a lot of people working on it. It comes with a set of great value propositions, but I don't think it will change the way



people deploy solutions: you'll still have to make business decisions. The fact that the OS is free, you still have to come up with support, service and services. I think it's more the same than it is different.

LXP: Moving on to a slightly different topic. We don't really want to go into the whole SCO debacle, but do you think that incident has taught the OS community anything? Is it something that we can learn from?

SC: Well, firstly the fact that SCO has sued IBM about some code – that's purely about business and money, totally separate from Linux. The way I describe it is that Linux is available to everybody, developed by everybody, so everyone can use it; kind of like the Internet. So if you're against Linux, you're sort of against the Internet too, because the properties are very similar. With that said, if SCO has something against IBM, by all means go against them with what you've got, but don't hold Linux hostage. The Linux development community don't like the fact they [SCO] said back in March: "You stole some of our code and we're going to come after you," – to paraphrase the argument. In the meantime, SCO won't say exactly what code, and so nobody can evaluate the validity of SCO's claims. It's changed from being purely a code dispute to a disagreement over the nub of FOSS philosophy as well, so the target has moved a little bit. The fact is, people have been accused of stealing something, yet they haven't been told what it is they are accused of stealing.

For the community, Linux is a very scientific project that has been put together over many years. It's a course of peer-review, and maintainers and evaluators in a hierarchy. It's not as though someone posts some code and the next day it's in the kernel – it doesn't end up there by accident. But Linux is always enhancing the process he uses, and I'm sure that he will enhance it even more as Linux gets more mature.

LXP: If there is a vulnerability to Linux, it could be from an IPR (Intellectual Property Rights) angle...

SC: Well, most of the code comes from people that the community knows, through contributing coders and companies having a history that is also subject to peer review. If it was code that had come from somewhere else, the chances are that someone would recognise it: it's hard to imagine it would never get caught. It's not like thousands of people write most of the code. Most of it comes from a very small group, I don't think it's a big concern.

LXP: Do you think you'll ever have to spend the money you've collected for the legal defence fund?

SC: I hope not. I hope SCO doesn't sue anyone. But you never know what the SCO management might do – they've proved that over the last year.

LXP: There was a definite suggestion that when the whole thing started people were a bit more cautious about Linux. Over time, that particular avenue of conjecture seems to have faded away...

SC: One of the reasons that we have the customer advisory councils is to continue to evaluate that kind of question. I am sure that people did a lot of thinking, and we published a few

white papers as well. I believe people must have evaluated it and came to their conclusions, and I don't think it has any effect on the market right now. It may have slowed down infinitesimally in Spring 2003, but not any more.

I think the indemnification announcements and the investment in Linux by organisations like Novell, IBM and HP have made a difference.

LXP: You mentioned earlier in this interview that you have your fair share of people working on the Linux kernel, including Linus Torvalds, of course. Does that mean you can direct what is prioritised?

SC: No. Linus is his own guy and his own organisation. Through our working groups and councils, we interact with a lot of people and identify a lot of issues that we think are inhibitors to Enterprise deployment. We set up projects for people to work on those, and we may develop code and participate in those projects; but at the end of the day it's up to Linus and it's up to Andrew what they are going to accept. We have no control: we might provide some insight and perspectives from users and vendors and the industry, but we have no control over what gets accepted. That's the way they want it, and the way we want it. We're thrilled to have them as part of OSDL, but we're not so naïve to believe we can in any way control the outcome. They like it because we're non-profit, vendor-neutral and we're independent. As long as they are doing those same three things, it's a great place for them to work.

LXP: We should perhaps return to the subject of OSDL Labs. Have they been successful?

SC: They've been going three years now, and have been hugely successful. They were started by the founders because the perception was that the development community had access to single CPUs, and this would give them an opportunity to run on 4-way and 8-way servers, looking to see how applications would scale. We have around 150 servers in the US running at about 80 per cent

capacity. Most of the projects are about testing and scaling.

LXP: Do you think that over time they become less relevant as technology increases in capability?

SC: It's probably changed. The need for access is less than it was, but the problems are bigger. The projects are getting bigger, the demands are getting bigger. I'm not worried about the equipment sitting there idle.

LINUX IS AVAILABLE TO EVERYBODY... SO IF YOU'RE AGAINST LINUX, YOU'RE SORT OF AGAINST THE INTERNET TOO, BECAUSE THE PROPERTIES INVOLVED ARE VERY SIMILAR

STUART F COHEN, CEO, OSDL



LXP: If you could summarise your objectives for the forthcoming year, what goals would you like to achieve in the next twelve months or so?

SC: Talking about specifics, I think ISV porting will be a big issue – porting applications to Linux and having a single image to maintain across distributions.

Also, I'd like to see us do a better job of balancing the membership – as we start 2004, we're probably a little vendor-heavy and should have more user representation. We definitely want to recruit governments and education centres. I would suspect we'll have more members from China and Asia in the near future, and expanding in Europe is a goal too. ■■■

MEMBERS OF OPEN SOURCE DEVELOPMENT LABS

As can be seen from this comprehensive list from the www.osdl.org/ website, OSDL is supported by many major organisations from the computing world in general, which simultaneously shows that Linux is far from being a niche market, and also ensures a high level of vendor-neutrality.

Alcatel www.alcatel.com/

Cisco www.cisco.com/

Computer Associates <http://ca.com/>

Co-Create www.ccross.com.cn/

Dell Computer Corporation
www1.us.dell.com/content/topics/global.aspx/alliances/en/red_hat

Ericsson www.ericsson.com/

Force Computers, Inc.
www.forcecomputers.com/

Fujitsu Ltd. www.fujitsu.com/

Hitachi Ltd. <http://global.hitachi.com/>

HP www.hp.com/linux

IBM www.ibm.com/linux

Intel www.intel.com/

Linuxcare, Inc. www.linuxcare.com

Miracle Linux Corporation
www.miraclelinux.com/

Mitsubishi Electric
<http://global.mitsubishielectric.com/>

MontaVista Software www.mvista.com/

NEC Corporation
www.nec.co.jp/english/product/computer/express/asia-pacific/index.html

NEC Soft www.necsoft.co.jp/index_e.html

Network Appliance
www.networkappliance.com/

Nokia www.nokia.com/

Novell www.novell.com/

NTT Group

www.ntt.co.jp/about_e/index.html

NTT Data Intellilink www.intellilink.co.jp/

Red Hat www.redhat.com/

Sun Microsystems www.sun.com/

SUSE Linux AG www.suse.com/

Timesys www.timesys.com/

Toshiba Solutions www.toshiba.com/

Transmeta www.transmeta.com/

Trolltech www.trolltech.com/

Turbolinux www.turbolinux.com/

Ulticom www.ulticom.com/

Unilever www.unilever.com/

VA Software www.valinux.com/

Wind River www.windriver.com/



EU COPYRIGHT BEGINS WITH A BANG!

COPYRIGHT & COMPUTER USE

Intellectual Property barrister **DAVID HARRIS** explains how the European Union Copyright Directive is *adding* to the legal minefield surrounding computing, research and business.

This article is intended to discuss the way the recent changes in copyright law are going to affect companies and individuals in the UK. At the end of October 2003, the law of copyright in the United Kingdom changed significantly with the *Copyright and Related Rights Regulations 2003* which modifies the *Copyright Designs and Patents Act 1988 (CDPA)*. It implements a European Union Directive passed in 2001 that harmonises copyright laws amongst EU member states: it was motivated by the Internet and American lobbying of the European Union.

The EU Directive has led to a number of important changes to UK copyright law; the most important changes remove or qualify some of the exceptions to copyright, such as fair dealing and the privileges for libraries. It also introduces new costs on any copying that is carried out for commercial purposes.

COMMERCIAL COPYING

Previously, a commercial motive for copying was irrelevant to the exceptions against copying. Now, when making extracts of works, it is the ultimate profit motive that determines whether the exception survives and a licence is required.

The amended law in section 29 of the CDPA states;

DISCLAIMER

THIS ARTICLE IS BASED ON UK LAW EXCEPT WHERE otherwise indicated. Substantial differences exist between UK law and that elsewhere. The consequence of all law varies greatly with individual circumstances and thus nothing in this article is intended as or should be construed as advice or acted on without seeking your own lawyers' advice. The author regrets that he cannot give personal legal advice.

(1) Fair dealing with a literary, dramatic, musical or artistic work for the purposes of research for a non-commercial purpose does not infringe any copyright in the work provided that it is accompanied by a sufficient acknowledgement.

Where a company is undertaking research, it is very likely that profit is the ultimate purpose: whether directly as a result of the research or as an incidental consequence of it, for example, 'blue sky' research by a software house researching academic papers on protocols for a new network product. The fact that it is vague research that may or may not pay off is immaterial, since the *ultimate purpose* may be a profit. The effects on the costs of research both in the public and private sectors are likely to be significant, since research groups rely particularly on collations of copyright material – which they

can now no longer make use of without a fee. It will be necessary for companies with research needs to assess whether their copying has such a purpose, and if it does, they will need a licence from the Copyright Licensing Agency. This won't always be so; for instance, where minimal extraction is done as part of a process undertaken for employee training or welfare, it is questionable whether this is commercial copying or is just private research with company assistance. As always, some companies will just ignore the rules and copy anyway, so only the honest are likely to be penalised.

COPYING

NON-COMMERCIAL & FREE SOFTWARE

The rules will not usually affect non-commercial bodies like charities, since their purpose is not commercial. Raising money to finance the running costs of a charity is most likely to be regarded as non-commercial, although there is increasingly some divergence of views on this. However, if the work is being done by a charity on behalf of a commercial organisation, then although the charity may have no profit motive, it may just be looking for help in its running costs, the end-user *has* a profit motive. It is likely that a licence would be needed since the ultimate purpose of the research is the prevalent motive to consider.

Although I talk here of charities, that is just an example: any non-profit group will be subject to this analysis. Therefore, some classes of Free software may be affected by these changes, depending on the timing of the copying. Where software is subject to dual-licensing – such as *MySQL* or *QT* – if the copying is done when the dual licensing is being contemplated, then there is probably a need for a licence, because at least part of the purpose is commercial. If however, the copying is within the normal scope of fair dealing; and at the time, the copying is done there is no commercial purpose contemplated, then even if dual-licensing is later undertaken, there is no liability. The intent is likely to be subjectively measured, *ie* what the copier intended; and timing and intent is all-important. The likely effect on the Free software portion of the project should ideally be nil, since any copying costs should be absorbed by the non-Free licence cost; but if the copying cost were significant, it might deter the project.

CRITICISM, REVIEW & NEWS REPORTING

There is an overriding requirement that for fair dealing to be permissible, the work must have been made available to the public; so long as it wasn't the result of an unauthorised act. The new provisions in s.30(1A) say:

for the purposes of subsection (1) a work has been made available to the public if it has been made available by any means, including –

- (i) the issue of copies to the public***
- (ii) making the work available by means of an electronic retrieval system***
- (iii) the rental or lending of copies of the work to the public***
- (iv) the performance, exhibition, playing or showing of the work in public***

(v) the communication to the public of the work, but in determining generally for the purposes of that subsection, whether a work has been made available to the public no account shall be taken of any unauthorised act.

EFFECTS ON THE COSTS OF RESEARCH ARE SIGNIFICANT, SINCE RESEARCH GROUPS RELY PARTICULARLY ON COLLATIONS OF COPYRIGHT MATERIAL

For rights holders, there is often a desire to sweep failure or embarrassment under the carpet. A dire novel or film might leave its publishers wishing the author had chosen to be a milkman. When this happens, they will file the work in the deepest vault and hope it is never seen again. Sadly, for *Matrix Revolutions* this is didn't happen, but Stanley Kubrick did manage it with *A Clockwork Orange*.

WHISTLE-BLOWING FOR FUN AND PROFIT

Employee conscience counts for very little...

WE ARE TALKING, OF COURSE, ABOUT one of the legal justifications of whistle-blowing. It should be said however that the English courts are not great devotees of this hobby. In the *Lions* case, it was considered permissible to publish this in a newspaper, since the option was to report it to the police – the very body who had a vested interest in covering it up. Generally, it is wrong to assume that genuine belief and moral outrage is automatically enough to justify copyright infringement; it may be, but it cannot automatically be taken for granted. Where the motivation is partially one of disclosure through the media for fame or reward, then the risk is a substantial one for the whistle-blower. For governments, whistle-blowers are a traditional sore point: the *Spycatcher* case was a grievous humiliation to the UK

government. The more recent case of David Shayler – the MI5 man who revealed MI5's plot to murder Colonel Muammar Gaddafi of Libya – emphasised the need to contain the (already rather limited) public interest defence to copyright infringement.

The worrying aspect of this new change in section 30 of the Act is that it may be over-strict in favouring those who wish to conceal their wrongdoing. On one argument, it is not clear from the Act's language that it permits the defence of public interest; and if so, it does not sit well with the *1998 Human Rights Act* and the *European Convention on Human Rights*. A countervailing argument is this defence should remain unaffected by the change as a result of s171(3) of the CDPA, which provides saving provisions that preserve existing common law, *ie* non-statutory, rules.

mostly. A criticism by Channel 4 of Time Warner's decision to withdraw the film used extracts of it to criticise the decision. It was held to be justifiable copying as an act of criticism. Under the new scheme, embarrassing reviews may be avoided because the journalist obtaining the review copy will have done so unlawfully. Quite what the urgent and critical underlying need was for this change is unclear, but touchy rights owners seem to have secured a

AN AGE-OLD PRINCIPLE OF COPYRIGHT LAW IS THAT YOU CANNOT USE COPYRIGHT LAW TO PROTECT YOURSELF FROM CRITICISM

change in the law favouring themselves – so no apparent change there then.

However, an age-old principle of copyright law is that you cannot use copyright law to protect yourself from criticism. In another case, *Lion Laboratories v Evans* (about the blood alcohol meters used by the police to detect drunken drivers), a copyrighted document was discovered by the press that showed that the meters may not produce

ISPS

The Act doesn't talk about ISPs *per se*, but of **Service Providers** (who under EC law are people who pursue an economic activity), and this includes ISPs as well as others providing Internet-based services. That said, ISPs are likely to be the high-profile targets for copyright organisations like the BPI, and it is this that is for many the focus of the new laws.

It is possible to obtain an injunction against any service provider, but it is necessary to demonstrate that the service provider has actual knowledge of the specific matter complained of. For a court to find **actual knowledge** under sections 97 and 191JA of the Act, a court:

accurate readings. The press revealed this and were accused of copyright infringement by Lion Labs. It was said that the infringement of copyright was acceptable so that wide publicity might be given to the issue.

THE VISUALLY IMPAIRED

The impaired who have an unusable copyright work can take advantage of the so-called **one for one** exception. This provides that if you already have a legal version, and there is no commercial version for the impaired that is useful to you, then you can have such a version made. There is also an exception whereby groups for the impaired can make these copies but only if the copyright holders have not created a licensing scheme – in most cases they will, again presumably they charge for the privilege.

This is all rather ironic, since much of the substantial UK opposition to the changes in copyright law was based on the adverse effects to those with impairments. Criticism of the DRM anti-circumvention provisions asserted that depriving someone of the right to use their own copies of works was wrong, and the government's response is to say that it is OK to make a useable copy, unless the owner wishes to make you pay them twice. Hardly the most generous of concessions, and a clear indicator of just how much the Patent Office is in hock to the interests of large copyright owners.

shall take into account all matters which appear to it in the particular circumstances to be relevant.

A factor in this is the receipt of any notice to the service provider by a complainant, including the details of the sender of the complaint, and the facts in the complaint. Oddly, this could be seen to be an argument for having a non-working abuse address if complaints are anticipated. If you don't have a working abuse address, you can hardly be said to have received notice of an abuse through it, and so that mitigates against liability. I suspect no court is likely to want to read the Act in that way, though.

OBSERVING, STUDYING & TESTING OF COMPUTER PROGRAMS

There has been some – welcome if limited – clarification of rights in the area of interoperability and reverse engineering. It's possible, under limited circumstances, to lawfully write code that inter-operates with another program for various purposes and to reverse-engineer the target program to do so. The full details are not something we will go into here, suffice to say that these rights are limited in operation, and don't necessarily support Free software development. For this purpose, it is necessary to study the program. The previous law was drafted to allow the user to use a **device or means** to do this, however there was some suggestion that while **means** ought to mean software – such as a software Information and Content Exchange (ICE) – to do so, there was possibility that it wasn't. It might have been possible to have an end-user licence that permitted reverse engineering as required by law, but which prevented study of binaries in order to do so. That would have defeated the legislation's intent, so the law has been clarified in the new draft:

50BA Observing, studying & testing of computer programs
(1) It is not an infringement of copyright for a lawful user of a copy of a computer program to observe, study or test the functioning of the program in order to determine the ideas and principles which underlie any element of the program if he does so while performing any of the [acts] he is entitled to do.

296A(1) Avoidance of certain terms.

Where a person has the use of a computer program under an agreement, any term or condition in the agreement shall be void in so far as it purports to prohibit or restrict –

(c) the observing, studying or testing of the functioning of the program in accordance with section 50BA.

This means it is definitely lawful to run, say, Microsoft Windows through a debugger just to see how it works, just for curiosity sake, so long as you have a valid Windows licence. Mr Gates must be unhappy about that.

COPYRIGHT PROTECTION MECHANISMS

Here we get to the pith of the controversy with the new copyright law. At heart, it is the fruition of the copyright owners' agenda. Or more specifically, the large copyright owners agenda, the agenda of the "IAA. Through their lobbying of WIPO in the early 1990s to the present day, their self-serving legislation is killing fair use and the public domain. This is at a time when they have had to concede the battle against *DeCSS*; first in Norway against 'DVD Jon', and also in California against Andrew Bunner.

A key objective of rights holders is to lock users into a controlled environment where the use of their works is based on maximising revenue, consumer control and lock-in. I have discussed copyright and copyright protection mechanisms at some length in *Linux Format* issue 38, and I'd recommend getting a copy for details. This area of law has become known as paracopyright, since it doesn't affect copyright directly, but rather to access to copyright works. In summary, the new rules apply a DMCA-style prohibition on circumvention of copyright protection devices. Unlike the older legislation – which discussed *devices* – a phrase that may or may not have applied to software *per se*, this latest reform addresses computer software directly:

Section 296 Circumvention of technical devices applied to computer programs:

a person (A) knowing or having reason to believe that it will be used to make infringing copies –

(i) manufactures for sale or hire, imports, distributes, sells or lets for hire, offers or exposes for sale or hire, advertises for sale or hire or has in his possession for commercial purposes any means the sole intended purpose of which is to facilitate the unauthorised removal or circumvention of the technical device; or
(ii) publishes information intended to enable or assist persons to remove or circumvent the technical device.

PHYSICAL CIRCUMVENTION DEVICES

For those providing physical hacks, like chipping, against devices such as the Xbox, there are equivalent provisions to section 296 in sections 296ZA and 296ZD. 296ZA that assists the device makers where:

(1) This section applies where –

(a) effective technological measures have been applied to a copyright work other than a computer program; and
(b) a person (B) does anything which circumvents those measures knowing, or with reasonable grounds to know, that he is pursuing that objective.

and in section 296ZD;

(1)(b) a person (C) manufactures, imports, distributes, sells or lets for hire, offers or exposes for sale or hire, advertises for sale or hire, or has in his possession for commercial purposes any device, product or component, or provides services which – [are promoted advertised or marketed for the purposes of circumvention; or which are designed, produced, adapted or performed for circumvention of the protection measures].

One argument makers of these hacks provide is that they provide a backup capability or some alternative functionality. That argument is addressed by the Act for devices that:

(ii) have only a limited commercially significant purpose or use other than to circumvent

It seems rather hard on the users of a legitimate hack that it should be forbidden because it has a dual use, and the proper one isn't as commercially successful as the improper one. The Act seeks to balance these alternate uses and the practicalities of enforcement: those cracking down on improper uses are often frustrated in their enforcement actions by the minor proper use.

CRYPTOGRAPHIC & SECURITY RESEARCH

Non-computer works Exceptions apply under section 296ZA for those undertaking cryptographic research where technical measures are applied to a non-computer work:

(2) This section does not apply where a person, for the purposes of research into cryptography, does anything which circumvents effective technological measures unless in so doing, or in issuing information derived from that research, he affects prejudicially the rights of the copyright owner.

If you can, spot the big catch in that section: crypto research is fine, just so long as you don't publicly publish your results or do anything that would compromise the interests of rights holders.

Computer works For technical measures applied to a computer work, section 296 rather than section 296ZA applies, and in it there are no equivalent comforting words about the acceptability of cryptographic research:

(1) This section applies where –

(a) a technical device has been applied to a computer program; and

(b) a person (A) knowing or having reason to believe that it will be used to make infringing copies

(i) manufactures for sale or hire, imports, distributes, sells or lets for hire, offers or exposes for sale or hire, advertises for sale or hire or has in his possession for commercial purposes any means the sole intended purpose of which is to facilitate the unauthorised removal or circumvention of the technical device; or
(ii) publishes information intended to enable or assist persons to remove or circumvent the technical device.

The wording of this section is different from that applying to non-computer works, and it is problematic because it doesn't make it clear whether any exceptions applies to cryptographic research. One would have to stretch the language somewhat to get an exception, and there should indeed be such an exception since that is required by *Article 6* of the EUCD and *paragraph 48* of the preamble to the text of the EUCD.

Although cryptographic research is mentioned by the UK legislation, the EUCD treats it as merely one example of **proportionality**; that is, balancing the objectives of the legislation against the proper needs of all users. Uses other than cryptographic research – such as security research – come into that proper use, but they are not mentioned at all in the UK act; and to be protected, security researchers would also need to rely on the widest interpretation of the section. 



◀ The DMCA was intended to protect research, but that didn't stop an RIAA threat against Professor Felten in relation to audio watermarking (strictly, he was involved in rights management, but the argument applies to all researchers). The EUCD was supposed to learn the lesson of that, but the UK Patent Office apparently didn't. The language still exposes researchers to a potential threat, which copyright holders will be bound to try and exploit eventually.

If you are Ross Anderson, perhaps now is the time to retrain as a driving instructor.

CRIMINALITY & DEVICE CIRCUMVENTION

How would 'DVD Jon' fare in the UK? Consider sec. 296ZB:

(1) A person commits an offence if he -

**(a) manufactures for sale or hire, or
(b) imports otherwise than for his private and domestic use, or**

(c) in the course of a business -

(i) sells or lets for hire, or

(ii) offers or exposes for sale or hire, or

(iii) advertises for sale or hire, or

(iv) possesses, or

(v) distributes, or

(d) distributes otherwise than in the course of a business to such an extent as to affect prejudicially the copyright owner,

any device, product or component which is primarily designed, produced, or adapted for the purpose of enabling or facilitating the circumvention of effective technological measures.

The short answer is that - unlike in Norway - here in the UK, he would risk going to prison:

Section 296ZB Devices and services designed to circumvent technological measures

(4) A person guilty of an offence under subsection (1) or (2) is liable -

(a) on summary conviction, to imprisonment for a term

not exceeding three months, or to a fine not exceeding the statutory maximum, or both;

(b) on conviction on indictment to a fine or imprisonment for a term not exceeding two years, or both.

Defences to criminal circumvention Against a criminal charge, there are no defences available of fair use, and no right to get access to the works you have bought if you use a non-monopoly platform. There are however some statutory defences, under section 296ZB:

THE EUCD WAS SUPPOSED TO LEARN FROM THE MISTAKES OF AMERICA'S DMCA, BUT THERE ARE MANY AREAS WHERE THE WORDING IS UNCLEAR

(5) It is a defence to any prosecution for an offence under this section for the defendant to prove that he did not know, and had no reasonable ground for believing, that -

(a) the device, product or component; or

(b) the service,

enabled or facilitated the circumvention of effective technological measures.

This boils down to "Oops, I accidentally designed a circumvention device," which doesn't seem a great defence. The more convincing defence that one did it to exploit one's otherwise lawful right to use one's own works seems to have been given short shrift by the Patent Office; presumably on the grounds rights holders wouldn't like it since it would make prosecuting *real* infringers harder.

The police and security services have the right of circumvention for law enforcement or national security purposes. Furthermore, there are rights given to the Secretary of State to authorise circumvention for people who would not otherwise be able to use a work, but only if the rights holders refuse to implement such a scheme.

DIGITAL RESTRICTION MANAGEMENT

Digital rights management is another important aspect of paracopyright, and forms a central plank in the permissions culture. 'Permissions culture' is the evil twin of the copyright public domain: it maintains that by locking up works indefinitely and extracting payments for all uses, an enhanced creative environment will thrive. Starving artists and musicians will cease their boycott of their unrewarding vocation, quit their 'proper' jobs at Burger Lord and return to productive cultural activity. As the Librarians Association said, **"the logic of this is that we we'll end up paying for the use of individual sentences..."**

Section 296ZG provides:

(1) This section applies where a person (D), knowingly and without authority, removes or alters electronic rights management information which -

(a) is associated with a copy of a copyright work, or

(b) appears in connection with the communication to the public of a copyright work, and

where (D) knows, or has reason to believe, that by so

doing he is inducing, enabling, facilitating or concealing an infringement of copyright.

(2) This section also applies where a person (E), knowingly and without authority, distributes, imports for distribution or communicates to the public copies of a copyright work from which electronic rights management information -

(a) associated with the copies, or

(b) appearing in connection with the communication to the public of the work,

has been removed or altered without authority and where (E) knows, or has reason to believe, that by so doing he is inducing, enabling, facilitating or concealing an infringement of copyright.

Liability for breach of this section will be dependant on the nature of the techniques used to implement it. Techniques that embed Remote Method Invocation (RMI) into the copyright work, such as digital watermarking, should be fairly well protected against tampering. Techniques that depend to



some degree on external meta information – or ‘phoning home’ – may be rather more susceptible. For example if an RMI-enabled work phones home to rmi.linuxformat.co.uk, then for the purposes of subsection (1)(b), that is information that *appears in connection with*.

However, if you redirect that network request to *localhost*, you are neither removing nor altering the information. Of course, whether such an RMI scheme could work would depend on its architecture, but that’s irrelevant to the legal

analysis. Similarly, consider what happens if external meta information is held on a file system by the protected work rather than being entirely embedded in itself; for example, if when downloading an RMI-enabled work, a key is put in */etc/rmi* or a Windows registry. If that part of the system is mounted using a network filesystem like *nfs*, you may be in breach of the Act if you unmounted it while using the work since that information is, according to a strict interpretation of subsection (1)(b), *Associated* with the work.

MAKING WORKS AVAILABLE TO THE PUBLIC

One of the rights of a copyright holder is to have the exclusive right to control a work’s dissemination to the public. The trend in copyright laws has not historically placed much emphasis on this type of language, but has tended to focus on words like ‘copying’. However, in a networked World, with its increase in information and file sharing, this copying assumes more of the quality – if not the technical detail – of broadcasting. For example, a file on a P2P network may be simultaneously and widely obtained by users all over the World. Similarly, the emphasis on the economic consequences of copyright work is now on rights obtained from broadcasters, rather than those making discreet copies, such as would have happened in the past with the publishing sheet of music. Additionally, the law has learned the advantages of technology-neutral language, and broadcasting has a specific legal meaning that may not transfer to other newer technologies.

The new regulations amend section 72(1) of the CDPA to make the inclusion of music in broadcasts an infringement so that now a licence is required. This covers the position one typically finds at gyms and bars that show MTV. Previously, the Performing Right Society (PRS) and Phonographic Performance Limited (PPL) operated a protection racket where they would threaten legal action against the owners unless they coughed up a licensing fee. Some resisted, saying that the showing was free to customers; it did not form any element of the price they paid for their drinks or membership and so fell within the old exceptions. Predictably, the cry went up from the music industry “*won’t anyone think of the poor starving musicians?!?*” and they lobbied for changes in the law which an eternally compliant Patent Office delivered. Other than music, the position remains that free broadcasts to a non-paying public remain an exception, so if you are filling your company’s lobby with broadcasts of news or documentaries, it remains cost-free.

The new *making available to the public* provisions affect several previous existing exceptions such as the *time shifting* exception. This remains the same in substance, with one change: whereas before, the copy was never infringing, now if it is used in an infringing way, it is held that the original copying was also infringing. The likeliest effect of this is in relation to Tivo-like services that might allow programs be shared over P2P networks.

CRIMINAL OFFENCE OF MAKING WORKS AVAILABLE TO THE PUBLIC

(1) Section 107 shall be amended as follows –
(a) after subsection (2) there shall be inserted –

(2A) A person who infringes copyright in a work by communicating the work to the public –
(a) in the course of a business, or
(b) otherwise than in the course of a business to such an extent as to affect prejudicially the owner of the copyright, commits an offence if he knows or has reason to believe that, by doing so, he is infringing copyright in that work.

MISCELLANEOUS RULE CHANGES

Libraries are also impacted by the new regime. By tradition, libraries always took a very conservative view of their obligations and imposed strict limitations on copying. The new rules under sections 38, 39 and 43 of the CDPA requires an additional assessment by the librarian of any commercial purpose for the copying:

(a) that copies are supplied only to persons satisfying the librarian that they require them for the purposes of –
(i) research for a non-commercial purpose, or
(ii) private study,
and will not use them for any other purpose.

The schemes for education are subject to similar changes, and so long as the instruction remains non-commercial, no fee is payable. ■■■

ABOUT THE AUTHOR

DAVID HARRIS IS AN IP barrister, practising at www.ukitlaw.com. When not suing people, he writes ‘depressingly poor code’.



RESOURCES

Licensing agencies

Educational Recording Agency (ERA)

The ERA Licensing Scheme permits staff at education establishments to record for educational purposes any radio or television broadcast output. For further details see www.era.org.uk

Newspaper Licensing Agency (NLA)

This license permits the photocopying of articles from a range of national and regional newspapers. For further details see www.nla.co.uk

Design and Copyright Society (DACS)

In the UK, DACS represents 40,000 artists from all over the world. The licence permits the copying of artistic works and the making of slides of artistic works within certain limits. www.dacs.co.uk

The Higher Education Licensing Scheme
www.uwic.ac.uk/itsu/copyright/digitisation.htm

The Copyright Licensing Agency
www.cla.co.uk/

Campaign groups

www.ukcdr.org/
www.eurorights.org/eudmca/index.html

The Act & relevant amendments

www.legislation.hmso.gov.uk/si/si2003/20032498.htm#15

The old computer regulations:
www.legislation.hmso.gov.uk/si/si1992/UKSI_19923233_en_3.htm

Changes to the above:
www.legislation.hmso.gov.uk/si/si1992/UKSI_19923233_en_5.htm#exnote

EMBEDDED FUTURES

Embedded Linux specialists MontaVista talk to **NICK VEITCH** about the challenges ahead.

Linux does scale, from the very large devices that we are probably more familiar with, to the very small that we often take for granted. The Embedded space is one of the niches where Linux has had a dramatic and rapid impact, and one of the leading lights is MontaVista www.mvista.com. Nick Veitch spoke to CEO Jim Ready about the embedded Linux market.

LINUX PRO: 2003 has been a big year for MontaVista, if for nothing else, huge growth in revenues

JIM READY: Well, that helps....

LXP: Is that evidence of long-term plans paying off?

JR: Well, we would hope so: there is pretty good evidence showing that is the case. We put in place some fundamental notions on how Linux was going to play out in the embedded market; we knew that in some cases it would take some time.

Secondly, a secret – that we learnt the hard way in the old business – was availability. It's very difficult to predict, for instance, which processors are going to be successful. We get situations where certain CPUs aren't widely successful, but are favoured by a key customer. Someone decides to pick something obscure, but which fits their needs, and you're going to have to support it if you want to go deep with them.

We have put in place a whole infrastructure to master the intricacies of supporting Linux on lots of different architectures. There's a lot of engineering that goes into that, and a fair amount of time and money to put it into place, but once you have it, you have almost a factory where you can produce lots of Linuxes on time at controlled cost. Then we have to build a virtuous cycle by delivering on time, delivering quality and having good partners.

LXP: One of the more famous quotes about Linux is that it's "the only OS that will run on hardware that hasn't been invented yet." I guess that kind of freedom to deploy it on different hardware is a great advantage in the embedded industry?

JR: It's key, because a lot of folks are going broke trying to bet on a particular architecture. It's a very dynamic

environment, and there's a lot of innovation on the processor side; because of that fact, it's pretty much established that if you bring out a new CPU of any capability, one way or another there's going to be a Linux available for it. In many cases, that might be because MontaVista is doing it, or it may be the manufacturer itself. In any case, that quote is very good and literally true – in many cases, we have brought up Linux on simulators. We've done that a number of times.

LXP: How many developers do you have who work full time on Linux?

JR: Right now, we have around about 200 people working for MontaVista, including contractors. Of that population, over half of those are technical – either working directly on Linux or supporting it. Our biggest expense is engineering. That will change over time, but that's where we should be.

LXP: In a certain sense, using Linux in the embedded space – while noteworthy – is not the core focus of Linux for a lot of the bigger technology companies who are pouring money into development. Do you find it difficult to get attention for the features you would like to see in Linux?

JR: What you say at one level is true. The only quibble I have is that there are also a lot of large companies pouring money into Linux, though less visibly. Corporate investment in the embedded side is ramping up very quickly, especially in the last year or so, with Sony and Matsushita and so on.

Core changes to the kernel is a good point. There you'd think, "Gee, the embedded guys need something but the enterprise guys don't care," so you'd never get it in. But in fact we've succeeded in precisely that, by improving the real-time aspects of Linux. It turns out that what people figured out – from Linus on down – was that it was a general improvement in the overall behaviour of Linux, it was useful for everybody. Fundamentally it was a good thing.

We've done a lot of work on power management and other things. I think what I'm getting at is that we've made

impact and had influence where it was fundamentally necessary but have also carried our products forward without any sort of conflict.

LXP: I suppose that one of the advantages with the way the system works is that, even though Linux works at OSDL, there is no question of certain things being prioritised because some people want them.

JR: That's right. It's a fascinating system, there's no doubt about it. And it does work.

LXP: It has so far. MontaVista is also active in a lot of consortia and for a around the world for different aspects of Linux use – the carrier grade group, and some of the more specific consumer electronics groups. Have you found that has helped your business to be recognised as a leader in the Linux space?

JR: Yes, it's part of the overall involvement with the community. I think it ranges from being fundamentally important, for example with OSDL and carrier Grade – it's so important to the business that we have to be contributors. CELF (CE Linux Forum – www.celinuxforum.org) may end up being the same way. There are parallels at the highest level; we're members of Eclipse too. It may vary in terms of the ebb and flow of influence, but it's part of the landscape of Linux. You can look at it as a kind of tax you have to pay: it does take time and money, but it's a two-way street.

LXP: Do you find that other companies that are involved with those groups, for example Sony are involved with CELF. Do you think that companies like that really 'get' Linux these days?

JR: Clearly, for all of us there is a learning curve. One of the benefits of MontaVista being on CELF is that there probably isn't anyone on the planet more experienced at doing this sort of stuff for embedded, in many dimensions – not just technically, but how Open Source development works. Sony will tell you, as an example, that it is an extremely capable company, and is coming up the curve on how all this integrates with its business, so we're a real asset to have as a supplier and a member of CELF. I'm sure it took IBM a long time to work out how it would integrate Linux with its business plan – and you have to give IBM credit for having got going on Linux a long way back

LXP: Let's talk about your products: it was last year that you launched the Carrier Grade and Consumer Electronics versions. Have they been working out? Before that you had one product, now you have several – how difficult is it to keep on top of things?

JR: It's harder than with just one product for sure, but Linux is remarkably scaleable. The same code base that resides in the Motorola A760 phone is the same one that's in carrier grade on 4-way servers running Oracle. No OS in the world apart from Linux does that! It's a great advantage to us. The backdrop is that with a disruptive technology like Linux, predicting where it may get the greatest attraction is hard. So we purposely had the strategy of one good general-purpose Linux – which is our professional version – but we knew all along that there would be segmentation.

When we got customers who were already shipping product, but were after specifics, we started to see patterns of requirements so we could target those specialisations. The classic strategy, or the temptation, is to super-verticalise: to produce a complete mobile phone stack for instance. That's not our strategy, we have a horizontal strategy with a better ecosystem – for example, *Openwave* runs on MontaVista. That seems to us to be a more stable way to go. Long-term guys like Motorola want the phone to be a Motorola phone, not something we cooked up.

LINUX IS “THE ONLY OS THAT WILL RUN ON HARDWARE THAT HASN'T BEEN INVENTED YET” – A GREAT ADVANTAGE IN THE EMBEDDED INDUSTRY...

LXP: Do you think you'll need to specialise again and further split the code?

JR: A good point. There may be some technical reasons. For example, we are doing really well on cellphones. Also there are STBs and other devices. So there may be a case for that. There may be specialisation from a marketing standpoint, and that's probably far more likely.

LXP: As we mentioned, 2003 was a good year. What goals have you set for yourself for 2004.

JR: If you believe the surveys about the rate of Linux growth, in the range of 30–40 per cent, we have to do even better than that. We believe there is good reason that we will grow and have to grow; we feel pretty good about our ability to do that. This will also be the year we establish MontaVista China and MontaVista Taiwan, to complement Korea and Japan which we already have. That will happen for sure...

LXP: That's very interesting. Obviously China has great market potential, and in a way the lack of penetration of technology so far is an advantage.

JR: There is an interesting parallel to that. In eastern Europe, where the phone systems were a bazillion years old, when they finally reinvested they went to the most modern and bypassed the intervening 50 years. There are a couple of aspects to that. Companies like Motorola and Ericsson and other telcos are already there; as they are our customers we can sell to them locally. We probably have about 50 companies we deal with there. On top of that, you have the political aspects of the Chinese government being favourable towards Linux, which bodes well for us.

LXP: How do you think MontaVista is faring against the competition in what must be a tough market?

JR: We've gone through a couple of rounds... When we started, there were other startups, and yet we managed to survive. Now we are in round two or three, and companies like Wind River – a proprietary company that used to deny Linux existed – are now very interested. I'm sure they'll generate a lot of noise, so we're preparing for that. ■■■

GRID COMPUTING

ORACLE

CONSULTING
THE ORACLE

Ever felt like your mainframe was more of a Zimmer frame? **PAUL HUDSON** thinks it's time you woke up to an age-old paradigm that's being revisited: grid computing.



Think of one big computer capable of processing hundreds of terabytes of data every second. In order to handle that much data, the computer has millions of nodes, each of which is an individual PC connected over a global network. Sound like reality? Well, no – it sounds more like a description of Skynet, the computer AI from *The Terminator*, but the overriding opinion of people when asked what grid computing continues to be the idea of a science-fictionesque supercomputer that will soon be sending Arnold Schwarzenegger back through time to cause more mayhem and destruction.

Newsflash: grid computing doesn't necessarily need to involve millions of computers, and in fact it doesn't even need to involve a hundred computers. Using commodity hardware means you can slash your costs and improve your performance *without* changing the way you work. For as long as most people can remember, Oracle added **i** to the end of its products. With version 10 of its product range, this has changed: say hello to Database 10g, Application Server 10g, and Enterprise Manager 10g. Grid computing is something Oracle is taking *seriously*, and Linux is central to its plans.

ON THE STARTING GRID

For such an important and successful company as Oracle to back grid computing so whole-heartedly is enough for most people to sit up and take notice. However, this is actually nothing new for Oracle – it has been researching grid computing and similar paradigms for years, and this latest move is really just the culmination of their efforts to date.

Oracle's Real Application Clusters technology swept the world off its feet by allowing machines to work together without requiring complex administration or maintenance. The announcement in August 2003 that Oracle and HP were able to 1.18 million transactions a minute on Linux and Itanium proves that Oracle leads the way when it comes to distributed computing, but the price/performance on that machine was \$5.52/tpmC. For the uninitiated, that means the overall machine costs over \$6 million, and for that you



Grid computing is "invincible and inevitable," and should be available to users "on tap the same way as electricity and gas."

get a 16-node cluster of 4-way 1.5GHz Itanium 2 machines, making a total of 64 CPUs.

If you use Oracle already, you probably already have alarms ringing: Oracle licences are charged on a per-CPU basis, which means you get the most value for money by implementing the fastest chips you can. Of course, there's an important trade-off between speed and price, and the popular view today is that AMD's ever-popular Opteron chip fits the gap perfectly. In this respect, you could trade in the 16-node 4-way Itanium solution for a 32-node 2-way Opteron system and watch your performance shoot up.

"32 servers?!" I hear you cry: yes, you read it correctly. If you want a 64-CPU server cluster without paying \$100,000 per chip, you need to use leverage the power of grid computing. Fortunately for you, this is what Oracle has been beavering away on all this while – over half of its developers are working full-time at making 10g easy to use, which means database admins are no longer the people who push buttons and follow arcane rituals to keep the server going.

In fact, in order to manage the naturally heterogeneous nature of grid computing – where machines are running on a variety of hardware and software – Oracle 10g does most of the hard work for you: gone are the days when adding an extra disk drive was risky business.

DEATH OF FIVE NINES

There was a big push towards 'five nine' computing at the turn of the 21st century, which meant that both the hardware and the software that powered a system had to be reliable enough to run on 99.999 per cent uptime – an effective downtime of an hour a year. The most popular way to accomplish this was through massive amounts of redundancy: a database server would have a hot backup sitting next to it ready to failover if the primary server goes down, then that small cluster would be mirrored again on

two other machines to provide cold backup. Although this makes five nines a possibility, it also means that at maximum load, only 25 per cent of the hardware is being used.

If those same servers were to be put together as a grid, even keeping their current geographical distribution, the system would have all four servers running live at the same time, giving much higher performance. Furthermore, having your data distributed evenly across the machines means that the system became much more reliable – if one machine fails, the other three automatically pick up the load. This is where the majority of the work in Oracle 10g has gone, because five nine computing is a dated concept.

Conversely, adding machines to the grid is as easy as plugging them in and installing Oracle; with just a few questions answered, Oracle can easily hook the machine up to your grid, copy relevant data to it, and add start using it almost immediately. Although this might sound childishly easy, you might be surprised to hear that it doesn't stop at five servers, that it doesn't stop at fifty servers, or even at five hundred servers – by plugging in more and more computers, you can scale your system upwards almost infinitely.

This gives a level of reliability that was previously impossible, which means your IT administrators spend more time solving problems that *matter*.

Previously, adding more and more servers to a task was subject to the law of diminishing gains – adding 100 servers to a 1000-server cluster would rarely bring about a 10 per cent performance improvement. With 10g, this is all handled internally as part of the system's self-management; you can transparently add 10 per cent more servers to the system without needing to tweak any settings: Oracle does it for you.

Of course, there *is* still scope for hand-tuning, which means if you have advance knowledge that you will need more resources at a certain time, you can forewarn the system in order to maximise resource usage. At slow periods, when the system isn't being used, you can easily re-use dormant machines for other tasks, thus eliminating even more redundancy. This is another interesting part of 10g – it has abstracted the actual hardware so far from the system that you can transparently migrate entire systems to and from your grid in order to make the most of your resources, which means that you are always getting the most value out of your hardware without end users noticing any difference.

Amazon.com, which is probably the most experienced e-commerce site in existence, already uses Oracle 9i on Linux and stands to benefit a lot from a move to 10g.

Rick Brauen, the senior manager of database services at Amazon.com, said, "The ability to manipulate and re-allocate database workloads within a cluster based on the peaks and valleys of our normal business cycles will be invaluable to us in maximising our IT investments while meeting ever-increasing service level requirements." Although clearly Amazon.com is a huge company with extensive IT demands, the same principles apply even to companies of under 50 people.

STANDARD EDITION

Perhaps the biggest misconception people have about Oracle is that it is expensive. Granted, no enterprise-level database server is cheap once you scale it high enough, but at the very low-end people don't want to be bothered with

complex licensing and pushy sales people. Starting with 10g, Oracle are making their database product much more accessible for small companies through Oracle Standard Edition One, which is essentially a cut-down release of the database better tailored for less-demanding users.

That's not to say any of the core features are gone – the system still works on a variety of hardware and software platforms, still includes all the legendary reliability of the

“MICROSOFT’S LONG-STANDING AMBITION TO MAKE WINDOWS THE DOMINANT OPERATING SYSTEM FOR SERVERS COULD BE THWARTED BECAUSE OF LINUX...”

LARRY ELLISON, ORACLE CEO



Enterprise Edition of Oracle, and still has all the advanced automatic management introduced in the Enterprise Edition. The key change, however, is the low cost of entry: at under £3000 per CPU, you get a lot of value for money. Of course, there are *some* limitations with Standard Edition, of which the most obvious is that it supports a maximum of two CPUs; however, if you ever find yourself needing to scale up further, you can perform an in-place upgrade to Enterprise Edition and keep your system running.

The target market for Standard Edition is clearly companies who previously would not have considered Oracle and want to dip their toe in the water. In this respect, the most likely users are people who need more reliability than Open Source databases offer without wanting to sacrifice performance. At this level, grid computing isn't really an option, particularly due to the CPU limitations built into Standard Edition. However, being able to upgrade to Enterprise Edition, and then scale upwards using that, is a big bonus and means you can't really go wrong.

THE ALL-ENCOMPASSING GRID

Along with Oracle Database 10g, both Application Server and Enterprise Manager have been grid-enabled and have adopted the 10g name. We'll be looking at these two over the next two issues of *Linux Pro*, seeing how they fit together to provide an overall solution that can benefit your bottom-line. For a detailed overview of The Grid in general, see *Linux Format* issue 46's *What On Earth* feature.

Switching to the grid isn't about adopting the latest industry buzzword. Utilising grid computing means you can start small – even with just four computers – and scale up smoothly whilst preserving your investment at all times. Not only is grid computing about making full use of your idle IT infrastructure, and not only is it about maximising reliability by distributing the points of failure, but it just makes good business sense. ■■■

VMWARE

THE VIRTUAL BUSINESS:

HOW VMWARE WILL SAVE YOU MONEY

PAUL HUDSON asks: Why have one Red Hat box for databases, a Debian box for web serving, and a BSD box for firewalling, when you can run all three on just one system?

Although many people still get horror flashbacks to the dot.com boom (or more likely the dot.com bust) when they see words like 'virtual' and 'business' side-by-side. Relax: we're not talking about two-man companies who have offices in the New York, Paris, and Peckham! Instead, this is about abstracting your IT goals entirely from the hardware you have available so that the sum of your compute power is no longer equal to the number of cubic metres you take up in a data centre.

Controversial as this might sound, particularly if you make and sell servers, adding another computer to an overloaded cluster is rarely the best solution. We were banging the drum about data centre usage just two issues ago, saying that, "despite the best efforts of sysadmins, the average resource utilisation hovers between the 15 and 20 per cent mark"; either because companies don't realise they have fallen into this trap, or because they don't know how to correct it. This means there's around US\$30 billion of computing hardware lying dormant every year. If only a tiny fraction of that is sitting on your bottom-line, I think you'll agree it's time to take action.

VIRTUAL HISTORY

VMware is a company that has made its name entirely through virtualisation. Most people associate the brand with its Workstation product, which is a helpful tool to help you simultaneously run multiple operating systems under one host. *VMware Workstation*, however, is there really for development and testing of products, and doesn't scale well when you start demanding real production use out of it.

Their product line goes up, however, to two server products: *GSX Server* and *ESX Server*, for low-demand servers and data centre use respectively. Both are aimed at businesses who want to build their IT systems on virtual infrastructure as opposed to the number of boxes they have, and both are fast

VMware Workstation 4.0 virtual machines have a modern architecture with improved graphics and sound capabilities.



and stable enough to be almost indistinguishable from real machines. Of the two, *GSX Server* is proving to be more popular, simply because it fits the needs of most people whilst still coming in at an attractive price. Both *GSX* and *ESX* allow you to map your IT resources to your business needs flexibly, and, although we hate to use such a common cliché, it works "at the touch of a button".

Now that VMware has been acquired by EMC, the company is not likely to change much – it remains an independent software subsidiary of EMC, and all its existing partnerships stay in place unchanged. For customers, this means their product just keeps on working; customers of *GSX Server 2.5* get a free upgrade to 3.0.

WHAT IT DOES

As the most flexible server virtualisation product on the market, *GSX Server 3* brings with a host of new features. More importantly, though, all the award-winning features that made it so big in the first place are still there if not improved – *VMware* clearly has no intention of changing its successful formula.

Perhaps the most impressive new feature in v3 is the automatic Virtual Machine start-up and shutdown. This can be scripted using standard tools, or also built into testing solutions such as IBM's *Rational*, so that *GSX Server* can be told exactly when to start up a VM, when to shut it down, what tests to run when the server is up, etc. As a result, you can write a series of tests for all operating systems your product supports, and have *GSX Server* automatically cycle through each OS, performing tests overnight with no intervention required. From a quality-

assurance perspective, that's a real time-saver.

However, this also has the impressive use of being able to dynamically adjust to your business requirements. For example, if you have a system running four virtual machines, two for your web server and two for your database, how can you have it scale up when more people start visiting your website? As all four machines are virtual, you have two choices: either shut down one of the database VMs and start up another web server VM, or just crank up another web server VM and shift the database VMs to a slightly lower priority. As no physical machines need to be cabled up, switched on or configured, this entire operation can be done automatically and behind the scenes.

Now consider the same thing in operation on a group of eight servers – four of which might be running one web server VM, and the other four running one database VM. In this situation, the database servers can be switched over to fully fledged web servers in seconds rather than minutes (or even hours), which means that your IT deployment is agile enough to cope with peaks and troughs as they occur rather than after the event.

One feature that was highly requested by end users was full VM PXE support. The *Preboot Execution Environment* allows computers to boot up using a network image rather than a local install, and is a very easy way to get a stock system install up and running. This should save a substantial amount of work for sysadmins, particularly when virtual servers are spread across a LAN or even a WAN.

For admins who have to deal with many virtual machines at once, *GSX Server 3* now comes with *VirtualCenter*. This is a remote UI system that allows you to administer all your virtual machines and host servers from one location, connecting over a VNC-like system so that you can actively remote control each of your servers. Like PXE booting, this is another time-saving feature that will score easy points with both efficiency-led IT managers and admins who are sick of trekking around with a sheaf of CDs in hand.

THE FULL RANGE

Sitting alongside *GSX Server* are both *Workstation* and *ESX Server*, for the desktop and high-end server market. All three are brought together by a common virtual machine interface, which means that a virtual machine created in *Workstation* works perfectly in both *GSX* and *ESX Server* simply by copying it across. Not only does this mean that developers can produce a VM test environment on a laptop while they are away from work then copy it to the company *GSX Server* when they get back, but also that companies can migrate from *GSX Server* to *ESX Server* if they need to scale their virtualisation even higher.

One common use for both *GSX* and *ESX Server* is to run legacy server software that is no longer suited to modern hardware. Because the VM environment never changes, the hardware can be upgraded as often as you want and the hosted operating system will never know the difference – it is always insulated from the changes. What this also means is that you can migrate a VM from one server to another as you see fit, potentially from a single-CPU Pentium 4 to a two-way Xeon and even up to the likes of an eight-way Xeon as performance demands. What's more

– you can do it transparently, meaning that, apart from perhaps running a little faster, end users will have no discernible interruptions.

All three VMs are able to act as 'bulletproof' test harnesses. For example, if you want to know what difference it would make if you migrated your database server to kernel 2.6, you simply take a copy of the VM, migrate it to a test environment, upgrade the kernel and run your tests. If the test fails, you can roll back the changes and try again on the test machine, or you can just delete the VM and carry on with the live server. In this situation, it becomes clear why VMware has gone to so much bother to make sure its VMs are independent of each other – no virtual machine is aware of any other VM running on the system, or indeed that it is actually a VM at all.

FREEDOM OF CHOICE

Just because you choose to run Red Hat on your server, does that mean you can't run other Linux distros – or even Windows – on top? Not at all. In fact, it is this freedom of choice that forms the core of IT virtualisation – why should your choice of hardware or software limit what you can do? The answer is that it shouldn't, because if it does, you find your IT systems dictate your response to new developments. The old adage: "when all you have is a hammer, every problem starts to look like a nail" applies perfectly to IT. If you ever find yourself saying, "well, we'd like to do that, but our servers/our software/our configuration means we can't", then you clearly have a problem on your hands.

In this situation, the first step to take is to stop letting your choice of hardware dictate your choice of software – take a look at where your system redundancy lies, and see whether you can start to group services together so that rather than having a web server at 30 per cent load and a database server at 30 per cent load, you have one server running two VMs and operating at 60 per cent load.

Making the first step towards virtualisation is always the hardest, but VMware can make it easy by letting you do it step-by-step. More importantly, you can switch over to

NO VIRTUAL MACHINE IS AWARE OF ANY OTHER VM RUNNING ON THE SYSTEM, OR INDEED THAT IT IS ACTUALLY A VM AT ALL

virtualisation *now* by having just one VM on each server – effectively keeping your solution exactly the same. Once this is done, you can plan further moves bit by bit until you eventually hit your goal of minimum redundancy.

While actually hitting zero redundancy isn't actually the Nirvana that it sounds like, you should at least aim to get most of the way there. Not only does it mean you can save significant amounts of money, but it also means you can make your existing hardware work harder, and also respond faster to immediate system requirements. Don't think of virtualisation as a dot.com pipe-dream: it's about restructuring your IT systems to make maximum use of what you already have. ■■■