

LINUX

THE UK'S BEST-SELLING LINUX MAGAZINE!

FORMAT

**THE OPTERONS
HAVE LANDED!**

64-bit servers straight from
the US – RackSaver dual
AMD64s on test
p54

BECOME A NETWORK GURU!

Increase your network skills with this essential 7-page guide **p47**

BURN TIME

The best CD-burning
apps head-to-head
in our fantastic
roundup **p34**

FAKING IT

VMware's latest virtual
machine on test – run multiple
OSes simultaneously **p18**

WINDOWS BEATER!

Incredible value at under £650
Evesham Quest LM – the latest
power-packed desktop system
reviewed **p20**

FIND OUT HOW TO:

- Choose Linux hardware **p64**
- Dual Boot with Windows **p66**
- Organise your money **p72**
- Share files with Samba **p76**
- Texture objects in Blender **p80**
- Master GIMP selections **p84**
- Create PHP extensions **p88**

"I personally manage more valuable IP rights than SCO has
ever held, and I take it damn seriously" – *Linus Torvalds* **p9**

DVD issue also available Printed in the UK
LXF43 AUGUST 2003 £3.99

9 771470 423019 0.8

(£3.99 outside UK and Republic of Ireland)

www.linuxformat.co.uk

Not working?

For something which seems on the face of it to be quite straightforward, networking is probably only second to software installation as a cause of grief. As our *Answers* section will attest, we get tons on mail on the subject, so this issue we have put together something of a networking special.

For a start, there is a great feature on Ethernet, which is actually what most people are thinking about when they talk about networking. Understanding the basics of what networking entails, and the parts that cables, hubs, routers and switches play in the grand scheme of things can be really useful when trying to iron out problems you may encounter.

Of course, usually you want a network so you can achieve something. One possible use is filesharing, and one of the best ways to manage filesharing, especially on a mixed network, is with *Samba*. This month we have the first of a short series on setting up and using *Samba* servers on your network (check out p76).

Also, you probably haven't overlooked the CD or DVD on the front of the magazine, which has the finest networking tools and software that we could find. There's something for all occasions there, from packet sniffers to file sharing utilities, firewall builders to monitoring software. We hope you find it useful (and do let us know of any other collections you would like to see).

We were also pleased this month to get some real-life AMD Opteron-based servers to test. These are the first units to be seen in the UK and we were pretty impressed with the performance in most areas – read our special extended review on page 54.

For those of you more concerned with the Linux desktop, we have reappraised the CD-burning applications currently available. Some of the old favourites still command much of a following in this area, but there are a few new contenders that you should be aware of too.

We continually try to improve the magazine, so if you have any suggestions, let us know!




Nick Veitch EDITOR

The world of networking needn't be as complicated as it sometimes seems **p47**

Opterons on test, and no sign of Captain Scarlet **p34**

The best CD-burning apps put through their paces **p54**



LINUX
FORMAT

AIMS OF THE MAGAZINE

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

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

MEET LINUX FORMAT'S TEAM OF WRITERS...



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



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



Richard Drummond
He may have moved to a different continent, but he can't resist writing for *LXF*.



Frank Charlton
Experienced journalist and longtime *PC Plus* contributor, Frank wants everyone to at least give Linux a try.



Paul Hudson
Paul is discovering that just because he works on *LXF* doesn't mean he can automatically handle his beer.

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

Mike Saunders
Now considered an *LXF* veteran, he's the only writer to have caused the editor to laugh out loud.

Chris Brown
A freelance Linux writer and Unix instructor. He has a PhD in Particle Physics, but hopes it doesn't show.

Jono Bacon
Jono's a busy polymath – core KDE developer, web developer, writer, musician and sound engineer.

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

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

LXF43 August 2003

Contents

LINUX
 FORMAT

 Welcome to another bursting-at-the-seams issue of *Linux Format*, your guide to all things Linux!

COVER HIGHLIGHTS

Ethernet Networking **47**

The world of ethernet unravelled - become a network guru!

18 Faking it with VMware

Run multiple OSes simultaneously without dual boot

20 Evesham's Windows beater

The best desktop box pre-loaded with Linux so far

34 Burn Time

CD-burning applications tried and tested

54 The Opterons have landed!

New 64-bit servers run rings around Intel's Xeon



'Not only has Linux usurped a MS monopoly, (Munich city government) but the mainstream press realise it's important...'


07 Celebrating a red letter day for Linux

'Again, Evesham has blown the competition out of the water with a Linux system to rock the home desktop world!'

20 Mr Hudson is beside himself with glee

LINUXPRO

Training Camp Learn Linux in a leafy locale
Linux Professional Institute Why 'distro-agnostic' certification is the future
Flex and Bison Compiler writing
Storage SATA Bridging ebttables explained



20 Evesham Micros Quest LM – power and value.



6 News

Analysis of the latest Linux stories

12 Mailserver

Got an opinion? Share it with us!

18 VMware Workstation 4.0 review

Don't like dual booting? Try this!

20 Evesham Micros Quest LM review

Linux desktop Windows-beater

22 MoneyDance review

Cross-platform Java cash-juggler

24 Win4Lin 5.0 review

MS Windows virtualization

26 Netvault 7 Enterprise Ed. review

Take the sting out of unwanted fscks

28 TextMaker 2002 for Linux review

Not just another word processor...

30 Book reviews

Top treeware for Tuxters!

34 Roundup

Linux CD-burning applications

40 HotPicks

The very best Open Source software

58 What On Earth is... /proc?

The virtual directory explained

64 Beginners' tutorial

Is your hardware up to it?

66 Linux and MS Windows tutorial

Novices' guide to partitioning etc

72 GnuCash tutorial

Balance your chequebook

76 Samba tutorial

Bridge the gap 'twixt Lin and Win

80 Blender tutorial

Adding textures to your models

84 The GIMP tutorial

Selections explained

88 PHP tutorial

Extension-writing to handle tar files completed and new version news

92 Answers

Your problems solved by experts

99 Back issues and subscriptions

LXF delivered direct to your door!


110 User Groups

Get help from your local LUG

114 Next Month

Sneaky peek at September's stuff

24



58



80



COVERDISCS

A DVD or 2 CDs packed full of the latest Linux goodies **102**



CDS A AND B

SpamAssassin rid your system of unsolicited emails; **America's Army** as played with by US forces; **Complete Networking toolkit** 98 programs to enhance your network, including: **GuardDog**, **AGT**, **Jwall** iptables front-ends; **Samba** stable and development versions; **Ethereal** a packet-sniffer to let you know what's going on; **Nagios** network & system monitor formerly known as NetSaint



DVD

GnuCash manage your finances; **3dDesktop** give your Linux setup the 'Wow' factor; **Slackware Live CD** create a distro that boots directly from a CD; **LDP** tons of documentation

Please read the coverdisc instructions from page 102 to page 109 before installing from coverdiscs!



SAVE MONEY!

SUBSCRIBE TO LINUX FORMAT

See page 100 or phone 0870 4448645



Newsdesk

● Neverwinter Nights Linux client ● LRP RIP ● Graphics Muse GIMP updates ● Linus leaves 'day job'
● HP's new dedicated Linux unit ● Big Medium CM ● Portable Ogg hardware ● 2.4.21 kernel released

SCO LAWSUIT

SCO vs the world

The SCO saga rages on and, it seems, the former Linux vendor is intent on spreading its litigious tentacles well beyond its first target IBM, hinting that Microsoft, Apple and almost every other operating system vendor could have their day in court. The exception is Sun Microsystems whose Solaris OS is said to be bullet-proofed against legal action, a fact that the company is attempting to exploit in its latest 'we're not IBM' marketing campaign.

Defying the accepted logic of IT news, the story is getting more complicated as time passes and the accusations from SCO become ever more insidious.

■ On June 13th (a Friday, superstitious folks...) SCO announced it had revoked IBM's UNIX licence, which forms the core of AIX, and is at the heart of the dispute. SCO alleges that by contributing AIX code to Linux, IBM has weakened UNIX and misused the source code it had licensed. IBM countered by assuring users that its UNIX license was "irrevocable, perpetual and fully paid up" and that it would vigorously defend the lawsuit.

■ A SCO insider (who elected to remain anonymous) said that portions of the Linux kernel were copied wholesale into the System V code tree to aid the

linuxformat.co.uk

Visit the Linux Format website at www.linuxformat.co.uk for an update on the latest events in the SCO saga as well as opinion from other readers of the magazine in the site's forums.



Demonstrators young and old took part in a recent protest at SCO's HQ.

company's Linux Kernel Personality project which allows SCO UNIX to use Linux applications. Use of the kernel is covered by the GPL, so changes should have been made publicly available. In response, one kernel hacker emailed SCO to say that he would begin legal proceedings against the company – and urge the other 400 main kernel developers to do the same – unless they ceased distributing his code and stopped their efforts to undermine Linux. "I've granted everyone the right to sell, distribute and use my work under the condition that they obey the restriction of the GPL...I've never authorised any other use of my work."

■ Around 70 members of the Provo Linux User Group (PLUG) visited SCO's offices in Lindon, Utah to protest against the company's actions. However, SCO staff attempted to hi-jack the demo by mingling with genuine protesters and

displaying banners like 'I Love Software Piracy', 'Try Communism – Use Linux' and 'Stealing Software is not a Crime'.

■ One of the 'players' behind the scenes is said to be The Canopy Group (www.canopy.com) an investment group whose Chief Executive, Ralph J Yarro III is also the Chairman of SCO board. Yarro was involved in Caldera's 1996 legal win against Microsoft. Other associated companies with Canopy involvement include Trolltech and LinuxNetwork.


■ Back in February 2000, *Linux Today* reported that SCO had made the source code for old versions of their UNIX OS (Mini UNIX, UNIX V6, PWB UNIX, UNIX V7) available to 'permit hobbyists and enthusiasts to have access to the source code of these historic releases, for personal and non-commercial use' (<http://makeashorterlink.com/?L5BF65205>) A day later, Bill Baxter, pointed

Myths vs. Facts in the case of SCO vs. IBM

Because of the technical nature of the SCO vs. IBM case, there are a lot of misconceptions surrounding the case. Let's set some of the facts straight.

- **Myth: SCO owns UNIX**
SCO owns the rights to System V UNIX – an ancient distribution no longer in production – and its own UNIX distributions, OpenServer and UnixWare. Based on the settlement of the 1992 case (SCS vs. University of California), other distributors of UNIX – including IBM's AIX, SCO's UnixWare, Sun's Solaris, and Linux – have all been so heavily modified from the original UNIX code that they can be considered entirely distinct products. SCO can claim no rights over any products not its own. SCO doesn't even own the UNIX trademark, UNIX is a trademark of The Open Group.
- **Myth: SCO's OpenServer and UnixWare are major UNIX distributions**
Even by the most liberal estimates, SCO's products cannot have more than a 5% share of the UNIX market.
- **Myth: IBM copied features from SCO's UNIX into Linux**
SCO has not enumerated which features IBM allegedly copied into Linux; they have used only the blanket term "enterprise features," which commonly means support for multiple processors, large amounts of memory, and a "journaling" filesystem. System V – due to its age – does not contain any of these features, but Linux and OpenServer do. In all cases, Linux out-does OpenServer in speed, range of features, and reliability. Why steal parts for a bludge from an Expert?
- **Myth: Linux would not have enterprise features without IBM's help**
Linux always had support for multiple processors and large memory long before IBM contributed any code to it. In fact, Caldera Systems – SCO's predecessor – had funded some of this development. IBM's contributions have accelerated the development of existing features more than they have added new features.
- **Myth: The Open Source development community – particularly Linus Torvalds – has a biased attitude towards intellectual property rights.**
SCO has not enumerated which features IBM allegedly copied into Linux; they have used only the blanket term "enterprise features," which commonly means support for multiple processors, large amounts of memory, and a "journaling" filesystem. System V – due to its age – does not contain any of these features, but Linux and OpenServer do. In all cases, Linux out-does OpenServer in speed, range of features, and reliability. Why steal parts for a bludge from an Expert?

For much more detailed information about the SCO vs. IBM case and the tangled history of UNIX development, please read the Open Source Initiative Position Paper, available at <http://www.opensource.org/isc-vs-ibm.html>



SCO staff attempted to raise the debate by infiltrating the protest.

out that these releases were not Open Source. "SCO retain rights to the source code," he wrote. "Maybe they even hope that some of their code will wind up in Linux, so that they can then sue, and render the Linux license terms invalid. Or would they be that spiteful? My guess = yes." Give that man a crystal ball!

If everything seems to be moving too fast the most readable and extensive *precis* of the issue, the Open Source Initiative's (OSI) positional paper at www.opensource.org, has recently been updated. It's a long read, but well worth it if you are concerned by the validity of the claims and the affect it may have on future Linux development.



Blender finally gets fully integrated audio support.

3D MODELLING Blender gets loud

Blender, the pro-quality 3D modelling package which was recently Open Sourced, has gained a new tool in its quest to become a fully featured production environment. *Instinctive-Blender* is an audio sequencing and syncing applet which brings facilities such as lip-syncing and effects compositing within the main application suite. The developers say it will improve efficiency, especially in

projects with pre-existing audio. The impressive feature list includes unlimited audio tracks, non-destructive cut and paste, output to CD (44.1kHz) or DAT (48kHz) quality, full video/audio synchronisation and preview scrubbing. Audio playback is also available from all views including 3D views and can resample all loaded sound files (stereo or mono) to the project's defined sample frequency.

BEOWULF CLUSTERS Penguins band together



Penguin Computing, one of the most well-known clustering hardware vendors, has acquired Scyld Computing, which develops high-availability Beowulf cluster management software.

Penguin CEO Sam Ockman said the acquisition would make the outfit the definitive provider of Linux-based high-performance computing solutions.

"Don Becker pioneered Beowulf clustering under Linux in 1994 and has been one of the industry's staunchest advocates," he said. "Over the past four years, Scyld has invested aggressively in the development of its software and has earned a strong reputation in the

Beowulf clustering market, reducing the complexity of provisioning and managing a Beowulf cluster to that of administering a single Linux workstation."

Don Becker, who will retain his position within Scyld, said the joint venture would soon be in a position to offer alternatives to the traditional monolithic supercomputers. "Our software addresses Beowulf users' dual concerns of cluster manageability and scalability and our technology brings both better performance and greater ease of use, whether your cluster is made up of eight computers or over 1000," he said.

NEWSBYTES

■ Nullsoft could be without its guiding light after Justin Frankle, developer of *Winamp* and the Gnutella protocol (loved by file sharers the world over), said he would be leaving the company which is owned by AOL/Time Warner, following the release and withdrawal of *WASTE*. *WASTE* was a private file sharing and chat application which limited sharing of groups with up to 50 users. AOL pulled the Windows client software very quickly, though it is still available on a number of mirror sites and a Linux port has already been completed.

■ IBM has updated its eServer Directory Server for i and p-Series systems with improved security and performance and advanced replication support. Product manager Eric McNeil said the release of v5.1 of the software is an important advance for Linux bringing a new level of performance and scalability to our favourite operating system. "We think we have made some significant advances there that should satisfy most companies' requirements," he said.

■ America's Department of Defense has become the first US government agency to officially put Open Source software on the same footing as its proprietary cousin. Linux and Open Source have been used in the background for a number of years – one study suggested that the DOD have some 251 active OSS projects underway – but this is the first time its use has been officially sanctioned.



■ The EnGuard secure Linux distribution from Guardian Digital has received an update. Featuring an industrial strength firewall, gateway facilities and intrusion detection, the software is available in a number of editions including a free community version.

Dean Townsley has brought Linux to his Acer TabletPC, detailing the process of installation and configuration on his website at <http://prometheus.physics.ucsb.edu/~dean/TmC100/AcerTmC100.html>

■ Linux kernel 2.4.21 has been released after a six-month wait. It's available at <http://kernel.org>

David Cartwright

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



COMMENT Linux? Ja!

“It's happened at last. Not only has Linux usurped a Microsoft monopoly (in the case of the city government of Munich's 14,000-seat Windows-to-Linux migration) but the mainstream press have got hold of it and realise it's important. I didn't spot this one in the various Linux news outlets; instead, I read it last night as I sat with a glass of wine and the latest edition of *Private Eye*.

Though I'm chuffed with Munich's realisation of the potential cost benefits of running Linux, I'm far more happy that Linux now seems to have expanded beyond 'techie' publications like this one into the media that the bean-counters and the guys who sign off the budgets actually read. This is a massive leg-up for Linux as a business tool.

The traditional impetus to move to Linux comes from the techies who've see it to be stable and a potential business benefit, and urged the management to go for it. Too many times, management believe sound business ideas and technical staff to be mutually exclusive. Even though the likes of IBM are great advocates of Linux as a real, stable business platform, there's always a healthy dose among management of "They would say that – they get it free, so their margins just got healthier".

That's all changed now, though. We're finally at the stage where the MD reads this stuff on the train in the normal publications, and is in a position to ask his fellow board members: "The city of Munich just saved millions on this deal, and their requirements must be more complex than ours, so shouldn't we be investigating it?"

The guys who lead the business decisions can't fail to have an "I want one of those" moment. And if that's not a watershed for Linux, I don't know what is. ”

MMORPG LINUX CLIENT

Neverwinter Nights now shipping

After a very long wait and accusations of vapourware, Bioware have finally released the native Linux client of *Neverwinter Nights*, the online roleplaying game (RPG). The client currently supports English, German, French, Italian and Spanish and requires the purchase of a valid licence for the game, which means the Windows version (approx £30), plus a downloadable Linux patch, available from www.fileshack.com/file.x?fid=2508. On the plus side, Bioware say, you do get three

Windows-branded shiny *NWN* coasters with every boxed game.

Set in a dungeons and dragons style mediaeval fantasy land, *Neverwinter Nights* allows players to build and populate their own worlds and comes complete with the tools needed to build new maps. The game was released last year and there was talk of this being the first game that would ship with Windows, Linux and Mac clients in the same box. However, the Linux version has been the subject of so many delays that it was feared it would never see the light of day.



<http://nwn.bioware.com/downloads/linuxclient.html> – are you brave enough?

Linux Web Watch/



Read the reasons for quitting – it's a common Linux problem.



And hello Coyote Linux, ascendant champion of the small-install.



Latest version of Knoppix includes a complete MAME installation.



Gentoo's Live CDs are available for various architectures.

The end of the Linux Router Project

Little distros – tiny, tiny ones that fit onto a single floppy disk.

Dave Cinage, the author of the Linux Router Project (www.linuxrouter.org) has killed his pet project, citing, among other things, a lack of financial support for his efforts from those who had based products and, in some cases, their whole enterprise on his project.

There are still other tiny distros out there though with similar aims. One which is based on the LRP and is still maintained is Coyote Linux

(www.coyotelinux.com). This has both Windows and Linux installation procedures and will run on a very basic machine with no local storage.

One of the more famous single disk distros is tomsrtbt (www.toms.net/rb) which is the perfect rescue system, though there is rather more to it than simple first aid!

If 1.44MB seems a bit limiting and you have a bootable CDROM drive, you

have a lot of space for a moveable Linux distro. The current king of 'Live' distros is the Debian-based Knoppix (www.knopper.net) which includes a full graphical installer, office suite, GUI and development tools. The most recent Knoppix project is a bootable arcade system using Linux and MAME. Carry your /home directory on a 64MB USB storage device and you can take your own Linux system anywhere!

All these distributions mentioned so far are fine if you're using x86 hardware, but if you have a PPC machine, you can put it to use with Gentoo Live distro. Just like Knoppix, this includes a full graphical installation routine, desktop environment and wide selection of applications providing the full range of functionality. The are Gentoo versions for many other architectures too.

NEWSBYTES

■ After news that the Thai government was subsidising the production of low-cost Linux laptops, Microsoft has launched a specific package including a limited version of Windows XP Home and a copy Office XP for the bargain basement price of 1,490baht (£22). Lindows founder Michael Robertson also claims that MS has been offering deep discounts on XP to buyers of Lindows PCs at Tiger Direct. Tiger disputes his claims.

■ Romanian AntiVirus vendor GeCAD has announced it will no longer develop its RAV range of products for either Linux or Novell. The company, which has recently been bought by Microsoft, will now be concentrating solely on creating a new antivirus engine for the next version of Windows.

■ The third release candidate for *Mozilla 1.4* has been released. The project's organisers say this will be the last 'unified' product release. Subsequent versions will be built around the spun-off browser and mail client projects.




■ The *Mozilla* side project *Jazzilla*, which implements the browser in Java, has achieved its first milestone release 10 months after it was rescued from development hell. The 100% Java browser uses *Mozilla*'s Rhino engine and is now compatible with standard *Mozilla* chrome.

■ Feeling the heat from Safari, the Mac OS X web browser built on the KDE project's HTML rendering engine, Operasoft has said it may discontinue its Mac browser. The news follows Microsoft's decision to abandon development of *Internet Explorer* for Mac, citing the difficulty of competing with a project with such ready access to the underlying operating system.

■ The patent which prevented Open Source applications such as *The GIMP* from using files based on Unisys' GIF image format has now expired. The patent for the LZW form of data compression, on which GIF is built, expired on June 20 in the USA. Users in countries such as the UK, Canada, Italy and France will have to wait another year before they can 're-adopt' the format.

■ NVIDIA has released a kernel patch (for kernel 2.4.20) finally bringing GART support to the nForce2 chipset.

■ SuSE has launched a Desktop distro aimed at corporate customers (review next issue) looking at alternatives to MS Windows. It includes *StarOffice*, *OOo* and *CrossOver Office* (allowing install of *MS Office*), as well as a licence to connect to SuSE Open Exchange Server.



GFXTrans goes beyond the standard 'Rotate' tools in *The GIMP*.

ART APPLICATION

Plug in a GIMP upgrade

Graphics Muse has updated its range of plug-ins for *The GIMP 1.2*. Distributed on CD, the package features tools which extend the functionality of the application in seven areas including layer management, business card creation and over/under exposure problems. There are also a number of additions to the suite including *GFXShapes* for drawing common shapes such as stars,

polygons, triangles and grids.

The \$24 CD also features a complete mirror of *The GIMP* plug-in registry, including 110 compiled plug-ins, 30 Perl scripts and 90 Script-Fu scripts, as well as 150 brushes and 130 patterns.

The previous version of the suite is also available for both Linux and Windows for \$15.

www.graphics-muse.com

OSDL'S FIRST FELLOW

Linus: full-time kernel hacker

Linux Torvalds has decided to take a leave of absence from his 'day job' at Transmeta to concentrate on kernel development. "Transmeta has always been very good at letting me spend even an inordinate amount of time on Linux, but as a result I've been feeling a little guilty at just how little 'real work' I got done lately," he wrote.

After spending over a decade managing the development of the Linux kernel in his 'spare time', Torvalds will now be paid a salary by the Open Source Development Lab (OSDL), a project set up by a wide ranging consortium to bring more high-end features to Linux. In a joint statement from Transmeta and OSDL, Torvalds said it felt a little strange to "finally officially work on what I've been doing for the last 12 years." But, he said, "with the upcoming 2.6.x release it makes sense to be able to concentrate fully on Linux. OSDL is the perfect setting for vendor-neutral Linux development"

Transmeta said that it was happy to release Torvalds on a undefined 'leave of absence' and that the company would

welcome him back when his work on the 2.6 series kernel was completed.

Torvalds has also been pulled into the SCO affair by accusations that his method for validating contributions and identifying the intellectual property origins of code led to 'a very significant amount of Unix protected code' finding its way into Linux 2.4.x / 2.5.x releases. Torvalds told CNET that, as the personal holder of "more IP than the average bear" he cared deeply about the issue. "I personally manage more valuable IP rights than SCO has ever held, and I take it damn seriously," he said.

OSDL was founded three years ago by a consortium including Hewlett Packard, IBM, Intel and NEC. Torvalds will become the organisation's first Fellow, guiding kernel development and setting priorities for OSDL's data centres.

Transmeta, meanwhile, announced a joint operation with China's Chinese 2000 Holdings to develop and promote the company's Midori Linux distribution in the Asia-Pacific region, especially China, Hong Kong, Taiwan and Macau. Midori Linux targets mobile / embedded apps.

Hoyt Duff

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



COMMENT

Wasting time

“ A recent convert to Linux said he was returning to MS Windows. Why? The learning curve was too steep, he was frustrated in finding answers, and discouraged by the rancour he found in various “religious” wars within Linux. He felt Linux was a waste of time.

Showing me how he tried to solve a problem he had, he had posted his question to a LUG where he was a member; no one answered. It was easy to see that he had done no research himself and had asked a poorly phrased question.

First, I suggested he read the essay by Eric S. Raymond *How To Ask Questions The Smart Way*. While some have issues with the tone of this missive, it contains many valid and useful points and is an insight into the culture wherein lie the answers we seek.

Secondly, I told him www.google.com/linux is the only true friend he will ever have. He was a newbie Googler, so a quick trip to www.google.com/help/basics.html started him on the path to Power Googling. It's up to us to separate the wheat from the chaff with advanced search techniques.

Here rises the Linux Conundrum: What do I look for if I don't know what it's called? One of my favourite tools, *kdict*, is a search engine front-end for words and phrases. It has often provided a clue for further Googling.

Third, I let him know that while some people find great personal solace in devoting their energy to unresolvable (and ultimately unimportant) debates, his path to Linux enlightenment needn't take those detours.

Finally, I reminded him that the smartest man is not the one who knows all the answers, but knows where to find them. That alone saves lots of time. ”

UNDERLINING COMMITMENT

Hewlett-Packard creates new Linux division

HP's commitment to Linux and Open Source development is boosted with the creation of a separate business unit to promote the company's expanding range of Linux products and services. Headed by Martin Fink (previously VP of HP's Business Critical Systems dept) it will be part of the Enterprise Servers and Storage group.


The new unit is said to be a response to HP's projected growth in Intel Itanium 2-based hardware and, said a spokeswoman, reflect the prominence accorded the OS. "This is just a way to make [Linux] more prominent both within HP and in the industry."

Hewlett Packard has also launched a joint venture with *VMware* to bring robust virtualisation solutions to the company's range of ProLiant servers. The tie up would, said HP's Ron Eller, lower total cost of ownership especially in production, testing and development environments.

"Enterprise customers need solutions to help them increase server utilisation, reduce IT costs and benefit from improved system and application management," said Ron Eller. "With *VMware* complementing our *ProLiant Essentials* management tools, customers can gain more agility, accountability and a better return on IT."



VMware and HP bring virtualisation to ProLiant servers.



It's not an emulator as such... See page 18 for our *VMware* review.

Embedded Linux News



PhotobaseTV will bring image editing to your television.

● **ArcSoft** has created a photo editing package aimed at users of 'Linux-based TV-centric' devices. PhotobaseTV can be used to do basic editing tasks such as rotation, cropping, colour adjustment and red eye removal and also features a full range of media management tools.

● **Radionet** has released a full set of Open Source tools to create a Linux based 802.11 wireless access point. The Radionet Open Source Environment (ROSE) has already been used to build a city-wide WLAN in Finland and is suitable for use with both embedded devices or standard PCs and laptops. Full source and documentation is available from www.rosewlan.com.

● **Prismiq** claims the WiFi upgrade for its Linux-based Media Player device is the first to bring DVD quality streaming over a wireless network. The device can be used to play video and audio wirelessly from a PC and also works as a web browser and general Internet access device. It is based on a customised Debian distribution, and you can find more details about it on the Web at www.prismiq.com

WEB MANAGEMENT

CMS for the masses?

The content management system (CMS) sector is very crowded, with many high-quality solutions available under the GPL. Global Moxie has updated its mid-range suite of Perl scripts (\$129) designed to simplify the creation and management of websites by non-technical staff.

Big Medium 1.1 is a fully featured CMS featuring simple installation and unusually modest system needs; no MySQL (or any other database) is needed, just the ability to run CGI scripts on your Linux/UNIX webserver. Despite this, Global Moxie claims the system will handle thousands of

articles and millions of visitors without putting too much pressure on the server, due to the 'elegance' of its Perl scripts.

Josh Clark said the software delivers the power and results that would have previously cost thousands of dollars in development time, while freeing web designers from the need to deal with the minutia of web content, so they can concentrate on designing. "Too many creative pros have devolved into glorified typists, as webmasters and designers are relied upon to make even minor editorial changes to websites. Big Medium empowers editorial and business staff to make changes themselves" he said.





OPEN SOURCE NEWS

The engine of development?

A conference supported by the United Nations has concluded that one of the key engines to growth in developing economies is Open Source software, which could be used to build and improve infrastructure without the costs associated with proprietary software. World Bank spokesman Bruno Lanvin said developing nations needed cheap and efficient technologies to make the "giant leaps necessary to catch up with the rest of the world." He added that many were mandating the use of Linux, which looks like it may become the number one OS in countries such as India and China very soon.

Meanwhile in Brazil, the government has announced its intention to migrate 80% of computers in the public sector (including state owned commerce) from Windows to Linux over the next three years. The project will begin with a pilot in one ministry and then spread throughout the administration.

Sergio Amadeu de Silveira, from the National Institute of IT told the country's leading financial newspaper, *Valor*, that the intention was to save money. "We are not just going to do a hasty migration," he said. "Our main concern is the security and the trust of our citizens. The biggest resistance to any change comes from the existing cultural inertia"

De Silveira said the adoption of Linux would promote its use among businesses and consumers in the country and could promote the development of local software and "the democratisation of knowledge."

The situation is also changing in Europe, according to a new report by Bloor Research. Report author Robin Bloor says that if governments adopt Linux in a big way, it will "boost the momentum for Linux everywhere." The analysis details the various migration programs taking place across Europe (like in Germany) which suggests that the UK is playing catch-up to its continental partners.

LINUX MUSIC LIBRARIES

Neuros gets Linux and Ogg friendly

Users of the Neuros Media Jukebox can now use a Linux PC as the host for syncing their music files, thanks to the beta release of Positron, developed on conjunction with Xiph.org. Announcing the availability of the second Beta, Neuros also said that full support for the Ogg Vorbis format, an Open Source rival to MP3, was in the 'final stages of development' and a testing program would begin before the end of the summer. The new releases are the result of a deal signed between Neuros and Xiph.org back in February.

NuRosseta, as the project has been christened, will make Neuros the first hard drive-based media player with support for .ogg files and the new



Neuros users will soon get an Ogg flavoured firmware upgrade.

functionality will be available as a simple firmware update.

Neuros has also launched a development site for Open Source

hackers who want to extend the capabilities of the device to buy devices at a reduced cost. www.neurosaudio.com/dev/dev_index.asp

Jono Bacon

Founder of UK Linux, KDE developer and all-round nice guy, Jono Bacon has finished Uni and is looking for new projects to work on...



COMMENT

Standards?!

“ Hey, who needs web technology standards? Who needs rules to be set out so web developers can ensure their code runs on a variety of browsers and not just a single, overrated commercial browser? Who needs people like the W3C doing this? Well...me for one.

Recently I have been messing around with CSS a lot; now my site www.jonobacon.org is entirely controlled by CSS. As I have been working on the site, I have had to consider the amount of browsers to be compliant with: *IE5, IE6, Moz, Konqueror, Lynx, Googlebot* etc. What has been making me increasingly annoyed is the sheer lack of standardisation in some browsers...most notably *IE*.

What is the point in the W3C developing standards if they are simply not implemented in every browser properly? The Internet by its nature is an open medium, and though I cherish this freedom, the standards forming the heart of the web should be mandated by the W3C. How this imposition would be implemented I don't know, but something should be done before the W3C loses its rightly placed authority on the technology of the web and have it replaced by a corporation and its agenda.

It's fascinating how difficult it is to be completely browser compliant so your site looks good in each browser. The problem seems to not be massive rendering problems, but small glitches here and there. IMHO, if I design a site in *Mozilla Firebird* (which is 'purely standards driven') it should work in every browser. It is difficult enough developing a quality, user-centric, efficient, secure site without having to mess around with CSS all day to ensure it avoids Browser A's quirk here and Browser B's bug there. ”

Mailserv

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

★ Letter of the month

This month's winner receives a copy of *The Red Hat RPM Guide* (ISBN 0-7645-4965-0)

Why upgrade?

I greatly enjoyed Paul Hudson's column in your June 2003 issue. A very common mistake in new GNU/Linux users is to believe that they NEED all the latest software irrespective of level of stability. I find it both amusing and irritating that your cover disc staff seem to make this same mistake month after month.

One of the things that attracted me to your June issue was the promise of a new version of *The GIMP*. I've not downloaded this program for a long time because of its huge size, so a cover disc version would be most handy.

I say *would be*, because in deference to Paul Hudson's advice

you've given us the, basically useless, development version (with all the latest features that don't work yet) instead of a nice USEABLE stable version.

Also in your *Essentials* section, you don't include an "essential" version of GTK, like 1.2.x - a version that many current applications are using, instead, again you give us the latest 2.x.x - a version that as yet very few stable programs require.

It seems now, that your magazine is the ONLY Linux mag that includes a cover disc, it would be nice if you could put just a little more effort into preparing it.

Andy Preston, *via email*
Blimey. LXF41 - GIMP plus 400 plugins... For Moz, GIMP, GTK etc we

have both stable and development versions on the disc. Usually the new (= unstable) releases are highlighted on the disks because people do want to try them - that's how they get to be stable versions after all. Most of the popular tools get featured as both unstable and stable releases, though not always in the same month.

We take the point about GTK, but I think the reason we included 2.x was because people are less likely to have it than the 1.2.x versions. There's nothing to stop us putting both on though, which we will endeavour to do in the future.

Perhaps with the help of this month's star Prize, *The Red Hat RPM Guide* (ISBN 0-7645-4965-0, £29.99) you'll be able to help contribute some packages too!

Although Mandrake pioneered the 'Pentium only' distribution, it's amazing what sort of Pentium you can get away with. I similarly used to use Mandrake on a very old Compaq P60-based laptop, and it ran fine. It may fly in the face of some of the stuff now being talked about the desktop market, but it is the huge range of choice offered by Mandrake that enabled you to set up a working system. Long may it continue.

In Sync

I have two comments about the *Take Control Of Your Life!* article in your April 2003 issue. First, Ximian Evolution seems to have a well-known problem with Palm synchronisation. On my Red Hat system, running the most current Evolution under the GNOME desktop, my contacts, memos, and to do items mysteriously duplicated every time I tried to sync. While there is a program that de-duplicates entries on the Palm side, that didn't help the Evolution files. After manually deleting entries one too many times (and in the process, perhaps deleting the original



MAD about MDK

I'd like to comment on the recent release of Mandrake 9.1. I have used this on a server machine at home since release candidate 1, upgrading at each new release with no problems. I was looking for a suitable distro to use with an old Toshiba Libretto 50CT I recently bought. The biggest problems here are that it has no CDROM drive and the floppy drive is connected via a PCMCIA card that is not recognised by Linux and therefore any floppy installation that uses more than one disk will fail. I was considering some of the cut-down single floppy distributions until I found that Mandrake had a number of single floppy install images. I tried the network.img and was able to run the installed from a single floppy, it recognised the D-Link DFE-650

PCMCIA network card and this enabled me to use the NFS install option. After copying the contents of the CDs to my original Mandrake server I was able to perform a NFS install. It may have been a lot slower than using a CD, but for a machine with limited hardware it was perfect.

With the Libretto being a fairly low spec machine (Pentium 75, 32MB RAM, 850MB HDD) I wasn't able to use KDE or GNOME, so *IceWM* was used as the window manager. I have now added a wireless network card for truly portable computing. The next plan is to do it all again but with a much bigger hard disk so that I can install the bits I wasn't able to before.

All I can say is, well done Mandrakesoft for a superb distribution that has been shown to be installable on limited hardware

systems unlike some of the other standard distributions!

Andrew Lindsay, *via email*



The LXF coverdiscs have some Linux-compatible PDA applications to try.

READER TIPS

FILE MANAGER

You didn't cover my *OwenShow* directory/file browser in your FM roundup a few issues ago. It's available at <http://codacentral.borland.com/codacentral/ccweb.exe/home> amongst their Kylix offerings under my name ("James Owen") including Kylix/Delphi source; the web page http://home.att.net/~owen_labs/ran17.htm#owenshow has some thoughts on the project. *OwenShow's* great talent is scanning whole directory trees at once – which can be good or sometimes bad, so the program provides a few modified scanning strategies.

James G Owen, *via email*

ADSL

I read your article on ADSL in June's issue with interest. I am due to get an ADSL connection installed this month. My local exchange is being activated on the 18th of this month.


How do I know this? Try pointing your web browser at www.adslguide.org.uk/availability/btprereg.asp. This is an excellent site, not only for figuring out what the trigger level is for your local exchange, but also other subjects, different ADSL modems for one. I was surprised that you did not have this URL listed.

Nigel Henderson, *via email*

GRAPHICS

On page 88 of LXF39, the question on support of the Radeon 7500 was completely wrong. While native support for the card in XFree86 4.1.0 is indeed lacking, support in 4.2.0 for both 2D and 3D are quite good. And in version 4.3.0 support is even better for the 7500, and newly added support for the 8500 and 9000 is coming along nicely.

I know for a fact that distributions with 4.2.0 included, such as Mandrake 9.0 and SuSE 8.1, automatically recognise the 7500 and configure the 2D and 3D support. And in 4.1.0, frame buffer support will give you the graphical 2D support, and may be good



OwenShow in action – download it and scan whole directory trees.

enough for the questioners needs. I hear that the Nvidia drivers are getting better and installation is getting easier, but they are still a closed source driver set to those where that matters. ATI has been much more forthcoming in information on their cards and make it a much easier proposition for open source driver development. Also, having more than one choice in video card vendors is good for the Linux community.

The right answer to the question should have been: use frame buffer support for the card with 4.1.0 where 3D support is not needed. Upgrade to 4.2.0 or 4.3.0, which is readily available for download at many sites, for native 2D and 3D support. But don't throw out that Radeon card as it will give good performance. Craig Nori, *via email*

KDE SCREENSAVERS

I searched out the screensaver on the net, and there were instructions to install it which only work temporarily! As usual, Mandrake overwrites things behind your back!

If you download any KDE Screensaver, just park the *.kss file in /usr/bin. Don't bother with the *.desktop file. Instead, run *MenuDrake*, and open the folder 'hidden' then 'screensavers'. Do 'Add Item' and put the name of the

screen saver into the top slot and /usr/bin/(filename).kss into the Command Line box. Press OK and Save. Now you can go to KDE Control Center and select Screen Savers and it will be listed and selectable. Phew!

Dave Spagnol, *via email*

GTK SKINS

On page 76 of LXF40 you state "skins are unavailable in the windows port of gtk". The following link will prove there is limited support available to GTK-based windows apps.

http://project_ut436.tripod.com/kgimp.htm

Adam Bryant, 'lemonii' from forum

DRIVERS

I just finished reading the question on the ATI drivers for the ATI Radeon 7500. I thought I would add a bit of information to help Doug before he buys a new video card. Have a look at the website <http://gatos.sourceforge.net>. This site is the project that works on supporting the ATI All-in-Wonder cards. However, the Gatos team have also worked on the 3D acceleration aspects of the ATI cards as well as improving the basic X86 support.


Have a look before buying a new card. Most ATI cards should work with these drivers including the Radeon series.

Kevin Foss, *Ottawa, Canada*

UNIHOST

Specialists in hosting dynamic web sites & e-commerce stores with secure SSL.

- Linux Hosting Specialists
- CGI / SSI
- PHP 4
- Perl / DBI
- MySQL
- PostgreSQL
- JSP / Java Servlets
- FTP / SSH Access
- Full Logs & Web Stats
- Huge Network Backbone
- 100 MB Storage
- Full Shell Access
- Monitored 24 x 7 x 365
- UPS
- POP3 E-mail Boxes
- Unlimited Aliases
- SSL Option Available



UNLIMITED BANDWIDTH

£250 per year (ex VAT)

SPECIAL OFFER:

20% off 1st year's hosting if booked before July 31, 2003. (£200 ex VAT)
Quote ref: LF603T

sales@unihost.net

0870 744 0721
UNIHOST.NET

Mailserv

« appointments, or perhaps that was yet another Evolution problem) I moved to *JPilot*.

JPilot, too, has its problems. It sometimes mysteriously does not want to copy memos to the Palm. But it's very fast in the sync process, and quick and responsive in use. Incidentally, the most current version does indeed have week and month views – and a nice touch is that the appointments list at the bottom of the screen whenever you move your mouse over a date.

James LaRue, USA

JPilot is pretty much essential if you want to sync data to a Palm, and far from resting on their laurels, the developers are adding new features such as those you mentioned.

Don't diss Windows

Everyday at work and home I do use Linux – Red Hat is my favourite distro – but at the same time I fully respect products like MS Windows NT, 2000 and XP. I use Windows

2000 (at work) as a normal client user (email, text processor, music etc). My problem with it is the philosophy – the closed unavailable source. But the product is definitely super for everyday use. I can not tell exactly how the situation is regarding high level servers in companies, but I believe that Linux does an good job in this area.

It annoys me so much that every second post I read on the Net it is being said that "Windows is so unstable and Linux is not", "don't mention the word Windows here", "Windoze" etc. I feel that those comments are very often from Linux newbies who want to be a part of 'the new world', but and the same time still have a dual boot with Windows running on their system and very quickly return to it as soon their email client shows blurred fonts... With this said I agree that sometimes Windows does crash. But one thing I can guarantee: If you install a Linux on your system with X windows, KDE

or GNOME, XMMS, *OpenOffice* etc I will bet that you'll get crashes as well. I see them, but not often but they are there, especially when you start to play with the system, USB, kernel etc. I agree that if you run with a very simple installation (without X just with an ftp server running) the chances for a crash are very minimal. (for Windows 2000 with ftp running, the chances for a crash are minimal as well).

So what is my point? No matter what system you have and use with different programs, you will have crashes... And if you don't "play" with it you won't get any crashes (but who doesn't wanna play...?) So please don't say that Linux doesn't crash – it does...

With all this said I still say hooray for Linux – keep the good work going, and let us get a free alternative to Windows which you can modify to fit your needs.

Kjeld Vang Larsen, Denmark

Here at *LXF*, we try to minimise the derogatory references to other OSes. Though having said that, it is hard not to poke fun at certain companies' marketing strategies!

As regards crashing, it certainly is possible to crash a Linux system, but usually you can tell when it's going to happen. Personally I run a Linux desktop here every day of the week,

with the most unstable and experimental software installed. I can honestly say the OS has only locked up on me once this year (I remember the date). Individual apps do hang or need to be killed, but rarely, and usually you know which ones they are going to be. Your *Mozilla 1.3* mail client isn't going to fall over, but the latest *Kmail* running on a beta version of KDE might be more suspect. I haven't used later versions of Windows, or any other OS, regularly enough to comment on how stable they are these days.

SuSE shenanigans

I hope this letter will stop people going through the hell I have suffered over the last 2 days. It started on Saturday after a visit to PC World to purchase a copy of SuSE 8.2 having recently installed the downloaded version on my test box after reading your review.


I arrived home with a box copy of SuSE 8.2 (2xDVD 5xCD and 2 massive books for the beginner and pro). Anyway the machine I was installing on already had a copy of Red Hat 9 on it, so I was confident installation should be swift and easy as it had been on my test box.

So after about 15 minutes the machine rebooted and came back up asking me for a root password which I gave it, then we moved on to

“No matter what system you have and whatever OS you use, you'll experience crashes at some point – even with Linux.”

Helpdex

shane_collinge@yahoo.com





Don't take it out on your hardware if your install is less than smooth.

setting up the network. Again this looked very simple as SuSE picked up the correct hardware and loaded the module I expected it to. The next section asks if you wish to test your network connection to the Internet to which I said yes. No luck. I then changed the settings to use a manual address, still no luck. I thought I'd finish the installation and look at this in a minute.

The installation finished and rebooted still no sign of a network connection that worked. The connection lights on the switch and the nic were showing connected and ifconfig showed the card to be up.

As part of the install I had done, *etherreal* was installed so I thought I would see what that thought was going on. It showed that packets were being sent out but not received. I then double-checked the hardware database on the SuSE web site using my working test box. Yep there was the card a realtek RT8139.

A reboot into Red Hat showed the card to be in perfect working order. Being at a loss as to what would cause this, I purchased a new NIC just in case, this time a Netgear FA310tx which was also supported. Exactly the same problem.

By this time I was getting bloody annoyed, especially as I had bought a copy of SuSE instead of just sticking with my downloaded version. I rebooted one more time before just in case and thought to try the safe mode. Well, there was life after all. I looked at the boot parameters for safe mode using *YaST* to see what safe mode had omitted to allow my network connection to work. It

turned out to be ACPI, as I found out by process of elimination. Which is something to do with power management I believe.

Two days nearly on and off wasted by a module I was never going to use anyway! I can only put it down to the motherboard or chipset on it, which is a Gigabyte GA-7avx as it worked with out any problems at all on my test box.

I am now using SuSE quite happily and I am very, very impressed with everything about it despite my moaning. Cheers for a great mag and I hope this helps someone if this letter gets published.

Trevor Collier-Moore, *via email*

It would perhaps be too much to ask the distro to know you were installing on a desktop rather than a laptop (though ACPI has other uses than battery monitoring). Unfortunately, the ACPI support in the kernel is far from perfect, mainly because of manufacturers interpreting things in different ways. **LXF**

SUBMISSION ADVICE

WHAT WE WANT:

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

WHAT WE DON'T WANT:

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

WRITE TO US AT:

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

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

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

»Linux is only for computer science students.«

Another myth disproved!



SuSE LINUX is for everyone – anytime – anywhere!

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

SuSE LINUX 8.2 PERSONAL

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

SuSE LINUX 8.2 PROFESSIONAL

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

@ info@SuSE.co.uk

+44(0)208 846 3918

+44(0)208 748 0538

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



Reviews

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

LXF VERDICT EXPLAINED

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

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

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

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

Value for money: Does it have a competitive price?

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



10 The close to perfect product.



8-9 Good, but has a few niggles.



6-7 Does the job, but needs work.



5-4 Average.



1-3 An utter disaster. Back to the drawing board.

THE TOP STUFF AWARD

If we really, really like something – we really think that a particular piece of software, hardware or any other sort of ware is the best stuff around – then we'll give it our Top Stuff Award. Only the very best will be chosen. It's not guaranteed to all products that score highly.



WHAT'S NEW...

VMware Workstation 4.0

Running a virtual machine is a great alternative to dual booting – if your system is fast enough to cope! **p18**

Evesham Quest LM

Powerful MS Windows-beating desktop box pre-loaded with Mandrake 9.1 and bundled with top-quality hardware **p20 >>**

MoneyDance

Linux lacked a personal finance app for ages, now three have arrived more-or-less at once. How does this measure up to GnuCash and Kapital? **p22**

Win4Lin Workstation 5.0

Its ability to run many 'key' Windows apps will make Linux more attractive to IT buyers in the corporate sphere seeking to keep costs down **p24**



Netvault 7 Enterprise Edition

Get into good network backup habits with this capable cross-platform tool **p26**

TextMaker 2002

Easier to use than OpenOffice.org, and has in a beginner-friendly MS Office look-and-feel... **p28**

Book reviews

MYSQL Enterprise Solutions, Teach Yourself C++, The Linux Process Manager, Apache Tomcat Bible and more **p30**

Opteron servers

Thinking of buying an Intel Xeon? Don't part with a single penny until you've read this! **p54**

LINUX FORMAT BENCHMARKS EXPLAINED

To provide objective performance comparison between machines running Linux, we run a set of in-house benchmarks. These are: *bonnie* and *hdparm* to test hard drive performance ('hd' in the benchmarks), MySQL *Super-Smack* to test how well a machine handles database serving ('mysql'), *ApacheBench* to test how fast a machine can serve web pages ('apache'), a *gcc* compilation of Linux kernel 2.4.19 ("compile"), and *oggenc* to convert a test .wav file to a .ogg file. These numbers are then averaged to produce an overall score, which may be adjusted

slightly now and then, if a machine has a particular high or low point that should be taken in to consideration. Combined, these tests really push hard drives, network cards, and CPUs to their limits, and so give quite a representative figure – a multiple of the performance our yardstick machine.

The LXF yardstick box attempts to represent an 'average' reader's setup: Debian 3.0 on an 866MHz PIII with 256MB of RAM. So, a machine which scores 1.5 on our *Apache* test served 50% more web pages than our yardstick, whereas a box that scores 3.0 for overall ran, on average, three times faster than our yardstick box.

BENCHMARKS

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

The blue bar in the example above represents the performance figure for the hardware, and the red bar is the benchmark figure. When a piece of kit performs lower than the benchmark, as in 'hd' and 'compile' above, the blue value will appear less than the red value. **LXF**

PC VIRTUALIZATION

VMware Workstation 4.0

Why dual- or triple-boot when you can go virtual? Richard Drummond examines the perfect way to cram all those operating systems into one PC.


BUYER INFO

No direct competitors to this virtual machine program, but try *Win4Lin* or *Wine* for MS Windows functionality.

- **DEVELOPER** VMware
- **WEB** www.vmware.com/
- **PRICE** \$299 (download edition), \$329 (boxed edition)

V VMware Workstation needs little introduction. Quite simply, it's a package of wizardry that lets you install and simultaneously use multiple operating systems on a single PC. It creates virtual environments in which 'guest' operating systems can live and run at near native speed on top of a 'host' operating system. Only Linux and Windows are supported as hosts currently, but you can install as a guest any PC operating system that you can on a real PC.

The hardware resources furnished by a *Workstation* virtual machine (VM) are conjured up by a combination of emulation and bridging to physical resources on the host PC. For instance, hard disk space can either be emulated – as a virtual disk – using a flat file on the host filesystem as a disk image, or you can access real hard drive partitions from the host; similarly, a VM's network



The new interface in *Workstation 4.0* isn't just prettier, it's easier and quicker to use too.

interface can be bridged to the host's Ethernet adapter, or *Workstation* can

create a virtual network between a guest and the host. It is this flexibility that is one of *Workstation's* strengths.

The latest release, *Workstation 4.0*, builds on this core functionality, but offers a revamped architecture and, on Linux systems, a brand new user interface for managing VMs. Other new features include support for Netware as a guest OS, a new snapshot function, and a shared-folder mechanism used between the host and any guests.

New and improved

Previously *Workstation* on Linux used Motif for its user interface, but was beginning to look rather dated when compared to the Windows version and was trailing in functionality. With version 4.0, Linux catches up with a brand new user interface built on the more modern GTK+ toolkit.

It's not just a question of style, though. The new interface makes a much more fluid environment for

managing virtual machines. Each VM that you open is now presented in its own tab-pane within the main *Workstation* window. If you have multiple VMs open, you can switch to another machine simply by clicking on its tab. When a VM is open but not powered up, a status screen will be displayed in its pane. This gives a summary of the machine and its state, and gives controls to power up that VM or open a configuration dialog to modify the VM's settings.

Along the top of the main window is a familiar menu bar with pull-down menus and, below that, a toolbar, which allows quick access to functions to power up, power down, reset, suspend or resume the currently selected VM. Here also are controls to take and revert to a snapshot and to control *Workstation's* display mode. Besides the default 'windowed' mode, a virtual machine can run in a full-screen mode where it takes over the whole

Upgrading

Backwards compatibility restrictions

VMware Workstation is easy to install and easy to upgrade. For instance, the installer script included with the tarball distribution will automatically recognize, back up and replace an existing install. The architectural changes in *Workstation 4.0* required various changes to be made in the way that a VM's settings are stored on disc which are incompatible with earlier versions. Handily for those who have used the previous iteration, *Workstation 4.0* includes a compatibility mode and can directly use VMs from version 3.0. *Workstation 4.0* has a built-in function to upgrade a virtual machine's architecture so that it can take advantage of the new features in 4.0, but an upgraded VM can't then be used on *Workstation 3.x*.

While testing, I encountered few problems with upgrading *Workstation* itself or with using version 3.0 VMs. The hair-pulling only starts when you upgrade a version 3.0 VM. The new VM architecture in *Workstation 4.0* means that such an upgrade is akin to replacing your motherboard, graphics card and sound card in a physical PC – all at once. How well a guest OS responds to such a wholesale change depends on the OS itself. Windows 98 coped in the tortuous way that Windows copes with such things, requiring half a dozen reboots while it replaces all the drivers. The Linux guests I tried, Red Hat 8.0 and Mandrake 9.1, responded rather more smoothly, although the guest toolkit for Linux guests under 4.0 has its own problems (see the main article).

Virtually fast

Just how fast is 'near native' speed?

VMware claims that guest operating systems hosted on *Workstation* can run at near native speed. To try and pin down this rather vague statement, I did some improvised benchmarking and compared certain software and tasks running on a plain Red Hat 8.0 install on the host machine, as a guest on *Workstation 4.0* and on *Workstation 3.2.1*. Rather than bore you with a big

table of statistics, I'll just give you some rough numbers.

For a series of compute-bound jobs including compressing an audio stream with the Ogg Vorbis codec and generating GnuPG signatures of large files, I found that both *Workstation 3.2.1* and *4.0* averaged an impressive 90% of native speed. Clearly, then, the overhead of *Workstation's* virtual environment

does not greatly effect processor speed.

Next, to try and gauge disk access speeds in *Workstation*, I used the *bonnie++* benchmark tool and measured filesystem throughput to the same *ext3* partition when used natively and as a *Workstation* raw disk and to a *ext3* virtual disk hosted in that same partition (both *Workstation* disks were in the persistent mode). I found that

Workstation could read the virtual disk at roughly 65% of native speed and could write to it around 45%; with the raw partition mounted under *Workstation*, those figures became 75% and 65% respectively. Again, any difference between *Workstation 3.2.1* and *4.0* was negligible. These figures perhaps show that the usefulness of virtual disks outweighs any performance disadvantage.

display – just as if it were running natively. New in *4.0* is a 'Quick Switch' mode combines the best of both worlds. The main *Workstation* window – minus window frame and tool bars – is expanded to fill the size of the whole display; only the tab buttons remain visible at the top of the screen so that you can still switch between VMs. To access any other controls, you simply need to move the mouse pointer to the very top edge of the screen, and the menu bar and tool bar will drop down.

The changes in *Workstation 4.0* aren't simply skin-deep: it also features a new virtual machine architecture. This is still based on a fake Intel PIIX4 chipset, but now an APIC controller and ACPI support make it look like a more modern PC. The graphics architecture has been overhauled too now featuring a virtual 16MB AGP graphics card with a VESA 2.0-compatible BIOS. Now when a driver isn't available for the guest OS, it can fall back on using colourful VESA-2.0 compatible modes rather than being stuck with the standard 16-colour VGA. This is great news for Linux guests, as it means that

many more graphical installers will now work. With an appropriate driver installed on the guest, such as the new XFree86 driver, this new virtual chipset now offers 2D hardware acceleration, giving a huge boost in performance for desktop use over the dumb framebuffer in previous versions. Sound has also been improved, with the old ISA Sound Blaster being replaced by a virtual Ensoniq AudioPCI128 chipset.

Guest features

Workstation 4.0 has two new features to save time when using virtual machines. The first of these is the snapshot function. *Workstation* already had the ability to suspend a running VM, so that you could later resume or power it up to the same state – in a way similar to which you can suspend and resume on a laptop. This snapshot function is a lot more powerful. It lets you save a VM's state at any point in time (including the 'physical' state of the VM, the VM configuration and the state of any disks) and then continue working with the VM. This a powerful feature for those that use *Workstation* for testing purposes.

Second is the ability to share folders on the host filesystem with guests. You simply nominate the folders you want to share in the VM configuration, and, when a suitable driver has been installed in the guest, these folders will be available for access. Under Linux, a special *vmhgfs* kernel module is supplied, and the shared folders are available as named directories under the mount point */mnt/hgfs*. This again is a powerful concept, as it provides a lightweight way of sharing data between the host and guest. Alas the implementation is currently somewhat lacking. No support for Windows 9x/ME guests is yet provided, and access control is poor in Linux guests. The contents of your shared folders are currently always owned by root when first mounted – despite who owns the files on the host – and there seems to be no way of changing that.


Many of the new features – shared folders, 2D graphics acceleration – as well as other features such as the ability to share a clipboard between guest and host, require additional support in the guest OS for them to work – in the shape of a package tools provided by *VMware* which you must install in the guest. Unfortunately, the Linux guest toolkit is one of the weakest aspects of *Workstation 4.0*. For example, the installation script installs the XFree86 driver for you and duly modifies your X config. However, it takes no notice of your existing X configuration, and replaces your font path, your module list and everything else with its own choices. Whether this configuration will work for you will depend on your particular X set-up. Another problem is that the *vmhgfs* kernel driver lacks any kind of license tag, so recent 2.4 kernels will report a 'tainted license' warning when inserting this module. The installation script inserts it for you and everything works as expected, but when you reboot your Linux guest at some later point, the



Workstation 4.0 virtual machines have a modern architecture with improved graphics and sound capabilities.

guest toolkit init scripts will fail because they interpret this warning as an error.

Such problems are easily corrected, but they do detract from *Workstation 4.0's* otherwise excellent presentation. The new features are well worth the upgrade, but with niggles like those mentioned here, I would recommend that you test the free evaluation version before parting with cash. It would be good to see more cooperation between *VMware* and the major Linux vendors, so that an appropriate build of guest toolkit could be made available with the distros themselves or via a distro's update channels. If this was done, then many such configuration problems would disappear. [LXF](#)



The new shared folders between host and guest make communication easier and lighter, but lack access control functions in Linux guests.

VERDICT

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

Good new features and managing VMs is quicker and simpler. More work needs to be done to ensure Linux guests can be installed and configured smoothly.

LINUX FORMAT RATING
8/10

Reviews Evesham Micros Quest LM



LINUX HARDWARE

Evesham Micros Quest LM



BUYER INFO

Feature-packed Linux-based PC bundle with no UK competitors.

- **SUPPLIER** Evesham Micros
- **PRICE** £645 (inc VAT)
- **WEB** www.evesham.com

Evesham Micros gained a lot of friends in the Linux community when they released their first Linux machine – the E-scape Li. We reviewed it in *LXF39*, and the conclusion was that it was a great machine – so long as you didn't try to push it too hard. This Evesham machine is powered by an Athlon XP 2400+ CPU, 512MB of RAM, and a DVD/CD-RW combo drive built in. The internals are powered by Nvidia's nForce II chipset, which means it has built-in networking, GeForce 4 graphics, and 6-channel Dolby-compatible surround sound. At the front there's two USB slots, and there's another four USB slots and three Firewire slots on the back – enough connectivity for *anyone!*

Also included in the package is an Epson Stylus C62 printer, and Cambridge Soundworks SW320 speakers, which consists of a wooden subwoofer and two satellites. The monitor in the bundle is an LG 17" Studioworks unit, which worked fine at the 1024x768 resolution the machine came configured with, and we can't really fault it at this price. There's also a good-quality Logitech optical mouse, a

Following on from the success of their budget Lindows-based machine, the champion of the UK Linux desktop PC market has released a high-end machine designed to rival their Windows boxes. **Paul Hudson** goes for a spin...

good-sized "Getting Started" book, and also a set of Mandrake 9.1 CDs.

The old E-scape machine came bundled with Lindows, but Evesham have now (wisely) switched to Mandrake 9.1, and the difference is immense – booting up takes you straight into KDE 3.1, with the latest software to hand. The engineers at Evesham has spent a great deal of effort to get this machine working so well, and its work has paid off – the machine is solid as a rock. In my "I'm an idiot who doesn't care about his computer" tests, I pushed the thing to the limits – tried deleting files, pulling the power plug out randomly, etc, and the machine was entirely unaffected!


Perfect preconfigure

As the machine comes preconfigured, the sound card and speakers worked first time, as did the graphics card, and the DVD player. I played back a few CDs on the speakers, and am happy to report they are of top quality – as long as you don't mind the fact that they aren't surroundsound speakers, these will do you well. The unit came bundled with Loki's *Soldier Of Fortune* demo game, which I duly started up and blasted my way through – the graphics card works like a dream, with incredibly

high frame rates and smooth animations abounding. With such power backing this machine, I doubt any games will be

able to challenge it for some time yet! I popped the *Lord Of The Rings* DVD into the drive, and, thanks to Mandrake's

You'd be hard-pressed to get value of this sort out of a High Street retailer, even *without* MS Windows.



supermount working smoothly, I simply had to double-click the Ogre icon to get started watching the movie – no hassle here, and the playback was perfectly smooth. The same goes for MP3s and Ogg files – *Kaboodle* started up in the blink of an eye and I was able to play back music, write in *OpenOffice.org Writer*, and have *The GIMP* in the background without the least stress showing on the system – again, this is a by-product of the first-class hardware backing the unit, and means that even power users will be hard stressed to get this thing to falter. I even tried compiling a new kernel at the same time as watching a DVD, and am happy to report the system didn't drop a single frame of the movie.

Bundled software

With 80GB of hard drive space inside, there's lots of room to install all the software you can handle. Evesham pre-install a selection of favourites, and supply Mandrake CDs for everything else. The default installation included KDE, Mozilla, *OpenOffice.org*, *KOffice*, *The GIMP*, and a selection of media players – more than enough to get home users started. One minor problem with default software choices is that a lot of command line "old favourites" are missing, and need to be installed from CD. However, this isn't a problem, as Evesham supply the unit with RPM pre-configured to read from CDs, so it's simply a matter of typing **urpmi flex** (for example) to have it installed.

The configuration of the software was excellent – font smoothing worked perfectly in *OpenOffice.org* (such a rare thing!), and I didn't spot any places where I thought "Ah, they could have enabled *that* to make things easier". The system comes set up to use Mandrake's default Galaxy theme, which is as pretty and easy as ever. One tweak that should help make the switch to Linux less scary for newcomers is the highly organised K menu – it's been kept very short, with programs placed sensibly under categories.

Bundled hardware

Making the USB mouse and printer work was simply a matter of plugging them in – the system was already preconfigured to handle them both, and I could start printing from applications immediately. The Epson printer came bundled with two ink cartridges (one colour, one black) and the print quality was excellent, and surprisingly quick.

As mentioned, the DVD player worked smoothly with no effort, and CD writing was a breeze. Everything comes under a 2-year on-site maintenance guarantee, which is admirable in these times. All my machines at home are Evesham, and the one time I had a problem (thanks to a freak lightning storm, no less), Evesham were round and solving the problem within 72 hours of my call – their warranty really is worth having, and is stunning considering the price of this unit.

Performance


This machine was designed to cope with the needs of home users, which means it has quite generic performance results – it doesn't have multiple RAID SCSI controllers for its hard drives like the servers we reviewed, neither does it have four Xeons in there, but neither is it supposed to. The oggenc score, which relies solely on CPU performance, returns an excellent 3.87 – one of the highest we've seen so far. The other results are also sterling for a machine of this class, so the Quest LM certainly is no laggard when it comes to shifting data. The GeForce 4 graphics card means that *Soldier Of Fortune* ran like lightning, and having so much RAM means you can configure MySQL to use large query and key caches. You can run multiple heavyweight applications side-by-side on this machine without much decrease in performance – we ran several *OpenOffice.org* programs, the *Lord Of The Rings DVD*, and *The GIMP*, then switched to the console and started compiling the kernel – the only time the music of the film stuttered was switching to and from X (a very intensive task), and even then it was only for about a quarter of a second. This machine is undoubtedly much more powerful than its E-scape predecessor, and suffers from none of the same problems of its cheaper sibling.

Conclusion

It's very fast, feature-complete for home and small business users, and comes at a brilliant price. We were hard pressed to find any faults despite giving the machine a much harder time than we normally would, which means this system should be stable even in the hands of complete PC newbies. It's an attractive unit, with a wonderfully tidy interior making for



Not a home cinema system, but still far better than most bundled speakers.



Smooth fonts and excellent performance all round means this machine will be a winner for users of all levels of Linux ability.

easy access to components, and, when you factor in the excellent warranty then it's easy to see Evesham have again pushed the envelope when it comes to Linux machines.

If you want a fast machine that's got RAM and hard drive space to spare, look no further. If you want a machine that's available at a great price and with no compromise on the components, you've just found it. If you're looking to try out Linux at home and don't want to worry about whether it will support your machine, your worrying is over – yet again, Evesham have blown the competition out of the water, and given us a system that will rock the home desktop world.

The only real drawback to this system is that must come with Mandrake, however that *will not* (and I really mean that!) be a problem for home users – in fact, that's probably an advantage, because Mandrake continues to be the easiest Linux distribution available. If they offered a selection of distributions (something

like Debian, for more experienced users, maybe), the machine will likely reach a somewhat larger market. However, at *this price*, I'd be happy to install Debian myself! **LXF**

BENCHMARKS

hd	1.32
apache	1.83
mysql	2.33
compile	1.46
oggenc	3.87
Overall	2.16

VERDICT

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

Evesham: 2, Competitors: 0. The above numbers speak for themselves. Leads UK Linux PC field by a big margin.

LINUX FORMAT RATING
10/10

PERSONAL FINANCE

MoneyDance

GnuCash and Kapital have shown that personal finance management and Linux are no longer mutually exclusive. Andy Chappelle sees how a cross-platform, Java-based competitor measures up.

BUYER INFO

Personal finance package in the *MS Money* mould. Competitors include *Kapital* and *GnuCash*.

- PUBLISHER Reilly Technology
- WEB www.moneydance.com
- PRICE US\$29.99

Once you introduced to managing your money by computer, it is easy to get obsessed with counting every single penny: putting in a fake balancing transaction at the end of the month just feels like failure. *MoneyDance* enables you to view exactly who gets what from your salary, though with the recent update from *GnuCash*, deciding whether US\$29.99 should go to *MoneyDance* is a tough choice.

This is a Java-based personal finance management package equally at home on Linux, Windows, OS-X and various other Unices. The main installation packages include an integrated Java run-time environment that pushes the download size to 17MB. Installation is as simple as unpacking the archive, though if you have Java already installed – which reduces the download to a more dial-up friendly 4.5MB – you'll need to create a symbolic link named **ire** to the correct location. Then it is just a case of **cd**ing to the *MoneyDance* directory and typing **moneydance**.

First run users are given the option of starting a new set of accounts from scratch, opening an existing set or importing a Quicken Interchange Format (QIF) file. QIF has become quite a standard format at the lower end of the finance sector, and *MoneyDance* imported our test file well. Choosing the first option takes you through a very simple wizard to pick the default currency and whether you want a full set of accounts or just a basic selection.

The first real screen shows a summary of your accounts, a wide range of currency exchange rates and a calendar showing scheduled payments. The plain design of this screen, which is indicative of the whole application, is clear and unfussy and does make it


easy to see what's going on. There is little in the way of customisation beyond changing a few colours and adding elements such as stock prices and liabilities. In terms of the competition, it comes somewhere between the austerity of *GnuCash*, which gives you a snapshot, and the bloated opening screen *MS Money*, which can be a little overwhelming. A good balance.

The register

The 'register' has become the standard way of entering data in finance packages and *MoneyDance* doesn't deviate from that method, though it does make things simpler initially than *GnuCash* by virtue of its use of easy labels: payments go out of your account, deposits come in. The benefit of the register look-'n'-feel is that it is familiar to anyone who's every received a bank statement. In fact, *MoneyDance's* UI is quite intuitive in a number of ways. For instance, creating a scheduled payment, which is useful for recurring items such as mortgage, salary or pension payments, just involves highlighting a manual entry and hitting the Memorize key. You then enter the frequency of payments, the day it is due and a few other details and confirm it; the next time you start up that entry will appear on your calendar. It makes planning ahead very easy. Furthermore, each entry can be split so if you were, for example, a freelance writer you could split each income entry into taxable and net income.

At the top left of the registry window is a quartet of icons. These provide a swift link to the home page, for a broad overview of your financial situation, and also to tools which offer a more detailed examination of where everything goes. The magnifying glass allows you to find entries, the pie chart generates charts that can be targeted on specific dates or categories and, finally, the report dialog.

Graphically, the reports are not particularly evolved, though they do the job. Hovering the mouse over each section of either the pie or bar chart calls up a bright yellow tooltip which tells you the amount the segment represents in pounds and percentage. The graphics



Setting up scheduled transactions is simple.

owe their somewhat primitive style to *MoneyDance's* Java core which gives me jagged, unattractive fonts when I've come to know and love the anti-aliased beauty of KDE. I know this is the price you pay for cross-platform use, but I want an good-looking interface!

On the money?

On the whole, *MoneyDance* is a good product, intuitive to use and powerful enough to cover home or small business finances. The reports and graphs, while not spectacular, give a broad outline of users' personal wealth. Though a whole range of currencies are available to you, users hoping to integrate *MoneyDance* with their UK online bank will be disappointed. Many banks have the facility to download account information in QIF format, so all is not lost.

The tools on offer, the loan calculator and budget manager especially, are generally useful and the 'extensions' system, which gives you access to a range of services such as Yahoo! stock quotes, is well integrated.

MoneyDance does score over *GnuCash* in one potentially important area: security. Encrypting your precious data involves nothing more than

selecting a radio button and choosing a password. Upon opening the file, you are prompted for the password before any detail is revealed. *GnuCash* stores its info in a readable XML file and while you can encrypt this manually, it is useful to have it integrated into the system.

If your needs are slight and you don't relish the prospect of wrestling with *GnuCash's* (some would say) fussy way of working, *MoneyDance* is certainly worth considering, though you should try the 30-day demo. Make a mistake and that payment, which will be quite near the top of your *MoneyDance* Register page, will taunt and haunt you for months! **LXF**

VERDICT

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

Competent, but *GnuCash* undercuts it on price and improves on the facilities. Try the demo before buying.

LINUX FORMAT RATING
 7/10

WINDOWS VIRTUALISATION

Win4Lin 5.0 Workstation

NeTraverse thinks Win4Lin 5.0 will appeal to corporate upgraders investigating Linux. Andy Channelle puts the software to the test.

BUYER INFO

Competition comes mainly in the form of VMware and CodeWeavers' CrossOver Office.

- **DEVELOPER** NeTraverse
- **PRICE** \$89.99 download/
\$99.99 boxed
- **WEB** www.netraverse.com

Win4Lin 5.0 Workstation is the latest iteration of NeTraverse's range of Windows virtualization software which also includes *Win4Lin Terminal Server* for allowing many users to access Windows applications from a central server. The *Workstation* edition itself is available as a downloadable or boxed product and, as the name implies, is built for use on a single machine. NeTraverse is suggesting that the advanced functionality in *Win4Lin* and the ability to run many of the 'key' Windows applications will help make Linux more attractive to corporate buyers seeking ways to reduce IT costs. Companies get the benefit of the low-costs and high security that Linux is famed for, without having to retrain their entire workforce in a new productivity suite or buy new Microsoft licenses. More



As it virtualises the whole OS, Win4Lin handles most Internet Explorer content – including Quicktime – out of the box. Linux digicam users rejoice!

expensive than *CrossOver Office* and a lot cheaper than *VMware*, *Win4Lin* also falls between these two products in functionality.

Installation

Previous versions of the software were criticized for the hoops the hapless installer had to jump through in order to get up and running, a process that

involved downloading and installing a special patched kernel. You still have to have a new kernel, but fortunately the whole installation routine is now – theoretically – graphical.

Installation is broken into three stages, beginning with acquiring the special kernel. For this to be really effective you'll need a supported distribution, which currently amounts to the big players, Red Hat, SuSE *et al*, though you can extract the patches from the source RPM to 'roll your own' *Win4Lin* kernel. The system correctly identified the default SuSE 8.2 kernel (2.4.20-4GB) and successfully installed the patch. Messing around with the kernel like this inevitably means a reboot and after restarting there was a new WIN4LIN entry on the GRUB menu, meaning you can choose to devote all the system resources to Linux if you know you won't need access to any Windows applications.

After the restart, you need to log in as root for the second stage of the install that sorts out the system-wide configuration. Again this is supposed to be fully graphical, unfortunately there was a snag during installation on my

system when it came to pointing the software at my Windows media – it didn't recognise it as bootable. This was anticipated by the excellent user manual though, and the solution involved nothing more arduous than a couple of typed shell commands to create a virtual boot-disk. NeTraverse assures me that in most cases the graphical system-wide install should run smoothly, and the command line alternative wasn't difficult, but the fact that the company has dedicated space in the manual suggests that this may be a common occurrence.

Once the second stage is complete you can remove the Windows CD – it's no longer needed – and logout of root.


The final stage is the User Install. This is completed under a regular user account, and you have to do it for every user that requires access to Windows. Just as an aside, I'm uncertain whether installing the same edition of Windows for multiple users on the same machine is within Microsoft's licensing restrictions, but it should be because the OS natively allows multiple users. If you have ever installed Windows on a standalone machine, this last step will be very familiar, though as all the files have been copied across to the hard disk it runs along about five times faster and the reboots take hardly any time at all.

Once that's done you'll have an icon on the desktop which, when clicked, launches Windows in a window, but you can edit the launch script, simply adding a **-f** suffix to launch in full screen mode.

Uncanny in use

Feeling brave, I backed up all the data on a Windows-running 500MHz Celeron, wiped the disk, installed SuSE 8.2, *Win4Lin* and Windows. The whole process took about an hour and, with the exception of the slight glitch as mentioned above, everything went smoothly. Let the testing begin...

The first bonus is that, as the Linux display drivers for my PC's graphics system are so much more evolved



DirectX support is regarded as 'limited', but is useable nonetheless.

Disk space

Make some room...

As each user needs his or her own installation, and Windows is not exactly noted for its svelte frame, *Win4Lin* has some pretty hefty space requirements. The first space hog is the requirement to have the contents of the Windows CD archived in `/var/win4lin`. For Windows 98 this is a manageable 40MB, but it goes up to 198MB for Windows Me.


Operating system Space required in /home

Windows 95	41MB
Windows 95OSR2	93MB
Windows 98	148MB
Windows 98SE	195MB
Windows Me	302MB


On top of this you may need to keep aside space for applications such as *Office 97* (116MB), *Office 2000* (400MB for a complete installation) and *Office XP* (577MB).

than their Windows counterparts, I could now run the OS at a higher resolution than the previous limit of 1024x768, which is great for working

on large *Adobe Photoshop* images or complex web designs. Sound support is also useable and doesn't, as I'd feared, interfere with the underlying



The range of supported applications is very good. From *Office XP*...



...to *Photoshop*, *Dreamweaver* and many more.

sound server. I tried a number of sound editing applications such as *Audacity*, which on Windows has the advantage of using VST plugins, with no problems.

This is all very useful, but NeTraverse has its sights set on getting Linux into offices across the world, so the most important weapon in the *Win4Lin* arsenal is the ability to running the productivity suite that offices across the world tend to use: and that means *MS Office*. Its claim to support all current *Office* versions is borne out by our testing. Everything from *Office 97* to *Office XP Professional* installed and ran as expected. While you can't cut and paste between OSes as you can with *CrossOver Office*, the shared directory structure – Windows' My Documents location becomes `/home/username/mydata` in Linux – making moving files around fairly simple. Within Windows, other parts of the Linux filesystem – root, `/home`, etc. – receive Windows drive letters, making the whole file system accessible so you could, if you so wish, open and save web pages in *Dreamweaver* directly from *Apache's* root directory.

Beloved CD-ROMs

Networking has also been expanded with the addition of Novell and Winsock II support, it's also now possible to create a virtual network between the Windows and Linux elements of your system. Set up here is a little fiddly, but basic Internet connection sharing (with Windows working through the Linux connection) is set up by default and, for most users, this will be sufficient.

Impressively the software also handled *Truespace*, which is DirectX based, and a range of educational and leisure titles including the beloved *Pingu* CD-ROM.

My 'old' Windows installation was old and bloated, so direct performance comparisons are probably not fair, but there is no noticeable performance degradation over a native installation. In fact, some things – application loading, navigating large images or sound files – seemed a little more responsive, and a reboot during software installation or after a Blue Screen of Death, while not as fast as *CrossOver Office*, is still performed in under a minute.

In fact the only significant usage problem I have encountered is

accessing audio CDs from within Windows. It just refused point-blank to recognise them, which is a shame as the MP3 ripper in *MusicMatch Jukebox* is so much faster than *GRIP*!

Conclusion

Unlike *CrossOver Office*, which doesn't rely on any Microsoft code, *Win4Lin* requires a valid Windows license so, unless you've got one of these hanging around the house/office, you'll have to factor in that cost. The difficulty you may face if you don't have a spare Windows CD, is that as the application only supports Windows 95, 98 and Me, you may have trouble getting hold of the media to install.

In comparison to CodeWeavers' *Wine* project, *Win4Lin* stands up rather better than you might have expected. It costs a little more, and requires a Windows license, but it handles a lot more software than *CrossOver Office*. For instance, it is possible to install all of *Office XP*, including *Outlook* and *Access*, *Photoshop's* sister product *ImageReady* works fine and, this could be the decider, *Dreamweaver* and *Adobe GoLive* work faultlessly. You also have other advantages, such as access to *Quicktime* video files and other *Internet Explorer* plug ins which otherwise would require the additional expense of *CrossOver Plugin*.

If you're in the 'just one application' camp (for instance, if your bank's online account system still refuses to sanction anything but *IE*) then *CrossOver Office* is the better option for your needs. However, if you live in a house with an OS schism (as I do), or need access to a wider range of Windows applications, the ability to switch from Linux to Windows and back again by hitting `<CTRL> <ALT> <F7/8>` will save a lot of wear and tear on your PC's reset button. And quite a bit of time too. [LXF](#)

VERDICT

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

Does exactly as it claims to – and does it well – but casual users will need to consider whether the pain of dual-booting and the features on offer are enough to warrant the extra expense.

LINUX FORMAT RATING
8/10

NETWORK BACKUPS

NetVault 7 Enterprise Edition

Paul Hudson tests the latest software designed to take the sting out of unwanted fscks...

BUYER INFO

Proprietary backup solution in a league of its own – for a cheaper alternative, try *Amanda*.


- **SUPPLIER** Bakbone Software
- **PRICE** £5995
- **WEB** www.bakbone.com

All of us wish that data loss would fsck off – it's no fun to lose your work, and it's even less fun if your company loses its work! However, precious few Linux users actually back up their data for one reason or another, which means they're unprotected if the worst happens. The usual reasons cited for not backing up include "I thought Linux was supposed to be stable?" or "I use a journalled filesystem – I have no need to back up". Sorry to tell you, but you're quite wrong – journalled filesystem or otherwise, you need to back up your data. It's rarely because of a problem with the kernel – the last file system problem in the kernel was some years ago. More often than not, it's merely an application error – sometimes it's a bug, sometimes the power goes out without a daemon having flushed its cache. Either way, data goes to information heaven, and it's unlikely to come back no matter how much you want it to.

The solution, unsurprisingly, is to back up your data on a regular basis. For home users with little data backing up can be as simple as copying their home directory to a CD once a week, however such a simple solution is rarely practical once you expand to multiple drives, computers, or even locations – you need something more powerful. Furthermore, in order to cut down on the size of the backup media, most people use incremental backups – the process of archiving only the data that has changed since the last backup – and simply copying large chunks of files won't allow you to do this.

The new competitor

In steps *NetVault* – one of the most popular cross-platform backup



A complicated GUI, but one that gives you all the power you could ever want.

solutions currently on the market. *NetVault 7*, the latest incarnation, comes in a sizeable box, that's jam-packed with two manuals and two CDs. The manuals, about two inches thick combined, are excellent – full coverage of installation on Linux, BSD, Windows, AiX, Solaris, and others, as well as comprehensive descriptions on how to use each part of the software. The second manual covers some of the many plugin modules that come with *NetVault* – tools for backing up databases such as Oracle, DB2, and MySQL, and also for groupware systems such as *Lotus Notes* and *Microsoft Exchange*. *NetVault* really makes it straightforward to get a solid backup solution up and running, irrespective of what software you're running.

Installing the product is simply a matter of running a binary from the CD as root, which means few should have problems. We tested it on Debian and Mandrake, and both worked smoothly with *NetVault*. The main software can run in either graphical mode (using X) or text-only, which is great for administrators who have to work through *ssh* or a serial connection.

Sometimes the GUI is a little too complicated for its own good – there

are option tabs within option tabs within option tabs. While this probably makes sense to users who have been at this thing for a while, it certainly left me constantly having to pick up the manual to save myself the time of having to search through the interface to find things. One advantage here is that the interface looks similar across all platforms – once you learn it on one platform you've learnt it for all platforms.

The software includes a variety of advanced functionality, from simple things like incremental backups to setting backup life, compression level, scripting, and HTML reporting, which means it almost certainly includes everything you need. Given that it is so easily extendible with plugins, anything that's not in the main program may well already be available as a module.

The one big down-side to the product is its price tag and complicated licensing. Licensing costs are decided on the version of *NetVault* you wish to use (*Workgroup* is the cheapest, *Enterprise* is the most expensive, and *Datacentre* fits somewhere in the middle), how many drives you have, how many clients you have, and more. The *Enterprise* edition, which is the version we reviewed, costs no less than

£5995, and supports backing up 25 clients. If this is a bit too rich for you, the *Workgroup* edition weighs in at £995, and *Datacentre* at £2995 – they have some features missing, though, and only support ten and five clients respectively.

Conclusion

Yes, it does everything you'll want from a backup solution, and yes it's cross-platform with excellent documentation. However, the price is fairly steep, and you might be best to consider other options. On the other hand, if you want a complete, supported backup application and aren't afraid to write fat cheques, *NetVault* leads the pack. **LXF**

VERDICT

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

A feature-packed, if somewhat pricey, backup solution geared towards medium to large businesses.

LINUX FORMAT RATING
8/10

ADVANCED TEXT EDITOR

TextMaker 2002 for Linux



Recovering from the shock of seeing yet another word processor application, Paul Hudson realises that this is in fact *not* just yet another word processor...

BUYER INFO

Not just yet another word processor, alternatives include *KWord*, *OpenOffice.org*

■ **SUPPLIER** SoftMaker

■ **PRICE** \$49.95

■ **WEB**


www.softmaker.de/index.en.htm

K Word, *OpenOffice.org Writer*, *AbiWord*, *Hancom Office*, and so on – is there really enough room for another word processor? I started up *TextMaker* with some trepidation – after all, it's hard to write anything original about a product once you've already reviewed ten others of its kind! The first thing to hit me was that it took about one-and-a-half seconds to load on my 800MHz PIII, which is faster than pretty much every other word processor out there, including something as simple (in comparison) as *Kate*!

TextMaker's appears to use its own graphical toolkit – fair enough, given that it runs equally as smoothly on Windows as it does on Linux, even managing to run on a Pocket PC system. Throughout the interface there are colourful menu icons, popup menus, tooltips, and other guides to help users get on track easily – SoftMaker really have done a *lot* of work to make this user interface as simple as it is, and should be commended for it.

Functionality

This product includes features we've all come to expect from word processors, such as customisable tabs, mail merge, drawings, and tables. The developers appear to have a modelled a lot of *TextMaker's* data flow – how users interact with the system – on a combination of *Lotus SmartSuite* and *Microsoft Office*, which are both venerable tools with a history of friendly and powerful user interfaces. Only once in *TextMaker* was I left wondering "how do I do this?", simply because the interface was so well thought-out. One minor drawback is



Mr Efficient? Meet Mr Feature-Packed – I don't think you two have met.

the menu option "Extras", where the developers pretty much dropped everything they couldn't work out where else to put.

TextMaker has excellent *Microsoft Word* import and export support, which will certainly count heavily with corporates looking to switch. One place it lacks is support for *OpenOffice.org* import and export – given that the formats are entirely open (and

remarkably easy to read and write), and that *OpenOffice.org* continues to rise in popularity, this is a bit of an odd exclusion and really ought to be rectified as soon as possible.

One place this program truly excels is with its fonts – they look beautiful by default, which is almost unheard of in the Linux world. Take a look at the example below to see how good *TextMaker* looks compared to

This tests how nice the fonts are
This tests how nice the fonts are
This tests how nice the fonts are
This tests how nice the fonts are

***TextMaker* does excellent anti-aliasing by default – no configuration needed.**

This tests how nice the fonts are
This tests how nice the fonts are
This tests how nice the fonts are
This tests how nice the fonts are

***OpenOffice.org's* fonts look quite ugly because they have no anti-aliasing. While this can be enabled, it usually takes a great deal of effort.**

OpenOffice.org Writer. It also has unlimited undo and redo capabilities, full mouse wheel support, background spell checking, a built-in thesaurus, index generator, and built-in document forms – this package is incredibly complete, and yet still manages to run like lightning.

Drawbacks

The one time I couldn't find what I wanted in the *TextMaker* interface was when I wanted to merge columns from a table. In *OOo*, this is simply a matter of selecting the cells you want, right-clicking, selecting 'Cell', then 'Merge' – but could I find this in *TextMaker*? Heck no. There are a few glitches in the system that could irritate – for example, changing the line spacing to less than 1 line makes the text highlighting go bananas. Furthermore, the spellchecking could use a little more work – "goodboy" isn't recognised as being two words lacking a space in between, and "Linux" isn't recognised at all. However, these are all minor things – the product is well rounded and mature.

Despite these problems, the product is much (*much*) easier to use than *OpenOffice.org*, and exponentially faster too. Add to that the fact that large amounts of it look and feel just like *Microsoft Word*, there's little reason people can't make the jump immediately – this is every bit as good as its competitors, and will almost certainly be a hit wherever people take the time to try it. **LXF**

VERDICT

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

Beats the pants of its competitors, and a price tag even the meanest user would be hard-pressed to complain about.

LINUX FORMAT RATING

9/10

MySQL Enterprise Solutions


Nick Veitch refutes that you need to be an expert to make use of this book.

BUYER INFO

- **AUTHOR** Alexander Pachev
- **PUBLISHER** Wiley
(www.wiley.com)
- **ISBN** ISBN0-471-26922-0
- **PRICE** £26.99

The title of the book may be a little misleading, as it may tend to suggest that you know a little about MySQL already, or that the content of the book is only useful to those looking to deploy huge business applications. The book actually does a good job of explaining MySQL and the pitfalls and shortcuts to good performance whatever you intend to use it for.

A preamble section explains what MySQL is, where it comes from, and how it is commonly used. There is also a fairly honest assessment of the strengths and weaknesses of the



software for enterprise use, though this is belied by some of the frankly gushing praise for the people behind MySQL ("Monty Widenius... deserves to be called Mozart of computer programming..").

There are plenty of tips on optimisation too. MySQL may have a reputation for being very fast, but that can easily be hampered by sloppy programming. A fair amount of space

is given to setting up your database, testing it and how to avoid unpleasant thrashing of the database tables – all good info.

When it comes to actually writing some code, the sections on C, PHP, Perl and Java explain the standard APIs and provide some sample code to show how they work. There is a lot of really good advice here on how to build a reliable application, and just as important, how to maintain it and keep it running well.

A very creditable effort is made to explain how to extend the functionality of MySQL by working with and adding to the source. This is a major undertaking with any large piece of software, and possibly doubly hard with something of the complexity of MySQL. The thought of meddling with this source would give me nightmares, but there are plenty of examples for those who wish to dabble.

This is a great book that covers a lot of ground for anyone considering using MySQL for any sort of serious task. I'm not a big fan of the paper or some of the typesetting, but it's a good read.

VERDICT

A good read for anyone considering a serious use for MySQL.

LINUX FORMAT RATING
 7/10

Teach Yourself C++

Nick Veitch gets some upper-body exercise reviewing this mammoth volume.

BUYER INFO

- **AUTHOR** Al Stevens
- **PUBLISHER** Wiley
- **ISBN** 0-7645-2644-8
- **PRICE** £22.95

There are many many books on programming, and a very many of those are about C or C++. To stand out in this crowded marketplace, a book has to be something special. But Wiley's *Teach yourself C++* is no ordinary book. Now in its seventh edition, and authored by the notable Al Stevens of *Dr Dobb's* fame, this aims to be as complete a book as you could hope to find on the C++ language.

The quality and depth of information is assured from the preface, which actually takes the time to explain the history of C and C++ and why they came about – it's always much easier to learn how to use a tool when you know exactly what it was designed for.

There are a few tips scattered about the book, but there could easily



be more. For example, post and pre-incrementing are covered in a page. The difference is explained, but without any idea as to when you might want to use one method over another. To be fair, the book is pretty huge at over 700 pages, so perhaps brevity won out.

The book covers everything you can probably think of, from simple C++ variables, through flow control, functions, operators, virtualising, overloading, class inheritance to

complex IO.

Throughout there are plenty of examples, though in parts it does seem a little too much like a textbook that you are being lectured on. Pointers for example, probably the most difficult part of C and C++ to get to grips with, are explained adequately but very tersely. Though some parts might be harder to get through than others though, this book will give you a complete understanding of how the language should work.

The last third or more of the book covers material outside of the core

language, but nevertheless very useful to get a complete picture. The Standard Template Libraries get one of the most thorough outings I've seen in a book. Great stuff.

The CDROM accompanying the book is mostly useful for Windows users however. Although it does include GPLed software and sources, and a complete GPL Integrated Development Environment for Windows called *Quincy*. All the source to the examples is available on the CD, though without Linux makefiles, – still, that'll be a bit more practice for you.

As a harsh judge of books (so I've been told) I could certainly see some room for some improvement here, but even as it is at present, it stands head and shoulders above many of the works in this field.

VERDICT

A thoroughbred in a field of also-rans.

LINUX FORMAT RATING
 9/10

The Linux Process Manager


With the codebase growing, it's becoming harder to learn how the Linux kernel works, says **Paul Hudson**.

BUYER INFO

- **AUTHOR** John O'Gorman
- **PUBLISHER** Wiley
- **ISBN** 0-470-84771-9
- **PRICE** £34.95

This is a very rare, and, seemingly at first at least, a very peculiar book. Large chunks of it – perhaps about 20% – are copied direct from the kernel. About another 60% is annotation of that code, with the last 20% being general background information. If that sounds like boring reading to you, then you're thinking the same thing as I was when I first picked it up – “who on Earth would read a book documenting source code?”

The key is that this is not just any source code – it's the source code from a major part of our beloved kernel, and the plain fact is that the book reads like a documentary. I don't know about you, but if I see a documentary about anything –



movies, animals, history, etc – I find myself unable to stop watching, because I always want to know more and more. I got that same feeling whilst reading the Linux Process Manager – it's non-fiction, yes, but it's about as close to a

“gripping page-turner” that we can have!

Owing to the fact that the kernel is now over twenty megabytes in size, it's too big to cover in one book. As a result, this book focuses solely on the process manager, covering scheduling, interrupts, and

signals – a complicated topic in itself, but one that fits well into the 800 pages of this book. The author, sadly now passed away, had long been a university lecturer on the subject of operating systems, and this shines through in the text – there's no ambiguity, and no doubt; you can simply dive in at the point that most interests you, and start reading from the authoritative text.

This book lives up to Wiley's ideal of providing timely, practical, and reliable texts. It makes for such an interesting read, and is about a topic so close to our hearts – this is sure to quite deservedly be one of Wiley's bestsellers for the year.

VERDICT

A classic is born, earning a place on bookshelves the world over

LINUX FORMAT RATING

■■■■■■■■■■ 9/10

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



EasyVserver solutions
Debian or RedHat O/S
True “root” access
4, 6 or 8GB raid space
1 IP address
Highly secure

p-p-p-pick up
a dedicated linux server

Make the most of Linux technology with the big name in Web registration and hosting packages. Our flexible, scalable, secure EasyVserver solutions start at just £39 per month. Back up by unrivalled support and know-how. If you want the best of Linux come along for the ride.

log on today at: www.easyspace.com

EasySpace 
your perfect partner for the web

Established in 1997 with over ½ million customers. Accredited ICANN registrar & nominet member. Prices exclude VAT.


Red Hat Linux 9 For Dummies

Richard Cobbett finds his level with this excellent trip up the learning Bluecurve...

BUYER INFO

- **AUTHOR** Jon Hall, Paul G. Sery
- **PUBLISHER** For Dummies / Wiley
- **ISBN** 0-7645-3990-6
- **PRICE** £20.95

Unlike many Linux guides, *Red Hat 9.0 for Dummies* understands that you're reading it because you actually want to use Linux rather than simply learn about it. When you're first starting out, you don't need to know or care what goes into etc, sit through a lecture on the politics of Open Source, and the only use for an fstab is to the chest, as a gentle reminder to the author that it's time to actually get started. Here, you're up and running within minutes – quite literally. The book comes with a two-disc Publisher's Edition of Red Hat 9.0, and using this you'll be on the desktop by page 53, online by 113 and be comfortably working in a productive fashion by teatime.



Following the normal *Dummies* pattern, the book is broken down into categories – installation, individual software packages and so on, and then into individual stumbling blocks that you might encounter, such as XMMS requiring a patch to play MP3 files, or the Anaconda installation program not being able to resize NTFS partitions. All steps are explained in

plain English, with a decent dose of humour to liven things up – although there is the occasional point (notably installing from RPM) where you have to wonder why the authors insist on opening a terminal instead of clicking on an icon. Of course, the advantage of doing it this way is that even the newest of users get some experience of the command line, making it much easier to perform more complicated tasks that are out of this book's remit.

While the range of topics covered is impressive, it seems odd that so much attention has been lavished on higher-end subjects such as building private networks and configuring servers. These sections are useful, but it's arguable that the space handed over to "So You Want To Build An Internet Gateway" and similar topics would have been more usefully spent on summing up packages like *GIMP* with more than "this program's good too", or offering larger tips and troubleshooting sections. Despite this, by the end of the book you'll be more than comfortable with Red Hat, and raring to take on the rest of the Linux world.

VERDICT

Entertainingly covers all the essential bases for getting started, but a bit more detail on the nitty-gritty of actually working in your new OS would have been the icing on the cake.

LINUX FORMAT RATING
 //////////////// 9/10


Apache Tomcat Bible

Paul Hudson reviews a new book claiming to cover the leading Java web server in its entirety.

BUYER INFO

- **AUTHOR** Jon Eaves, Rupert Jones, & Warner Godfrey
- **PUBLISHER** Wiley
- **ISBN** 0-7645-2606-5
- **PRICE** £34.95

By itself, Java is a big topic, and Java Server Pages (JSP), the section of Java allow you to create your WWW site in Java, is no different. *Tomcat* is an official *Apache* project that provides a complete reference implementation of both the Java Servlet Specification and the Java Server Pages specification, and makes up a major part of *Apache's Jakarta* project. This book aims to cover all of *Tomcat*, as well as several other parts of Jakarta, including *Ant* (the XML-based build tool), *Struts* (the model/view framework builder), *Log4j* (the logging system), and *Cactus* (the test framework), which all in all is quite a task



in the 800 or so pages in this book. About two-thirds of the book is dedicated to teaching how to install, configure, and use the system. Uniquely, the book also covers choosing an IDE, handling complex tasks such as load balancing, and how to customise

Tomcat. The book isn't perfect, though – Wiley continues to persist with its policy of not distributing CDs with books, which means you'll need easy access to the web to get most use out of this book. Secondly, the book is Windows-centric for many topics, particularly when it comes to installation. UNIX readers like us will need to ignore the Windows-only

instructions and figure it out ourselves, which might cause a few hiccups for some. When it's good, it's good, but it's fair to say that the book lets itself down quite badly by not straying outside the Windows environment, which is a shame because other parts (particularly the *Ant* parts and 80+ pages on "Creating a real application") are excellent.

Overall, it's a fair read that can give you a solid introduction into how *Tomcat* works and why you might want to make the jump from *Apache*. However, it is likely to leave you with more questions than when you started, which, at the price, is a bit disappointing. **LXF**

VERDICT

A good read, but you'll need to skip over the bad bits to feel satisfied.

LINUX FORMAT RATING
 //////////////// 8/10

Roundup

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

CD-burning apps



OUR SELECTION AT A GLANCE

- X-CD-Roast
- CD Bake Oven
- K3b
- Simple CDR-X
- webCDwriter
- Arson

Whether you like to master, to burn or to rip, you'll find something of interest in **Richard Drummond's** comparison of CD writing tools.

We last looked at the crop of CD writing suites available for Linux about 18 months ago, so we were interested to see what changes had been made since then. Would the previously top-rated packages still rule the roost for features and ease of use? Or would some brash young up-start sneak in and grab the laurels?

What we are testing here is GUI-driven, general-purpose CD writing

tools. These applications all provide broadly similar features for mastering and writing data tracks to disc, mastering audio discs, and ripping audio tracks from disc. Some even have tools for creating VideoCDs. While there are many other useful applications for working with CDs under Linux, they tend to be specialised or single-purpose applications. If you simply need a quick tool to rip and encode audio

tracks or perhaps to make filesystem back-ups to CD, then pop along to <http://freshmeat.net> and have a browse. Here we are looking at the jack-of-all-trades.

All of the applications on test are merely front-ends for a commonly available set of command-line tools which do all the real hard work, such as *mkisofs* for mastering data tracks, *cdrecord* and *cdrecord* for the actual writing, *cdparanoia* and *cdada2wav* for

audio ripping. I say 'merely'; but the devil is in the details. Building a good user interface is not easy, as a look at the range in quality of the applications on test will reveal. But there's not much point in using a friendly front-end if it's not actually friendlier or quicker to use than the tools it is abstracting away from. GUI design, then, will be a key factor by which we will rate these applications, and is just as important as the features on offer.

X-CD-Roast

Stalwart CD-burning suite steadily accrues features.


■ **VERSION** 0.98alpha14 ■ **WEB** www.xcdrtoast.org

X-CD-Roast has been around, it seems, for ever – yet it is still tagged as alpha status. While this is technically true (it isn't feature complete yet) it's actually very stable and dependable. *X-CD-Roast* has a simple GTK+ 1.2 interface. Configuration is easy, and it automatically detects any CD readers or writers that you have attached. There's even experimental support for using the new ATAPI interface for ATAPI writers (rather than the SCSI generic interface via the SCSI emulation layer). Support for networked drives is also included.

The alpha14 release adds much-needed features, though its interface remains broadly constant. The same double list-views are employed when mastering data tracks – one display is the ISO tree you are building and the other contains your local filesystem.

While this has some useful features, such as being able to filter out files using shell-style patterns, it was always cumbersome to use. This has been improved markedly with the addition of drag-and-drop. Many other interface tweaks make *X-CD-Roast* much more pleasant to use all round, eg it no longer forgets your ISO tree when you switch from the Master Track window. *X-CD-Roast* provides a fairly complete set of tools for configuring your ISO build, and usefully, warns you when your twiddling breaks strict ISO compatibility. The only significant features missing are support for Hybrid/HFS filesystems and the ability to hide files when building *Joliet* or *Rock Ridge* discs. No big deal if you're not doing cross-platform work.

Audio support is comprehensive. You can retrieve and write CD-Text



X-CD-Roast's interface won't win any design awards, but with drag-and-drop now supported it's certainly quicker to use than it used to be.

information from and to audio discs, and you can do CDDB lookup. Only CDDA-format WAV files are supported when ripping and writing, though.

An annoying aspect of *X-CD-Roast's* interface detracts from its usefulness for audio work – you may only write tracks to and read tracks from folders that you have designated in the *X-CD-Roast* settings as image folders – awkward for data tracks, but a positive pain when it comes to dealing with audio tracks.

VERDICT

Ease of use	7/10
Audio features	7/10
Data features	8/10
Performance	8/10

No fuss, no frills. *X-CD-Roast* is the dependable workhorse of CD burning.

LINUX FORMAT RATING
8/10

Arson


It won't set the house on fire, but this KDE-based CD-writer is a pretty capable package.

■ **VERSION** 0.9.7 ■ **WEB** <http://arson.sourceforge.net/>

Arson began life as a front-end to cdrdao for creating audio discs, but, as development has progressed, support for mastering and writing data tracks and VideoCDs has been added. *Arson's*

strength still lies in its audio handling abilities, however.

While it lacks the panache of *K3b's* interface, *Arson's* GUI is mostly well laid out and easy to use. Negative points are



With support for a variety of encodings, ID3 tags and path templates *Arson* excels at ripping and encoding audio tracks.

the clumsy and cluttered configuration dialogs and virtually non-existent error detection and reporting. There's also little documentation included.

Arson's main screen is divided in to four tab panes, one each for creating audio CDs, VideoCDs and data CDs, and the last for ripping audio tracks. Tools for mastering data tracks are limited: you build a filesystem tree by adding files and folders as you wish and pruning the entries you don't want. You cannot rename entries or create new empty directories, however. The ISO Settings dialog supports *mkisofs's* main options to control how your filesystem is created, although this is inconvenient to use. Most of what you need is here, such as support for Rock Ridge and Joliet extensions and El Torrito boot images. You cannot hide individual files from the Rock Ridge or the Joliet filesystems, and support for Hybrid/HFS filesystems is not included. A greater deficiency is that you cannot create multi-session or mixed mode discs. On the plus side, a quick function for generating a data track directly from a directory hierarchy is provided, useful for doing back-ups, for instance.

When it comes to ripping and writing audio tracks, *Arson* is much better equipped. It supports an

impressive array of sound formats, including WAV, MP3, Ogg Vorbis, FLAC and AIFF. CD-Text retrieval and CDDb lookup can be used to obtain track info when ripping, although the former appeared broken on testing. Author, track, etc. templates are supported when generating the output path of ripped files, and ID3 tags can be attached where appropriate. Audio tracks cannot be written on the fly – they must be decoded to a temporary files first. Normalizing of audio tracks is supported by *Arson*, but not the writing of CD-Text information. Basic VideoCD support is included, but MPEG2 streams must already be appropriately formatted.

VERDICT

Ease of use	7/10
Audio features	8/10
Data features	5/10
Performance	7/10

Great for ripping audio discs, but *Arson* lacks the features its rivals offers for burning data tracks.

LINUX FORMAT RATING

6/10

CD Bake Oven

Another rewrite for version 2.0, this is now even more tightly integrated with KDE desktop.


■ **VERSION** 2.0.0 beta 2 ■ **WEB** <http://cdbakeoven.sourceforge.net/>

CD Bake Oven 2.0 is still at the beta stage, but, like its predecessor, it has some unique and powerful features. *CD Bake Oven* is a KDE-based application, and heavily makes use of the KDE infrastructure. Its settings dialogs are all *KDE Control Center* modules; it relies on KDE's *cdaudio* IO slave for ripping and encoding audio tracks; and it features its own IO slave for constructing ISO filesystems with *Konqueror*, accessible via the new **cd://** URL.

The main interface of *Bake Oven* is structured similarly to version 1.8. Thus, the main window is split into two main areas, the top half has a filesystem browser, while the bottom contains the audio discs or data tracks that you are building, with each in a separate tab

pane. You can simply drag-and-drop from the top to the bottom to add files and directories to an ISO filesystem or audio tracks to a audio CD. The interface is clean and clear for the most, and some context-sensitive help is provided. There's little documentation besides.

Features for constructing ISO filesystems are good, but sadly the filter tool from previous versions has vanished. Still, the *Bake Oven* GUI allows you to fairly rapidly construct a directory tree, by dragging and dropping, either from the integrated browser or from the desktop. *Bake Oven* supports Rock Ridge, Joliet and Hybrid/HFS tracks, and lets you select in which of these filesystems an individual file or folder will appear. There's also El Torrito support.



With its own *KDE Control Center* modules and *Konqueror* plug-ins, *CD Bake Oven* integrates well with the KDE desktop.

Audio ripping and encoding is done via a separate dialog, accessed via the Tools menu. CD-Text retrieval is currently unsupported by the KDE *cdaudio* IO slave, so is not usable from *Bake Oven* either, though CDDb lookup and ID3 tagging does work. When burning, the writing of CD-Text is supported with extensive controls. On-the-fly decoding of audio isn't supported, nor is normalization. WAV, MP3 and Ogg formats can be used when ripping and burning audio tracks.

VERDICT

Ease of use	7/10
Audio features	7/10
Data features	8/10
Performance	6/10

CD Bake Oven is well equipped for mastering data tracks, but is significantly less capable, but still fairly competent at audio work.

LINUX FORMAT RATING


7/10

RoundupCD-burning

K3b

Combines ease-of-use with powerful features.

VERSION 0.9pre1 WEB www.k3b.org



A polished interface and a bucketload of powerful features make *K3b* a fantastic all-round performer for CD ripping and burning.


Although relatively unheard of until recently, *K3b* has risen from obscurity to leapfrog over the competition both in terms of usability and functionality. It boasts a well-designed and intuitive user interface that is simple to use, yet gives access to some advanced features. The GUI's prettiness and uncluttered layout are just part of the story: *K3b* exudes quality, from its fully automatic configuration, its helpful warnings and error reporting, and extensive context-sensitive help (although no real user manual is provided). It offers a level of features not matched by its peers, including fine-grained control over ISO generation and support for ripping and reformatting of video streams from DVD (in conjunction with the transcode tool).

Integrates well

K3b is KDE-based and like *CD Bake Oven*, it integrates well with the KDE desktop. Similar to its rival, you can drag-and-drop files from the desktop when building CDs, and, similarly, *K3b* includes an ISO-image MIME type, which launches *K3b* to allow quick writing. It also adds entries to *Konqueror*'s context menu to quickly build audio CDs from audio tracks and data CDs from folders and files. Though it is KDE-based, *K3b* doesn't go for *CD Bake Oven*'s tight coupling with KDE, however.

K3b's main window adopts the two-area approach, with the top half containing the source file browsers and the bottom half any disc-building projects that you are working on. A *K3b* project can either be an audio CD, a VideoCD, a data track or a mixed-mode disc, and you can have multiple projects open at once, displayed in separate tab-panes. The top-left pane of the main window shows a directory hierarchy. Clicking on folders here will list their contents in the browser window in the top right – and, not surprisingly, you can drag and drop files and folders onto a *K3b* project to add them. This directory browser also lists entries for any CD and DVD drives or burners you have attached to your system. When a disc is inserted, click on the corresponding drive, and its contents will be displayed in the browser. Filesystem tracks will be mounted and can be browsed like a normal filesystem; the contents of audio discs will be displayed with optional CDDB lookup; and, if you have transcode on your system, the contents of video DVDs will be listed as well (UDF support in your kernel isn't required). When viewing audio or video content, ripping tools are launched from here via a context menu.

When creating data tracks, you populate an ISO filesystem by dragging and dropping files and



K3b is unrivalled by other general-purpose CD writing applications when it comes to video ripping and encoding tools.

folders from the source browser or the desktop, renaming entries and creating new folders where appropriate. *K3b* supports Rock Ridge and Joliet extensions (but not Hybrid/HFS filesystems) and you can control on a per-file level in which filesystem each entry appears. Support for adding El Torrito boot images has just been added, and *K3b* allows you to include multiple boot images per disc.

Audio discs

K3b is less well-equipped for audio work. When ripping, although CDDB lookup is supported, CD-Text retrieval isn't. *K3b* can currently only encode audio tracks as WAV files (so, naturally, there's no ID3 tagging). You can use path templates to give the generated audio streams meaningful file names. Until support for other formats is added (it's on the to-do list), you'll probably want to go elsewhere for audio ripping, though. CD-Text output is supported when burning audio tracks, and *K3b* can decode and burn on the fly from WAV, MP3 and Ogg encodings. When decoding to temporary files, rather than on-the-fly, audio may be normalized.

When it comes to video, *K3b* is in a class of its own – none of the other programs here can compete. Not only is it the only application on test which supports DVD ripping, it includes tools for transcoding DVD video streams. With its built-in transcode support you can recompress MPEG2 streams using DivX4 or DivX5 and whole host



of options and output the result as an AVI file. It even automates picking a bit rate. You simply choose how many CDs or the file size you'd like to cram the video into, and it will select the best bit-rate given your chosen codec and options. Scaling and cropping of the video stream is also possible. Support for burning VideoCD 1.0 and 2.0-format and SVideoCD discs is included, plus rudimentary support for Playback Control (PBC) and for CD-I discs. However, you must pre-format the MPEG2 streams appropriately yourself, because *K3b* currently has no built-in tools to do this for you. *K3b* cannot decode from other video formats such as AVI when burning VideoCDs, either.

VERDICT

Ease of use	9/10
Audio features	7/10
Data features	8/10
Performance	9/10


K3b is quickly becoming the standard by which all other CD writing applications for Linux will be judged.

LINUX FORMAT RATING
9/10

SimpleCDR-X

A lightweight CD writer that aims for simplicity but falls somewhat short of its target.

■ **VERSION** 1.3-pre2 ■ **WEB** <http://ogre.rocky-road.net/cdr.shtml>



Building ISO filesystems with SimpleCDR-X is something of a struggle.

Developing a CD-burning application that is lighter and simpler than some of those on offer here is a worthwhile pursuit. *SimpleCDR-X* takes such a tack. Built on the GTK+ 1.2 toolkit, it provides

only the basic functionality that you will need 90% of the time. Unfortunately, as far as simplicity goes, it misses the mark by having a cluttered, confused user interface and poor stability.

SimpleCDR-X's main interface is split into five task-oriented tab panes. The first is for constructing data tracks, the second for burning audio CDs, the third for audio ripping, the next a stand-alone tool for tagging MP3 files, and the last is for program settings. The settings dialog is messy. You can call *cdrecord* to scan for SCSI generic CD drives and writers but must enter the details yourself manually. Non-SCSI generic CD readers are also supported – you simply enter the drive's device file – but the GUI doesn't make this obvious.

The tools provided by *SimpleCDR-X* for ISO mastering are basic and tortuous to use. You are presented with a list of actions down the left-hand side of the screen. Next to that is the directory hierarchy that you are building, and on the right is the contents of the currently selected directory. You can add files and folders, but this must be done with the primitive GTK+ file dialog. You can create directories, but only at the root of the filesystem. No renaming of entries is allowed. Mixed-mode discs are not supported, but multi-session discs are. Rock Ridge and Joliet

extensions are supported, but you cannot choose in which filesystem files appear. There's no facility for creating bootable discs.

Audio ripping has improved with version 1.3. You can now actually scan a CD to get a list of tracks and do CDDb lookup. (In 1.2.2 and earlier you couldn't list a disc's contents.) You must choose whether you want to encode to MP3 or Ogg streams in the main settings dialog though, and ID3-tagging is supported only with MP3s. You can decode and burn WAVs, MP3s and Ogg to disc, and *SimpleCDR-X* can normalise audio tracks. On-the-fly decoding isn't supported, nor is CD-Text.

VERDICT

Ease of use	5/10
Audio features	7/10
Data features	6/10
Performance	6/10

The meagre feature set and unpleasant interface doesn't compare favourably with the competition.

LINUX FORMAT RATING
5/10

webCDwriter

It's good to share, so why not let anybody on your network use your CD or DVD writer?

■ **VERSION** 2.5.1 ■ **WEB** <http://joerghaeger.de/webCDwriter/>

A client-server solution – built on the same familiar tools, such as *cdrecord* and *mkisofs* – that lets a machine's CD burner or multiple CD burners be

accessed from any host on your network. Now that CD-writer prices have fallen, you may think that's a waste of effort, but install the appropriate tools

and *webCDwriter* will work equally well with DVD writers – which are still pricey enough to make this technology useful. The *webCDwriter* server takes care of all of the resource sharing required, caching of mastered images, and, if necessary, user authentication.

Two client-side interfaces to the CD server are provided: a simple command-line client and a Java-based graphical interface. The latter is a true thin-client solution. Since it's packaged as an applet, nothing needs to be installed at the client end; the user just needs to point a Java-enabled web browser at the right port on the server.

Installing *webCDwriter* from source is a bit of a pain, but RPMs are provided which simplify things greatly. The applet itself is pre-compiled and signed, even in the source packages. Previously, it was signed using a Certificate Authority untrusted by most Java distributions. Thus for clients to be able to download and run the applet you either had fiddle with Java's key store on each client, which was a headache, or disable Java's security manager. *webCDwriter* is now a lot more transparent to install and use. The CD server's config file may need some tweaking, but it can largely

discover most of the info it needs and use defaults for the rest – just like the other standalone packages being tested.

The Java interface lets the user master ISO filesystems (with Rock Ridge and Joliet extensions, Hybrid/HFS support, multi-session discs and even boot images), burn pre-mastered ISO images, and create audio discs from WAV and MP3 streams. The interface is simple, but functional, and online documentation is included. On-the-fly CD copying from a local (to the server, that is) drive is also supported, but the CD server will only recognise drives with a SCSI-generic interface. No audio ripping tools are supplied.

VERDICT

Ease of use	8/10
Audio features	5/10
Data features	8/10
Performance	8/10

Features don't compete with the likes of *K3b*, but ideal for cross-platform, network-transparent CD writing.

LINUX FORMAT RATING
8/10



Simple interface, but there's a surprising amount of functionality packed in.

CD-BURNING APPS THE VERDICT

The usual question after we do one of these roundups is 'So which is best?' No matter what area of hardware or software that we have covered, the response has to be an inevitable but evasive 'Well, it depends on your requirements.' Well, for a change, we're going to be decisive – but don't expect this to set any sort of precedent for LXF roundups in future! *K3b* is clearly, unequivocally and without doubt the best CD burning application currently for Linux.

The reason *K3b* is so good is not simply that it has some of the most advanced features, it's because its features are easy to use. When you're talking front-ends, interface is everything. *K3b* features a clear, easy-



Ease of use and power features make *K3b* the clear winner.

to-navigate GUI design, automatic configuration, good error detection and reporting, helpful online

documentation, and a sheer level of polish that is not matched by any of the other packages. After using open-source software for so long, I'm actually taken aback by how good *K3b* really is. Sure, it's not perfect. Software never is. But the project seems to have enough momentum to address and overcome any shortcomings it currently has.

Having said all that, you may still want to use other packages. Ease of use is largely a personal judgement, after all. Or perhaps you don't want the overhead of running KDE. Each of the other CD writers on test have their own particular aspect of merit. *Arson*, for instance, is really handy for ripping and encoding audio discs and noticeably better equipped than *K3b* in this regard. *X-CD-Roast* and *CD Bake Oven* are both good at mastering data tracks. In fact, if you need to create Hybrid/HFS discs for use on

Burnt out?


CD writers that've faded...

If you remember the last time we did a roundup of CD burning apps, you may notice that some of our then-favourite packages are not met mentioned here. Most of these reside soundly in the 'Where are they now?' pile. Some, like *KonCD* and *KreateCD*, seem to have disappeared entirely. Others, like *GnomeToaster* and *GCombus*, have seen little if any improvement in the intervening 18 months. *GnomeToaster*, for one, is still a fairly capable package, even though it has received no development recently. If you can put up with the non-intuitive interface, it compares favourably with the packages on test here in terms of features.

the Mac, then *CD Bake Oven* is probably your best bet. Even *SimpleCD-R-X* is useful, if you need a lightweight application.

The future?

How will these applications fare in the next 18 months? What changes can we expect? Well for one, Kernel 2.6 should be well and truly out by then. Support for packet-writing with CD-R/W media hasn't yet been merged into the kernel development source tree, but it'll probably make it in time for 2.6. Packet-writing, along with the UDF filesystem (as used on DVD Video discs), enables the use of CD-R/W media just like regular hard drives, and is a much-awaited addition to the Linux kernel, especially for those that use CDs as a cheap back-up medium. Packet writing won't make traditional CD-burning obsolete, though, particularly since support for CD-R media has yet to be implemented. And, besides, there's plenty of scope for the applications on test here to grow. Expect more and better features, especially for video processing and burning. Hopefully, the developers of these CD writing applications can spend some time to add some polish to their software. It's a more complete overview of form and functionality that will make them increasingly more attractive to users, not just extra features. [LXF](#)



X-CD-Roast lacks *K3b*'s polish and tools, but its still a worthy package.

Table of features

How the different CD burners in this roundup measure up.

	Arson	CD Bake Oven	K3b	SimpleCD-R-X	webCDwriter	X-CD-Roast
Version	0.9.7	2.0beta2	0.9pre1	1.3pre2	2.5.1	0.98alpha14
Platform	KDE 2/3	KDE 3.1	KDE 3.1	GTK+ 1.2	Java 2 (Swing)	GTK+ 1.2
El Torrito support	yes	yes	yes	no	yes	yes
Hybrid/HFS support	no	yes	no	no	yes	no
Multi-session support	no	yes	yes	yes	yes	yes
Mixed-mode support	no	no	yes	no	yes	no
Audio formats	WAV,MP3,OGG, FLAC,AIFF	CDA,WAV,MP3, Ogg	WAV, MP3,Ogg	WAV, MP3,Ogg	WAV,	WAV
CD-Text	no	write-only	write-only	no	no	yes
ID3 tagging	yes	yes	N/A	MP3 only	N/A	N/A
Audio normalisation	yes	no	yes	yes	no	no
VideoCD writing	yes	no	yes	no	no	no

Hot Picks

The best new open source software on the planet!



Mike Saunders

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

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

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

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

Hot Picks at a glance

Firebird	40
WeirdX	41
ManEdit	41
4est's Weblog	42
ELinks	42
Middleman	43
Gallery	43
Tux Typing	44
Frozen Bubble	44

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!



GECKO-BASED WEB BROWSER

Mozilla Firebird

■ **VERSION** Version 0.6 ■ **WEB** www.mozilla.org/projects/firebird/



Firebird shows its tabbed-browsing and annoyance-eliminating features.

Six months ago, back in issue 37, we briefly looked at *Phoenix 0.5* in *Hot Picks*. This *Mozilla*-derived standalone Web browser aimed to remove some of the weight and bloat from its big brother, while adding new features and improving performance. Since then, a lot has happened in the *Phoenix* community; complications with the *Phoenix BIOS* company (who were producing their own small browser) led to a name change, but even that didn't go smoothly. *Firebird* was settled on as the new name – after much debate – but complaints from the *Firebird Database* developers have caused further problems.

More significant, though, was the *Mozilla* project's decision to adopt *Firebird* as its browser component. Following the release of *Mozilla 1.4* (which will replace 1.0 as the new stable branch), development activity will be based around *Firebird* and its mail/news client equivalent, *Thunderbird* (see LXF41 *Hot Picks*). As always, this has ignited more

flamefesty in the community, but the general belief is that it's a good route to take. *Mozilla* received a lot of criticism for being one large monolithic app, so splitting it up will combat that particular problem.

Firebird 0.6 has been a long time coming, but a vast amount of work has been done in the last half year and it's steadily maturing into a featureful and stable browser. As a sort-of 'fork' from the *Mozilla* codebase, it shares many similarities and is still built around *Gecko* (the rendering engine) and XUL (user interface language). Large amounts of cruft have been chucked out though, and more attention has been paid to a friendlier front-end.

She's on fire

Mozilla's coders never intended their app to be an end-user product – that was the job of *Netscape et al*. *Firebird* wants to bring control back to the user, and not the website developer. For the first time, *Firebird's* source is available as a tarball with no need to do a CVS

pull; we've provided the source on our coverdisc, but the simplest way to get up and running is by using the binary archive. Simply extract it in any place, change into the directory produced and run **./MozillaFirebird**. Providing you have a reasonably up-to-date distro, it should start up without any trouble. Note that it still uses **~/phoenix** to store its settings, and it's best to create a new profile for this release (as so much has changed).

Firebird's interface is simpler and cleaner than *Mozilla's*, with less clutter, fewer buttons and no sidebar enabled by default. Regular *Phoenixers* may be surprised to see a theme overhaul – the decent but garishly yellow *Orbit* skin has lost to *Quite*, which slightly resembles *WinXP*. It doesn't occupy as much space as *Mozilla's* default *Netscape*-ish theme, and the ability to reconfigure the toolbar with drag-and-drop is a superb and warmly welcomed bonus over *Moz*.

Other major developments since *Phoenix 0.5* include a snazzy new Preferences dialog, which groups various aspects of configuration into panes, along with smooth scrolling, automatic image resizing (if an image is larger than the viewing window), a context menu for the bookmark menu and the usual assortment of tweaks and bugfixes. Undoubtedly, the wait has been worth it.

Should *Mozilla* users switch? *Firebird* has a few nice touches, such as the small icon that appears bottom-left when a site tries to open a popup window (easier to permit it on a site-by-site basis), and it is noticeably snappier. It has *Moz's* excellent support for Web standards, a plethora of privacy and security options, advert filtering and tabbed browsing. Unfortunately it's not quite as stable, the form auto-complete function is particularly crashable.

Overall, *Firebird* is a great browser, jam-packed with goodies and ready for even more through extensions. It's fast, attractive, easy-to-use and mostly reliable, and it's guaranteed a prosperous future. Give it a test run – you might just replace *Mozilla* for good.

JAVA X SERVER WeirdX


■ VERSION 1.0.31 ■ WEB www.jcraft.com/weirdx/

One of the X Window System's most famous and popular strengths is its network transparency. Rather than fiddling with inelegant hacks (VNC) or expensive addons (*MS Terminal Server*), X users can enjoy simple and flexible remote application display built right in; enormously useful for both home and corporate Linux installations alike. XFree86, being the *de-facto* Linux X server, naturally offers this, but working in a mixed environment isn't so easy. X servers are available for Windows and other non-Unix-like OSes, but *WeirdX* stands out by being free and portable through its Java codebase.

WeirdX is a fork of JCraft's *WiredX*, the latter being a similar Java X server but not Open Source. In order to

generate interest in their product and as a goodwill gesture to the free software community, JCraft maintains *WeirdX* under the GPL. It's available as a standalone Java .jar application, and also as an applet for *MS Internet Explorer* – both are on our coverdisc. It needs JDK 1.1.x, Swing 1.1.x or Java2.

Running as a .cab applet under *IE* is clearly of huge value; no software installation is required on the Windows side, and it's a cinch to get working. With the weirdx-IE.html and weirdx-1.0.31-rsa.cab files dropped in place on a webserver, the Win32 user simply needs to fire up *IE*, point it at that HTML doc and begin working. A display is opened on the machine, ready for apps to be displayed on it (eg **xterm -display <IP>:2** or similar). Settings



IE on Win98, with Linux Mozilla and an xterm displaying via the applet.

like display dimensions can be set in the HTML file, and there's an option to display windows individually (out of the browser window).

WeirdX suffers from the occasional glitch with apps which make heavy use of the X protocol, but we ran *Mozilla*, *OpenOffice.org* and *FluxBox* without any major niggles. It's not all that fast though, so rendering highly

complex graphical software can be troublesome, but for most apps it does the job competently. In all, it's an impressive tool for mixed Win32 and UNIX environments; Windows desktops which occasionally need a Linux application don't need dedicated X server programs, and even for home use it's extremely handy to have around.

GUI MAN PAGE EDITOR ManEdit

■ VERSION 0.5.10 ■ WEB <http://wolfpack.twu.net/ManEdit/>

Writing good documentation is an essential (and often tedious) job in any open source project. Users need instructions, and developers need docs before they feel happy contributing. UNIX manual pages (typically shortened to 'man pages') are the best quick reference source for info on running a program. Despite the efforts of the GNU Project, which regards them as obsolete in favour of its 'info' format, man pages will be around for the foreseeable future.

Unfortunately, writing a man page isn't as straightforward as typing away in a text editor; in order to provide the the headers, footers, sections and bold/underline etc. attributes, the 'roff' formatting system is used. roff is complex and powerful, but takes time to master; thankfully, then, *ManEdit* is here and supplies a familiar and

graphical front-end for working on manual pages.


Built around GTK 1, *ManEdit*'s main window presents an approachable document editing layout, with a tree list of sections down the left and text pane

on the right. After the program reads in a man page, it converts it to XML – this tag-based format is easier to work with, and it's explained in thorough detail in the online docs. A handful of templates are also provided for those starting from scratch on a page.

ManEdit sports a preview window, which produces a quick render of the page (although direct editing through it isn't supported). Syntax highlighting can be toggled on, but the author warns of possible crashes with it; still, fonts and colours can all be modified

through the Prefs box, and on the whole it's a solid app. Cut and paste, search and replace, a tag insertion toolbar and context menus all make it a breeze to pick up and use.

Even though some features have yet to be implemented, and some already in place need more polish, *ManEdit* does its job very well – it didn't fall over with the huge and complex man pages we threw it at. If you're working on a project and need a quick way to prepare documentation, give *ManEdit* a look.



ManEdit's main editing window in action, along with the instant-preview box.

BLOGGING TOOL

4est's Weblog

■ VERSION 2.00 ■ WEB <http://www.semithere.com/?pro>

Web-logging, or 'Blogging' as it's come to be known colloquially, is a very interesting phenomenon that's taken off wildly recently – so much so that O'Reilly even has a book on it. A blog is a webpage or site where an individual posts his or her musings, like a diary, and friends can make comments; it's another communication system. Prominent industry pundits, hardware and software developers, and millions of others are all contributing to this 'opinion journalism' revolution, and naturally there are masses of tools to assist in it.

Forrest (4est) Jones' *Weblog* is a small package which makes setting up a weblog a walk in the park. With the archive extracted, it's just

a case of copying the relevant files over to a webserver, setting some permissions and altering the visual style to your preference. The information in the docs/ directory explains all, and it's pleasingly easy to get started with. We tested on *Apache 1.3*, but the Perl-based script should work fine on any webserver that supports CGI.

The *Weblog's* default graphical style is a minimal affair, with dotted orange on white lines and a table of links down the left hand side. Entries are listed newest-first, with links to read and write comments beneath; additionally, the message composer is notably accomplished, with its own brand of 'pseudocode' (HTML-ish tag system) for adding spice to comments. Sadly, though, the 'smiley' clickable

images only seem to appear as plain-text emoticons when viewing comments – no replacement is done.

Changing the visual theme involves hand-hacking the templates.html file, but it's easy enough to navigate around. Similarly, the articles and comments are stored in plain text files for easy editing, and there's even a swearword filtering system in place – ideal if family frequently visit your

blog, and you've become the target of a troll.

4est's Weblog tools may not provide the total flexibility that power-bloggers demand, but for most situations it's ideal, being small, fast, configurable and easy to use. If you're thinking of keeping a small blog, but want something simple yet easily manageable, it's definitely worth a trial run.



Is my boring life really blog-worthy? Answers on a postcard, please.

TEXT-MODE BROWSER

ELinks

■ VERSION 0.4.3-rc1 ■ WEB <http://elinks.or.cz>

Back in the early days of the World Wide Web, browser programmers had more freedom in rendering a page. , for instance, could be drawn as bold, underlined or even flashing – pretty much a total contrast to today's Web, where pixel-perfect rendering is crucial in image-drenched sites. Still, some programmers stuck to their guns, and *Lynx* has been one of the most popular purely text-based fringe browsers. The more recent *Links* put greater effort into table and frame rendering; now *ELinks* is an enhanced version.

Due to the author of *Links* being conservative with changes or updates, some coders gathered patches and began work on a more experimental codebase. *ELinks* takes the small but solid featureset of its reliable cousin, including the basic character-cell UI, FTP and SSL support, and adds all manner of extras: Lua scripting language

support, tweakable keybindings, HTTP authentication, proper cookies, and stacks of further fixes and touches.


ELinks is built against the *ncurses* text-drawing library, and will use the mouse at the console (GPM) or in an

xterm if required. As with *Links*, hitting F10 brings up the main menu, while S launches the bookmark manager. *ELinks'* comprehensive options box uses a handy tree of settings – virtually all configuration can be achieved inside the program itself. Navigation is done with the cursor and Ins/Del/PgUp/PgDn keys, a system which works well.

With speed being its main selling point, how fast actually is *ELinks*? We tested *Firebird* against it by rendering a 615K Slashdot page (offline, no images) in both: 1.8s for *ELinks*, 5.2s for *Firebird*

(P3 800). Impressive. Of course, its rendering isn't so precise, but we've been using *ELinks* on all sorts of sites and unless they require JavaScript, they're mostly fine.

ELinks maintains the simplicity and speed of *Links*, but improves upon the usability of its predecessor with lots of smart features. It's outrageously fast, very stable, easy to configure and satisfying to use. Naturally it isn't ideal for all sites, but for general reading and quick browsing jobs, it's highly recommended by us.



ELinks renders text-heavy sites like Slashdot with such immense speed it leaves even the likes of *Firebird* standing.

HTTP PROXY

Middleman

■ VERSION 1.5.3 ■ WEB <http://sourceforge.net/projects/middle-man>


While the Web can be full of annoyances, such as popup windows, adverts and unpleasant material, solutions exist to make life on the Net somewhat easier – *Firebird*, as reviewed this month, boasts plenty of features for this. Another option is to use a proxy server; this sits between your browser and the Web server you're connecting to, caching files for speed and performing operations on the incoming data, so it's highly useful for filtering out dodgy content if kids are using a machine. *Squid* is by far the best-known HTTP proxy, but *Middleman* is an interesting newcomer.

Middleman should compile without hassle – there are no esoteric dependencies, and the resulting file is just a single binary. Scripts are

supplied in the tarball for installing in a system's startup dirs, but the easiest way to get up and running is with `/mman -c config.xml`. As this suggests, the main configuration file is in XML format – it's readable but long, and it's much easier to configure the proxy through the browser itself.

Before a browser can be used, it needs to be configured for the proxy; most major browsers support this. *Eg* a browser on the same machine as the proxy would use `127.0.0.1:8080` for *Middleman*. Then, viewing `http://mman/` brings up the HTMLised interface – it's neatly separated into categories, with connection tracking and logging features.

Middleman supports HTTP/1.1, full caching of DNS info and HTTP-served content, filtering based on URL,



The web-based configuration system busy at work in Mozilla.

header and keyword, along with external parsers which can be hand-written. This flexibility is great: you can easily replace the User-Agent header to make stubborn sites cooperate, for example, and a bunch of notorious advert servers are blocked by default.

With a plethora of configuration choices, *Middleman* can be configured

to simply make the Web appear less annoying as you browse, or make it safer for children to surf. Equally, the caching functions noticeably improve performance (especially if there's several people sharing the same connection), and the config system is generally easy to work with and the included documentation is brilliant.

WEB PHOTO GALLERY CREATOR

Gallery

■ VERSION 1.3.3 ■ WEB <http://gallery.sourceforge.net>

Continuing the Web theme of this month's *Hot Picks*, *Gallery* is a utility to help in creating another type of popular site: the photo gallery. Thanks to the huge proliferation of cheap-but-decent digital cameras, and with

more easily available webspace than ever, it's the latest trend to have a vast online archive of snaps. Just uploading .jpegs and writing some HTML around them isn't feasible with a large picture collection – software like *Gallery* is essential to manage them properly.

Being written in PHP4, *Gallery* needs *Apache* for serving and the *NetPBM* tools for image manipulation, and it's also fairly simple to install. After placing the gallery directory somewhere in the DocumentRoot and running the `configure.sh` script, a browser needs to be pointed at the directory and the web-based setup system begins. It helpfully checks for all required components and makes sure `httpd.conf` is OK, and then moves on to general options. Very smart and friendly, even if you're not a particularly confident Linux user.

Once in place, the administrator account (or any other you add to the settings) can begin adding files to the album through the browser's file upload dialog, and a thumbnail and caption is then generated. Nested albums (ie an album within an album) can be created, while the slideshow option automatically steps through each photo within a defined time limit. Users can even add their own comments to each picture.

Gallery's default black-on-blue is easy on the eye, but the colours (or background image) can be defined through the properties box; similarly, pictures can be sorted by date, caption or number of views, and they can all be resized on-the-fly. Just about every option useful in a web gallery is present, and it's commendably easy to use.

All in all, *Gallery* excels in its simplicity and tweakability; it's ideal for running a small community site where others want to upload their snaps, and the online assistance in checking that PHP and other components are working correctly is top-notch. We'd like to see support for drop-in themes, but otherwise it's a solid and well-constructed package.



A simple album being worked on, with the image-upload window ready.

TYPING TUTOR

Tux Typing

■ VERSION 1.0.3 ■ WEB <http://tuxtype.sourceforge.net>

Remarkably, in face of revolutionary layouts and RSI-battling designs, computer keyboards have remained pretty much the same for the last few decades. Microsoft's fairly successful middle-split Natural model has been the most radical mainstream development – even so, it takes a great deal of use before one can become a fast typist. Recognising the small but vacant Linux market for 'edutainment', Tux4Kids has looked to the great Mavis Beacon for inspiration and has produced *Tux Typing*.

Buildable on Windows as well as the usual Unix flavours, *Tux Typing* needs only the virtually ubiquitous SDL libraries to compile. If you receive a missing image error after installing, just run it from the

`tuxtype/ dir` in the source. We've provided an RPM on the coverdisc as well, which should work with most recent distro releases.

Following the 'we-mean-business' splash screen, *Tux Typing's* main menu offers a choice of play modes: key cascade, word cascade or practice. A brief tutorial and demonstration explain how the simple but effective game engine works, and it's basically this: hungry Tux, lettered fish falling from the sky, and pressing the appropriate key moves Tux beneath them to eat. Exactly why they sound like crisp iceberg lettuce is never made clear, though.

For the first play mode, fish fall down individually, so it's a test of finding the right key. In the second, the fish are arranged into full words, and the user needs to type in the



Sad, sly pop pie? Even hideous green fish don't escape Tux's appetite.

whole word before Tux moves over for his grub. Apart from the limited few-frame animation of our favourite mascot, the letters are clear and the green fish never accidentally blend into the backdrop.

With several difficulty levels, some very picturesque postcard-from-Antarctica backdrops and a short but perfectly looped classical sample, *Tux*

Typing has been well thought-out and our test infant enjoyed it. Work is underway on version 2.0, and we hope that Tux4Kids' other play-and-learn projects (maths, drawing, alphabet etc.) are equally well-received. Linux's assault on the desktop needs a wide range of applications, and *Tux Typing* just moved us that little bit closer...

PUZZLE GAME

Frozen Bubble

■ VERSION 1.0.0 ■ WEB www.frozen-bubble.org

Helicopters, pinball machines, spaceships and battlefields – we've looked at a number of action games in the past few *Hot Picks*, so it's time for something more cerebral. If you remember the *LXF Awards 2002*, you may also recall a certain cutesy puzzler being runner-up in the New Games category – yes, it was *Frozen Bubble*, and with version 1.0.0 now out it definitely deserves a closer look. Regular gamers glancing at the screenshot will spot the main influence immediately: it's *Bust-A-Move*, but on Linux.

On the coverdisc we've provided RPMs for Red Hat and Mandrake, along with the source tarball. As it's built around the popular SDL media library you should encounter no major hassles when installing, but you'll also need the `SDL_Perl` bindings <http://sdlperl.org> on your box too.

`SDL_Perl`? It does seem odd that a interpreted language focused on text manipulation would fit nicely with a game toolkit, but Perl is undoubtedly powerful and works well here.

Frozen Bubble's well-crafted cuteness remains true to the *Bust-a-Move* spirit; colourful little penguins abound, although at least the background tunes aren't filled with overbearingly hyperactive Japanese children's voices. Indeed, *FB's* in-game themes are strangely unfitting – they're reminiscent of a point-'n'-click futuristic RPG. Still, the visual polish is applaudably good, with crisply drawn sprites and attractive static images.

FB's game mechanics are based on a kinda inverted-Columns system – three or more touching spheres of the same colour will disappear on contact, but instead of rotating stacks, you have to fire individual balls at a pre-built



Cripes! That lone green one is going to make life hell later on...

layout (even bouncing them off the walls). It's a very simple and, on later levels, wall-headbuttingly frustrating idea. Point your arrow with the cursor keys, use Up to fire, miss your target and watch your problems grow. And then the roof gets lower as time passes...

With 100 1-player levels (and a thoroughly respectable level editor to do a bit of DIY), along with a 2-player mode, *Frozen Bubble* is ready to eat up hours like they're seconds. Silly, pretty, tremendously playable and fun – splendid stuff. **LXF**

Become a NETWORK GURU!

cover feature



If you've ever wondered what makes some of the largest networks in the world tick, **David Coulson's** look at Ethernet networking will explain...

Almost everyone will have looked at creating a network at some point, even if it is as simple as needing to connect two boxes to share the Internet connection. Whether it be a wired LAN or wireless, chances are it will utilise IEEE 802.3, more commonly known as Ethernet. While there are many Ethernet standards out there with varying throughput levels, the underlying technology is consistent throughout and Ethernet networking is particularly compelling as there is far more to understand than simply putting a cable into the hole that fits.

Even if an individual's networking experience is limited to stringing a crossover cable between two boxes, or building a basic home network, there are numerous advanced features of Ethernet networking which can be done using Linux. Just as how a Linux system can replace an IP router and firewall, it can also be a direct replacement for many components in an Ethernet network, although hardware requirements may often make this impractical. One rarely has enough PCI slots to create a 24-port switch out of an Athlon system and a selection of 4-port network interface cards.



“If you buy an Ethernet device in England or Germany, it will work with other Ethernet devices in Japan, the USA or Africa.”

www.linuxformat.co.uk



A brief history of Ethernet

The first Ethernet packet was transmitted at Xerox Palo Alto Research Center (PARC) back in 1973. The devices were invented by Bob Metcalf and David Boggs, who by 1976 had further refined their lab network to support in excess of one hundred devices. The original specification for their task was to create a network allowing hundreds of systems to be connected, in addition to providing bandwidth capable of making use of the (then) new-fangled laser printer. One should bear in mind that at this point half-a-dozen computer systems, much less one hundred, had not been in the same building before, much less connected together.

The first commercial interest in Ethernet was by Digital Equipment Corp. (DEC) in 1979 and, along with Xerox and Intel Corp. wanted to create a commercial product from Ethernet. The original idea was that these three companies would control Ethernet and anyone who wanted to use Ethernet would have to purchase hardware from them. As Microsoft have found out over recent years, the US government is not always supportive of monopolies, and rather than face the prospect of anti-trust lawsuits at a later date, the decision was made to work with the Institute of Electronics and Electrical Engineers (IEEE) to create an open industry standard for Ethernet. The IEEE is a non-profit group to allow open standards within the Electronic community.

The first commercial Ethernet device was released by 3com in 1981, following the creation of the company by Metcalf in 1979, and was the first use of Ethernet out of the lab. The next year 3com released the first Ethernet Network Interface Card for a personal computer and little over a year later the Etherlink ISA adaptor was released for the IBM PC. Both of these devices used Thin Ethernet over co-axial.

By 1983 the IEEE released the 802.3 specification for Ethernet and the patents for Ethernet were given to the IEEE, allowing anyone to obtain the specifications to build and sell their own Ethernet devices for \$1000. Further to the IEEE specification,


```

david@macha:~ (pts/17)
macha:~# ifconfig eth0
eth0
Link encap:Ethernet HWaddr 08:00:46:18:05:04
inet6 addr: fe80::a00:46ff:fe18:a604/64 Scope:Link
UP BROADCAST MULTICAST MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:14 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:100
RX bytes:0 (0.0 b) TX bytes:1396 (1.3 KiB)
Interrupt:9 Base address:0xa000

macha:~#

```

ifconfig shows the basic configuration of an interface, even if it's not Ethernet, but we can use it to find out the MAC.



The original diagram to describe Ethernet. While Ethernet has grown from the co-axial-based 10Base2, the actual functionality remains identical.

Ethernet obtained international approval through the International Organization for Standards (ISO) in 1989 making it a global standard. If you buy an Ethernet device in England or Germany, it will work with other Ethernet devices in Japan, the US or Africa. Having such a great international standard allowed Ethernet to become the basis for the Internet through the ease of integration between existing networks.

Practical Ethernet

As exciting as the history of Ethernet is, it has little impact on the ease of getting a real Ethernet network up and running in the year 2003. To start with the basics, every system that needs to be connected to the Ethernet network requires a Network Interface Card, or NIC. NICs are available for practically every system out there, from Apples and IBM-compatible systems, to high-end workstation and server equipment including Sun and DEC. With the advent of Cable and DSL for Internet connectivity, an increasing number of systems come complete with an Ethernet NIC and the vast majority of motherboards for Athlon and Pentium

systems now come with on-board Ethernet. Even if your system does not have Ethernet already, it can be added with a simple PCI card that can cost as little as £10 (or less if you know where to look). For older systems, ISA network cards are also available although they will be unable to get close to 100Mbit on newer networks due to bus limitations.

PCMCIA network cards are also available for notebook computers and PDAs with expansion slots, and just as with everything small and mobile, they cost significantly more than the PCI cards available for regular systems. Another option for laptops and PDAs is a USB-based network interface card. While USB1 only supports 15Mbit per device, USB2 can easily handle the 100Mbit from a current network and a number of manufacturers have USB NICs on the market, a number of which work perfectly with Linux.


Almost all current Ethernet cards will use either 10BaseT or 100BaseT running at 10Mbit and 100Mbit respectively. Earlier network devices utilised the 10Base2 standard over co-axial, but this was limited to 10Mbit and two metres between devices, so is reserved for legacy systems now. However, as all of these are 802.3 Ethernet they can communicate between each other through a device which supports the required standards.

With two devices, one can simply connect them with a crossover Ethernet cable, which is used to connect two similar devices. A crossover cable is a standard 4-pair Cat5 cable where the RX on one end


is connected to the TX on the other, and visa versa. This is the simplest Ethernet network one can create. With a 10Base2 network, as it is a continuous chain, the physical structure of a two node network is the same as a fifty node network, as it still requires two terminators and two T-pieces to connect it all together. If one attempts to connect two 10Base2 cards directly with co-axial, then functionality should not be an expectation.

The alternative to a crossover cable is to use a simple hub-based network and straight-through cables. A straight-through Cat5 cable has both ends wired exactly the same way and is used to connect dissimilar devices. Each system would have an individual straight through cable between its NIC and the hub, as opposed to a 10Base2 network where there is a continual chain structure. A hub is the most basic method of creating a three-node Ethernet network, as all it does is receive an Ethernet frame on a port and broadcast it out all of the other ports. It's up to the individual systems to decide which frames they should process and which they should not.

Ethernet RJ45 plug – while it looks like a standard connection used by modems and phones, it is actually slightly larger and has eight pins instead of six.



Each Ethernet node device has a Media Access Control address, otherwise known as its MAC. This is a unique number for the device on the local Ethernet network, which is assigned by the manufacturer. When one system wants to send an Ethernet frame to another it will create the Ethernet frame containing the destination MAC address, as well as its own as the source MAC. A hub does not have a MAC address, as it does not accept or send any packets, just blindly relays them between ports. Before an Ethernet node will send a packet it will 'listen' on the media to find out if another device is sending a packet. If there is another packet on the network it will wait for a while, otherwise it will send it immediately. It is quite possible for two devices to send packets at the same time, assuming the network is clear of traffic, which causes a collision. When a collision occurs, the devices which sent the packet will see nonsense coming back and will keep quiet for a reasonably random period of time, wait for the network to be empty, then try to send it again. On a fairly low-traffic network collisions occur, but they are not really a major problem.



However, once traffic gets heavier then collisions can actually cause a significant effect to the overall throughput of the network.

To avoid collisions, we need to split up the collision domain. The collision domain is the area of the network where two devices talking at once will cause a collision. By chopping this in half we can reduce the number of collisions by fifty percent, making far more bandwidth available for useful traffic. We can do this by replacing our hub with two hubs and a bridge. A bridge forwards Ethernet frames between two collision domains and two devices can talk at the same time, assuming they are not on the same side of the bridge. Essentially, the bridge will 'hear' traffic on both sides of the network and will learn what devices are available on which side of its connection. If it knows which side the packet is supposed to be heading to it will either forward the frame

The Linux kernel supports tens, if not hundreds, of Ethernet NIC chipsets, but www.google.com is useful in finding which one works with your NIC.

“Even if your system does not have Ethernet already, it can be added with a simple PCI card for as little as £10.”

between the two collision domains, or it will simply drop it.

A bridge is critical to making the most use of the available bandwidth on the system. Rather than using a bridge along with hubs, the modern equivalent of a bridge is a switch. A switch is a bridge with a large number of ports, allowing each system connected to it to function as if it was on its own collision domain. With a





switch, each port can make use of the maximum amount of bandwidth available, rather than having to fight with collisions for the bandwidth. All production networks are built now with switches and given the current cost of switches vs hubs it's hardly worth using a hub at all anymore.

A standard Ethernet network allows for 100m between switches and 500m overall. Segments are connected with either hubs or switches, although they obviously do not require a linear topology. All that is needed is that any two systems are less than 100m from each of the hubs or switches they are connected to, and the two systems furthest apart on the overall network are less than 500m apart. We can extend the network further using various methods, but if we have a really huge Ethernet network the signal strength will be reduced over the long distances involved, plus the system will have to wait a rather long time period to actually receive the packet sent to it.

Ethernet and Linux

Obviously Ethernet isn't a whole lot of use without an operating system which supports both the Ethernet specifications, as well as devices which allow for connectivity to the physical Ethernet network. Linux has support for Ethernet networking 'out of the box', and supports the vast majority of Network Interface Cards on the market. Ethernet device support can be either compiled into the kernel, or made available as a loadable kernel module. Compiling the kernel module is done in the standard way by issuing:

```
# make menuconfig
# make dep clean
# make blmage
# make modules
# make modules_install
```

All Ethernet devices are available under *Network device support* and *Ethernet (10 or 100Mbit)*, and there certainly are quite a few of them in the list. Devices are listed by the specific chipset used by the NIC, rather than the NIC manufacturer and model number, so it's well worth checking the box or physically looking at the card before attempting to compile in support for the device. Most distributions come complete with

support for most NIC devices, but if they are only available as a module then they will not help in figuring out which one is needed for the NIC in the system.

Network devices are registered with the kernel by the driver to create an *ethX* device, where X is an integer which increments with each detected NIC. With drivers compiled into the kernel, the *eth* interfaces are allocated based upon the order the drivers were compiled in, but in the case of a module, they are allocated in the order the modules are loaded. Assuming the modules are always loaded in the same order, then the NICs will be allocated the same *eth* interface all the time. We can also use the *nameif* utility to name NIC interfaces based on the MAC address, so that the NIC always appears in Linux with exactly the same address, independent of the order of module loading. We can give the interface any

name, so our external NIC connected to our upstream provider can be called *inet0* and our internal network connectivity can be called *lan0*. This is all done with the */etc/mactab* file which associates the MAC with the interface:


```
# cat /etc/mactab
inet0 08:00:46:18:A6:04
lan0 00:FF:13:A5:6C:D7
```

Running *nameif* at anytime will rename the NIC interfaces within Linux, although it's best to have it run by the *ifup* script so that it correctly renames interfaces when they are brought up.


The *ifconfig* command is used to configure the NIC and a simple *ifconfig eth0* will show some useful information on *eth0*.

```
eth0 Link encap:Ethernet HWaddr
08:00:46:18:A6:04
    inet6 addr:
    fe80::a00:46ff:fe18:a604/64
    Scope:Link
```

When we don't know what eth device a NIC is going to get, we can use *nameif* to allocate interface names based on MAC address



Etherapp helps determine what systems are causing traffic on the network, making it easier to establish connectivity issues or compromised systems.



Snort is a wonderful intrusion detection system, which can be thrown on a box attached to a hub making the interception of traffic transparent.

```
UP BROADCAST MULTICAST
MTU:1500 Metric:1
RX packets:0 errors:0
dropped:0 overruns:0 frame:0
TX packets:0 errors:0
dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:100
RX bytes:0 (0.0 b) TX bytes:0
(0.0b)
Interrupt:9 Base address:0xa000
```

The most important part of the current output from **ifconfig** is the **HWaddr** value, which contains the MAC of the NIC. Should the eth interfaces not be useful when deciding which NIC is which, the **Hwaddr** value will aid in identifying the NIC, as the MAC is generally labelled on the chipset or on the panel next to the jack. With **ifconfig** we can give the device an IP address, bring it up and various other things. To 'bring up' an interface simply means that the kernel will begin accepting frames on that interface, rather than the actual physical Ethernet link will start functioning. Generally an Ethernet device will immediately start trying to get a link with whatever it connected to, if anything, when either the system is booted and the NIC is powered, or when the kernel detects the device and tells it do something useful.

Ethernet Debugging

Once we have our Linux box connected to the Ethernet network, we can do a number of useful things with it, the most basic of which is packet sniffing. As standard, the kernel will drop any Ethernet frames that are not destined to the MAC of

the local Ethernet NIC, so we won't see packets being sent to other systems even if we are on the same Ethernet collision domain. When connected to a switch we only see our traffic anyway, but on a hub we get everyone's traffic broadcasted to us so we can easily find out what is happening on the network.

To sniff the traffic on the network we need to first tell our kernel to accept all frames coming in from the network, even if they are not being sent to us. This is called 'promiscuous mode' for obvious reasons, and is setup with **ifconfig**:

```
# ifconfig eth0 promisc up
```

All we have to do now is run some software that will interpret the packets coming in from our network and give them to us in a halfway useful form. A common packet sniffer is the **tcpdump** program, which will extract the IP information from an Ethernet frame and tell us which IP address and port the frame is going to and where it comes from.

A slightly more graphical variant of this is **Etherape**, which draws nice pictures of the traffic on the network, showing which specific protocols are passing traffic between systems, the hosts they are going to and from, as well as allowing us to pull the data out of the frames so we can intercept the traffic on the network. As one would expect, if we ever forget our POP3 password or want to know what passwords people are trying to use via telnet to get into our router or switch, then we can sniff the traffic from our network and figure it out. >>>

User-space Ethernet

You don't need masses of Cat5 and switches to build a network




The OpenVPN tunnelling daemon is one of many user-space processes that use TUN/TAP to have the same functionality as an Ethernet device.

While Ethernet is most commonly used over Cat5 or co-axial cabling between systems, it is perfectly possible to utilise the advantages of Ethernet networking within user-space, either within a single Linux host or between hosts over another transport mechanism.

User-space Ethernet functionality is provided by the TUN/TAP kernel module, which provides Ethernet and IP virtual devices which pipe the frames into userspace via the **/dev/net/tun** device. The most common use for such a system is for Ethernet or IP tunnels and the TUN/TAP driver was written specifically for the **vtund tunnelling daemon** (<http://vtun.sf.net/>). Originally the Ethertap module was used for similar networking, but this was replaced with the TUN/TAP device in the 2.4 kernel. Once one has registered a TAP device, the device is available via **ifconfig** as if it was a regular network device like a NIC. We can give it an IPv4 address, include it in a bridge, firewall it, and indeed do pretty much everything else you can do to the traditional **eth0** device. However,

something needs to be running in user-space to make use of the frames flying out of the **/dev/net/tun** device, otherwise they will simply disappear into nothingness. In the case of the **vtund** system, the frames are sent over either a TCP/UDP, PPP or character-based connection to a remote **vtund** server and leave from the appropriate TAP device on the other end.

Another popular use of the TUN/TAP device is with the ever-popular User-mode Linux (<http://user-mode-linux.sf.net/>) kernel patch. To provide Ethernet connectivity to the user-space Linux instance, UML can use a TAP device to communicate with the host, which will either route the packets out onto the network, or can create a simple bridge containing the host's **eth0** device and the tap devices belonging to the UMLs, so they can directly communicate with the LAN and even use LAN services, such as DHCP. While UML does support other methods of network connectivity, the TUN/TAP device is the only one that allows the UML to directly access the LAN which the host system is connected to.



User-mode Linux relies on the TUN/TAP transport to allow the user-space kernel to connect to physical Ethernet devices via a bridge.

Networking



As everyone in the network field knows, the top of the list is our good friend *Snort* (<http://snort.org/>). *Snort* will intercept traffic on an IP network, look for buffer overflows in various protocols, check for CGI exploits, portscans and numerous other things to improve the security of the network. *Snort* is a great tool to put on a firewall, or have listening on a NIC which is connected to the same segment as a router, as it gives the opportunity to check Ethernet packets

```
david@macha:~ (pts/17)
macha:~# mii-tool
eth0: no link
macha:~# mii-tool
eth0: negotiated 100baseTx-FD flow-control, link ok
macha:~#
```

***mii-tool* can be used to find out the status of the MII on a NIC or to have it attempt to negotiate a particular speed and duplex with the other end.**

externally from the machine which will be receiving the packet. With a box setup with *Snort* without an IP

address and with a promiscuous *eth0* device, it can perform traffic analysis totally transparently. For more on *Snort*, see this month's *Linux Pro*.

“For dealing with more complex network topologies, Ethernet is a wonderfully customisable and adaptable standard.”

Advanced Ethernet Configuration

While *ifconfig* is useful, it doesn't tell us a whole lot about our NIC. It'll tell us about the IP configuration and specific details of the NIC status with

Ethernet Bridging

You don't need to buy hardware to create a bridge if you've an old 486 or Pentium system with a couple of NICs

While most people think of a switch when bridges are discussed, it is possible to make use of a simple bridge with only a couple of ports to connect a number of collision domains together. A common use of a simple bridge is to connect a hub to a router, or to link a cable modem to an internal LAN.

The Linux kernel supports bridging using the *bridge.o* module. Once this module is either loaded or compiled into the kernel, we can start to create our Ethernet bridge using the standard NICs in the system. Using the NICs for a bridge does not neglect the need to have the correct kernel drivers available, as the bridge functions above the actual driver layer.

We can start to create a simple bridge using the *brctl* command. This can allow us to create our bridge and allocate network devices to it as ports. To make the bridge device we do:

```
# brctl addbr br0
# brctl stp br0 0
# brctl setfd br0 0
# brctl sethello br0 0
```

This will add a *br0* interface to our *ifconfig -a* output, along with a *br0* device in our *brctl show* output. The *br0* network interface serves no other purpose other than to allow the Linux box that is providing the bridge to function as a node on that link-level network. To actually make the bridge do something we need to add some ports to it, and we can use regular *eth* devices as ports. We start by bringing up *eth0* and *eth1* in promiscuous mode without an IP address and add them to our bridge:


```
# ifconfig eth0 up promisc
# ifconfig eth1 up promisc
# brctl addif br0 eth0
# brctl addif br0 eth1
```

The bridge will not start to forward frames between the two ports until we bring it up:

```
# ifconfig br0 up
```

Now the bridge is up and running, we now have a little two port switch running on our Linux system. There is no limit to the number of ports on a bridge within Linux, although we are obviously limited by the number of NICs we can fit into one box. One can also include tap devices, used for user-space networking, into the bridge, allowing a user-space application to communicate with our LAN at the link level.

The bridge support does add some extra capabilities that are more commonly found on high-end switches. The most useful of these is the Spanning Tree Protocol, which is an inter-bridge communication method to maintain a functional bridge topology. When we have switches connected together we can easily create a loop, causing the packets to be passed continuously from one switch to another. STP solves this by allowing the bridges to figure out who is connected to what and to stop some bridges from forwarding frames. The result is that some switches may not forward packets between certain ports as there is a higher priority route via another switch somewhere else. The side effect of STP is that it provides automatic failover capabilities for an Ethernet network, as when we disconnect one



With a quad-NIC system, creating a redundant bridge configuration is child's play.

bridge from the network, the others will communicate with each other and rearrange their forwarding configuration to allow frames to flow between all of the segments.

With two Linux systems, each with four NICs, it is fairly trivial to set up a resilient uplink to the Internet, as if either box fails, or one of the NICs is disconnected, the bridge will switch another port or bridge into forwarding mode to allow continual connectivity. The *sethello* option in *brctl* tells the bridge how frequently it should send out Bridge Protocol Data Unit (BPDU)

packets to all other bridges on the network. Along with *sethello*, *setfd* specifies how long the bridge should wait before forwarding packets, as obviously we don't want it forwarding packets if the other bridges have not yet decided if they need to be in the forwarding or blocked state. Depending upon the complexity of the network which the bridges are connected to, along with the amount of other data on the network, it is possible to create a network configuration which will very quickly fail-over between NICs and bridges when a failure occurs.

the kernel, but in terms of the actual link from the NIC to the hub or switch, **ifconfig** is completely oblivious.

To check the status of the NIC at a lower level than the kernel, we use a tool called *mii-tool*. This will query the Media Independent Interface, or MII, on the NIC, which allows the kernel to talk to the NIC in the same way independent of the actual link it has to the outside world. An Ethernet 'link' is when two connected devices negotiate the best Ethernet standard to use, based on the capabilities of the devices at either end. There are a number of different standards, such as 100BaseT-FD, 10BaseT-HD and so forth. Essentially each standard has a maximum throughput in Mbit, as well as a duplex status. The duplex status of a NIC can either be half or full, meaning that the NIC can operate in one of two ways. A half-duplex device can only send or receive at any time, which is what one will experience when connected to a hub. In full-duplex mode, the NIC can both send and receive frames at the same time, which is certainly preferable on a high-end network. It's not uncommon for two devices to get confused, so one functions as half-duplex and the other full-duplex, which can cause problems. Using *mii-tool* we can force a particular mode at our end of the link. Lots more information can be found in **man mii-tool**.

Conclusion

While Ethernet networking may look simple on the outside, once one starts dealing with more complex network topologies it spirals into a wonderfully customisable and adaptable network standard. There is certainly far more to Ethernet than we have had room to cover in this short article, including VLANs, Cost of Service and bridge-based frame filtering with *eatables* (<http://eatables.sf.net/>). See this month's *Linux Pro* for more on *eatables* too. The simplest of Ethernet networks can be used to experiment with some of the huge range of the technology's capabilities without investing in high-end hardware.

More info on the basics of Ethernet under Linux, along with example usage involving IP networking, can be found online at www.tldp.org/HOWTO/Net-HOWTO/. 

Glossary of Ethernet Terminology

Know your MAC from your VLAN

ARP The Address Resolution Protocol is used by IP networks to associate the IP address with the MAC address on a system. When an IP device wants to send a packet over an Ethernet link, it needs to know which MAC address to send the frame to. To do this, the IP device will broadcast an ARP packet to the network, and the system with the appropriate IP will respond with another ARP packet. Once the system receives the response it will cache the result for a period of time to limit traffic on the network.

BRIDGE If you've got multiple Ethernet segments which need to share traffic, then a bridge will join them together. Linux supports bridging using the 'bridge.o' kernel module and the *brctl* utility. <http://bridge.sf.net/>

Cat5 Category 5 is the current specification for the cable used to run 100BaseT Ethernet. While Cat5 contains four pairs, only two are actually used by the 100BaseT standard, although 1000BaseT4 requires all four pairs to be wired correctly.

COLLISION If two systems send out Ethernet frames in the same collision domain, these frames will 'collide' and the senders will send a jam signal across the domain to keep the other devices quiet while it sends its frame. As one would expect, a collision will occur frequently on a hub and reduces throughput.

ETHERNET The IEEE 802.3 standard for Ethernet is the most common form of networking used for LANs within educational establishments, corporations and private homes. There are a number of different wiring standards used for Ethernet, including 10Base2 over co-axial, 10Base-T over Cat5, 10Base-F over fibre, 100Base-T over Cat5 as well as Gigabit Ethernet over either fibre or Cat5e/Cat6.

HUB The cheapest way to link multiple Ethernet links together is to use a hub, otherwise known as a repeater. A hub simply takes a packet in through one of its ports and broadcasts it out of all the others. Unfortunately, this causes collisions, reducing throughput. All devices connected to a hub are limited to half-duplex mode, limiting them to sending or receiving, but not both at the same time.

LINK When two Ethernet devices are connected they will negotiate a link, so that they are both talking at the same speed and are set at either full-duplex or half-duplex. Once this negotiation is complete, the little light will come on at each end.

MAC Each Ethernet device is identified by its unique Media Access Control, or MAC, address which is a 48bit value written in the form 00:FF:13:A5:6C:D7. The first three 8bit segments identify the manufacturer and the last three the individual device.

MII The Media Independent Interface is used by network cards to negotiate the capabilities of the link, including speed and duplex settings. This component can be controlled on a Linux system using the *mii-tool* utility to force a particular speed/duplex or to check link status of a device.

RJ45 10Base-T and 100Base-T standards, as well as Gigabit Ethernet over copper, which use Cat5 cabling have a small plug of the type RJ45. While the RJ45 plug and socket look very similar to those used by modems and telephones for connecting devices, RJ45 has 8 pins and is slightly larger.


SEGMENT When multiple Ethernet devices are connected, the entire network is known as a 'segment'. Generally a segment is defined as being an Ethernet broadcast domain, in that ARP requests and other broadcast traffic does not go any further. An Ethernet segment is usually attached to an IP router, or created by the use of a VLAN on a switch.

STP Spanning Tree Protocol is used by high-end switches or bridges, so they can communicate and avoid loops in the network. If 3 Ethernet switches are connected in a loop, they will negotiate with each other and one will stop forwarding frames between the two ports which have the other switches connected. STP is a very useful protocol, as when the network topology changes, it will automatically calculate the new topology of the network allowing for topologies which support fail-over to be created.

SWITCH Most Ethernet devices are connected through a Switch, which utilities transparent bridging. As traffic passed between the ports on a switch, it will learn which MAC addresses reside on which port, so that it does not have to send the frame to each port as a switch does.

VLAN When one can't afford lots of switches, it is possible to segment a single switch into multiple broadcast domains using VLANs. Each port on a switch belongs to a particular VLAN and when a port receives a broadcast frame, or the switch does not know the MAC of the device to send the frame to, it will only broadcast it out of the other ports in the same VLAN.

WIRELESS Ethernet will happily operate over Cat5 or co-axial cable, but can also be used over the airwaves, using the 2.4GHz band which is unrestricted. There are a number of wireless Ethernet standards operating between 2Mbit and 54Mbit, although they all operate on a similar principle to wired Ethernet. Of course, every packet from a wireless network is broadcasted out into the open, so appropriate security measures must be in place when using wireless networking technologies.



Opteron



The Opterons have landed

After a long wait, **Paul Hudson** finally has two Opteron machines under his desk. *Linux Format* presents: the UK's first in-depth Opteron-on-Linux review...

BUYER INFO

Two 1U 64-bit servers emerge to dominate a new market and leave the competition trailing in the dust.

- **SUPPLIER** RackSaver
- **PRICE** Price: \$4000 for the 242 machine, and \$2500 for the 240 machine.
- **WEB** www.racksaver.com

Despite launching back on April 22nd, not all that many Opteron servers have managed to sell so far. This isn't because they're bad, mind you – simply because supplies have been short, as usual, particularly because many of the early releases go to the media for review purposes. And so here we are – RackSaver, the company who made such a storm by producing the highest-performing x86 servers around when Opteron launched – has been kind enough to send us two machines from its new Opteron range. Opterons, it seems, are like buses and Little Chef restaurants. The first machine is an Opteron 240 (that's 2 x 1.4GHz CPUs for non-AMD marketing

droids), 2GB ECC RAM, and a 40GB hard drive, and the second machine is an Opteron 242 (2 x 1.6GHz) with 4GB ECC RAM, and the same 40GB hard drive.

Both machines came with SLES8 / UnitedLinux pre-installed, however the installation certainly wasn't done very well, as the system was a *mess*. `/usr/bin/gcc` or even `/usr/bin/cc` didn't work, as, for some reason, `/usr/bin/cpp` was being used – naturally this managed to screw up most default compilations. To make things worse, 80% of the standard include files were missing, which meant that even after we'd sorted out the *gcc* mess, we still couldn't compile anything. As we had no CDs supplied, and the machine was set up to read from an

NFS store we didn't have, we were unable to use *YaST* to correct the problem. To be fair, these *were* review machines – if you buy one of these, they will almost certainly come properly configured. We got in touch with our friends at SuSE, who were more than happy to provide us with some SLES 8 for Opteron CDs.

Performance

Normally the first thing we do is to go into all sorts of detail about the specifications of review machines, but just this once it's important to make an exception – I think everyone reading this wants to know how well Opterons perform!

The first test we ran was the *GCC* test, because this is where the key changes between x86-32 and x86-64 lie. As you may know, we use a compile of kernel 2.4.19 in our tests to get our benchmark figure, however, because that kernel version has no support for x86-64, we used 2.5.10 instead. 2.5.10 was chosen as it was the nearest in code size to 2.4.19, which, while not making it a direct comparison, is at least fairly close –



hyperthreading technology in the Xeons. As CPU scores scale linearly, users would need dual 3.4GHz Xeons to match Opterons running at less than half the speed. This incredible speed difference is again highlighted in the *ogenc* benchmark results, with the 1.6GHz Opteron returning 4.68 compared to our dual 2.8GHz machine from last issue returning “just” 4.04 – a 15% lead for AMD, again making their CPUs twice as fast as Intel’s Xeon, clock for clock.

So, on raw CPU performance it’s clear that Opteron is game, set, and match ahead of Intel. However, it doesn’t stop there – with 2GB and 4GB of PC2700 RAM behind them respectively, the 240 and 242 machines have more than enough capacity to load vast amounts of data into RAM and serve it up without touching the hard disk. This is reflected in the *Apache* benchmark, where the 1.6GHz CPU returned a result of 5.5 – that’s 5.5 times faster than our test machine, or 50% faster than the twin 2.8GHz Xeons. An incredible result, certainly, but it hides the only gremlin in this system – hard disk access.

The hard disk part of our test uses *hdparm* and *bonnie++* to measure overall hard drive performance. When just using *hdparm*, which tests buffered disk and disk cache reads only, returned a score of 3.19 – very respectable for a non-SCSI drive. However, the *bonnie++* test returned a measly score of 0.17, which is entirely unacceptable for departmental servers. We got in touch with both RackSaver and Arima (the makers of the motherboard) who did some research into the problem and said they would produce a fix for new

machines in the form of an updated BIOS, which, once installed, would correct the problem.

One minor peculiarity we noticed was in the MySQL test – the machine returned a score of 0.51, which is awful. However, when we took SLES off and put Red Hat Enterprise Linux WS on (a 32-bit OS) the score jumped to 2.99 – in line with what we would expect based upon the other scores. To add to the confusion, putting SuSE Pro Linux 8.2 (their 32-bit OS) on the machine caused the benchmark to *crash*. SuSE have yet to respond to what’s causing the problem, however it’s likely to be a fairly minor hardware support glitch.

Hotter than hot

The machines both come with two network interfaces at the back, plus a standard floppy and CD-ROM at the front. Of interest, though, is the lack of a serial port at the front, which means that administrators who fix faulty rackmounted servers such as these by plugging a serial cable into them are out of luck. However, this could be considered as just an overdue step in the right direction – serial connections aren’t really used all that much any more, particularly as SSH is so widespread now. With so few ports on this thing, it’s basically a power machine – not designed for easy connectivity, but it crunches numbers faster than a speeding bullet.

One particularly noticeable thing about these RackSaver machines is that they run incredibly hot – running one of them for 30 minutes or so made the top of the unit almost uncomfortable to the touch, which made us wonder how hot the machine would get if it were busy for a longer period of time, or indeed if there were more of them running at the same time.

Another very evident sign that the machines were turned on was the sheer amount of noise they make. With two fans on the front and another three to the right of the machines, there’s certainly a lot of noise when powered on. In fact, these machines were so loud that, with only one of them turned on, we even had *Mac Format* staff (the nearest magazine to us here at LXF Towers) complaining that the noise drowned out their iPod tunes!




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

RICHARD SEIBT, CEO
SUSE LINUX AG

if anything, it will be slower, owing to the new code added in the experimental branch. Let me remind you that last issue we reviewed a dual 2.8GHz Xeon machine, which returned a score 1.74 for compilation, and that this Opteron machine has dual 1.6GHz CPUs – hardly a fair comparison, you might think!

You’d be wrong: the 242 machine returned a compilation score of 2.14, which blows the Xeons out of the water by a margin of around 22%, despite the much-vaunted



Sales of Opterons haven’t been huge so far, but it is early days yet.



Opteron




Because these machines are destined for large racks out of the way of staff, noise certainly isn't an issue. However, I must admit to being a little worried about the general heat level a rack of these machines would put out – if you intend on deploying several of these machines in one room, you would be wise to get a powerful cooling system installed.

Docs and support

The two RackSaver machines came with nothing else in the box other than the usual rackmounting kit. No documentation, no CDs, not even an A4 sheet describing what the hardware actually was. This is quite a peculiarity, because RackSaver is known for having first-class support and help for customers. Again, this is likely to be down to the fact that these are review machines, but it would've been nice to have some documentation regardless.

Getting support from RackSaver was admittedly tricky because it doesn't have a UK office – when I sent emails off with questions, I would generally get responses the following morning or a day later. The problem lies in the fact that RackSaver's only office is in



In case you were wondering what all the physical parts of the Opteron chips actually do, this should make it fairly clear.

California, which puts it 8 hours behind us, and even more for those of you who are further east than the UK.

SLES 8

We reviewed SuSE Linux Enterprise Server 8 back in issue 40, and it scored 8 out of 10 – a great distro with a few niggles. At the time of writing, SLES 8 is the only working Opteron-optimised Linux distribution available – we've yet to be able to get Mandrake Corporate Server 2.1 for AMD64 to work, and Red Hat's effort is still a long way off. As such, there is little choice for an operating system if you want to take full advantage of the Opteron hardware, but it's far from a sorry state of affairs because, excluding the performance problems already mentioned, SLES 8 for Opteron is as good in every area as its IA-32 cousin.

Despite these computers using the very latest hardware throughout, SLES 8 detected everything without bother, and very little was different from SLES for x86 – other than the big jump forward in performance. Using the 64-bit GCC, everything compiled with optimisation for the Opteron, which means that for all intents and purposes the CPU is transparent. This reflects very

well on SLES as a product – you really can use it on whatever hardware you've got, from aged 486s to top of the line Opterons, using the same skill set.

Naturally we'd prefer to have other working Opteron OSes on the market, to provide competition if nothing else. With SuSE selling thousands of copies of SLES thanks to its excellent Opteron support, Red Hat must be kicking itself for being so slow off the starting blocks.

Conclusion

The Opterons we tested, despite being dragged back by sluggish hard disk performance before the BIOS patch, utterly crush the competition. Intel's Xeon CPU (its leading x86 workstation and server chip) doesn't come close to the Opteron when it comes to performance – our tests show that clock-for-clock the Opteron is twice as fast as the Xeon in many tests, and that's despite the fact that our Xeon test rig had a much higher-end storage solution behind it.

When it comes to price the situation is a little more confused. AMD has traditionally based its market position on "we're as good as Intel, and cheaper too". Naturally this mantra has been dropped with AMD64 – it's not too

Kernel 2.4.21

From Kernel 2.4.21 onwards, the AMD64 platform is now supported fully. There have been a large number of fixes and optimisations made to the AMD64 kernel since SLES 8 was released, which means that you're likely to get even better performance than we did if you load the new kernel.

What now for Opteron?

Market-leading products should be reaping rewards...

Are these RackSaver machines the beginning of a flood of Opteron servers? Well, with this stunning price/performance comparison, one would think so. However, here at LXF we're certainly not being inundated with offers of hardware, which points towards most suppliers still being a little hesitant about the platform. When RackSaver announced that its QuatreX-64 server had taken the lead in the TPCC 4-way scores despite coming in at under half the price of the nearest competitor, few seemed to notice. Even now that AMD's claims of supreme performance have been confirmed by independent tests,

there does still seem to be a notable lack in uptake from manufacturers.

Whether this is due to a continuing lack of parts from AMD, FUD being spread by AMD's competitors who realise they stand to be ousted from their most lucrative market, or something else – it's hard to say. While it's good to see that we're getting one Opteron review request for every one Xeon review request, the platform still deserves a lot more.

With Intel's Gallatin Xeon and the new Itanium 2 chip having been launched at the end of June, Intel has got one last chance to hold back the AMD64 tide – only time will tell how successful they will be.



Not many ports compared to some of the competitors, but you can always expand connectivity using some sort of external hub.

In line with expectations?

LXF double-checks the performance figures

In *LXF41* where we tried to estimate the performance improvement we'd see when switching to 64-bit. The conclusion was that the extra registers would likely give a 10-15% speed up, with even more being possible for computationally intensive code and code that can take advantage of SSE2 instructions.

At first, you might think "Those *Linux Format* monkeys were completely wrong – the Opteron is twice as fast as the Xeon, not 15% faster!" However, that's essentially comparing chalk and cheese – after all, the Opteron is largely the same as the Athlon processor. Remember that, although AMD may print its latest CPU as being 3200+, in reality it only runs at 2.2GHz, so it's only 400MHz faster than the fastest Opteron currently available – a difference of about 20%.

The performance difference between Opteron and Athlon is at its least where computationally intensive code is few and far between – particularly the *Apache* benchmarks. However, wherever the Opteron gets the chance to tear away – particularly in *MySQL*, where there's a lot of 64-bit integer mathematics – it surges forward.

We have to admit we were surprised when we totted up the figures at quite *how much* difference in speed the Opteron gives, particularly when concerning intensive code. The extra registers, the new flat memory model, the integrated memory controller, the innate 64-bit operation, and the other tweaks AMD made to the Opteron mean that this CPU isn't just a little faster than the next best thing – it's a lot faster, *and then some*.

smart to compete on price in a marketplace where the main requirement is quality. As such, the Opteron CPUs are quite a long way from cheap, but that's more than justified by the performance they earn you. At the time of writing we found 242 Opterons on sale for \$726 each (in quantities of one) with 3.0GHz Xeons on sale for \$688 each (also in quantities of one). As you can see from the results above, a 242 Opteron runs a great deal faster than a 3.0GHz Xeon, so when it comes to price/performance AMD still has Intel beaten hands down.

However, when it comes to the price of a complete machine, the difference becomes quite noticeable – the Sharq server we reviewed last month was £4640, whereas the cheaper of these two RackSaver machines we tested is \$2500. Once converted to sterling and import tax is added, the 240 machine is likely to come in a little over £2000, which means you can get two of the 240 machines for the price of one Sharq, and *still* have change left over,

despite the fact that just one of the 240 machines beats the Sharq server on everything but hard disk performance.

When it comes down to your choice of operating system, the way is very clearly led by SuSE Enterprise Linux 8 – SuSE have done a lot of work to make the system suitable for people at all technical levels whilst also making sure they take advantage of the new technology offered by the CPU. The RackSaver machines we tested came with a very poorly installed set up, but it's almost certain that that's because they were review machines – RackSaver has worked very closely with both Red Hat and SuSE to make sure that its machines are compatible.

Neither of these machines came with a high-performance storage solution, which means that you should consider other options if you want SCSI or RAID – here at *LXF* we'd be very interested to see these same machines in a 2U format, backed by a set of high-powered hard drives, as we think

there'd be another big performance jump. In fact, based upon these prices it seems quite likely that a 244 machine with a set of SCSI hard drives in a 2U machine will still come in at under £4000 including tax.

It's clear to see that Opteron is a superior platform right now, and not just because of its price. On the performance front, no x86-compatible machine comes close to Opteron, despite the fact that we didn't even review the most powerful Opterons on the market. Furthermore, the mere fact that Opteron is x86-compatible is a huge boon in itself, because there's an increasing number of proprietary programs for our OS which won't come out in an Opteron-optimised form just quite yet. We tried Red Hat Enterprise Linux WS on the machines (look for a full RH WS review next issue) to test out 32-bit performance and it was just as impressive as the 64-bit performance – these Opterons take whatever you throw at them, regardless of the bit-size.

AMD took a huge risk with Opteron, and it seems to have paid off better than any AMD exec could have hoped for. RackSaver, for its part, is showing full commitment to the new platform, and are producing machines that are set to take the world by storm. If you have recently splashed out £5000 on a new Xeon server, or, worse, £20,000 on a new Itanium server, don't worry – I imagine its bulk will prop doors open remarkably well. [LXF](#)

“On the performance front, no X-86-compatible machine comes close to Opteron, despite the fact that we didn't even review the most powerful Opterons on the market.”

BENCHMARK: 240

hd	1.66
apache	4.9
mysql	2.6
compile	1.88
oggenc	4.11
Overall	3.03

BENCHMARK: 242

hd	1.68
apache	5.5
mysql	2.99
compile	2.14
oggenc	4.68
Overall	3.4

VERDICT

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

Two very capable, very affordable machines. If you were thinking of buying a Xeon server, you'd be crazy not to reconsider and plump for these.

LINUX FORMAT RATING

9/10

What on Earth is... /proc?

Proc is simultaneously one of the most powerful and one of the most confusing aspects about Linux, but what is it, why is it useful, and what is its future? Paul Hudson gets all virtual and takes a look...

» OK, give me the lowdown – what is /proc?

If installed, and it probably is, **/proc** is a virtual directory on your system that contains information about your machine.

» A virtual filesystem? As opposed to what?

A virtual filesystem is one that doesn't really exist physically – the "files" in there aren't really files, which is why it's referred to as virtual. If you run the command **ls -l /proc**, you'll see most of the files in there don't actually appear to have any content – they're mostly 0 bytes. Furthermore, they have a date set to the current date and time, which means they're constantly changing.

» So the files in there aren't really files?

Well, in Unix *everything* is a file, so in **/proc** what we have is system information represented as files. One of the files in there, for example, is *cpuinfo* – it

appears to be 0 bytes in size, but if you run **cat /proc/cpuinfo** you'll see it will output a lot of information about your processor(s).

» cat /proc/cpuinfo just seems to give me a lot of numbers!

Well, yes – that's one of the downsides of using **/proc**! It is designed to provide raw data, and doesn't make much effort to format it all nicely. For **cpuinfo**, you'll see that there's information about the model number, actual speed in Megahertz, specific instructions supported by the chip, as well as other various bits and pieces. Think of **/proc** as containing information about your running system – it's not pretty, but it's very detailed.

» So *cpuinfo* contains information about my CPU – what other files are there?

There are lots of virtual files in **/proc** to provide all sorts of information about your PC – here's a quick list of the important ones:

/proc/apm contains information about the power management state of your computer, as well as version numbers and configuration options. The program *apm* presents the same information as **/proc/apm**, except it's nicely formatted, which makes it easier to understand.

/proc/filesystems lists the filesystems supported by your kernel, and this info is used by mount when no particular mount type is specified.

/proc/meminfo is as useful as **cpuinfo** – it contains lots of information (a lot of it is quite complex and can be ignored) about the status of your RAM and also your swap space. Combined with **cpuinfo**, these two are good to read when you're connecting to an unknown system, as they let you know what kind of system you're connected to.

Another very helpful file in **proc** is **modules**, which tells you what modules are loaded into your kernel, how big they are, and whether they will be automatically unloaded or not. Modules listed with information in square brackets show that a given module is dependent on others to load.



»» So what's in `/proc/sys/dev`?

In the `/proc/sys/dev` directory, you're likely to have a `cdrom` directory, which contains information on CD-ROM devices in your system. Look in there for a file, "info", and `cat` it – you'll see that it contains information about your CD-ROM devices such as how fast it is, whether it can write CDs, whether it can read DVDs, and such. Again, you'll find this particularly useful when you're connecting to an unknown system, but also very helpful for debugging purposes.

»» And `sys/kernel`?

The `/proc/sys/kernel` directory contains lots of information about your kernel, and most of it should really be left alone. If you really want to tweak stuff, you can echo a 1 into `ctrl-alt-del` (execute the command `echo 1 > /proc/sys/kernel/ctrl-alt-del` as root) to make the kernel reboot very quickly (but potentially dangerously!) when you press Ctrl-Alt-Del as root, or you can echo a 1 into `sysrq` to enable the System Request key on your keyboard. Generally `/proc/sys/kernel` is best left untouched, though.

»» I can guess that `sys/net` contains networking information!

Yes indeed, and there are quite a few files in there that you change easily to get helpful results. For example, echoing a 1 into `icmp_echo_ignore_all` will stop the kernel from replying to ping requests from other computers. Many of the other files are also just as useful in preventing attacks on your system, although care should be taken before making any changes – be sure you know what you're changing, and what effect it will have.

»» How long do these changes last?

As mentioned earlier, your `proc` directory, being virtual, gets wiped as soon as you reboot, which means any custom configuration changes you make are lost. However, there is a tool that makes reading and changing data in `proc` much easier, called `sysctl`.

»» What does `sysctl` do?

Well, for a start it makes changing values in `proc` much easier (and a little safer!) because you don't have to echo numbers into files any more – the syntax is a little clearer. Secondly, it can list all the information you're able to change in your `proc` directory, as well as current values. Finally, it can be given a list of changes to make on every boot up, which means that your custom `proc` configuration is restored when you reboot.

»» How do I use it to set variables in `proc`?

Use the `-w` parameter to set a variable, then provide `proc` with the name of the file you want to change, an equals sign, followed by a value in double quotes. So, to stop replying to pings, use this:

```
sysctl -w net.ipv4.icmp_echo_ignore_all="1"
```

Note that each comma-separated word is a subdirectory inside `proc`, and also that we don't include `sys` at the beginning because it's implied.

»» So that lets me change anything inside the `sys` subdirectory?

No, only things that *can* be changed – some variables, for example, require a recompile to be changed. You can get a complete list of what can be changed inside `proc`, as well as the current value of each key, by typing `sysctl -a`.

»» So I've got a list of all the changes I want to make permanent – how do I make `sysctl` redo these changes each time I reboot?

Chances are you're using a distro that has a `rc.sysinit` file in `etc/rc.d` (Mandrake/Red Hat). You can add to this script your own calls to `sysctl` to make it set up your custom configuration options at boot time. Be sure to check for any existing `sysctl` calls – you're likely to have several. Alternatively, your system might be set up to call `sysctl` with the `-p` parameter, which means it loads up settings from a file. This file is usually `/etc/sysctl.conf`, and contains the same list of changes that you would have entered from the command line (minus the `sysctl -w`). For Debian-like systems, check the file `procp.h` to see where `sysctl` is being called.

»» What are all the numbered directories for?

There are lots of directories in `proc` that have numbers as their name, and these directly correlate to a process you've got running.

»» Does that mean that each process I run has a directory of its own?

Precisely, and each process directory has lots of information about that process – inside each virtual directory for a process in `proc` you'll see a variety of virtual files, such as `cmdline`, `cwd`, `environ`, `exe`, and `status`. `cmdline` contains the command line used to start a given process; so, for example, as I'm writing a review of RackSaver Opteron servers right now (see page 54 of this issue), process 26286 on my machine is `Kate`, so running


```
cat /proc/26286/cmdline
```

will return:

```
kate /home/paul/lxf/43/rev_opteron/rev_opteron.txt
```

`cwd` is a link to the current working directory of the process. For my process **26286** the working directory is my home directory, so if I type

```
cd /proc/26286/cwd
```



I can see the contents of my home directory.

environ contains a list of environment variables set up for a given process, which is a remarkably clever thing. For example, I have *bash* running as root in process 25772, and it has its own individual set of environment variables – if I change a setting elsewhere, that copy of *bash* remains independent, and this is how it's accomplished. If I try and read **/proc/25772/environ** as a non-root user, I get a permission denied error, because it's impossible to “spy” on what the root user is doing. Trying to read the same file as root naturally works fine, and outputs a list of environment variables and values that process has used.

exe is a link to the executable itself, and **status** contains various information about each process, such as whether it's active or sleeping, which UID it's running as, etc. So, as you can see, there's a huge amount of process information available if you really want to find out what your system is up to.

» Why is self a symlink to a directory?

/proc/self is always a symlink to a process directory, and, if you open two shells and look at **/proc/self**,

you'll see it changes to reflect the process that's looking at it. What this means is that if you ever need a process to examine itself, it can just look at **/proc/self**, which will in turn point to the correct process directory.

» Is proc ever going to be replaced?

Hopefully, yes. Right now there's an awful lot of information in there that isn't terribly useful, or at least shouldn't be in there. In order to clean things up, a lot of work has gone into the new 2.5 kernel to shift a lot of the information from **proc** into a new **sysfs** filesystem that would work much like **proc** does, except would also include much more detailed (and more organised) information. The goal is that **sysfs** will be compiled into the kernel by default, making it as ubiquitous as **proc**. If you haven't looked at kernel development for some time, you might remember **sysfs** as *driverfs*, its original name.

» If proc is going to be replaced, is it still worth spending the time learning about it?

Definitely! **proc** won't be replaced entirely, at least not in the foreseeable future. Current plans have

The best online reference to proc is the Red Hat 9 manual – be sure to check it out if you want to learn more.

Learn more about proc

Information about the ins and outs of **proc** is remarkably thin on the ground, but perhaps the best reference is in the Red Hat manual. If you don't have a copy close to hand, you can read it online at www.redhat.com/docs/manuals/linux/RHL-9-Manual/ref-guide/ch-proc.html. *Linux Gazette* also touched on **proc** very lightly some time ago, and you can read what information is archived online about it at <http://www.linuxgazette.com/issue46/fink.html>. Or Google may turn up something new for you...

proc remaining to take care of process information at the very least, so the numbered directories will remain. How much stuff – if any – is taken out for good has yet to be firmly decided, and may differ amongst distributions. All you learn about **proc** will be applicable elsewhere and help your overall understanding of Linux, and it's a good way to see exactly how your system is functioning. **LXF**

Tutorials >>>

Our experts offer help and opinions on a whole host of Linux applications

YOUR GUIDE TO GETTING THINGS DONE!

Whether you are just starting out in Linux, or an experienced veteran, there's always more to learn. Every issue of *Linux Format* is packed full of practical advice, and nowhere is it more concentrated than in our tutorials pages.

Here you'll find expert guides to all sorts of things, from Basic Linux usage to understanding and deploying network solutions, from simple script coding to the complexities of Perl regular expressions, Java server apps and more. We aim to bring a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact us by post, or at linuxformat@futurenet.co.uk or log on to our website and post your suggestions in our special forums? (www.linuxformat.co.uk). Hope to hear from you soon!

Nick Veitch EDITOR

HOW CODE IS REPRESENTED

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap
between lines:
begin
mniWordWrap.Checked := false
end;
```

Usually, you'll find the code on our CD/DVD too.

THIS MONTH TEACH YOURSELF...

Beginners' guide

Last month you chose your distro – now it's time to see whether your hardware is up to it **p64**

Linux and Windows

Want to run Linux alongside, on top, or instead of MS Windows? Here's the novices' HOWTO **p66**

GnuCash

If you're missing *Quicken* or *MS Money*, this is the Free financial management answer **p72**

Samba

Bridges the gap between Linux and MS Windows for file and print services. Know your **smbd** from your **nmbd** **p76**

Blender >>

Making your models more realistic by adding a range of textures to their surfaces **p80**




ILLUSTRATION COURTESY - WWW.BLENDER3D.ORG

The GIMP

Selections are a fundamental to every art application – here's how they work in *The GIMP* **p84**

PHP

Making an extension to handle TAR files – the conclusion; plus news of the new iteration **p88**

TIP OF THE MONTH! LOSE WEEKS OF YOUR LIFE

All too often, us Linux users are thought of as being geeks with no interest in games or having fun, but that just couldn't be much further from reality, right? The truth is that Linux geeks like to play games as much as anyone else, if not more – the difference is that we often eschew FP-shootery and prefer “classic” games such as *Net Hack* and *Empire*.

Well, if you weren't aware, there's a fantastic open-source project available called *ScummVM*, which, so long as you own the original disks/CDs, will play a great many of the old LucasArts puzzle games, such as *Monkey Island 1 and 2*, *Loom*, *Sam &*

Max, Day of the Tentacle, and *Indiana Jones* – the best point-and-click puzzle games of their era bar none.

All the games work and play as well as they did originally back when DOS was considered to be a good platform for games, including all the original graphics and music. To accomplish this wizardry, the game uses its own SDL-based rendering engine, so the graphics are as smooth as originally intended. You can even opt to push graphics rendering onto your video card using OpenGL, which allows you to play games in a window alongside other applications.

DVD readers will find *ScummVM*

on their coverdisc. Installation is easy if you're using a popular distribution – Debian uses, for example, can **apt-get scummvm**. Alternatively, CD readers can just go to the *ScummVM* homepage at www.scummvm.org and download/compile the GPLed source code. Just think – in under ten minutes you could be enjoying the likes of *Zak McKracken* and *Maniac Mansion* all over again!

One word of warning, though: prepare to lose whole weeks of your life when you start playing as Guybrush Threepwood and wade through the strange and unreal worlds of LucasArts' classic games!

HARDWARE COMPATIBILITY

Beginners' Guide to Linux

PART 2 Linux is very versatile, capable of running on an enormous range of hardware from smart watches to universe-simulating supercomputers. Most installations, however, fall somewhere in between...

Purchased your PC in the last couple of years? If so, and you haven't messed with internals to any great degree, you should have no problem installing and running any of the latest distributions. As a guide, the following configuration will give you many happy, frustration-free hours of Linux computing. As always though, bigger, faster, better and more are

Even if you've got a Microsoft keyboard, it should still work with Linux on your PC.



good watchwords, especially if you plan on editing your holiday videos or recording an orchestral performance.

PROCESSOR There are versions of Linux available for all the major architectures – that is the very core of the computer – though the most common versions are based on Intel's x86 specification – Pentiums (II, III and 4)

and Celerons from Intel, and Athlons or Durons from AMD. While Linux can be coaxed on to even the most basic hardware, a processor better than 300MHz will ensure you don't spend too much time waiting around for things to happen, and anything above 500MHz will be perfectly adequate for all but the most demanding applications.

HARD DISK You can run a PC on a tiny amount of hard disk space, but if you want to install KDE, GNOME and a decent selection of applications you really need something in excess of 3GB of free space. The latest Red Hat, SuSE and Mandrake distributions claim a minimum of 2GB for a full graphical install, but if you scrimp here, you may spend a lot of time looking at worrying messages saying you're about to run out of disk space. If you have a fairly large disk and you are dual-booting with Windows, 3-5GB for a Linux partition (see box below) will usually be enough for the system and basic applications, while the Windows partition (or a third Windows partition) can double up as a data storage area that is accessible from both operating systems. Many removable storage devices such as Zip disks are recognised and mounted automatically by the latest generation of Linux distributions.

RAM As with hard disk space, you can run a Linux PC on a tiny amount of memory – but why would you want to when RAM is so ludicrously cheap? A 256MB memory module can be had for less than £25, but if you don't have the cash or capacity, 64MB

Linux basics

What On Earth is a PARTITION?

COULD THIS BE SOMETHING TO DO WITH DIY?

No, not really, though you've hit upon a handy analogy. Think of your hard disk as a room. Chances are that, at the moment, it's either empty – in which case you can move all your Linux furniture in without needing to rearrange things – or it's full of Windows stuff. If you're in the latter situation and you want to keep Windows, you'll need to create some space in order to let new Linux furniture have some space alongside the existing stuff.

AND HOW DO I DO THAT?

Without getting technical, partitioning a disk splits it up so the computer 'sees' your single hard disk as two, three, four or fifty different devices. And as with our room, we are not really making anything new, we're just moving stuff around and building walls between what is already there. Think of it as Feng Shui for your PC's hard disk.



SO WHERE DOES ALL THE WINDOWS STUFF GO?

You obviously need space to make these partitions. Your distro's installation routine (or third-party software) will usually 'resize' the Windows bit of the room before creating the new partition. To see the effect of partitioning in Windows, you can right-click on the C: icon in 'My Computer' and choose Properties from the list of options that are presented. There will be less free space on the drive. If you're resizing a Windows

partition, it is a very good idea to defragment your drive first. Look in Start > Accessories > System Tools for the 'Defrag' tool, and run it before you partition.

THERE'S STILL ONLY ONE HARD DISK IN MY COMPUTER!

That is because there is a one way door going between the two rooms. While Windows is unable to 'see' Linux partitions, you'll be able to access C: from Linux, meaning you can share data quite easily.

or 128MB should be suitable for pretty much anything you'd care to try. Anything less and you may spend more time waiting around than actually being productive!

CD-ROM A CD-ROM drive is almost essential, and a bootable one is very handy if you're installing one of the more recent distributions. You can check whether a drive is bootable by slipping a Linux CD into it and rebooting the PC. If it boots into Linux, fine, if not all may not be lost; reboot again and before the operating system begins to load, hit the **** key, this should take you into the BIOS (or look to see the boot up message which will tell you which key to press. Using the arrow keys, highlight Advanced CMOS Settings or the equivalent and press **<ENTER>**. You should now see a list that includes entries for 'primary', 'secondary' and maybe even 'tertiary' boot disks. By default the primary disk is the floppy disk, with hard disk second. Highlight the Primary entry and hit the **<ENTER>** key to bring up a list. If there is a CD-ROM entry, select it and Save and Exit.

If, on the other hand, your CD-ROM is not bootable, you'll need to create a boot disk. This is quite simple. In Windows (or on another PC), you need to access the first CD of your distro set and look in the /boot directory (SuSE) or /images (Mandrake and Red Hat) where you'll find a boot image which can be copied to a floppy disk and booted. This will then give you access to the CD for a full installation.

DVD drives and CD/RW (ReWritable) drives are always handy to have, and accessing them shouldn't pose a problem for most Linux distributions.

MODEM Modems, or a particular type of modem at least, are not a happy subject for Linux users and can cause the most difficulties for virgin installations. The reason is that many off-the-shelf PCs use a number of cost-cutting measures to keep the prices low, and one of the most prevalent is the 'Winmodem' which offloads much of the MODulating and DEModulating onto Windows. Unfortunately, the majority of Winmodems are defiantly closed source and are incompatible with Linux (there are a few exceptions, see www.idir.net/~gromitkc/winmodem.html for an exhaustive run down of compatibility), but all is not lost as

you can pick up an internal PCI modem for less than £15 and an external one for not much more.

We'll return to the modem problem in the near future.

NETWORK If you will be plugging into – or building – a network, or you're lucky enough to have broadband Internet access, a network card (NIC) is a vital component. Linux support for network cards, and integrated network hardware, is very good and will, in most cases, be

configured automatically during the installation process.

GRAPHICS AND SOUND Whether your PC is equipped with an integrated graphics and sound system, such as Intel's ubiquitous i810 chipset, or a high-end AGP graphics card and pro-quality sound card, you should not encounter any problems during installation. The fact that graphics card developers such as NVIDIA are making serious efforts to get the best from their cards under Linux is a sign of how much impact the operating system is having in the computer industry.


PERIPHERALS Keyboards, mice, monitors and printers are all well supported. In fact when it comes to printers, many of the older models have better drivers for Linux than Windows, thanks to the efforts of Open Source developers trying to get the most of their existing hardware.

There are a number of issues with scanners and digital cameras that we'll address in the future, though if you have a recently purchased USB scanner or camera, you shouldn't suffer any insurmountable problems.

Get with the common people

There are thousands of graphics cards, hard disks, CD-ROMs, Mice, etc. available to you; and the number of combinations that your average Linux distribution may encounter is staggering.

Problems, you may think, are inevitable, but Linux vendors – in house and through supporting the Open Source community in general – have spent an awful lot of time making sure both cutting-edge and legacy equipment are well supported. If your PC is a 'standard' build, there is a good chance that your installation will run without a hitch. If you have a more esoteric feature-set, don't despair. There's still a very good chance that if hardware support is not available out of the box, somebody somewhere has stumbled upon the same issue and would be winning to share their solutions. [LXF](#)



There are stripped-down Linux distros coded specially for much older machines, but for most home computing tasks on the main full-featured desktops like Mandrake or Red Hat, a minimum of 300MHz processor speed is essential.

Most essential peripherals are supported by Linux these days.



NEXT MONTH

Now we've decided on a distribution, and confirmed our hardware is Linux-friendly, it is time to bring them together to create a functioning system. Don't worry, many people have done all the difficult stuff, so we can just take it easy with the installation!

LINUX & WINDOWS

Linux for Windows users

An introduction to disk partitioning and running Linux alongside or on top of Microsoft Windows.



Frank Charlton

Frank has worked with PCs since the early 1980s, and has been using and writing about Linux since he discovered the first versions on Atari and Amiga.

He is a regular contributor to *PC PLUS* where some of this material has appeared previously.

This six-page beginners' tutorial looks at the work involved in preparing hard drive partitions to accept a fresh Linux installation, both on stand-alone machines and PCs where Linux is forced to share disk space with Windows. However, if the idea of plonking Linux onto your new Windows machine is too scary to contemplate, the third section deals with variants of Linux which run at the same time as Windows in ways that don't require any hard drive manipulation whatsoever.

Despite the inclusion of partitioning utilities either alongside or within modern Linux distros, the thorny issue of drive partitioning is still something of a hot topic for many Linux novices. Back in the mists of time, performing a new installation was not a job for the faint-hearted. The original installation programs for early distributions like Red Hat and Slackware were completely text driven, and there were none of the concessions toward ease of use and friendliness that we take for granted these days.

Original Linux installation programs simply assumed two basic things about potential users: that they knew how to manually partition a drive and make space available, and that they had already done it before going anywhere near the Linux installer itself. Most operations were carried out using arcane MS-DOS tools like *FIPS*, a tool that split and resized existing Windows partitions to make room for Linux. Thankfully, the Linux market has changed a lot in the last couple of years: every Linux distribution worth a Linux novice's attention has managed to come up with some way of making the task of partitioning drives as painless and transparent as possible.

Before we begin, there's one thing we can't stress clearly enough – back up your system before you begin if you're intending to install on a system that already has Windows running on it. Read exactly how and why you should back your system up before you proceed in the boxout on the right. It's important, and *Linux Format* cannot be held responsible if you fail to act on the advice and lose important data.

Back up, back up, back up! Make a copy of all your old data



An image back-up tool such as *Norton Ghost* is useful – it can even back up your Linux partitions once you've finished!

Things can – and at some stage, probably will – go wrong. Even a perfect installation, which is carried out with the utmost time and care, can fall victim to circumstances beyond your control. If you're in the middle of performing a critical operation such as resizing and restructuring your Windows partitions and you suffer a simple power failure, the end result is likely to be the complete loss of your data. Add the human error factor – we all click the wrong button sometimes – and you can see the potential for tears when installing a new OS.

Don't let this put you off! All this talk of disaster and data loss can easily be avoided by making a backup of your system before you even take the Linux installation CDs from their cases.

Backups come in two basic flavours – file and image. A normal file backup is one carried out by a program like *Microsoft Backup* or *Dantz Retrospect*, where each and every file is copied and saved. This is fine for normal disaster prevention, but won't always protect against the sort of situation we're discussing here, where complete partition loss can occur. A better option is a tool such as *Norton Ghost* or *PowerQuest DriveImage*. These tools take a complete snapshot or image of the contents of your drive, complete with all partition information. If the worst happens, you can restore your PC to its exact pre-disaster state. If you have a removable media drive then back up to that – even a CD writer will do if you have enough blank discs.

Partitioning your hard disk Making room for Linux

Generally, partitioning solutions fall into two camps – those which are executed outside of the Linux installer itself, and those which are internal routines carried out as part of the installation process. The former usually takes the form of a well-established Windows utility, often a cut-down version of best sellers like PowerQuest's Partition Magic. If your chosen distro follows this route, we can assure you that there is very little chance of

anything cataclysmic happening to your data should problems arise. First of all, these programs are commercial quality, and generally undergo rigorous testing before they reach the release stage. Secondly, without being too patronising, their Windows heritage means they're usually extremely well protected against the curse of user error, so that plenty of warnings are given even if you try to overwrite your main Windows boot partition.

Essential links

The Linux Partitioning Mini-FAQ <http://kmsself.home.netcom.com/Linux/FAQs/partition.html>.

Slightly out of date, but still contains useful information.

Experts Exchange www.experts-exchange.com/Operating_Systems/Linux/Q_20173134.html

Some excellent user-contributed advice concerning practical partitioning issues.

The Linux Filesystem Explained www.freeos.com/articles/3102/
Slightly technical, but a good primer on what the various parts of the filesystem actually are.

The LILO bootloader www.linux-tutorial.info/cgi-bin/display.pl?68&0&0&0&3


A terrible URL to try typing correctly, but an excellent explanation of the default Linux boot manager.

The second method can always be thought of as an unknown quantity to some extent, and we certainly can't blame anyone for feeling a little cautious when allowing an unknown program to run riot across precious hard drive partitions. All we can say to offer an element of reassurance, is that we've happily allowed many a Linux distribution to carry out the necessary partition manipulation without a single disastrous occurrence, both on our own machines and those belonging to friends and colleagues.

When the time comes to actually carrying out the partition jiggling, there are a number of things to think about. If you're planning on installing Linux as the sole operating system on a PC, then your worries amount to exactly zero. Wipe the entire hard drive, and sit back and watch as your chosen distribution's installer happily hogs every acre of disk space as the new home for Linux. For that matter, this approach is often the best way to go if you're nervous about tinkering with the default settings of a shop-bought PC which came with Windows pre-installed – dig that old PC out of the loft, or pick up an older model at a ridiculously cheap price and let it become your dedicated Linux workstation. Linux is a lot happier with older hardware than Windows will ever be, and even a knockdown old Pentium will happily run most distros without the slightest hint of a problem.

Dual booting

For most of us though, the dual-boot option is the method of choice when we decide to investigate the wonders of Linux. Dual booting is simply a term used to describe the situation where a




Many Linux distributions come with their own internal disk and partition manipulation utilities

PC plays host to more than one operating system, and can boot into either without one interfering with the other.

It's likely that your PC manufacturer has already partitioned the entire hard drive surface, formatted it and handed it on a plate to MS Windows. In cases like this, there are two ways to proceed. The first and probably most sensible is to install a second hard drive mechanism, and give that entirely over to Linux. In this case, the Linux installer doesn't need to touch your existing Windows partitions at all, and will partition and format the new mechanism for you. The pros are obvious – keeping the operating systems completely separate reduces the chance of misfortune either during installation or later on. In addition, it takes away the grief of interfering with your Windows partitions.

The second method is only slightly more risky, and involves resizing your existing partitions to make space available for Linux. We'll be looking at this in more detail in the next section, when we'll take you through the procedure involved using a typical Linux distribution, providing practical advice on how much space you really need to install Linux on a Windows system. We'll also



If you want to dual or multi-boot more than one OS, you'll need a boot manager

Mount points explained

'Fake' directories = more advanced functionality than MS Windows!

Windows and Linux handle disk drives and partitions differently. Since the early days of MS-DOS, Windows has always used a specific method of referring to and accessing hard drives and partitions, which we know as the drive letter method. When Windows loads, it assigns a letter to each drive on the system, whether it's a physical hard drive, partition, CD-ROM, floppy or other removable media drive. The hard drive the PC boots from is always labelled as **C:**, and the floppy can always be accessed via **A:**. It's an easy system to get to grips with, and while it's not necessarily that efficient in

advanced terms, it gets the job done and ensures an easy life. For most of us.

Linux doesn't use this drive letter access model. Rather, it uses something called **mount points**. This is an issue that confuses many of us when we're making the first tentative steps into the Linux world. To put it simply, a mount point is essentially a fake directory in the Linux file hierarchy that actually leads to a physical disk or partition. When you create a shortcut to a drive in Windows and place it on your Desktop, it's a similar arrangement.

While Windows places itself and just about everything else including the Program Files

directory on the boot drive, Linux doesn't. The Linux filesystem is much more advanced, and allows for complete configuration. With a normal Linux installation, the installer will recommend that you place some parts of the filesystem on different partitions – at the very least, **/usr** and **/home** should be separate. When you're looking at the root directory of a Linux system, you'd see **usr** and **home** as two directories. Moving into these folders actually takes you to separate partitions, but Linux does this with such ease and transparency that you're never aware of it – nor need you be.

TutorialBeginners'Linux

◀◀ be discussing the various boot managers available that handle the task of choosing which operating system to boot when you switch on. Every Linux distro comes with one or more included, but there are some very good commercial equivalents available which can make the experience a lot more pleasurable.

Finally, there is another option to consider – using a version

of Linux which doesn't require any kind of partition manipulation at all, and actually runs from within Windows, making use of the standard Windows filesystem. There are dedicated distros available that do this, as well as solutions like *Virtual PC* and *VMWare* (see page 18) which act as a host for the alien operating system of your choice. We'll examine this more in the last section.

The file `/etc/fstab` is where the mount points are actually configured – but don't edit by hand unless you know what you're doing.

```

Konsol
File Sessions Options Help
UW PICO(tm) 3.7 File: /etc/fstab
/dev/hdb1 /boot ext2 defaults 1 2
/dev/hdb2 swap swap defaults 0 2
/dev/hdb3 / ext2 defaults 1 1
proc /proc proc defaults 0 0
usbdevfs /proc/bus/usb usbdevfs defaults 0 0
devpts /dev/pts devpts defaults 0 0
/dev/cdrom /cdrom auto ro,noauto,user,exec 0 0
/dev/cdrom1 /cdrom1 auto ro,noauto,user,exec 0 0
/dev/fd0 /floppy auto noauto,user 0 0
/dev/hda1 /mnt/win_c vfat noauto,user 0 0
/dev/hda2 /mnt/win_d vfat user 0 0
/dev/hda3 /mnt/win_me vfat user 0 0

[ Read 12 lines ]
^G Get Help ^O WriteOut ^R Read File ^Y Prev Pg ^K Cut Text ^C Cur Pos
^X Exit ^J Justify ^W Where is ^V Next Pg ^U UnCut Text ^I To Spell
  
```

Linux alongside Windows

Despite what Microsoft might like you to believe, installing Linux beside Windows is simple and safe.

Before you begin to worry about the potential problems

involved with performing radical surgery on your existing Windows installation in order to get involved with the world of Linux, think carefully. Do you have an older PC stashed away somewhere, an ageing machine built from what Microsoft and most of the computer press would consider 'obsolete' components? If the answer is yes, then dig it out. Clean it up, switch it on and make it come alive again with a Linux installation.

Using a stand-alone PC

Linux in its most basic form will run on the oldest machines – well, almost. The original versions of the operating system were designed for much older processors than the 2GHz monsters you see in shops these days, and Linux will still run perfectly well on older hardware. Many of us who've been using PCs for years have older machines using 486 or even 386 processors gathering dust in lofts and storage boxes, and these can make ideal test-bed machines for the start of a Linux hobby. As long as they have the basic required hardware – a working base unit with processor and hard drive, a monitor and input devices – then there's no reason for them to sit dormant. If you're not going to use your old PC for Linux, your local Linux User Group can probably put you in touch with a charity who will be grateful of your hardware donation!



A KVM switch like this one from Belkin (www.belkin.com) lets two PCs run from one keyboard, mouse and monitor

If you're short on desk space, then pick up a KVM switching device – with one of these, your old machine can share keyboard, mouse and monitor with your new one, switching over whenever you like.

If you don't have an old machine but like the idea of leaving your shiny new Windows PC well alone, then pick one up second-hand. As well as the obvious approach of finding an old machine at a computer fair, many outlets now refurbish and sell older machines at very low prices, advertising them as 'Internet Ready PCs'. Generally, you can pick up an older first-generation Pentium machine complete with refurbished 14-inch monitor, keyboard and mouse for around £150. Buy a KVM switch and just the base unit and you can expect to pay even less.

Armed with a stand-alone PC, you can happily devote the whole thing to Linux. Whichever distribution you opt for, the installer will provide an option to format the entire hard drive and claim the space for itself. In this case, you can go with the recommended partition options all the way, since there are no concerns about leaving space for Windows.

Using an existing Windows PC

For those of us without the resources, space or desire for setting up a separate Linux machine, the only option is to install Linux on the same machine as Windows – unless you're feeling brave enough to do away with Microsoft completely, that is!

Running more than one operating system on the same PC is of course entirely possible, and no matter what Microsoft would like us to think, it's neither especially difficult or detrimental to Windows itself. The process is known as 'dual booting', since the two operating systems exist side-by-side and you choose which one to boot your PC into whenever you switch on. While both operating systems will to some extent be aware of the other's existence, on a correctly set-up dual boot system you can happily use either without fear of one interfering with the other.

Most modern off-the-shelf PCs come with a single hard drive installed, which is usually very large and offers much more space than you actually need for Windows. Whether Windows needs it or not is neither here nor there at this point though, since it will have claimed it all for itself.

Reclaim hard disk territory

To install Linux on an average system like this, the first thing you need to do is fight with Windows a little and reclaim some of that idle disk space. You can't just tell Windows to back off though; instead will need to physically push it out of the way in order to get some space back. Basically, there are two ways you can achieve this – by using a stand-alone third-party application such as *PartitionMagic* from PowerQuest, or by entrusting the task to your chosen Linux distribution.

Both approaches have advantages and disadvantages. Using a partitioning tool yourself gives a greater degree of control and flexibility, but allowing the Linux installer to do it makes for an easier life if you're not sure exactly what you're doing. From the standpoint of a complete novice to both Linux and performing surgery on disk partitions, we'd go for the latter option and use the tools built into your Linux distro's installation program. These have been tested to death, and offer just as much safety as the proprietary third-party application route, so you can trust them.

Most Linux installers will offer different degrees or classes of installation, depending on what you want your new operating system to do. These range from a bare minimum console-only installation to a full-blown graphical Windows replacement OS,

Essential links

POWERQUEST www.powerquest.com

Home of the company behind back-up tool *Drivelmage* and partition tool *PartitionMagic*. If you have an older version, you may need to upgrade, especially if you use Win XP.

SYMANTEC www.symantec.com

Stop here for information on *Norton Ghost*, the award-winning image back-up and restoration tool.

OPEN WATERS www.openwaters.net/linux/

For information on dual booting with Windows XP.

ZDNET www.zdnet.com.au/itmanager/technology/story/0,2000029587,20265546,00.htm

For info on using Linux on a machine hosting Microsoft's newest OS.

complete with almost every application you can imagine. Which one you opt for is entirely up to you, but remember this – you can always add more applications to your Linux system at a later date, and it's best to make sure you have some spare space on your hard disk to play with when the time comes.


Partitioning for Linux

At the most absolutely basic level, you will need a minimum of two partitions to install Linux. One will hold the OS itself, and the other will act as swap space. While Windows uses normal disk space to hold the swap files used by virtual memory, Linux doesn't and requires a dedicated partition. Your installer will recommend how much space to use, but don't go for anything less than twice the amount of the physical memory your PC has – it's better to have too much swap space than too little.

Ideally, you'll go with more than one filesystem partition, though. If you can afford the space, go for three – swap, main filesystem and a home partition. This system will be suggested by your installer anyway, since keeping the /home directory hierarchy separate from the system files improves reliability and performance. If anything should happen to the main filesystem, your user data files are safe on a discrete partition and you can even reinstall the OS without losing personal files. Again, it's pointless us giving exact figures here, as it will depend entirely on factors including how much drive space you have to play with and how many users will be requiring accounts on your Linux system.

Go with the recommendations of your chosen distro's installation program and you won't go far wrong.

A tool such as *PartitionMagic* is essential if you want full control over how much space both Windows and Linux can use



Linux on top of Windows

Safe solutions where Linux requires no disk manipulation

So far, we've looked at the work involved in preparing hard drive partitions to accept a fresh Linux installation, both on stand-alone machines and PCs where Linux is forced to share disk space with Windows. However, there are other ways that you can investigate, learn about and use Linux – ways that don't require any hard drive manipulation whatsoever. In this section, we're examining these alternative solutions – variants of Linux that run at the same time as Windows.

It's standard practice for a computer – any computer, and not just the PC – to run one operating system at a time. After all, the OS is the control program that provides the entire working environment, so why would you possibly want to run more than one concurrently?

There are actually many reasons for wanting to do this, but for the purposes of this article, the answer can be reduced to a single word – convenience. As we've seen, installing Linux onto a PC that already hosts Windows is a process fraught with potential problems. Even a smooth installation can mean time and drive space invested on an OS that you might not even like – after all, there's a vast choice of Linux distributions out there. So, surely you shouldn't have to commit to a full install every time you want to try one out for size. In many cases, you don't have to.


Tinkering with Linux

Essentially there are two ways to tinker with Linux without carving up your drive partitions or installing an extra hard disk. The first is to use an OS hosting device – a virtual machine, effectively – to run another OS on top of Windows. The second is to use a customised Linux distribution that's built specifically to run on top of your existing Windows system. There are pros and cons to both of these approaches, which we'll be examining here.

Customised distros

A while ago, a few clever programmers cottoned on to the concept that not everyone is willing or able to go through the process of re-partitioning their hard drive and installing an alien operating system.

Topologilinux is more flexible than WinLinux; it will even happily sit on top of newer versions of Windows, such as XP.



WinLinux 2003 might look like an ordinary Linux desktop, but this is Linux running full-screen on top of Windows 98.

To the sheer horror of many hardcore Linux aficionados, a new breed of distribution was born – a version of Linux that ran on top of Windows, and didn't require specific drive partitions formatted with the Linux file system.

Most of these Windows/Linux variants operate under the same principle. A large file is created on the Windows drive, which acts

Why go all the way?

Other ways of running Windows apps on Linux

In an ideal world, applications written for specific operating systems would happily run under any other OS. If you read *LXF* issue 42's six-page feature, you'll know about **WINE**, the project that enables you to run Windows programs directly from your Linux desktop. It's not an emulator as such (as the recursive **WINE Is Not an Emulator** name affirms), but rather an implementation of the Windows API, which sits on top of Linux.

Amazingly enough, a project to perform exactly the opposite is also underway. **LINE** – guess what the acronym stands for – is dedicated to making it possible to run Linux programs directly from Windows without requiring Linux itself to be installed. Like **WINE**, **LINE** is what's known as a 'compatibility layer', which intercepts system calls from Linux applications and translates them into the Windows equivalent, if possible. Also like **WINE**, it's still in development and won't run any Linux commands or applications

directly out of the box.

At the moment, **LINE** is alpha software, meaning it hasn't even reached the less flaky beta test stage yet. Unless you're technically proficient with Linux, we wouldn't recommend you tinker too much with **LINE** just yet. It's fresh but unstable, and likely to fail unless you have the knowledge to pursue a particular solution.

So why feature **LINE** in a Linux tutorial aimed at the beginner? In a word, potential. We first featured **WINE** many moons ago in its infancy, and even getting a simple application like Notepad to run was a chore. **WINE** has now evolved into a mature product, and comes ready-to-use with many Linux distros. With support from the active developer community, we expect **LINE** to progress along the same lines. Operating system convergence and cross-compatibility is a hot topic for the future, and you can read more about **LINE** at <http://line.sourceforge.net>.

as a virtual hard drive, and the customised distro is installed inside that 'container' file. Once this is done, the version of Linux is booted-up as easily as any Windows application, by simply launching it from the Desktop. Depending on the distribution, it may then run inside a window, take over the entire screen, or offer the possibility of both. Either way, you have a full version of Linux, which thinks it's running from dedicated drive partitions but is easily accessible without the usual boot manager.

The advantages of this approach are obvious. Firstly, you don't need to worry about dividing up or re-partitioning your hard disk. Secondly, you don't need to leave Windows during the initial installation process, or use custom boot disks or CDs to install. Finally, you don't ever need to leave your Windows environment to enter Linux, since the whole thing is done within the confines of Microsoft's OS. For many novices, these reasons alone are enough to investigate this type of user-friendly Linux distro, and a number of them have appeared over the last couple of years.

Running one OS on top of another isn't without its pitfalls, though. To begin with – no matter what the manufacturer's blurb may tell you – this type of Linux is never as fast in use as a real installation to dedicated partitions. The Linux filesystem is effectively emulated inside the large container file, which lives on a drive that uses the standard Windows filesystem – it isn't hard to see that writing a file to disk effectively means two sets of operating system calls. Having said that, as long as you're using a reasonably fast PC the difference is not huge – and if you've never experienced a genuine Linux installation you'd be none the wiser anyway. With these points in mind, it's time to look at some of the commercially available distributions that are able to run on top of Windows.

WINLINUX 2003 Probably the best known of all the Windows-hosted Linux distributions is *WinLinux*, now in its 2003 version. WinLinux has been around for quite a while, and has proven itself to be very stable indeed.

The latest release is a variant of Red Hat 7.3 using kernel release 2.4.18, and comes with the latest version of KDE 3. Installation is ludicrously simple, since it all runs from a standard Windows install program. There are no difficult technical questions to be answered, and you can have the entire thing up and running as quickly as you would any normal Windows application – it really is that simple. Hardware detection is performed automatically, and your Windows settings cloned as accurately as they can be.

As long as your hardware is supported by Linux *per se*, *WinLinux* will identify and configure it without you having to.

Essential links

WINAXE www.xwinpro.com/x-windows-on-windows.html

WinaXE is another implementation of Linux that sits on top of Windows, and offers full support for all versions of Windows.

WINE www.winehq.com


See how work on running Windows apps on Linux is progressing.

VMWARE www.vmware.com

Download an evaluation copy of this superb virtual machine and request a demo key.

CONNECTIX VIRTUAL PC www.connectix.com

Another very smooth virtual machine from the company sued by Sony for its PlayStation-on-PC software a while back. Solid, flexible and well worth a look if you're willing to break the Open Source tenets and spend a little money on this worthwhile application.



In a screenshot taken by the LINE author, you can see numerous Linux applications running directly from a Windows 98 desktop.

WinLinux requires a Pentium II processor or higher, and will only run under Windows 95, 98 or Me. Unfortunately, NT, 2000 and Windows XP aren't yet supported, which is a fairly significant negative point. For more information about the system, take a look at www.winlinux.com.

TOPOLOGILINUX A lesser-known (not to mention almost unpronounceable) name, this Linux variant is based around the venerable *Slackware* distribution. *Topologilinux* comes with kernel 2.4.20, KDE 3.05a and X Windows 4.2.1.1. Again, installation is extremely simple thanks to the Windows-based approach, and *Topologilinux* has a big advantage over *WinLinux* – it fully supports NTFS partitions and will happily run on top of any flavour of Windows, including XP. *Topologilinux* is available for download from: <http://hem1.passagen.se/svto/tlinux/index.html>

VIRTUAL MACHINES For increased flexibility, you could always opt to use a so-called 'virtual machine', which is effectively a PC within a PC. Using a system like this, you can play host to just about any operating system you like, and can have as many as you need sat alongside each other, disk space permitting. This approach offers supreme flexibility, giving you the chance to investigate more than one Linux distribution at the same time before making your mind up. You can also choose to use any distro you like, and you're not locked to one particular vendor, as you would be with proprietary solutions such as *WinLinux*.

VMWARE The original and most successful commercial virtual machine, *VMWare* has been around for a long time. And justifiably so, since it's a superb product. Install *VMWare*, set up a couple of virtual machines using the wizards,

and install your chosen distros just as if you were sat at a completely separate PC.

However, you really do need a fast PC with lots of memory to get the best from *VMWare* (or indeed any virtual machine). The system will run on slower machines, but you won't get an awful lot of work done, as even simple tasks like clicking between apps can take more venerable systems an age to process. *VMWare* isn't cheap either, but you can request a demo key, which gives you free access to the full program for a limited time. If your hardware is up to the task, give it a try. [LF](#)

PERSONAL ACCOUNTANCY SOFTWARE

GnuCash 1.8



Using your PC to manage your business or household finances can really help you get much more control over your cash-flow. Andy Channelle lays bare his inner accountant.

For users who have previously relied on a Windows application such as *Quicken* or *MS Money* to handle their finances, the migration to *GnuCash* can be a little disorientating. In fact, *GnuCash* is more similar in its structure to something like *Sage* with its adherence to real accounting principles than the home-orientated packages. Fortunately it retains the intuitive 'register' style interface for most common tasks and, if your needs are simple, all the difficult stuff goes on in the background.

For the non-accountants out there, the application works on double entry procedures, the basis of which is that, like physical matter, money doesn't simply appear and disappear, it has to come from - and go - to some other place. Every transaction entered into *GnuCash* has to have at least two accounts. For instance, money paid for your shopping moves from your account to the supermarket, and your salary simply moves money from your employer to you. Every transaction must have both a credit and debit element. *GnuCash* uses five account types (though there are a few subdivisions):

ASSET This type of account includes things like cash in your accounts, stocks and shares, your home and anything else you own. With Asset accounts, a debit entry increases the value of the account, while a credit entry decreases it.

LIABILITY This is the opposite of an Asset account, so a debit entry here will decrease the value of the account while a credit entry increases it. Asset accounts may include entries such as mortgage, loans, credit cards and Hire Purchase agreements.

INCOME This account contains details of salary, interest share dividends etc. Credit increases, debit decreases the income account balance.

EXPENSES This tends to be the most numerous account type. It details your daily expenditure and may include food, petrol, rent, utility bills, etc. Debits increase Expense account balances while credit decrease it.

Tip

It may be becoming apparent that the words 'debit' and 'credit' don't, as you may expect, have the same fixed meaning when it comes to accounting. Either may increase or decrease an account balance. It helps to remember that, for the sake of consistency, debits are left column entries and credits are right column entries in the register.




Fig2 Creating an account with the druid.



Fig1 The opening screen first time you load *GnuCash*.

EQUITY The Equity account is used to give you a snapshot of your net worth. A simple equation for this is Assets (what you own) minus Liabilities (what you owe) equals Equity.

Double entry bookkeeping seems complicated at first, but it just follows the movement of cash in the real world, making it the perfect method for financial monitoring and reporting. Moreover, *GnuCash* does most of the hard work, allowing you to get on with looking after the pennies.

New accounts

The first time you launch *GnuCash*, the set up druid appears (see **Fig1** and **Fig2**) offering the opportunity to configure the application, set options such as default currency and choose which type of account, or accounts, you are setting up. By selecting more than one option here, for instance Business Accounts and Simple Checkbook, you can manage both home and work accounts in a single, easy to back up environment. It also simplifies the transferral of money from one account to another. We have chosen the Common Accounts option to begin with, which is a minimal set of accounts suitable for home users.

Once you have decided to what extent you need the software and chosen your accounts, you move on to the accounts list,




Fig3 Subsequent launches start you at the accounts list.

where you can set an opening balance for most of the default accounts. This is where a copy of your most recent bank statement comes in handy; you can simply enter the amount (credit or overdrawn) in your current account, petty cash, savings account etc. Accounts in credit should be entered normally, though there is no need to include a currency (£) sign, and overdrawn accounts should be enclosed in brackets. These will then show up in the register highlighted in red. You can, of course, enter an opening balance just like a regular entry in the register later.

On subsequent launches the Accounts List is where you'll start, as shown in **Fig3**. The Accounts List expands and collapses to reveal more detail by clicking on the + or - sign to the left of each entry. On the right of the window, each element has a balance, which is a summary of everything within that branch of the accounts tree, every sub-menu within the tree will also have a balance, which makes it quite easy to see, for instance, how much you have earned from investments or spent on childcare.

Click the Expenses entry and take a little time to examine the list. Are there elements, such as Taxes>Medicare, that are irrelevant? If so, either highlight it and select the 'Delete' option from the toolbar, or right click and select 'edit' to make it more meaningful for your circumstances. It may seem a little long-winded to go through the menus this way, but it actually makes things a lot less complicated later on, it is also a good place to keep a record of account numbers, Customer Service phone numbers etc. Once you're satisfied with the list, do 'File>Save As..' and save your file in your /home directory.

It is time now to move on to the very heart of the application, the register. When dealing with the register, shown below in **Fig4**, it is important to remember that entries always need a source and a destination.

Below the menu bar, toolbar and Account overview (which offers a running total of your financial state) there is the register, which looks a lot like a bank statement. The easiest way to start is by making an entry. So, open 'Assets>Current Assets and double click on Checking Account - this is *GnuCash's* default current account name - so your screen looks like the image below.

The date is automatic, though you can adjust it by selecting the small arrow on the right of the date box. Give the first entry a description, such as Salary or whatever, and then select the drop-down list and find the 'Income:Salary' entry. Finally enter an amount, £1,000 seems a nice round figure. What we have done here is transfer £1,000 from a source (Income:Salary) to a destination (Assets:Checking), and if you go back into the Accounts List, you'll see this reflected in positive balances for both accounts.




Fig4 The register - entries need a source and a destination.




Fig5 Creating a scheduled transaction.

Chances are the salary is going to be a consistent entry in our accounts so it makes sense to save some time and effort, and create a scheduled transaction. That way *GnuCash* does all the remembering and tedious data entry. Right click on the Salary entry and select 'Schedule..' to bring up the scheduled transaction dialog, as shown above in **Fig5**. Here you can select payment frequency (daily, weekly etc.), number of payments due and a start and end date for the payments to begin. As this is a salary payment, we select 'Monthly' and 'Never end'. The Advanced tab gives you the opportunity to set extra options such as having entries put into the register automatically a number of days before they're due and set reminders.

Money comes in, but it also goes out, and we need to set up a scheduled transaction for the mortgage payments. Not created an account for your mortgage? Go back into the Accounts List, select New Account and add a child account to the Liabilities parent.

Back in the register for your Checking Account, the actual entry procedure is the same as before, but with a few small differences. Firstly, as our mortgage tends to go out at the end of the month, the first change is the date. Clicking the small arrow in the date box launches the calendar where we can choose the appropriate date for the entry. Next we input the description and, in the Transfer box select the previously created 'Liabilities: Mortgage' account. Finally you put the actual payment in the Withdrawal column and set up the scheduled payments as outlined above.

At this point you may notice a blue line separating our first two entries. Think of this as a line dividing the present (entries above the line) from the future (those below). It's also worth noting just above the register, the 'Summary bar' has been updated to reflect future transactions.

It is possible, and in many cases, desirable to do all your day-to-day management within this one account, and in that sense, *GnuCash* is actually very similar to *Quicken*, *MoneyDance et al*. However, the use of a real double entry system makes for a very flexible application capable of running both home and small business accounts.

In with the old

If you're migrating away from another application, *GnuCash* makes it quite easy to import your data, but you must make sure **➤**

TutorialGnuCash




Fig6 Importing is easy if you follow the druid's advice.

◀◀ it is exported from *Quicken*, *Money* or whatever in the correct format. The application can handle either OFX/QFX (Open/Quicken Financial Exchange) data or, more commonly, QIF (Quicken Interchange Format) files. The latter can be exported from most personal finance packages or – in many cases – can be downloaded from your bank and is actually what we are importing, as in Fig6 above.

The procedure is very simple, but it pays to read the import druid's advice thoroughly. Select File>Import>QIF. We are then prompted to specify the file and click 'Next'. After the data has been converted, we can either pick another file to import (if, for instance, you're taking data from more than one source) or move on. Click 'Next'. As we've already created current and saving accounts, *GnuCash* attempts to match up the imported files. The next stage converts the categories of our old application to accounts for *GnuCash* and lets us specify an account for

Equity Accounts

Knowing your worth at any given time can be either useful or disheartening, depending on whether the final numbers are black or red. Equity Accounts allow you to take a snap shot of your wealth, but you don't make entries here in the same way as you would with the other accounts. There are a number of ways you can use these special sections, but it is common practice to move all your income and expense balances into an equity account at the end of the financial year. This allows you to reset the balances in ordinary accounts to zero, so the books don't become too unwieldy, while still retaining the ability to see the big picture.

uncategorised entries. As before, diligence here pays dividends in the future.

One thing the import process doesn't recognise is scheduled transactions so, if possible, get a print out of your direct debits and standing orders from *Quicken* or *Money* to reconcile with the new file in *GnuCash*. The process of setting up Scheduled Transactions (detailed on previous page) isn't too arduous, and if you don't remember to do it, you could even miss a mortgage payment...

Where does it all go?

It's arguable that, even above convenience, the real value of managing your money on a computer is the ability to see, both visually and numerically, where all your money goes. It's great for budgeting, but it can also be quite frightening when you realise that 70 per cent of your disposable income goes on alcohol, broadband and cat5 cabling.


Tax friendly

Making payment to the Inland Revenue less taxing

Two years ago a small victory was won by the Linux and Macintosh communities, when the UK government abandoned plans to offer, essentially, a self-assessment website for filing tax returns that was available only to Windows users. So if you're recently had a tax return thud itself down onto your doormat (or you filed online last year and just received a slip of paper with a web address on it), now is the time to sort out your finances. The first thing you'll need to do is get registered with the Inland Revenue at www.inlandrevenue.gov.uk/sa/index.htm. Once you've put in your details, you'll be sent a user ID, a randomly generated selection of 12 numbers and letters, which can be used to login to the Government's online Gateway. You then have to request an activation PIN, another 12 digit 'password' for online Self-Assessment. As both of these transactions involve the postal service, it's better to get this out of the way early in the tax year to avoid a rush in December (and the £100 penalty for late filing).

Once you have all this information to hand, you can login to the website and use the Inland Revenue's own online software to file your return. This uses standard web technologies and so can be completed with *Konqueror* or *Mozilla*.

The other bit of good news is that if you've used *GnuCash* to maintain your business accounts for the past year (and we have to remember that these things can only be as accurate as the information we provide), filing in much of the self-assessment forms will involve nothing more than a little judicious cutting and pasting! If, for instance, you were a freelance writer your incoming payments may have been



entered under 'Income:Sales' while outgoings may be coralled under 'Expenses' (which may also have subdivisions for capital expenses and so on. And of course, if you've been really diligent in your record

keeping, splitting payments into tax, National Insurance etc. (see above), the final outcome won't be a shock – indeed, you can give yourself a smug pat on the back for a job well done.

Date	Num	Description	Transfer	R	Deposit	Withdrawal	Balance
12/06/01	switch	B&Q	Garden	y		3.96	-1,007.83
12/06/01	switch	Tesco	Groceries	y		66.02	-1,073.85
12/06/01	switch	M&S	Groceries	y		7.08	-1,140.93
12/06/01	switch	Mothercare	Tom	y		64.07	-1,205.00
12/06/01	switch	Boots	Healthcare	y		9.87	-1,214.87
12/06/01	switch	Tesco	Petrol	n		26.98	-1,241.85
15/06/01	dd	SWEB	Bills:Electricity	y		54.33	-1,296.18
15/06/01	dd	Telewest	Bills:Telephone	y		29.27	-1,325.45
15/06/01		Tesco	Groceries	n		54.05	-1,379.50
15/06/01		Tesco	Petrol	n		28.27	-1,407.77
15/06/01		MVC	Birthdays	y		4.06	-1,411.83
15/06/01		Balance Record	Unspecified	n		64.95	-1,346.88
16/06/01		Egg Card	Expenses:Miscellaneous	y		27.26	-1,374.14
16/06/01	Num	Boots	Groceries	n	Deposit	14.70	-1,359.44
16/06/01		Tesco	Groceries	n		4.08	-1,363.52
17/06/01		Tesco	Groceries	n		24.97	-1,388.49
18/06/01		Britannia	Bills:Mortgage Payment	y		208.00	-1,596.49
04/06/03		Future Publishing pay	Income:Salary	n	1,000.00		-496.49
26/06/03		Mortgage	Liabilities:Britannia Mortgage	n		268.00	-764.49
04/06/03				n			

Fig7 **NOTE:** scheduled transactions are not recognised on import, so make sure that you check new accounts against their sources when importing new data to avoid any later pecuniary embarrassment.

From the register, reports are limited to an Account Report, which just regurgitates the raw data in a printable form and a Transaction Report which does the same for individual transactions.

More useful reports and charts are available from the accounts list. The key to getting a useful report or chart out of any finance package is setting the right criteria: there's no point in looking at four years of data if your interest only covers the last quarter of 2002. The reports within *GnuCash* are more functional than pretty, but are no worse for their simple presentation, and are available via the 'Reports' menu entry. Home users will probably find the 'Account Summary' option and the various charts in 'Income and Expenses' most useful.

To get an overall picture of your financial situation, select Reports>Income and Expense>Cash Flow. You'll notice that this will give you information from the 'Current year start' to 'today'. We need to hone that down a bit, so select the 'Options' icon to bring up the configuration dialog, as shown in Fig8. Here we can select important details, such as which accounts to include in the report (use the Accounts tab to include or discard specific accounts), but we can also set a more useful time period to monitor. This dialog uses the standard Calendar widget, so select the 'From' and 'To' radio buttons, and then define your dates by either typing them into the space, using dd/mm/yy format, or clicking the small arrow which will bring up the clickable calendar. The arrows on either side of the month and year labels move you back and forward in time (not literally!) and today's date will be highlighted in blue.

Once you've set your variables, you can hit 'Apply' and the new report is created. [LXF](#)




Fig8 The big picture: just how skint are you?

Splitters!

Divide and conquer your payments

There will be occasions, especially if you are using *GnuCash* to manage the finances of your business, that simply entering figures and printing reports is not enough. You may need to split purchases and sales figures for VAT purposes, for instance, or view the tax and national insurance payments due on income. *GnuCash* calls these entries Splits, and creating them is as simple as setting up scheduled transactions.

As an example, we can right click on the previously entered salary payment and select Split transaction. This will add a new line below the original entry which can be filled just like

any other transaction. In this case, we select 'Expenses:Tax:Income Tax' as the destination for the payment, and put a figure into the 'Deposit' column – if you go back to the Accounts List, you'll see this amount in the relevant account. Unfortunately *GnuCash* does allow entries as a percentage figure, so you may find it handy to have a calculator to hand.

Splits can be really useful at home too. For example if you set up your Mortgage you can split payments into Capital and Interest payments to get a real snapshot of your equity/debt ratio.

FILE AND PRINT

First Steps with Samba



PART 1 Bringing the gap between MS Windows and Linux for file and print services, with **Dr Chris Brown**.

My wife sounded puzzled. "Someone from *Linux Format* called. They want you to do a piece about the samba". And no wonder. Why should someone like me, with less physical grace than a JCB, be writing a tutorial about a Brazilian dance, and why might a computing magazine print it? So I explained that this was one of those cute Unix names like *Apache* and GNU and GNOME, and she said this was clearly some new use of the word 'cute' she hadn't come across before. Will she ever see the light?

Samba, to set the record straight, is a suite of services and programs that runs on Linux and that provides file and print services to Windows clients. It began as a one-man project in 1992 (the one man was Andrew Tridgell) and has since grown into one of the world's flagship open source projects, with a large, active, international development team. (See www.samba.org) The name *Samba* comes from the original name of the protocol which Microsoft used for file and printer sharing – Server Message Block, or SMB. (Later, in a typical Microsoft marketing upgrade, the name was changed to "Common Internet File System" or CIFS. The new name is misleading. For a number of reasons, SMB (or CIFS) is really only suitable for deployment across a local area network, not across the Internet. And despite the use of 'common' in the name, it remains primarily a Microsoft-centric protocol.)

Interoperability between Linux and Windows is an essential part of any Linux migration plan. The commercial importance of *Samba* is that it allows file and print services which would otherwise be provided by Windows NT or Windows 2000 machines to be provided on Linux at zero software cost and, probably, with a smaller hardware footprint than Windows. Replacing a few NT file servers with Linux and *Samba* is one of the easiest ways for a company to dip its toes into the Linux waters. Because the Windows client machines can't tell the difference between a share exported by a *Samba* server and a Windows file share, very little reconfiguration is required, and the change is essentially transparent to the end users.

In this two part tutorial, we'll first have a look at how file and printer sharing works in the Windows world and then move on to see how *Samba* fits in. We'll look at some simple *Samba* configurations before progressing to the slightly thornier problem of user identity and authentication.

A Short History of NetBIOS & SMB

It feels a bit odd writing about Microsoft's technologies for a Linux readership, like doing an article on Richard Branson for BA's inflight magazine. But a little background is necessary to understand what *Samba* is all about.

The history of the SMB protocol goes back a long way. In the beginning, (soon after the heavier elements had started condensing out after the Big Bang,) IBM invented the PC, DOS, and a piece of




Fig1 Browsing the network neighbourhood.

firmware that handled the low-level I/O within the machine called the BIOS (Basic Input/Output System). When PCs started being put onto networks, an extension to the BIOS, called NetBIOS, was developed to handle machine to machine communication. With NetBIOS, each machine had a name, up to 15 characters long, which was used directly to address network packets to their recipients. NetBIOS was a lightweight protocol requiring very little memory, and worked fine for the small LANs for which it was designed. However, it had two major drawbacks – it had no mechanism for routing, and it relied heavily on broadcasting. These features made NetBIOS unsuitable for use either on WANs or on point-to-point links. To overcome these limitations, a way was developed for NetBIOS traffic to be encapsulated into TCP/IP datagrams. This scheme, known as NetBIOS over TCP/IP (NBT) is defined in RFCs 1001 and 1002. The scheme retained the concept of NetBIOS names, but uses IP addresses for routing and delivery. *Samba* uses NBT.

NBT brought with it a requirement for NetBIOS name resolution – that is, it needed a mechanism to convert the NetBIOS names to the underlying IP addresses. Just as there are several ways to do hostname-to-IP lookups, there are several ways of performing NetBIOS name resolution, too. Clients can consult a local file (on Linux it's usually `/etc/samba/lmhosts`), or they can use a NetBIOS Name Server. Microsoft's implementation of a NetBIOS name server is called WINS (Windows Internet Name Server). As we will see, *Samba* can act as a NetBIOS name server, too. As a last resort, clients will broadcast requests to resolve NetBIOS names.

Resource Sharing

There are a couple of concepts central to the Windows resource-sharing model: workgroups and shares. A workgroup is simply a group of machines, identified by a workgroup name, which share resources amongst themselves. A share is a named, advertised resource on a specific machine, usually a filesystem or a printer. For example, my Windows 2000 machine BASHFUL is advertising a file share called projects which corresponds to the directory C:\J2SDK_Forte\projects.

The 'projects' share on the machine BASHFUL is identified by the name '//BASHFUL/projects'. This is an example of what is rather grandly called the Universal Naming Convention, or UNC. (I find the claim of universality hard to take, and have a strong suspicion that its use does not even extend throughout our own galaxy. But I digress.) Note that BASHFUL is the NetBIOS name of the machine, which is not necessarily the same as its TCP/IP hostname – *ie* the name DNS knows it as.

The collection of resources made up by all these shares is known in the Microsoft world as the Network Neighbourhood. You'll almost certainly have browsed a network neighbourhood using *Windows Explorer*; if not, **Fig1** shows what it looks like. Here, we're looking at the machines in the workgroup Westwick, and have expanded the view to show the shares on one specific machine, Dopey.

How does *Windows Explorer* know what machines are in the workgroup? Is there a configuration file somewhere that makes all this explicit? No, what happens is that all machines register with one specific machine on the network called the master browser. The master browser is not pre-determined, it is chosen by holding an 'election' across the network. Potentially, any machine that speaks SMB (including a *Samba* server) can take on this role. The list of machines held in the master browser is called the browse list. (Note that browsing, in this context, is not related to web browsing). In **Fig1**, we see that machines Bashful, Dopey and Petal are in the browse list for the workgroup Westwick. And although you can't tell from the figure, Dopey and Petal are both *Samba* servers running on Linux. The actual list of shares on Dopey, which is expanded in the figure, is not obtained from the browse list; it is obtained by contacting the SMB server on Dopey directly.

For those of you who prefer a command line interface, even on Windows, the command

```
net view \\dopey
```

will show the shares on dopey, and the command:

```
net use m: \\dopey\chris
```

will attach the share chris as drive letter m:

If you're familiar with native Linux file sharing (using NFS), there are a couple of differences between NFS and Microsoft's file-sharing mechanisms to keep in mind. First, SMB has an extra level of naming for exported shares. For example, a server can export the directory C:\images\vintage under the share name 'oldcars'. Clients access the files via the share name; they never see the name of the underlying directory. Using NFS on the other hand, a server might export the directory /images/vintage and the client would use the self-same directory name in an NFS mount command to attach the filesystem; for example:

```
# mount dopey.example.com:/images/vintage /mnt/oldcars
```

The second difference is that NFS doesn't support the notion of workgroups or browse lists. It's true that you can ask a specific NFS server about its exports, for example:

```
# showmount -e dopey.example.com
```

will list the exported file systems on **dopey.example.com**, but




Fig2 SWAT's global configuration screen.

you have to just know, somehow, that **dopey** is out there in the first place. There is no command that says "show me all the NFS servers on this network". [See the following note below.]

[Note: OK, I guess that's not strictly true. You could run a command such as `rpcinfo -b mountd 3` which would tell you how many machines on the local subnet have a mountd daemon listening. But that's not really the same as having an actively maintained browse list.]

Where does Samba fit in?

OK, enough of the Windows networking lecture. Where does *Samba* fit in? Well, *Samba* can act in the role of an SMB file and print server, as a NetBios name server, and as a master browser. The 'server' side of *Samba* actually consists of several daemons, utilities and configuration files, the most important of which are shown in the box immediately below. >>

Key server-side components

Samba's utilities and configuration

FILE	DESCRIPTION
smbd	The main daemon providing file and print services
nmbd	Daemon providing NetBIOS name resolution (equivalent to Microsoft's WINS) and network neighbourhood browsing
smb.conf	Samba's main configuration file, usually in the directory /etc/samba
testparm	A program which performs a basic syntax check of the smb.conf file
smbpasswd	A tool for managing samba accounts and passwords

« Configuring Samba

Samba can be installed from binary RPMs, or from source using the traditional command sequence of **./configure; make; make install**. Or, easiest of all, it can be included in the package selection when Linux is initially installed. Rather than dwell on any of that, we're going to take installation for granted and move straight on to doing some basic configuration of the server.

On my Red Hat system, the *Samba* config file is `/etc/samba/smb.conf`. Your mileage may differ. You can either hand-edit this file, or use a web-based graphical configuration tool called *SWAT* (*Samba Web Administration Tool*). *SWAT* is prettier, and we'll look at a couple of screenshots later, but it's easier to discuss the configuration options by examining the file directly.

There are potentially a lot of options that can be placed into `smb.conf` (the manual page describing the file's format runs to 6,500 lines!) but you can get basic services up and running with quite a simple file. The file as shipped with the *Samba* distribution contains a lot of helpful comments which may be sufficient to get you going. Here's an example, minus the comments. (The line numbers are for reference; they are not part of the file.)

```

1 [global]
2 netbios name = DOPEY
3 workgroup = WESTWICK
4 server string = Samba Server on Laptop
5 printcap name = /etc/printcap
6 load printers = yes
7 printing = lprng
8 log file = /var/log/samba/%m.log
9 security = user
10 encrypt passwords = yes
11 smb passwd file = /etc/samba/smbpasswd
12 hosts deny = ALL EXCEPT 192.168.0. 127.
13
14 [cdrom]
15 path = /mnt/cdrom
16 read only = yes
17 guest ok = yes
18

```

```

19 [homes]
20 comment = Home Directories
21 browseable = no
22 writable = yes
23 valid users = %S
24 create mode = 0664
25
26 [printers]
27 comment = All Printers
28 path = /var/spool/samba
29 browseable = yes
30 public = yes
31 writable = no
32 printable = yes

```

The section headed `[global]` (lines **1-12**) define settings that apply to the server as a whole. Lines **2-3** specify the NetBIOS name for this server, and the workgroup it will be part of. Lines **5-7** tell the server to make available as shares all the printers listed in the machines `/etc/printcap` file. (This is the standard configuration file for the traditional *lprng* print system on Linux.)

Line **8** determines where *samba* will do its logging. The `%m` in the file name is a *Samba* variable. *Samba* maintains a number of these and replaces them by the appropriate value. The variable `%m` is replaced by the NetBIOS name of the client machine. So this line is a Cunning Trick to cause *Samba* to log each client in a separate file. For example, entries for client `bashful` will be logged to `/var/log/samba/bashful.log`. The 'log level' directive (not shown in the example) controls the amount of detail written to this file.

Lines **9-12** are concerned with access control. The setting **security = user** means that users have to establish their identity when they connect to the server by supplying a user name and a password; access to shares is then controlled on the basis of that user identity. There is another, fundamentally different, mode called **security = share** in which clients gain access to shares not by establishing their own identity but by supplying a password associated with the share. This security mode was used by early versions of Windows such as Windows 95 but is not recommended for current use.

Line **10** says that passwords sent from windows clients will be encrypted, and line **11** says which file *Samba* should look up the encrypted passwords in. We'll talk more about this in the second part of the tutorial when we discuss authentication in detail. Finally, line **12** says that we'll only accept connections from clients on our local network `192.168.0.0` or from the local machine via its loopback address.

Having completed the global configuration of the server, at lines **14-17** we get to define an actual share. The share name is **cdrom** and the Linux path name corresponding to this share is `/mnt/cdrom`, on which, presumably, a CD is mounted. Line **16** stops users from trying to write to this share and line **17** allows users to connect to this share without a password. On the Linux machine, these unauthenticated users will assume the identity of the "guest user" which by default is usually the Linux account "nobody". Since nobody doesn't own any files, (sorry about the grammar, but I said what I meant), such users will only have access to world-readable files in the share.

The magic [homes] share

The `[homes]` share, defined in lines 19-24 of the file, is treated specially by *Samba*, and provides an easy way to automatically export all users' home directories. It works like this. When a

Fig3 *SWAT*'s share configuration screen



```

File Edit View Terminal Go Help
[chris@dopey chris]$ smbclient '//bashful/My Images' -U administrator%adminpw
added interface ip=192.168.0.48 bcast=192.168.0.255 mcast=255.255.255.0
Domain=[BESHFWICK] OS=[Windows 3.0] Server=[Windows 2000 LAN Manager]
smb: \> dir
.                D          0   Mon Jun 9 16:49:53 2003
..               D          0   Mon Jun 9 16:49:53 2003
beetle_green.png A 1574888  Sun Jan 19 11:25:58 2003
beetle_orange.png A 2099415  Sun Jan 19 12:13:29 2003
beetle_pic.bmp   A 4133974  Sun Jan 19 10:56:25 2003
beetle_pic.png  A 2309828  Sun Jan 19 11:40:46 2003
beetle_spotted.png A 2221458  Sun Jan 19 12:37:46 2003
beetle_stripped.png A 208572  Sun Jan 19 11:36:19 2003
chip.bmp        A 2387758  Sun Jan 19 12:38:20 2003
fish.bmp        A 46062   Thu Feb 20 20:38:20 2003
Steve_and_Liz_wedding.zip A 568868  Sun Oct 27 21:57:58 2002

40005 blocks of size 262144. 22947 blocks available
smb: \> get fish.bmp
getting file fish.bmp of size 46062 as fish.bmp (977.9 kb/s) (average 977.9 kb/s)
smb: \> quit
[chris@dopey chris]$

```

Fig4 Typical smbclient dialog.

windows user connects to the *Samba* server and authenticates, *Samba* will dynamically create a share named after that user and associate it with that user's home directory on the Linux machine. So if user Mike logs in, he'll see a share called *mike*. If Jane logs in, she'll see a share called *jane*, and so on. This is obviously very convenient when large user populations are being served. Line **21** of the config file stops the **[homes]** share itself from showing up in the browse list; it does not stop the user's individual home share from being visible.

Line **23** uses another of *Samba*'s variables. **%S** will be replaced by the name of the current share, which in this case is the user's name. So this line ensures that the only user who can access this share is the one whose home directory it corresponds to. It would be possible, for example, to add the administrator into the list of users who was allowed access to this share.

The **create mode** directive at line **24** is used to limit the access permissions that newly created files in this share will have. This is probably not the place to start explaining how the octal number 0644 is interpreted (you probably know already if you've used the Linux **chmod** command to set file permissions), but the result is that any files created in this share have (at most) the permissions **rw-r--r--** (The whole business of mapping the windows access control model for files onto the UNIX model is very tricky and is something I'd rather not get into right now.)

The last section of the *smb.conf* file (lines **26-32**) establish printer shares on this server. There is no need to export each printer individually, although you can if you want. As configured, *Samba* will automatically create shares for all printers defined by the underlying print system.

There is a useful little tool which comes as part of the *Samba* suite with the rather generic and unimaginative name of *testparm*. This tool will check the syntax of your *smb.conf* file and will report all the options in effect, including all the defaults which are not explicitly set in the config file. It's useful to check that you haven't mis-spelled a directive, or tried to specify an illegal value for an option.

Configuring with SWAT

If you'd prefer not to have to learn the syntax of the *smb.conf* file, you can configure it using a graphical, browser based tool called *SWAT*. Having been (to quote Eric Raymond) "marinated in UNIX" for many years, I'm not a great fan of graphical configuration tools, but as they go, *SWAT* is pretty good. It's a part of the *Samba* suite but may be separately packaged (on my Red Hat system it's in the RPM package *samba-swat*). It's started via *xinetd* and will need to be enabled, either by editing the file */etc/xinetd.d/swat* or by running the command **chkconfig swat on**.

Once running, *SWAT* listens for HTTP requests on port 901 and you can connect to it with a web browser by entering a URL such as <http://localhost:901>. *SWAT* offers several configuration screens allowing you to configure the globals section (see **Fig2**), individual shares (see **Fig3**), or the printers section of the file. There are other screens which allow you to view the text of the config file you've generated, and the status of the running servers. On its home page (not to be confused with the page for the **[home]** share) there are convenient links to all the *Samba* documentation.

As you'll see from **Fig2** and **Fig3**, *SWAT* just provides a fill-in-the-blanks alternative to hand-editing directives into *smb.conf*. It does, however, ensure that directives which can only take one of a specific set of values are not set to an illegal value. For example, the boolean option **guest ok** can only be set to 'yes' or 'no'. Also, the 'Help' links down the left-hand side of the screen take you directly to the entry in the *smb.conf* manual page which describes that directive. This is a nice convenience.

Client-side Samba

Although *Samba* is mainly known as a server to provide file, print, name resolution and browsing services from Linux to Windows clients, there is a client side to *Samba*, too. The client-side tools can be used to access shares on a Windows machine, or simply to run tests against a *Samba* server without having access to a Windows machine. One of the key client-side tools is a command line utility called *smbclient*.

For example, in order to list the shares on the machine *bashful*, use the command:

```
# smbclient -U administrator -L bashful
```

Note that **bashful**, in this example, is the NetBIOS name of the server, which as we've seen is not necessarily related to the DNS name of the machine (assuming it has one). In the example, *administrator* is the user you're logging in as; you'll be prompted for the password. Alternatively you can enter the password on the command line like this:

```
# smbclient -U administrator%adminpw -L bashful
```

although I'm uncomfortable with this approach because I prefer echoing to be turned off when I type a password.

You can also use *smbclient* to connect to a specific share. If you do this, you'll get an interactive prompt at which *smbclient* offers an ftp-like command repertoire. **Fig4** shows an example of connecting to the share 'My Images' on the server *bashful* and retrieving the file *fish.bmp*. The dialog should look familiar to anyone who has used ftp from the command line.

The Linux kernel also supports SMB as a filesystem type, so if you want to do more than just copy a few files across from an SMB share it's probably easier to mount the share onto a local directory. The command would look something like this:

```
# mount -t smbfs -o username=administrator
//bashful/projects /mnt/bashful
```

In this example, **administrator** is the name I'm logging in as on the SMB server, **//bashful/projects** is the UNC share name, and **/mnt/bashful** is the local mount point. Again, you'll be prompted for the password. Now you can simply **cd** to the directory */mnt/bashful* and access the files there as if they were local. Let's just repeat that to be perfectly clear: here, we're using Linux as a client to access shares exported from an SMB server, presumably a Windows machine, although there's no technical reason that you couldn't share files between Linux boxes this way, too. **LXF**

NEXT MONTH

We'll wrap up this two-part tutorial by taking a detailed look at the tricky business of authentication in *Samba*.



ILLUSTRATION COURTESY - WWW.BLENDER3D.ORG

3D MODELLING

Blender textures and materials

PART 4 Jono Bacon shows how to make our models more attractive by giving their surfaces a more lifelike patina.

So far in this *Blender* tutorial series, we have been looking at some basic modelling. Although this is good general introduction to the *Blender* interface, so far all of our models have been grey, lifeless objects that look neither realistic or appealing. This month we plan to rectify this situation and explore how we can add materials to our scenes to add life and colour to our objects.

Material girl

For those of you with no experience with 3D modellers such as *Blender*, the concept of materials may be new to you. A material is basically the colouring of an object to make it look more convincing within its context. This context for example may be a real-world scene, so the material must look as realistic as possible to the viewer. It is important to remember that a great looking scene with poor-quality materials will really let the final quality of the scene down.


Although the concept of a material sounds like a physical thing, a material is actually just an image or colour that has been designed to look like the desired material. A classic example is moulded plastic. If you think of an actual moulded plastic item (such as one of the stupid boxes you might have had to make in Design/Technology classes at school!) you will be aware that the

material plastic looks quite shiny, dense and has a very characteristic look. If you now consider a material like steel, a similar box made out of that will look far different from the plastic. Although the actual object in the scene is no different (in terms of how you created it) the material can vastly change how it is finally viewed.

Blender not only has support for materials, but has extensive support for Textures. "What is a texture?" I hear some of you cry. Well, if you missed the great explanation in Michael J Hammel's *GIMP* tutorial (*LXF39*), texture is an image that has a particular material's textual design on it to make your object look more realistic using a graphical representation of materials that have these textures. An example of this is Wood. If you look at a lump of untreated wood, you will see grains, knots and other small details. If you wanted to create a wooden table in your scene for example, you would need a 'Wood' texture, as a normal material would not give this detail.

Understanding materials

Materials is a large and complex subject, and there's certainly not enough space here to discuss all of the theory behind these complex beasts. I will however cover some of the theory that is important to knowing how to create materials effectively.



The concept of a material is simple: we apply it to an object. Where a material gets complex however, is in how the environmental surroundings affect how it is displayed, and the biggest environmental entity to affect a material is lighting. We can prefix this with a simple example. Just look at your copy of *Linux Format*. Look at how the light is cascading onto the pages. Now, some of you will be relaxing in bed reading this with a soft lamp lighting the magazine, some of you will be outside and some will be under a strong artificial light reading the magazine. Although the material is the same (paper and ink), the light greatly affects how that material looks.

The key to understanding the lighting is how the light reflects off the magazine. In *Blender* there are two types of reflection to be aware of – Diffuse and Specular reflection. Diffuse reflection will take the light from the light source and when it hits the material, the light source will be scattered out at a number of angles, in a similar way that when you throw a stone in a pond the ripples move out. The practical implication of using a Diffuse reflection is that a matte finish is the product of the reflection. If your object needs a smooth solid material, this is a good choice.

Specular reflection, on the other hand, takes the light from the source and the reflection is a single reflection which departs in the direction of the Incident Viewing Angle. This angle is the angle in which the viewer of the scene (the camera) is viewing the object. What happens is that the light comes from the light source onto the material, and bounces off towards the viewing angle. As such, with Specular reflection, the angle is critical to

seeing the scene properly, in the same way that viewing the magazine directly will look different to viewing the magazine from the opposite angle (assuming you are using a specular light source such as a spotlight for example).

Getting started

To get started playing with different levels of reflection and materials, it is good to use a simple object. In Top view, scale the default plane up so it is much bigger and add a UVSphere to the scene. Scale it up so it is bigger and ensure that it is sat on the

Fig1 Our scene to play with materials.




Fig2 A ball with our first material.

TutorialBlender




Fig3 The ball gets a wall to go with the floor.

◀ default plane (this will possibly mean moving the ball up a bit in the Front view. When you have created the ball, add a spotlight to the scene so it is positioned on the ball. Your scene should look similar to that of **Fig1** on the previous page.

When you hit **F12** and render the scene, you will see the grey ball in view. To add a material to our ball, ensure you are not in edit mode. Select the ball with the right mouse button and then click on the Materials button or press F5. To create a material we need to make use of Datablock button which is like a [-] sign on a button to the right of the Render button on the button bar. Press in the button and select ADD NEW. You will see MA:Material to the right of the Datablock which can be renamed by clicking on it. With this new material, a stack of new buttons will come into view in the main button area.

Although we have a new material, we have not done anything with it, and the first thing we should do is set the colour of the material. We can do this with the R, G and B sliders in the button area. These buttons adjust the amount of Red, Green and Blue that is in the material. Adjust the sliders until you have a colour that you are happy with.

When you have adjusted the sliders, we can look at the other sliders that involve some of the theory we just discussed:

SPEC SLIDER Deals with the amount of Specularity that is




Fig4 The ball gets coloured walls.

Texture Resources

Equip Blender with more




Fig5 There are many different textures that you can use to make your modelling more realistic – here from left to right are examples of Brick, Granite and Wood textures.

TEXTURE LIBRARY

<http://textures.forrest.cz/>

SOME BASIC TEXTURES

www.cc.jyu.fi/atk/blender/blacksmith/textures/

GRSITES TEXTURE ARCHIVE

www.grsites.com/textures/

DR ZEUS' TEXTURES

<http://drzeus.best.vwh.net/Art/Textures/Textures.html>

emitted from the material. If you adjust the slider higher you will see the increased specularity on the ball in the sense of an increased size of the reflection on the ball.

REF SLIDER This slider adjusts the reflection on the material. If you increase the value of this slider, you will see how the overall light is reflected off the ball and hence adjusts the matte shade of the ball. Rather than use this slider on its own, it is often a good idea to make use of this slider in conjunction with the colour sliders (R, G, B) to get the desired effect.

HARD SLIDER Controls the amount of hardness that is applied to the reflection on the ball. If you reduce the slider value you will see how you can soften it to the point of overwhelming the ball with reflection. Once again, play with it and experiment.

Spicing things up

So far we have looked at some of the materials options available. We will now spice our scene up just a little bit and create a feeling of space in our scene by boxing our little ball into a room. To create a pseudo room we are only going to actually add two




Fig6 The wooden textured ball.

walls, and even then, they are not walls; merely two planes that will act as walls for us.

First, go to front view and add a plane. Resize the plane and stretch it so it becomes a rectangle. When you are happy with the size position it so it sits on top of ground plane. Do the same now in Side view to create another wall. When you have created both walls, put them together. You may also need to move your camera and light over a little bit, but your scene should look similar to that in **Fig3**.

When you render the scene now you will see how the shadows will cascade off the wall. Remember to make sure that you have clicked the SetSmooth button in the Edit buttons for each object in the scene. Next, add a material to each new wall, and while we are at it, add a material for the floor. Remember that when you create a new material when clicking on the material Datablock, name the material so you don't mix them up. When you have added your materials, your scene should look similar to that in **Fig4**.

Adding Textures

Textures are in many ways the life that adds realism to your scenes. The world is certainly not a world just made up of matte colours, and there no reason why your *Blender* world should be.

To add textures to your scene you are going to need some texture files. You have a few options here. One is to create your own textures in something such as *The GIMP*. Unfortunately we do not have the space to cover that here, so we will use the easier option of using pre-built texture files. There are various sites available to get these textures (see the *Texture Resources* box for details), but I have got a Teak texture from www.cc.jyu.fi/atk/blender/blacksmith/textures/teak.jpg. The texture in its original form is shown in **Fig5** opposite.

To apply this texture to the ball, first right click the ball to select it. Next, click on the Materials button and look towards the right of the block of buttons. Within the buttons you will see another Datablock. Click on the Datablock, then ADD NEW. You can click on the Texture button (the one next to the Materials button) or press **F6**. Along the top of the Texture buttons there is an Image button. When you click on this button you can use the Load Image button to locate the texture file and load it. You will see the loaded texture in the preview pane on the left side of the Texture and Materials button groups. When you render the scene now you will see the texture applied to the ball.




Fig7 More textures added to the scene.

Creating an item of furniture

IKEA use this type of software to design theirs, so why don't you...?

To continue developing our modelling skills, I will show you how to create a stool (stop sniggering at the back there) and we will look at a couple of new techniques while creating the stool. This stool will be similar to a bar stool and will essentially comprise of a seat and three legs. To make things a little interesting, we will make our stool a little bit funky in its design. Hey, that is what creativity is all about!

First, start with a new *Blender* scene and scale the default plane so it is much bigger to act as our floor. Next, in the Top view create a circle and scale it up a little. We will now squash our circle into an ellipse by going out of edit mode, pressing **S** to scale the circle and then press both mouse buttons together. Your will see the circle squash and you can now resize the ellipse. Next, go to Front view, extrude it down one square and you now have a seat. Although the seat is there, it is still hollow, as the circle only draws an outline. To make the seat solid we need to create a face for each side of the ellipse. Go into edit mode, ensure that all vertices are pink (unselected). Now, press the **B** key to select the top row of vertices and press **Ctrl-F**. This will connect the vertices to form a face.



Fig8 It's a little basic so far, so why not add your own textures to make it more realistic?

Repeat this for the lower vertices.

Before we continue to make the legs, we will make use of a new button. Right-click the seat when not in edit mode so it is pink, and then go to the Edit buttons. In there is a SubSurf button. This will create a smooth surface between in-between vertices instead of a solid connected surface.

To create a leg, first leave edit mode and then go to the top view and add a circle. In front view extrude the circle down a little bit.

Extrude again a little further and scale the vertices down. Extrude these smaller vertices

further a little bit and then extrude further, again scaling the vertices but this time scaling up. Finally extrude a few more times to create a fatter bottom part of the leg. Now click on the SubSurf button and the leg will look smooth. Duplicate each leg by right clicking it and press **Shift-D**.

Next do the usual procedures such as moving the camera. Adding lights and you will end up with something similar that of **Fig8**. My stool is still quite limited and does not make best use of the lighting and materials, but it gives you a good start to refine the stool to perfection!

One thing important to notice is how the texture is applied to the ball. Notice how *Blender* has applied the woodgrain in quite a realistic fashion, and the light effects discussed earlier have still been applied to the new texture. You can see my render in **Fig6**. When adding more textures you can see the final result of my scene in **Fig7** on the left.

Conclusion

This month we have examined some of the issues that really bring life to our scenes. Although we have only taken a basic look at adding materials and textures to our scenes, it really gives an impression of how important they are. In future issues we will take a look at some more advanced uses of these facilities to create even more authentic scenes.

Until we resume next month, I recommend that you practice these modelling skills. Start creating some imaginative shapes in *Blender*, and try and model some household objects such as your TV, computer, washing-up liquid bottle and other everyday items. The practice of these skills will make the complex topic of animation far easier to deal with if you are fluent in the skills that we have covered so far. 

NEXT MONTH

We will begin taking a look at one of the key areas in *Blender* – animation; looking at simple key frame animation and then look at more complex animation with the use of virtual bones in our scenes. These concepts will form the core skills that we will use in future issues to develop our own animated shorts in *Blender*.




IMAGE EDITING

GIMP Selections

PART 6 Every project in whatever art app you choose will use Selections. Michael J Hammel selects the *GIMP* essentials...

Selections are the foundation upon which much graphic artwork is built. From object isolation to mask generation, selections are often the first tool pulled from the *GIMP* Toolbox when an artist starts a new project. But selections are more than just a tool. Combined with a scanned photo, channels and bucket fills they can be artwork unto themselves.

To cut flame from a candle an artist might reach for the Bezier Selection tool or perhaps the Freehand Selection tool. Each will suffice, though in different circumstances. But there are far more ways to isolate a shape than what you might find in the *GIMP* Toolbox. The flame burns red and blue and might be a candidate for Select by Color. Perhaps those colours aren't adequate, but their composite channels are – then a Fuzzy Select may be the answer. But what if these are too imprecise or time-consuming? Can't we just pluck that flame from the page?

Indeed we can, using a paintbrush and the Quick Mask. Selections are powerful tools for finding related pixels within an image. While most users will use a selection to isolate a shape or colour a box, complex selections can be used on existing images to create stylistic impressions that resemble hand paintings or coloured drawings. The complexity of the selection can be reduced for style or increased for detail, the latter pulling more pixel information while the former leaves more to the viewers' imagination.


The GIMP provides multiple selection options from the Toolbox: Elliptical, Rectangular, Freehand, Bezier, and Fuzzy Select tools. Outside of the toolbox you can find the Quick Mask on the bottom of any Canvas Window and Select by Color in the Canvas Window Select menu. Options for selections are found in the Select menu and in the various tool options dialog windows. Combining selections – adding, subtracting and creating unions – is done by holding the Shift, Control (Ctrl) and Shift+Ctrl keys, respectively, down before starting the selection. Constrained selections (perfect circles or squares) can be made by holding the same keys down while dragging.

This issue we'll look briefly at combining selections, since this is often one of the hardest processes to understand, along with using selections to create simulated drawings from real photographs. In the last tutorial we'll look at using the oft-overlooked Quick Mask.


Adding, Subtracting, and Unions: Selections as building blocks

Rectangular and elliptical selections can be useful on their own, but often it is a combination of two or more of these that gives the user the outline they need. A simple example of this is to place

stripes over a ball. With a few simple keystrokes added to the typical click and drag of the mouse, this process becomes another no-brainer.



1 We start by adding a transparent layer to a new Canvas window with a white background. Choose the Elliptical Selection tool from the Toolbox. Hold down the Shift key, click in the upper left of the Canvas window and drag to the lower right. This creates a circular selection. Double click on the Gradient Tool to get its Tool Options dialog box. Choose 'FG to BG' for the Blend option and 'Radial' for the Gradient type. With their defaults set (black and white, respectively), swap the foreground and background colors by clicking on the swap arrows. Click inside the selection toward the upper left and drag to the lower right. This creates a simulated 3D sphere in the transparent layer. In the Layers and Channels dialog, choose 'Alpha to Selection' from the Layers menu for this layer. We now have a selection around the sphere.



2 Add some vertical guides by clicking on the left side ruler and dragging into the Canvas. Make an even number of guides – here we added six in order to add three vertical stripes separated by gray area. Choose the Rectangular selection tool from the Toolbox. With the Shift and Control keys held down, click near the left-most guide outside of the sphere. When you start to drag, holding the mouse button down, release the Shift and Control keys. Clicking near the guide will cause the selection to snap to the guides. By holding the Shift and Control keys down when you start to drag the mouse you have told GIMP to create a union of the existing selection with the selection you are creating with the Rectangular selection tool. By releasing the two keys while you drag you allow your rectangular selection to be drawn with the lower left corner following the mouse. If you didn't release the keys, you'd end up with a perfect square instead, which is not what we wanted in this case. Once you release the mouse button, a selection is created where the rectangular selection and the existing selection overlap – the area of the sphere where the stripe will be added.



3 Bucket fill the selection from the previous step with red, setting the Bucket Tool Options blend mode to 'Color'. Repeat this process a couple of times to add more stripes. There are actually easier ways to do this – and less destructive – but this shows how to use the union option for combining selections. In practice, you might fill the selections in their own transparent layer and change the Layer Blend mode to color instead in order to prevent changing the sphere while still generating the desired composite effect.

Select by color: Pencil drawings made quick and easy


The artistic side to selections includes using them as the basis for simulated pencil sketches. The trick here is to use multiple selections to pull in just the right amount of the original image. The Select by Color dialog is perfect for this sort of work.



1 Start with any decent photo. While the photo doesn't have to be the best quality, it helps to have good content. We've chosen a balloon launch with several balloons rising above a lake. Desaturate the image. Add a white layer but turn off its visibility. Add a transparent layer above that. Open the Select by Color dialog. The dialog will hold a black preview showing that no pixels have been selected. Make sure the Selection Mode is set to Replace.



Select by color: Pencil drawings made quick and easy (continued)



2 Click in the image and check the preview in the Select By Color dialog. Adjust the Threshold and try again till you get an image that looks good in the preview, but does not include the entire original image – we're looking for a sketch here, not a painting. Note that the white areas in the Select by Color dialog will be black in our sketch.




3 Close the Select by Color dialog with a selection still active. Make the transparent layer active by clicking on its name. Be sure that the Keep Transparency button is not set for this layer in the Layers and Channels dialog. Now bucket fill with black in the selection in this layer. Use Ctrl-Shift-A to turn off the selection, then turn on the visibility of the white layer. Viola! Instant pencil drawing. You'll have to experiment a few times and might need to repeat the selection/bucket fill steps a few times to get the image you want. The trick here is to find a selection that pulls out the interesting parts of the original photo.


Channel selections: Going from photo to drawing

We can expand on that simple drawing and create a painting with only a few added steps. Working with the Fuzzy Select tool will


help with creating selections. We'll finish by adding a few blurs and applying a canvas to the final image.




1 Start with the same photo, but this time don't desaturate it. Open the Layers and Channels dialog and go to the Channels page. Turn off the visibility (click on the eye icon) in two channels and look at the image using only the third, let's say the blue channel. Now try a different channel, let's say just the red channel. Look at what parts of the image stand out in each channel. We'll be making Fuzzy Selections in different channels to grab different parts of the image.



2 We've turned off the Green and Blue channels by clicking on their channel names. With the Fuzzy Select tool active, click and drag around the image - in this case we clicked and dragged around the balloon to select part of it. We don't want the whole balloon, however. So we use the Fuzzy Select's Tool Options and reset the Threshold, then hold down the shift key and click and drag in another part of the balloon. We did this several times to get a selection that outlines the balloon but does not grab it completely. The bottom right image shows the selection in the red channel and the upper left image shows that selection copied from the background, pasted into its own layer and placed above a white layer. The painting is taking shape.



3 Repeat the selection process in different channels, targeting other key components of the original image. Copy each selection and paste as new layers each time. Make sure all channels and the white layer are visible to preview your work. Turn off the pasted layers and the white layer and make sure the original layer is active before trying to make your copy from the selection. The results vary with your selections. Here we made selections of the balloon, the shoreline and basket, the sky and the water each as separate layers. The image is good as it is, but we can add a little artistic flair to it.




4 Save a copy of this as an XCF file just so you can come back and make changes later if you want. Then flatten the image into a single layer. Apply a very slight Wind (Filters > Distorts > Wind), with strength and threshold set to 1 and the both edges affected. Next, choose a Gaussian Blur RLE (Filters > Blur > Gaussian Blur(RLE)) using width and height of 1 pixel. Finally, apply a canvas texture (Filters > Artistic > Apply Canvas) of no more than 2 pixels in depth. Our results are pretty good. We've zoomed in to show the detail of the painting. Adding wind provided a light sense of brush strokes. Blurring took out any remaining photographic resemblance.


Quick Masks: Paint and isolate with the "Unknown Tool".

Enough of this artsy stuff – what about real selections? What is the easiest way to isolate a complex shape? The answer comes


from a seldom-mentioned little red box found at the bottom of every Canvas window: the *GIMP* Quick Mask.



1 The problem with some selection tools in any art application is their imprecise nature. Selecting a single orange in this image using 'Select by Color' would get lots of oranges. We can trim this down to just one orange by making a rectangular selection and then setting the Selection Mode to 'Intersect' in the Select by Color dialog. The result isn't bad, but it just doesn't get the whole orange. Multiple attempts are necessary to really get close to getting the whole orange and nothing else.



2 Fuzzy Select can be used to select nearby pixels that are close in colour to the first pixel you click on. Dragging across the orange from the brightest to darkest areas also gets close, but when you get near the edge of the orange you start to pick up the neighbouring orange as well. A tip while working with Quick Masks – zoom in on the area to select. This will make it easier to be precise while painting your quick mask.



3 The fastest way to get that one orange is the Quick Mask. Start by clicking on Quick Mask button on the lower right of the Canvas window (the little box with the red outline). This causes the image to be overlaid with a reddish tint. Next, choose a brush. Hard-edged brushes will work better for Quick Masks. Make sure the Quick Mask channel is active in the Channels page of the Layers and Channels dialog. Now you're ready to start painting over the orange to make your selection. Paint with white to expose the orange (and create the area to be converted into a selection) and paint with black to replace the reddish tint (re: the mask). This will make the orange show through the red veil created by the Quick Mask. The area that shows through will become the selection. When you're done painting, click on the selection box to the left of the quick mask. Your selection is ready. [LXF](#)

FILE MANIPULATION

Practical PHP programming

Paul Hudson finishes production of a PHP extension to handle tar files – at last!

Welcome back, PHPers! If you've missed the last two issues, we've been creating a PHP extension to handle tar files. Creating extensions is a particularly tricky topic as far as the PHP language goes, which is why we've been going over it in such depth. If you recall, part one of this mini-series saw us creating a basic extension using `ext_skel`, modifying the `m4` configuration file, and finally compiling smoothly with the rest of PHP. In part two, the extension was improved to include support for the first new function, `tar_list()`, which returned an array of the names of all files contained inside a tar archive, and we also looked at what made up a `zval`. In this concluding part to the mini-series, we'll be adding the final two tar functions, `tar_add()` and `tar_extract()`, then reviewing how the PHP extension process works as a whole, and finally looking at ideas for other extensions.

If you didn't quite grasp how `tar_list` works from last issue, it is strongly recommended that you go back and re-read last month's tutorial, as we'll be building on much of the same knowledge in this issue. If you've made it this far you're onto the home straight – you're already able to create and compile your own extensions, we're just going to be looking at a few more advanced topics this issue as well as rounding up the subject of extensions.

As with the prior two tutorials in this mini-series, knowledge of C is a requirement – don't worry, next month's PHP tutorial will be much easier!

Tarred and feathered

Of the two functions left for us to implement, `tar_extract()` is definitely the easiest, so we'll be starting with that. The function

prototype we'll be working to is:

```
bool tar_extract(string tarfile, string location)
```

which means that `tar_extract()` takes the name of a tar file as parameter one, where it should extract the tar archive to as parameter two, and returns boolean true if it was successful or boolean false if there was an error. As always with PHP, the first changes to make are to declare the functions inside the extension, so we'll get that out of the way first.

Open up `php_tar.h` (you may want to take backups of all your existing files, in case things go wrong!), and look for the line `PHP_FUNCTION(tar_list)`. Below that, you need to add two more lines to define the new functions we'll be adding, so it should end up looking something like this:

```
PHP_FUNCTION(tar_list);
PHP_FUNCTION(tar_add);
PHP_FUNCTION(tar_extract);
```

You'll need to edit `tar.c` also, so that the `tar_functions` list array looks like this:

```
function_entry tar_functions[] = {
    PHP_FE(tar_list, NULL)
    PHP_FE(tar_add, NULL)
    PHP_FE(tar_extract, NULL)
    {NULL, NULL, NULL}
};
```

With those changes, we're all set to go ahead and program the `tar_extract()` function. You can use `tar_list()` as your base to save yourself quite a bit of typing, because large amounts of the functions are the same. Enter this next function into `tar.c` below `tar_list`:

```
PHP_FUNCTION(tar_extract)
{
    TAR *t;
    char *tarfile = NULL;
    int tarfile_len;
    char *location = NULL;
    int location_len;

    if (zend_parse_parameters(ZEND_NUM_ARGS()
        TSRMLS_CC, "ss", &tarfile, &tarfile_len, &location,
        &location_len) == FAILURE) {
        return;
    }
}
```

New PHP version

From 4.3.1 to 4.3.2

PHP 4.3.2, the first "proper" patch to the 4.3 tree is now available. 4.3.1 fixed just one bug, which had been fixed in earlier releases and broken in 4.3.0. So, 4.3.2 is the first meaningful patch, and it's a "big" difference – so big, in fact, that this is the first time since

PHP 4.0 that the upgrade has been described as "strongly recommended" on the PHP homepage. This new release incorporates over 150 (yes, one hundred and fifty) key bug fixes in the language, which should make everything that much more stable and reliable, while

at the same time not affecting backwards compatibility in the least. "Strongly recommended"? 100% – this thing went through five release candidates before being declared final, and so is a massive improvement over all existing versions.

```

if (tar_open(&t, tarfile, NULL, O_RDONLY, O, TAR_GNU) == -1)
{
    php_error(E_WARNING, "%s", strerror(errno));
    RETURN_FALSE;
}

if (tar_extract_all(t, location) != 0)
{
    php_error(E_WARNING, "%s", strerror(errno));
    RETURN_FALSE;
}

if (tar_close(t) != 0)
{
    php_error(E_WARNING, "%s", strerror(errno));
    RETURN_FALSE;
}

RETURN_TRUE;
}

```

If you compare that to `tar_list`, you'll see the changes are two new variables to hold the extra parameter, a minor change to the call for `zend_parse_parameters()` also to handle the extra parameters, and the `"while ((i = th_read(t)) == 0)"` code being replaced with an `if` statement. The two new variables are the `char*/location/` and the integer `/location_len/` – these are to store the second parameter string (where to extract the tar file contents to) and its length. The third parameter (not the second, remember that `TSRMLS_CC` is a macro that has a comma in) is now `ss`


Memory Leaks and You

Grokking mallocing, easy emallocing

One of the most commonly cited "killer features" of C is that it allows you complete control over your memory usage – you assign and free memory as you please, and can be very exact about the position your program is in. However, entirely contradictory to this state of affairs is PHP, where one of its killer features is that you don't have to worry a whit about what happens to memory when your script runs – every variable, object, or array you create is automatically freed at the end of your script. These two are absolute polar opposites – how does PHP handle making sure errant extensions don't haemorrhage memory?

Well, the answer is simple: whereas you and I are used to using `malloc()` and friends to grab memory for our uses inside a C program, *Zend* has its own memory assignment functions that are abstracted from `malloc()` *et al.* It's owing to this abstraction that the *Zend Engine* gets full control over all memory allocations, and is therefore able to determine whether a block is in use, which means it can automatically free unused blocks and blocks with lost references. By keeping such very close track of memory in scripts, the *Zend Engine* can entirely prevent the memory leaks that often plague standard C programs.

Naturally, functions such as `malloc()` and `realloc()` can be called from inside your extension whenever you want data stored in memory to outlive execution of your script, but on the whole you should use the *Zend* replacements, `emalloc()` and `erealloc()`. If you desperately do need to use the non-abstracted functions, and you're aware of all the potential issues that might arise from you working in a fully multi-threaded environment, then you should step very carefully – memory leaks can be catastrophic in a web server environment because `mod_php` will end up taking up hundreds of megabytes of memory and may well result in a denial of service!



rather than `s`, because we're expecting two string parameters being passed in – hence the addition of `&location` and `&location_len`.

The key change, of course, is the new `if` statement: `if (tar_extract_all(t, location) != 0)`. Previously we called `th_read()` to loop through each file in the archive, but this is not necessary as *libtar* provides `tar_extract_all()` function to extract all the files. `tar_extract_all()` takes two parameters, which are the `TAR*` file that is an opened archive, and the location to extract the files to – we use the `TAR*` we declared earlier, and the `char*` location that, hopefully, our user passed in. If there was a fatal error extracting the files, `tar_extract_all()` will return `-1`, and if this happens we use the PHP macro `RETURN_FALSE` to return boolean false from the function and exit.

If the tar file is opened, extracted, and closed without hitch, we use `RETURN_TRUE` to return boolean true from the function, and we're done – that was easy! It won't compile just yet, because we've added the definition for `tar_add()` but have yet to add the implementation. Let's do that now...

Adding to the tar pit

We've now got `tar_list()` and `tar_extract()` working, which just leaves one function: `tar_add()`. If you recall last issue, you'll remember the prototype given was this:

```
bool tar_add(string tarfile, string file_to_add)
```

We'll be changing that just slightly, to this:

```
bool tar_add(string tarfile, array files_to_add)
```

This change is for two reasons: firstly, most tar files contain more than one file, simply because there's no point having it otherwise. Secondly, it gives us the chance to look at how arrays are handled in extensions. Thirdly, if users pass an array of one item, it works the same way anyway – everyone's a winner! Here's the first draft of code we'll be working with:

```

1 PHP_FUNCTION(tar_add)
2 {
3     TAR *t;
4     char *tarfile = NULL;
5     int tarfile_len;
6     zval *arr;

```

Syntax highlighting IDEs make the world go round. Not that it ever was any other shape...

```

7  zval **entry;
8  HashTable *harr;

9  if (zend_parse_parameters(ZEND_NUM_ARGS()
TSRMLS_CC, "sa", &tarfile, &tarfile_len, &arr) == FAILURE) {
10     return;
11 }

12 if (tar_open(&t, tarfile, NULL, O_WRONLY | O_CREAT,
0644, TAR_GNU) == -1)
13 {
14     php_error(E_WARNING, "%s", strerror(errno));
15     RETURN_FALSE;
16 }

17 harr = HASH_OF(arr);

18 for (zend_hash_internal_pointer_reset(harr);
zend_hash_get_current_data(harr, (void**)&entry) ==
SUCCESS; zend_hash_move_forward(harr)) {
19     convert_to_string_ex(entry);

20     if (tar_append_file(t, Z_STRVAL_PP(entry),
Z_STRVAL_PP(entry)) != 0) {
21         fprintf(stderr, "tar_append_file: %s\n",
strerror(errno));
22         return -1;
23     }
24 }

25 if (tar_close(t) != 0)
26 {
27     php_error(E_WARNING, "%s", strerror(errno));
28     RETURN_FALSE;
29 }
30 }

```

As you can see, this code is quite a bit more complicated than both `tar_list()` and `tar_extract()`, because we need to deal with PHP's array objects. Internally, a PHP array is a hash – a set of data with a key and a value each. If you were curious, an object is also a hash internally, and works in much the same manner as arrays. The functions `zend_hash_*` all, unsurprisingly, manipulate hashes inside C code, and there's a massive collection of hash-related functions available for you to use – check out `zend_hash.h` for a list.

On lines **6** and **7**, two `zvals` are declared to hold the new information that will be passed in. The `zval` `arr` will hold the array `zval` in its entirety (see the box in last issue, *Anatomy of a zval*), `entry` will hold an individual value from inside the array, and `harr` is declared to hold the hash table of our array. The hash table holds the actual array data – keys and values – of a `zval`, which is the key information we want.

LXF website redesign

At last! And using PHP of course

As I'm typing this, the grand redesign is still being planned – we want to get it "right". All being well, the whole project will be Open Source (and community-developed) pretty much from the off, and your thoughts and input are most welcome –

check out the forum "This Website" on www.linuxformat.co.uk to keep up-to-date with announcements. *PostNuke* is going to go, and hopefully we'll be able to increase performance by several notches while we're at it!

Questions from the forum

Choosing a PHP editor

There has been some discussion on the LXF forums regarding choosing a decent PHP editor – it might sound easy at first, but there are lots of options. *Quanta* is a good program, but it has quite a way to go – particularly with regards to performance. *Zend Studio* is the best out there, but unfortunately it's not free – if you don't mind laying out some money, you can't do better.

I always recommend people give *phpmole* a try, because they're usually surprised that a PHP editor written in PHP works so well. *vim* and *emacs* work quite well, particularly for those who are seriously used to them. I myself have a free licence of *Zend Studio*, but (and I think this probably says a lot) I still use *Kate*!

Line **12** is slightly different to our previous calls to `tar_open`, because this time we open the file as `O_WRONLY` (write only) **OR**ed with `O_CREAT` (create if non-existent), and we also pass in `0644` as parameter five, which sets the file to read/write for us, and read for others. If you find you have no permissions to modify the tar file created by PHP scripts, make sure you don't have `O` for this parameter!

On line **17** we use the macro `HASH_OF` to take the hash from our array, `arr`, and place it into `harr`. `HASH_OF`, found in `zend_API.h`, simply returns the hash table of a given `zval`, whether that be an array or an object. From there, we hit the complicated part of the function: iterating through the array. We use three functions to form our loop:

```

zend_hash_internal_pointer_reset(),
zend_hash_get_current_data(),
and
zend_hash_move_forward().

```

As you know, *for* loops take three arguments – definition, condition, and action – which means that this loop starts by resetting `harr`, sets `entry` to the current value the internal pointer of `harr` is at, and moves the pointer forward one place, iterating through the array, as its action. The end result is that the loop goes from the start of the array to the end of the array, picking out the values as it goes. Inside the array, the first thing we do is call `convert_to_string_ex()`, which is another *Zend* macro (although not prefixed with `zend_` like many of the others) that makes sure a given `zval` has its type set to `IS_STRING`. It takes a pointer to a pointer to a `zval`, which is how we declared `entry`, so all is well.

Lines **20** to **23** handle the tar-related code, and uses a new function, `tar_append_file()`, to do the hard work. `tar_append_file()` takes three parameters – the TAR* file to write to, the name of the file to add, and the name the file should have within the archive. As with the other two tar functions, we're using `t` for our TAR* file, so parameter one is a doddle. Parameter two and three are the same as each other – we read files in and store them using the same name. `Z_STRVAL_PP` is a macro to a macro that returns the string value of a `zval`. In last issue I mentioned the macro `Z_STRVAL`, which returns the string value of a `zval`. `Z_STRVAL_PP` uses the `Z_STRVAL` macro, but also dereferences the pointer pointer, saving two characters of typing – sometimes programming with PHP can be quite confusing thanks to the extreme use of macros!

If `tar_append_file()` fails, it will return `-1` – we catch this and return an appropriate error message. So, the loop as a whole iterates through the array passed in, adding each file in the array to our tar file – quite straightforward once deciphered. Lines 25 to 30 run the usual `tar_close()` routine to clear up before the function exits.

That's the new function complete, so we can now go ahead and recompile PHP to take advantage of the new functions `tar_add()`, and `tar_extract()`. Once it's recompiled, you can test the new functions out by creating a few dummy files, and running a basic script along the lines of this:

```
<?php
$array[] = "index.html";
$array[] = "index.html2";
$array[] = "index.html3";

tar_add("mytar.tar", $array);
tar_extract("mytar.tar", "foobar");
?>
```

As you can see, we add three dummy HTML files to an array, then pass it into `tar_add` with parameter one being "mytar.tar" to create a new tar archive, then extract the new archive to the directory `foobar`. Quick, easy, and powerful – once all the hard work is done, it's easy to re-use your code.

Extensions Round up

OK, it did take three issues-worth of explanation, but the end result is that we've managed to create an all-new PHP extension that adds something valuable to the language, and it really wasn't that hard after all. Using `ext_skel`, it's a cinch to get started with extensions because all the drudgery is done for you – to get going, all you really need is an idea and a few man pages discussing how a library works. Once `ext_skel` has done its work, it's merely a matter of planning out the functions you want to write – most of the things that sound complicated, such as making sure variables are passed in in the right order, are done automagically by the *Zend Engine*. The more you write for PHP, the more you appreciate the power and freedom the *Zend Engine* gives you – as long as you're smart, there's little need to worry about memory leaks, user errors, and such.

Once you're able to get around the PHP source code easily, you'll also find it gives you a greater understanding of the language in general, increasing your skill level at writing PHP scripts. So much goes on behind the scenes in PHP that it's nearly always worthwhile checking how it works in the source code to make sure you're not calling code that's more complicated than you realise. Another good reason to be competent with the PHP source code is that you no longer have to worry about manual pages being out-of-date. Some pages still

Make your mark

Brainstorms 'R' Us

Would you like to get your name in the mag and learn about stuff you're most interested in?

We're looking out for ideas for new *Linux Format* Practical PHP tutorials, and where better to look than to you, the reader? If, while reading past issues of *Practical PHP*, you've thought "I wish they'd covered XYZ in more depth...", or "I really want to know how to use...", then now's the time to get your voice heard!

Send me your ideas to paul.hudson@futurenet.co.uk – all the good suggestions you send in will be covered in future issues. So far, the topics we have covered in some depth include MySQL, XML, CLI, GUIs, media generation, templates, extensions, and more. If you're short of ideas or want to know more, you're welcome to write in with comments about prior issues – we're always looking to improve the overall quality of tutorials.

PHP 5...?

Big changes over the horizon

PHP 4.3 may well be the last minor version from the PHP 4.x branch, which would be no surprise given that 4.0 was launched way back in May 2000 – it's served everyone well, to say the least! Between now and 5.0, there will be a 4.3.3 release to add features that were left out in the long 4.3.2 feature freeze, and probably even a 4.4/4.5 as a stop-gap release before 5.0. In the meantime, extensive amounts of work are being done on PHP5 to make it as significant an upgrade as PHP 4 was to PHP 3. If reading the php-dev mailing lists aren't your

thing, try visiting www.zend.com/zend/future.php to read a selection of features written by Zeev Suraski regarding the new functionality in PHP 5 – there are some *big* changes coming, so it's best to get on the ball now so you're *au fait* with all the new stuff before it happens!

We'll be covering PHP 5 in depth once it gets close to release. For now, it's best that you at least make sure you know what will change, how it will affect you, and whether it's best for you to try to bring your code in line now to save hassle later.

have "no description available" or flat-out incorrect data – once you're able to jump into an extension and see for yourself what it does, this becomes a non-problem.

If you're stuck for an extension to write, just take a look through the lib directory on your PC to get some ideas. Debian users can type `apt-cache search ^lib` to get a good list of possibilities, or, if you're an RPMer, try something like `rpm -qa | grep ^lib`. As you can see looking through the PHP documentation, people have written extensions for all sorts of weird and wonderful things – from writing GUIs using GTK to writing Java apps with, well, Java.

You might find it easier to try making the *libtar* extension more powerful – there are several other functions in *libtar* that we haven't looked at at all here, such as using *gzip* to make *tar.gz* files, and this might be a good starting point for the slightly less adventurous. There's a lot that can be added to the extension to make it more flexible and user-friendly – adding files that don't exist generates an unhelpful error message, there's no checking to make sure a tar file doesn't exist before it's created, or, worse, that it's been created by other users and we're likely to get a "permission denied" error back. Try to be flexible – *libtar* is capable of lots, yes, but there's even more you can do (adding entire directories in one call, anyone?) just by playing around with new ideas.

Conclusion

Extension writing isn't hard, as I'm sure you can now see. The PHP system comes bundled with a massive array of macros that are there to make life easier for those wanting to jump into the language. Sometimes they make life a little hard – particularly when you get compilation errors that are nigh impossible to track down thanks to macros being used so extensively. However, on the whole, adding to the PHP language can be fun and rewarding. You get to extend your favourite language with new features otherwise only available to locally compiled programs, and you get to learn a lot about the internals of your system to boot. Marvellous!

Delving into extensions is probably the hardest thing we'll be covering in this series – partly because there are few things more difficult than writing extensions, but partially because we have only received a couple of tutorial suggestions so far, which suggests that I may have taught you all you need to know, young grasshoppers! Read the box on the left, and send in your ideas... [LXF](#)

NEXT MONTH

We'll be looking at optimising PHP and MySQL to make your websites fly. On the whole, this is a much easier topic than extension writing, and I think you'll all be thankful for the break! If your SQL theory is a little weak, don't worry – we'll be starting off with some very basic "gotchas", and working our way up to more complicated techniques.

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



Fetchmail doesn't

I am wrestling with this email lark... I think I am sending emails OK, and I seem to be receiving emails in my ISP's mailbox. However, when I try to collect them with *fetchmail*, they disappear into thin air! What am I doing wrong?

Tony Bradley, via email

When *fetchmail* downloads a message from your ISP, it will send it through the local MTA. It would appear that your MTA is either not configured, or it's sending your email into a black hole.

The first step is to check the mail server logs. This would be `/var/log/mail.log` or `/var/log/maillog` depending on the distribution and the specific *syslogd* configuration. Once you know which logfile is for your MTA, you can `tail -f /var/log/mail.log` and then run *fetchmail*. Hopefully, you will be told in the logs what your MTA is doing to the messages.

You can also run `fetchmail -vvv` to get debugging output from *fetchmail*. This will tell you exactly what it's doing, in terms of contacting your MTA and who it is delivering the message to – as we don't have your *fetchmailrc* configuration, it's difficult to establish if there is a fault with your settings or not.


Dependency blues

I run SuSE Linux 8.0 Pro and bought your magazine yesterday, eager to upgrade my KDE3.0 to 3.1. So, after `cd`ing to Desktop/KDE3.1/SuSE on the CD-ROM, I typed

```
rpm -Fhv *.rpm i18n/*.rpm
```

as stated on page 99 of the magazine. However, nothing was installed, and instead I got:

```
error: failed dependencies:
libvorbisenc.so.2 is needed by arts-1.1-49
libvorbisfile.so.3 is needed by arts-1.1-49
```



fetchmail pulls mail from a remote system, but if the local MTA is misconfigured then it will disappear.

`libvorbisenc.so.2` is needed by `kdebase3-3.1-51`

`libvorbisfile.so.3` is needed by `kdebase3-3.1-51`

`libXm.so.3` is needed by `kdebase3-nsplugin-3.1-51`

`libvorbisenc.so.2` is needed by `kdelibs3-3.1-46`

`libvorbisfile.so.3` is needed by `kdelibs3-3.1-46`

`libkateinterfaces.so.0` is needed by `kdeaddons3-3.0-5`

When I try to install RPMs that contain these files, other programs start complaining they need the (older) versions of them, that are already installed on my computer. How should I solve this?

Diederick de Vries, the Netherlands. Ideally one would solve this issue by upgrading the packages which require earlier versions of *libvorbis* and *libXm*. By upgrading these packages to the same version as required by KDE3, then the errors which you are seeing would disappear.

You can either force the installation of a new *libvorbis* and the symlink, the more recent libraries to the names which existing software will look for. Of course, this assumes that the two *libvorbis* versions are backwards compatible and a piece of software can use either.

Without upgrading the packages which require the older version of *libvorbis*, then every solution is going to be nasty and break the SuSE package system somewhere, either by installing software which is not from a package, or by forcing a package to install without checking dependencies.

TCO Nothing?

Our small firm has a network with Windows and a server running Linux (Red Hat 8). Works great! There are about 15 of us here, all using different computers during the day and some 10-15 casuals during the peak times. The server runs *Samba* as a PDC that makes



The `make xconfig` command creates the `.config` file for the kernel

passwords easy as they are on the server. I intend (hope!) to upgrade the Windows workstations (Win9x, Me) for Linux using *OpenOffice.org* or KDE's offering. Now, is there a util available so that all the new Linux boxes log onto the server *aka* Win? I can't see myself having to set up all those passwords on each and every Linux box! So, can I give the boot to Windows (and the licence fees if we upgrade)? Your help is much appreciated!

Steve Coleman, via email

Unix has a system called *Yellow Pages*, which is now known as NIS. This is a system which allows certain configuration files on a server, typically

`/etc/passwd` and `/etc/group` to be distributed over a network. When you login into a NIS client, it will contact the NIS server for user information, rather than checking `/etc/passwd`.

Setting up NIS is fairly straightforward and documentation can be found at www.linux-nis.org/nis-howto/

Mandrake

Just over a month ago I decided to buy a copy of Mandrake 9.0 from PC World. Installed it – no problems apart from my ADSL connection (not very happy with that) so I spent the next two days trying to find drivers for my

Conexant Access Runner (USB) and came across <http://sourceforge.net/projects/accessrunner> where I downloaded all the necessary files, then I extracted them to a local dir where I typed `make` and this is what I got:

```
[root@localhost root]# cd
/root/drivers/cxacru
```

```
[root@localhost cxacru]# make
```

```
cd init && make && make install
```

```
make[1]: Entering directory
```

```
~/root/drivers/cxacru/init'
```

```
make[1]: Nothing to be done for
`all'.
```

```
make[1]: Leaving directory
```

```
~/root/drivers/cxacru/init'
```

```
make[1]: Entering directory
```

```
~/root/drivers/cxacru/init'
```

```
install -c -m 755 -p cxload cxioctl
/usr/sbin
```

```
make[1]: Leaving directory
```

```
~/root/drivers/cxacru/init'
```

```
cd firmware && make
```

```
make[1]: Entering directory
```

```
~/root/drivers/cxacru/firmware'
```


```
install -c -m 644 -p cxinit.bin
/usr/sbin
```

```
make[1]: Leaving directory
```

```
~/root/drivers/cxacru/firmware'
```

```
cd module && make && make
```

```
install
```



`usbview` easily establishes what devices are connected to the USB controller.

A QUICK REFERENCE TO: Patching

Anyone who has used Linux for more than a couple of months will have heard of *patches* – Small updates to code to fix bugs, add functionality, or simply to avoid users having to download the entire new version of a piece of software if only a few dozen lines of code have changed.

patches deal with text, so they are mainly used by developers and those of us who delve into the source code now and then to build a bleeding edge piece of software. Simply speaking, a *patch* tells us which lines have changed, and where to find those lines in the existing text file. *patches* are usually referred to as *diffs* as they contain the differences between two files. A basic *patch* looks something like:

```
--- net/ipv4/tcp_ipv4.c.orig
Mon Dec 11 00:49:44 2000
```

```
+++ net/ipv4/tcp_ipv4.c      Thu
Dec 14 17:11:59 2000
```

```
@@ -113,7 +113,7 @@
```

```
* For high-usage systems, use
sysctl to change this to
```

```
* 32768-61000
```

```
*/
```

```
-int sysctl_local_port_range[2] = {
1024, 4999};
```


```
+int sysctl_local_port_range[2] = {
1024, 9999};
```

```
int tcp_port_rover = (1024 - 1);
```

```
static inline int
tcp_hashfn(__u32 laddr, __u16
lport,
```

It may appear cryptic at first, but it's only the first character of each line which tells us something is happening. In this example, it removes one line and inserts another in its place.

Of course, we don't manually *patch* files from a *diff*. Instead, the *patch* utility




The `patch` command allows source trees to be updated easily using a compressed *diff*.

does it all for us. By feeding *patch* a *diff* on *stdin*, it will attempt to patch the files described in the *diff*. *patch* is usually quite intelligent and will figure out where to put a *diff* if the line number changed or something is a little different, but if it fails it will dump

the failed patch to a *.rej* file allowing us to attempt fix it manually.

```
# cd /usr/src/linux
# patch -p1 < ~/tcp-port-range.diff
or
# bzcatt ~/tcp-port-range.diff |
patch -p1
```

FREQUENTLY ASKED QUESTIONS X NETWORKING



FAQ When my system boots up the display is very small and scrolls around as I move the mouse. How do I stop it scrolling AND set the display at 1024x768.

X supports a number of different resolutions, so your display may know about 640x480, 800x600 and 1024x768. As X can't dynamically modify the size of the X display, only switch resolution of the video card, it will set the size of the desktop the maximum size, then set the video display resolution to the first one in the list.

If the first listed mode is '640x480', then X will allow you to scroll around the desktop to experience the full 1024x768 display. The solution is to list the highest most first in the 'Modes' list for the appropriate bit-depth in /etc/X11/XF86Config-4

Section "Screen"	
Identifier	"Default Screen"
Device	"Generic Video Card"
Monitor	"Generic Monitor"
DefaultDepth	16
SubSection "Display"	
Depth	16
Modes	"1024x768" "800x600" "640x480"
EndSubSection	
EndSection	

FAQ I have two systems, but only one has a monitor. Can I run X applications over the network on the headless box and control them on my workstation?

X is a totally network dependent system, and every X program which

runs, from xterm to gnome-panel communicates with the X server via a network connection. For efficiency purposes, local clients connect via Unix domain sockets, but remote clients can connect via TCP sockets over IP networks.

If the workstation has the IP 10.1.1.1 and the other system has the IP 10.1.1.2, then we need to permit 10.1.1.2 to connect to our X server, which is running on 10.1.1.1. This is done with xhost:

```
$ xhost +10.1.1.2
```

'xhost' must be executed from within X, otherwise you will not have permissions to modify the xhost list. You can now shell into 10.1.1.1, using telnet or ssh, then run the appropriate client with a modified DISPLAY variable pointing at our 10.1.1.1 X server.

```
$ DISPLAY=10.1.1.1 xterm &
```

Alternatively, from within an

The X desktop supports a wide range of network capabilities, but it often likes bandwidth.

existing X session, such as xterm, we can ssh to 10.1.1.2. ssh supports X forwarding, which allows us to run X clients transparently on a remote system and have them appear locally.

FAQ Is it possible to connect to an existing X server, so I can control applications from two systems?

As X is assigned to a video card, there is no easy way to view the X display on a remote box. However, one can run an alternative X server which does not rely on a video card. The tool VNC (www.realvnc.com) runs a virtual X server, which one can connect to over an IP network using a VNC client,

```
<< make[1]: Entering directory
~/root/drivers/cxacru/module'
Makefile:7: /usr/src/linux/.config: No
such file or directory
make[1]: *** No rule to make target
~/usr/src/linux/.config'. Stop.
make[1]: Leaving directory
~/root/drivers/cxacru/module'
make: *** [CX_MODULE] Error 2
[root@localhost cxacru]#
```

So I decided to go to the forums where I was told I needed the following RPMs: *atm-2.4.0-1mdk.i586*, *libusb0.1_4-0.1.6a-2mdk.i586* and *libusb0.1_4-devel-0.1.6a-2mdk.i586*. I don't know why - I installed the full package when Mandrake installer asked. Luckily I still have XP to access the Net, but I

would like to scrap it altogether. Until my ADSL is sorted for Linux, this does not look like it's going to happen. So what is it that am I doing wrong?

Kai, via email

It would appear that you are missing a '.config' in your /usr/src/linux source tree. You will need to install a RPM containing the kernel source tree, so that you can compile the kernel module against the correct kernel source. The package will be called something similar to *kernel-source-2.4.19-32mdk.i586.rpm*, although you must ensure that the kernel version matches that from `uname -r`, otherwise the module will not be compiled against the correct kernel.

Red Hat startup

My situation is that I want to have Red Hat 9 running on a machine. The machine is a Compaq Armade E500, 600 MHz Pentium III and 128 MB RAM. Processor is enough but the RAM falls a little short when I use GNOME and the computer starts stuttering.

So, now I am trying out some of the window managers. One of the things that always happens when installing the window manager from source is that it doesn't make an entry to the GDM. Sometimes the RPM install does make an entry which of course helps, but not all of them. How do I add an entry to GDM? Or will it be the best to boot

too a terminal and then startx? But how does startx know which to start? I've heard suggestions to add stuff to .xinitrc and .xsession but where exactly so that it doesn't screw up any of the stuff that is already there?

Here's some ideas for a tutorial. *Fluxbox*, *Blackbox* etc covering installation and configuration; adding some *wmapplets*. The screenshots of the users are really great but the documentation to get it looking like that is not easily available or I find that what is covered goes over my head.

Louis Botha, Finland

Rather than mess with *gdm*, it's much easier to simply set *.xsession* to run

such as *vncviewer*. Multiple VNC clients can connect to a single VNC server, so it's easy to switch between systems without causing disruption to the X display.

As one would expect VNC is rather bandwidth-hungry, so there is an alternative known as *TightVNC* (<http://tightvnc.sf.net/>) which provides particularly effective compression, allowing VNC to function over a low-bandwidth connection. It's never very clever to run highly graphical or quickly changing applications over VNC, such as *The GIMP* or *Xine*, but with a fairly plain desktop VNC is functional over low-bandwidth connections.

FAQ I have two workstations, each with a monitor. Can I control them with a single keyboard and mouse?

There is a great utility called *x2x* (<http://handhelds.org/z/wiki/x2x>) which allows a remote X display to be controlled with a local keyboard and mouse. The pointer can be warped at either side of the display, so two displays can behave like a multi-head system. Windows can't be dragged from one side to another, as they can be on a system running *Xinerama*, but everything else works the same, including the clipboard. For example:

```
x2x -to 10.1.1.3:0.0 -east
```

This would warp the pointer to display 10.1.1.3:0.0 when the mouse moves to the right, or east, of the current display.

the desired window manager when you login. *.xinitrc* is used when X is started via 'startx' and *.xsession* is used when you login via *xdm* or *gdm*, so the former can be ignored.

As default *.xsession* should not contain much of anything, but you can simply create an empty file and put in a single line:

```
windowmaker
```

Note that the file must be granted executable permissions with **chmod +x .xsession**, otherwise it will be ignored. If you want to run more than just the window manager, then you can include other commands:

```
kicker &
enlightenment
```

Any command before the window

manager must be followed by an **&** symbol so they run in the background, but the window manager must be the last line in *.xsession* and *not* have a **&** at the end, otherwise X will simply exit when you login.

SuSE slipping

I have two computers connected through a KVM switch box to one screen, one mouse and one keyboard. Before, I could switch from one computer to another without any problems running SUSE 8.0. Now I've updated to SUSE 8.2, when I switch over and back to the SUSE 8.2 computer, the screen is working more-or-less by itself. Menus pop up and close, the mouse is dragging icons instead of opening them, when I move the mouse and pick it up the pointer jumps to the bottom left-hand corner of the screen – control over the mouse is more-or-less lost.

I have to reboot the computer to make it right. I've found on different forums the same problem but unfortunately without solutions. Can you help me?

Marcel Bosch, *The Netherlands*


The first thing to check would be to ensure that the mouse is configured correctly under SuSE 8.2. The symptoms you are experiencing are typical of an incorrectly configured PS/2 mouse. For example, it may be configured as IMPS/2 but not be an Intellimouse, or *vice versa*. Check */etc/X11/XF86Config* to see how exactly your mouse is configured, as your problem should hopefully be overcome with a fairly straightforward change of the 'Mouse' device.

```
Section "InputDevice"
Identifier "Mouse"
Driver "mouse"
Option "CorePointer"
Option "Device" "/dev/misc/psaux"
Option "Protocol" "PS/2"
Option "Emulate3Buttons" "true"
Option "ZAxisMapping" "4 5"
```

Scanner struggle

I'm running Mandrake 9.0 and I am having some trouble getting my Canon N650U USB Scanner to work. I've looked on their website and there they state that the Plustek backend can be used with this scanner. Following their instructions I've added the following to *plustek.conf*:

```
[usb] 0x04a9 0x2206
```



There are many window managers available for X, including *Window Maker* (www.windowmaker.org).

```
dev/usbscanner
I have also added the following to
/etc/modules.conf, as shown in the
VueScan article (LXF35):
options scanner vendor = 0x04a9
product = 0x2206
After making these changes, if I
then run
cat /proc/bus/usb/devices;
it gives:
```

```
C:* #lfs = 1 Cfg# = 1 Atr = 80
MxPwr = 500mA

I: lfs# = 0 Alt = 0 #EPs = 3 Cls =
ff(vend.) Sub = 01 Prot = ff Driver =
(none)
E: Ad = 81(l) Atr = 03(Int) MxPS = 4
lvl = 5ms
E:Ad = 82(l) Atr = 02(Bulk) MxPS =
64 lvl = 0ms
E:Ad = 03(0) Atr = 02(Bulk) MxPS
= 64 lvl = 0ms
T: Bus = 01 Lev = 01 Prnt = 01 Port
= 01 Cnt = 02 Dev# = 3 Spd = 12
MxCh = 0
D: Ver = 1.00 Cls = 00(>ifc) Sub =
00 Prot = 00 MxPS = 8 #Cfgs = 1
P: Vendor = 04a9 ProdID = 2206
Rev = 1.00
S: Manufacturer = Canon
S: Product = CanoScan
```

Am I right in thinking that the kernel module still hasn't loaded, and if so what else do I have to do to make it load? Any help you can give to point me in the right direction will be greatly appreciated.

John Knops, *Australia*

The only way to really check that a module hasn't loaded is with **lsmod**, although none of the above will actually load the kernel module. You can do this

by performing **modprobe scanner** as root on the command line.

USB scanner devices have a major number of 180, so you can simply put the following in */etc/modules.conf* to have the scanner module load when something tries to read from */dev/usb/scanner*:

```
alias char-major-180 scanner
```

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

Answers



Flippin' students...

Q I have a question about restoring the Linux booting screen. I installed Windows 2000 and Red Hat 8.0 on the community college computers where I work. One student typed the `fdisk /mbr`

command in Windows. Since then, the Linux or Windows option screen does not display at boot time. If I use a boot diskette, I still get to the Linux screen. I do not want students to be forced use the boot diskette to get to Linux. I want students to be able to choose the OS (Windows or Linux) at boot time. How can I restore the screen?

Steve Youn, via email

A Depending on which bootloader you are using, the problem could be resolved by issuing one of two commands:

1. If you are using GRUB, the default bootloader for Red Hat 8.0, you should log into your server using the boot disk you have and issue the following command:

```
grub-install /dev/hda
```

replacing `/dev/hda` with the device from which you are trying to boot.

2. If you are using lilo then you will just need to type:

```
lilo
```

at the command line. Both these commands rewrite the master boot record with information from their respective config files.

Attack of the clones

Q I have recently taken over a project that is deploying a customised version of Red Hat (7.3) in a number of academic sites across the UK. We use the *Systemimager* suite to customise before release and effectively 'clone' machines on deployment. Access to these machines is limited to SSH for security (SSH is the only open port as they are Internet-facing). Many of our users are first-time Linux users, and we are encouraging people to deal with

their own systems administration through training courses, however, in the meantime I'd like to be able to monitor their machines remotely (via SSH) – basically log file checking, keeping tabs on disk usage etc. I'm used to managing, at most, 3 Linux/Unix systems at any one time that can be done "by eye". How can I scale up to keeping tabs on 80 machines whilst at the same time I train people to perform their own admin duties?

Dan, via email

A Dan, Looking through log files is certainly a crucial aspect of a system administrator's role, whether you're working for a business, academic institution, public sector or a charity. On single systems that are not too busy, it is quite feasible to look through your important log files by

hand looking for any unusual messages. I would, however, recommend a log parsing and reporting application such as *logwatch* to provide you with much more meaningful and concise information as to what is happening on a Linux server. With reference to your situation, *logwatch* will be useful in one of two ways:

1. *Logwatch* is designed to be able to run on a regular basis as a cron job and can either mail the results to a user, or alternatively, as in your situation, you could log into the server via *ssh* and run *logwatch* from the command line to print the results to *stdout*, which you can redirect to a text file. *Logwatch* will also give a report on disk space per partition at the end of its report.

Logwatch should be installed on your Red Hat 7.3 system, however, I would download a newer version from www.logwatch.org rather than rely on any older version that you might have. The configuration file resides in `/etc/log.d/conf` and is named `logwatch.conf`.

2. To view output in a console just run:

```
/usr/sbin/logwatch -print
```

Hacked or gremlin?

Q Hello Rackspace! I'm in the unenviable position where I'm starting to think that my server has been hacked! :(

Recently I've taken over a server from another company and in order to consolidate the data on the two servers I have taken the hard disk from the second server and connected it to my server. I've mounted it under `/mnt/othercompany` and moved all the required data around. This process took several days. Now, when I try to unmount the drive I get the following error:

```
umount: /mnt/othercompany: device is busy
```

I've made sure nobody else is logged in and I'm not using the mount at all. I cannot see anything `ps -auxw` which I believe might be accessing the drive. Rebooting did fix the problem, but I wanted to leave the server as it was to investigate, and have hence mounted the disk again. Now I cannot unmount it (again, there's the same error).

Tom, via email

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

The screenshot shows the Rackspace Managed Hosting website. At the top, there's a navigation bar with links for Home, Bookmarks, The Mozilla Organ..., BT Com - 400, Google, and MDR Retail - P... Below that is a search bar and a 'Search' button. The main content area features a large banner with a woman's face and the text 'It's not about SERVERS, it's about SUPPORT.' Below the banner, there are several sections: 'Fanatical Support Defined', 'Service Level Agreements', 'MyRackspace', 'How We Do It', 'Customer Awards', and 'Alamo'. On the right side, there are three smaller boxes: 'Did You Say 100%?', 'Happy Birthday Rackspace', and 'Move to Rackspace and Save'.

WIN A SHARP ZAURUS PDA FROM RACKSPACE MANAGED HOSTING

www.rackspace.co.uk

Every month, the best question related to Systems Administration that a LXF reader sends in wins a Sharp Zaurus SL-5500 running Linux.

With industry support from the likes of IBM, SAP, Oracle, and Sybase, the Zaurus features a QWERTY keyboard behind a sliding cover. The 65,536 colour screen is 240x320 resolution. An Intel® StrongARM® 206MHz processor and 64MB memory makes the Zaurus ideal for downloading software or playing MP3 and MPEG1. The docking station that fits into its USB connector, allowing users to synch with their desktop, schedules, address

book and ToDo lists. The email app supports POP3 or IMAP email servers. Write and send emails and browse the Internet with a compatible modem card, or using the infra-red port and a mobile phone; or any replied, forwarded, or composed emails are transferred to your outbox the next time you synch. SD and CF slots allow upgrades and addition of software.

Email all your sysadmin questions and problems to Hans Huberland: sysadminqa@rackspace.co.uk Every published query on these system admin pages wins a Rackspace T-shirt.

It's only 74mm wide by 138mm long and a svelte 18mm thick – so the Sharp Zaurus really is Linux in your pocket!



★ Star Question – PDA winner!

This issue's lucky winner is **Andrew** – your new Zaurus PDA will be with you shortly!

Thanks for a great mag. This may seem like a petty request but I hope you can help me. I've been trying to convince my company to let me use Linux on more of the internal servers, and they've given me permission to migrate our mail server to a Red Hat 8.0 system. The migration went smoothly and the mail server is working very well. There's just one problem: I set up the system using a Mylex AcceleRAID 170 controller using RAID 5 and one of the disks has crashed (the LED on the disk does not light up). Being on RAID 5 the system is still up and running and now I need to repair it. This is where the catch comes in. I do not want to take the server down to

rebuild the array as one of my main arguments explaining why Linux would beat our old Microsoft platforms was incredible uptime. In Windows we would be able to rebuild this disk locally but I have been unable to find any utilities to do this from within Linux.

Andrew, via email

There are two approaches you can take to resolve your problem. The first (and easier approach) is to manage the array using the *Global Array Manager* from Mylex. You can download it from the official site at www.mylex.com/pub/support/current_raid/linux-gam_510.tar.gz – extracting this file should reveal several RPMs. You'll need to install the *GAM server* on the server and

the *GAM client* on the workstation. This will obviously not install on distributions that don't support the RPM package format. Furthermore, this application is not open source, however, they do work very well.

For the purist or those who cannot install RPMs it is also possible to issue commands directly to the RAID controller by echoing strings into the proc filesystem.

I would recommend verifying that the RAID array is really in a deteriorated state by checking the contents of `/proc/rd/status`, this will report whether any of the RAID arrays on any installed Mylex controller are OK or FATAL. If you need more information you can look at `/proc/rd/c0/current_status`

(assuming `c0` is the channel in question). This should tell you the state of each disk. If the disk is in fact dead, you should remove it and add a new disk before issuing the following command:

```
echo "rebuild 0:3" >
/proc/rd/c0/user_command
```

This command will rebuild redundancy data on the disk with ID 3 on the primary channel. **Cat**'ing out `/proc/rd/c0/user_command` should verify that the rebuild process is taking place. Mylex has provided quite a bit of information about these user issuable commands and this documentation should be installed in `/usr/src/linux-2.4/Documentation/README.DAC960` as part of the DAC960 module.

A Before assuming that the server has been compromised, let's make sure that nothing has been overlooked. To find out who has a file or directory open and is not allowing you to unmount you can use the `lsdf` command (short for **ls open files**). If your mount point is `/mnt/othercompany` then issue a `lsdf /mnt/othercompany` and check the output. The first column of the output tells you which application has the file open. If `lsdf` shows no output and still will not let you unmount then you may want to check the integrity of `lsdf` as

the server could indeed be compromised.

Un-blacklist

Q I have a Red Hat 7.3 server that runs *Qmail*. A while ago the server got hacked and relaying was turned on – this resulted in some pretty nasty mails being sent through my server. I now find that my server is blacklisted at several blacklist sites. My ISP has provided me with an additional IP address, which I would like to send mail out from. I do not want to

change my server's primary IP address, just change the address which *qmail* uses. I cannot get *qmail* to use this new address without changing the primary address though.

Greg, via email

A I've done a little investigating regarding your issue and I have not been able to find a way to change the IP *qmail* sends mail out on. The easiest way around this problem is to set up *iptables* to force all outgoing traffic on port 25 to use the new additional IP address. For example, if

your primary IP address is 192.168.1.5 and the new IP address is 192.168.1.10. In order to allow the traffic to be forwarded you'll need to make sure that `/etc/sysctl.conf` contains:

```
net.ipv4.ip_forward = 1
```

The *iptables* rule you would need to `/etc/sysconfig/iptables` is:

```
-t -nat -A POSTROUTING -p tcp -
dport 25 -j SNAT --to-source
192.168.1.10
```

Basically, this tells the *iptables* to rewrite any packet sent out on port 25 as coming from the new address and then sends it on its way.

Answers

RH RPMs encore

I'm a very experienced software developer and since all my work is with the various flavours of Java, I moved to Red Hat 7.2 for my OS last year and the Java development has not been interrupted :-). I really like Linux (most of the time), but it seems that some of the simplest things seem to cause so much grief!

I run Red Hat 7.2 on a server and a laptop and it's just about holding together. I have a "fairly" reasonable understanding about how Linux works (I've re-built the kernel with moderate success, my network is OK but I have to `ssh` to the server to control the ISDN link). My printer (HP Officejet G85) doesn't work – I'm completely lost with the printer management (queues, spoolers, drivers, back-ends etc.) so I keep going back to 'the dark side' to print :- (I could probably fill your entire *Answers* section with questions, but here's an example of one of the annoying issues I have at the moment.

I would like to experiment with Python. It's neither urgent, nor totally necessary, but it looks interesting. Patrick O'Brien's first article recommended getting the latest version of Python so that the tutorial could be followed. Here begins the usual tale of woe.

Red Hat 7.2 has Python v.1.5.2, so I downloaded `python-2.2.2-7.i386.rpm`. Then I downloaded another six packages that it may or may not rely on (`devel`, `tools`, `tkinter`, `wxPythonGTK`, `sip?`).

Trying to `rpm -uvh` (or `-ivh`) provides the first hint of problems. Apparently my `rpmlib` is too modern?!? It seems that `rpmlib` (PartialHardlinkSets) should be `<= 4.0.4-1`.

1 I can usually identify these dependency problems using:

```
rpm -qa | grep
```

or

```
rpm -q --provides | grep
```

but in this case, I have no idea what `rpmlib` is actually associated with.

2 Checking my version of RPM (`rpm -q rpm`), it's at 4.0.3-1. Shouldn't this mean that `rpmlib` is at the same release?

OK. Not much good so far, but I'll continue with the other dependency problems. No trace of `libdb-4.0.so`. I managed to find this as part of `db4-`



Our Python series with Patrick K O'Brien will return after a brief hiatus.

4.0.14-14.i386.rpm. However, in this case `rpm -Uvh` was a bad choice. Loads of broken link warnings, but `rpm -ivh` seemed to work. This is the first time I've come across the idea that different "versions" can be installed simultaneously, rather than having to upgrade.

It seems my version of `glibc` isn't new enough (Red Hat 7.2 comes with 2.2.4-13 and I need `lic.so.6(GLIBC_2.3)`). First I got the source tarball for `glibc 2.3`. Then I needed to install `pkg-config` (this must have gone without any problems because I recall doing it, but it was very quick). But... the `glibc` source attempts didn't go anywhere, so I opted to search for the RPM instead.

3 If you build things from source, how do you get RPM to understand what you've done? I seem to recall there's a way to get RPM to "rebuild" its database, but if I do this, will it include the new (from source) builds? And by the way, how exactly do you do this?

4 What potential damage may occur by upgrading something as fundamental as `glibc` (or should a new version be installed (`-ivh`) instead of upgraded (`-Uvh`)?)

5 When RPMs (or tarballs etc.) are labelled for Red Hat 8.0, can they be used for 7.2 or should I only use 7.3 (or even only 7.2)? How about Red Hat 9.0?

6 So. `rpm -Uvh` for `glibc-2.3...` and `glibc-common-2.3...`

Argh! `binutils < 2.13.90...` conflicts with `devel-2.3`! Let's try without `devel-2.3` then. This looks better(ish). Now it only complains about `rpmlib`.

I saw a newsgroup message somewhere on the internet that said a very similar problem could be cured with the `--nodeps`. Well, this is how I entirely stuffed my system before (although that was with KDE3.0 and I've seen your advice on that since).

Should I use `--nodeps`? When is it safe to do this? `glibc` is just one of many elements of the Linux system that seems to be "fundamental". Is there somewhere that describes how to safely upgrade this kind of thing (the `glibc` homepage was sorely lacking in any advice)?

Thank you for your patience (if you've read this far). My apologies for the rambling nature, but, in a way, this reflects my frustration and experience with Linux updates. Maybe, one day, if I finally understand this process, I might consider updating my desktop(s) to KDE3.1 and GNOME 2.2.

I seem to have dug myself into a hole. My systems that run Red Hat are essential for my business and in most cases, the standard system is adequate, but I live in fear and trepidation of the time when I have to do an update in case I screw things up.

One quick final question. If I were to purchase Red Hat 9.0, would this involve a "simple"

upgrade to my existing system or would it require a fresh installation? – the specific issue I have in mind here is that I have a license for some software that was awkward and time-consuming to set up and uses the hard disk ID, so reformatting means great pain.

Jim McLachlan, via email

Dependency problems from RPMs with Red Hat seem to be becoming a familiar trend in the *LXF Answers* pages of recent months. If there was a single issue with Red Hat that they need to address, then it doesn't really take a genius to establish what that would be.

Short of installing Debian, which handles packages much better than Red Hat due to `apt` and `dpkg`, I would suggest that if they haven't done so already, all Red Hat users should take a look at a wonderful tool called `apt-rpm`. Just like `apt` in Debian, it will download packages on the fly, handle dependencies properly and update the system with a single command.

`apt-rpm` can be found at <http://apt-rpm.tuxfamily.org/>, along with plenty of documentation and example `apt-rpm` sources to use with a whole array of Red Hat versions.

As with your final question – Yes, it is perfectly possible to upgrade to Red Hat 9.0 without reinstalling the system. Some packages may require reconfiguration, or you might need to look at obtaining upgraded packages for certain programs elsewhere if they are not part of Red Hat's standard distribution. If you were to reinstall, there is no reason why any hard drive or partition identification would change, assuming that you did not repartition the drive. The installer would simply write a new filesystem to the partition. [LXF](#)

Posting to the forum

The LXF online community

Got a technical question? Other *LXF* readers may be able to help!

The forums at www.linuxformat.co.uk have a section dedicated to technical queries, hardware, programming languages and general help. As well as being able to call on `lxfadmin` (when there's no deadline!) and the ever-present 'anonymous', the forums are also frequented by Linux heroes like **Jeremy, Nelz, Fingers99, Rhakios, Erin** and many others brimming with knowledge and experience of using Linux in a wide variety of situations.

missed one?

LINUX FORMAT BACK ISSUES

Every month *Linux Format* brings you the best tutorials, the essential reviews and the latest news. But if you've missed out on a must-read feature or a vital 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).



July 2003

Product code:
LXFB0042(cd)
LXFD0042(dvd)

DVD HIGHLIGHTS:
Opera 7i beta, GAIM, ABCDE, Pingus, KDE-Sounds/Aqua Fusion, LDP, Alsaplayer, KOffice, RUTE, Seclude, Freevix, DamnSmallLinux

CDs HIGHLIGHTS:
Morphix, GuardDog, NVidia XDrivers, GCC 3.3, TORCS, Linux Cookbook, SETI

MAGAZINE FEATURING:
All about X, MySQL review, Sharq server review, WINE explained, Steganography, Distributed Computing, Red Hat Enterprise AS, Choosing a distro, GCC 3.3 review



June 2003

Product code:
LXFB0041(cd)
LXFD0041(dvd)

DVD HIGHLIGHTS:
Slackware 9.0, Mutt, OpenOffice.org, Xine, TakeAJoint, Peanut Linux, FCKeditor, KDE 3.1.1a, mirror of Linux Documentation Project

CDs HIGHLIGHTS:
Mandrake 9.1, AC3D, KDERadio, The GIMP – including 400 plugins from The GIMP Registry

MAGAZINE FEATURING:
AMD's 64-bit Opteron CPU, Linux Laptop Megatest, The Hurd, JCodeBox review, Build your own RPMs, Red Hat 9 vs SuSE8.2 vs Mandrake 9.1, OpenAuth Linux PC review



May 2003

Product code:
LXFB0040(cd)
LXFD0040(dvd)

DVD HIGHLIGHTS:
Vector Linux, Dropline GNOME, Lopster, jEdit, QtParted, QuickRipDVD, BaitAndSwitchHoneyPot, AdiosBootCD, zTappy, MondoRescue

CDs HIGHLIGHTS:
Mozilla 1.3, Opera update, IMP server, Kino, IceWM, ECLIPt Roaster, VariCAD

MAGAZINE FEATURING:
SCO v IBM \$1 billion lawsuit, Blender tutorial, LXF Awards 2002 results, Linux on Apple Mac, Gigabyte SR113 server, EFI to replace BIOS, VariCAD & IEMS mailserver reviews



April 2003

Product code:
LXFB0039(cd)
LXFD0039(dvd)

DVD HIGHLIGHTS:
CensorNet, K3b, VIPs, FrozenBubble, Linux Open Media Box, IDMS Linux, WhiteDune, OBM, AstaroSecurityLinux

CDs HIGHLIGHTS:
SME Server, TightVNC, GNOME 2.2, VOCP, CoyoteLinux, UT2003 update, Blender, Qt, BootEverywhereLinux

MAGAZINE FEATURING:
Customising Linux, Evesham Linux PCs, Open publishing with WIKI, Linux PIMs and PDAs, SME Server review, File Managers Roundup, Serial Networking, GnuCash



March 2003

Product code:
LXFB0038(cd)
LXFD0038(dvd)

DVD HIGHLIGHTS:
Python, The GIMP, Lycoris, Yellowdog, Knoppix, Cygwin, DansGuardian, Senken, AbiWord, X Northern Captain, Eclipse

CDs HIGHLIGHTS:
KDE 3.1, UNIX Tools For Windows, Konstruct, phpMyAdmin, X Net Strength, AutoUpdate

MAGAZINE FEATURING:
You, Copyright & The Law, X apps on Mac OS with Fink, SVG versus Flash, Audio Editors roundup, Python tutorial, GIMP secrets, SuSE Office Desktop examined



February 2003

Product code:
LXFB0037(cd)
LXFD0037(dvd)

DVD HIGHLIGHTS:
IPCop, Postnuke, Lillypond, Squid, Allegro, Cocoon, Fandango, gPhoto, Xrmap, Ted, AdvancedPHPDebugger, Flightgear updates

CDs HIGHLIGHTS:
Yoper, KOffice 1.2.1, Parted, DOM Tooltip, Guarddog, GL-117, PythonImagingLibrary

MAGAZINE FEATURING:
Lindows on test, Serial ATA, Window Managers roundup, Arkeia review, Xft2, build your own Linux backup box, Java compiling, Jool Kwartz review, New Larger Linux Pro mag

Other back issues also available, please call for further information, giving the issue number and/or date required

To order, call the back issues hotline on **0870 4448645**

or **+44 1458 271178** for overseas orders

Remember to quote the issue code shown above and have your credit card/switch card details ready.

Prices (inc post and packing): CD – UK£5.99, US & Eur £6.99, rest of world £7.99

DVD – UK £6.99, US&Eur £7.99, rest of world £8.99.

LINUX

www.linuxformat.co.uk **FORMAT**

SUBSCRIBE AND

ON A YEAR'S SUPPLY OF LINUX FORMAT

That's 3 issues free!



Title _____ Initials _____ Surname _____

Address _____

Postcode _____ Country _____

Email address _____

Daytime telephone number _____

Instruction to your Bank or Building Society to pay Direct Debits

  Originator's Identification Number **768195**

Media with passion
Future Publishing Ltd, Cary Court,
Somerton, Somerset, TA11 6BR

Please send this form to address shown

1. Name and full postal address of your Bank or Building Society branch

To: The Manager _____ Bank/Building Society

Address _____

Postcode _____

2. Name(s) of account holder(s)

3. Branch sort code (from the top right-hand corner of your cheque)

4. Bank or Building Society account number

5. Instruction to your Bank or Building Society

Please pay Future Publishing Direct Debits from the account detailed on this instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Future Publishing and if so, details will be passed electronically to my bank or building society.

Signature(s) _____ Date _____

Ref No (Office use only) _____

Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

You may photocopy this blank form and then complete it. You cannot fax or email Direct Debit forms.

QUOTE ORDER No. **LXFP43**

LINUX THE UK'S BEST-SELLING LINUX MAGAZINE

FORMAT

Please choose your method of payment **1** or **2**

- 1.** Direct Debit (UK only – double CD Edition) – I wish to pay £27.45 every 6 months
or Direct Debit (UK only – DVD Edition) – I wish to pay £29.95 every 6 months

Please complete the form on the left

I understand that I will receive 13 issues during the next 12 months.

2. Cheque/credit card 13 issues

UK	<input type="checkbox"/> £54.90	<input type="checkbox"/> £59.90
Europe (including Eire)	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
North America	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
Rest of the World	<input type="checkbox"/> £91.90	<input type="checkbox"/> £96.90

Cheque (payable to Linux Format. Sterling cheques drawn on a UK account.)

Visa MasterCard American Express

Switch Issue No/Valid date _____

Card No _____ Expires _____

Signature _____ Date _____

UK READERS RETURN THIS COUPON TO:

Linux Format, Subscriptions, Future Publishing Ltd, Freepost BS4900, Somerton, Somerset, TA11 6BR.

OVERSEAS READERS RETURN (POSTAGE PAYABLE) TO:

Linux Format, Subscriptions, Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6TB, UK

WEB: www.linuxformat.co.uk then click on **Subscribe**


Please tick this box if you would prefer not to receive information on other offers.

ORDER HOTLINE ☎ 0870 444 8645
OR FROM OVERSEAS ON +44 870 444 8645

OFFER ENDS: 12.08.03

SAVE £24.47!*

– THE LEADING LINUX MAGAZINE



= SAVING OF £24.47*

- **THE ESSENTIAL** Linux read for all Linux users – can you afford to miss out?
- **EVERY ISSUE** comes with the ideal supplement for IT professionals – *Linux Pro* covers trends, technology, security issues, industry comment and thought-provoking case studies
- **GUARANTEED DELIVERY** direct to your door so you never miss an issue!

RISK-FREE GUARANTEE!

You can cancel your subscription at any time and receive a full refund on all unmailed issues, no questions asked

SPREAD THE COST OF SUBSCRIBING

For under £15.00* every 3 months you can pay by convenient Direct Debit. Simply complete the form opposite and return it back to us today.

Call now on 0870 444 8645 to make sure you receive the next issue

* Prices and savings based on DVD edition.

www.linuxformat.co.uk

LXF43 AUGUST 2003 **101**



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 xvzf --bzip2 /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in .tar.gz or .tgz, and the second for Bzipped files, ending in .tar.bz2 or .tbz2. Naturally, you change the paths to suit the location and name of the archive. and replace /mnt/cdrom with whatever is applicable to your system (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 do not have to use Linux to burn the ISO to a disc. All the Linux-specific bits are already built into the image file. Programs like *cdrecord* simply dump it to the disk. If you don't have a CD-writer, find someone who does have one, and a DVD drive, and use the CD burning software on their computer. It can be Windows, MacOS, AmigaOS whatever.

No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions it is also possible to mount the images and do a network install, or even a local install from another disk partition. The methods often vary between distributions, so check on the distro vendors website for more information. [LXF](#)

Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format DVD*. We include requests, so why not send us your wish-list?

As you'll see from the CD pages, one of this month's *Linux Format* CDs is dedicated to networking. The contents of this networking CD, and more besides, can be found in the Networking directory of the DVD. The "more besides" consists of some packages that were too large to fit on the CD.

NETWORKING-ROUTERS ASTAROSEcurityLinux

The Networking CD contains the ever-popular SmoothWall, that turns an old PC into a dedicated router and firewall. Naturally, the DVD also has this, but we have a couple of extra router/firewall projects too. *Astaro Security Linux* is, according to its

website "an integrated software solution that provides superior performance in an all-in-one Firewall. Its hardened operating system, stateful packet inspection, content filtering (virus and surf protection), application proxies and *IPsec*-based VPN provides a powerful solution to today's security issues". In other words, it is an all in one network router, gateway and firewall that runs on a dedicated PC. As with the other such packages, *Astaro* is supplied as an ISO image. Burn this to CD-R, following the instructions on the Essential Disc Info page, and boot from it. In common with *SmoothWall*, *IPCop* and *ClarkConnect*, *Astaro* takes over the computer it is installed on, erasing any previous information from the hard disk. You cannot install it alongside another operating system.

NETWORKING REMOTEACCESS

Thin clients, or diskless workstations, are becoming an increasingly popular means of deploying and maintaining a large number of desktop computers. Instead of each computer being a standalone system, with its own installation of operating system and software, which all needs maintaining and updating, thin clients boot and run software from a central server. Not only does this reduce the hardware costs for the workstations, as each one needs less processing power and no storage, but it greatly simplifies system administration.

There are several projects to run thin clients using Linux, we have two of them on the DVD. The *Linux Terminal Server Project* and the *PXES Linux* thin client. Setting up such a system is obviously more complex than shoving the DVD in the drive of a computer and clicking an RPM file, so you will have to read the documentation for each of these. If you have a need for such a system, or are simply interested



DVD users get two extra gateway/firewall distributions, *Astaro Security Linux* mentioned on the left, and *ClarkConnect*.

in how these things work, this is the place to start looking.

DISTROS TRUSTIXSECURELinux

It seems that the majority of new distribution releases are aimed at the desktop user or are general purpose, jack-of-all trades distros. The latest release from *Trustix* goes against that trend. This is a distribution aimed squarely at server use. As such, the priorities are security and stability, rather than hand-holding installation and configuration tools. Because security is so important for a server, *Trustix* comes with an automated software update tool, *swup*, so that security fixes can be applied as soon as they become available. We have a beta of the forthcoming *Trustix 2.0*, it is supplied as an ISO image, ready to create a bootable CD-R.

OFFICE GNUCASH

When we reviewed *GnuCash* four issues ago, we commented on the improvements made over previous versions. Since then, the developers

have released several updates, adding even more features. If you haven't come across *GnuCash* before, there's a tutorial in this very issue. While accounts packages may not be the most exciting of uses for a computer, (you may well prefer to shoot people in *America's Army* than look at columns of figures) it is a very practical application.

I use *GnuCash* myself and find it a very straightforward way of keeping track of my overdrafts. Unfortunately, installation of *GnuCash* is not always so straightforward. This program has a lot of dependencies on the cover discs is a bit of a balancing act. Many of them are included with most distributions, so putting them on the CD or DVD would be a waste of space. We have included the more non-standard requirements. It is possible you may have to download a Perl module or two from www.cpan.org, depending on what your distro provides. If you are building from source on an RPM-based system, you may also need to install some devel packages for it to compile. Even though I already had version 1.8.2



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.

Date	Num	Description	Total Deposit	Total Withdrawal	Balance
4/21/1997	xy123	paycheck	7,000.00		7,000.00
		net deposit into bank acct	Bank:ABC Savings n	7,000.00	
		Income Tax Withholding	Expenses:Taxes:FICA n	2,500.00	
		Medicare Insurance	Expenses:Taxes:Medicare n	500.00	
		gross income	Income:IGOR Consulting n	10,000.00	
4/22/1997	101	Open Stock Trading Account		5,000.00	2,000.00
		We are going to get rich!	Bank:Mega Brokerage Settlement n	5,000.00	
		We are going to get rich!	Bank:ABC Savings n	5,000.00	
5/1/1997		Interest at 3.5%		5.67	2,005.67
		Int	Bank:ABC Savings n	5.67	
		Int	Income:ABC Bank Interest n	5.67	
6/1/1997		Interest at 3.5%		6.78	2,012.45
		Int	Bank:ABC Savings n	6.78	
		Int	Income:ABC Bank Interest n	6.78	
7/1/1998		Interest at 3.5%		7.89	2,020.34
		Int	Bank:ABC Savings n	7.89	
		Int	Income:ABC Bank Interest n	7.89	
1/1/2001					

GnuCash can be a little tricky to install, but once running it is easy to use, yet quite powerful. There's an introductory *GnuCash* tutorial on page 72.

installed, from a Mandrake cooker RPM, I had to install several devel RPMs from the installation DVD before the source would compile.

GnuCash isn't a particularly difficult program to build, but it can be somewhat time-consuming if each run of `.configure` throws up another package you need to install. Having said that, it built perfectly once the dependencies were met, and `checkinstall` built an RPM that updated my old 1.8.2 installation to the latest 1.8.4. If you need a good accounting program, do persevere with the installation, the end product is definitely worth a small amount of hassle.

GAMES SCUMMVM

A lot of people enjoy playing old games. Many games have been ported or cloned for Linux, like *Pingus*, a clone of *Lemmings* that was on the DVD in LXF's last issue.

However, many old games have not been converted, for various reasons. *ScummVM* is a classic gaming

interpreter and emulator that plays Lucas Arts adventure games and several other formats. The games themselves are still copyrighted, so don't expect to find a fully playable copy of *Monkey Island* or *Day Of The Tentacle* on the DVD. You need the original disks to be able to play these games. The program on the DVD is your excuse to brave the spiders and dig your old adventure disks out of the loft or cellar for a nostalgic session of retro-gaming.

MOBILE SHARPROMV3.10

Last month's DVD included the latest *OpenZaurus ROM*, a 3rd party ROM image for the Sharp Zaurus PDA running Linux. Now Sharp has upgraded the official ROM for the Zaurus 5000D and 5500. This adds much of the functionality of the Zaurus 5600 to the older models. The 3.10 ROM image supplied here is for US models, no UK image has been released yet. Some people have reported problems using the US ROMs



Dig out your old adventure games to re-live the past using *ScummVM*.



High-end video editing and compositing, including render farm support, with *Cinelerra*.

with the UK hardware, so install this to a UK Zaurus at your own risk. The upgrade process is a little convoluted, so the first thing you should do is read *5500v3UpgradeInstructions030512.pdf* thoroughly. The next thing to do is read it again.

One change from previous upgrade procedures is that you will not be able to read files made with the Backup/Restore program after upgrade, so you must backup using either the Windows synchronisation software, or by copying or archiving directories to a linked computer or a Flash memory card. You can backup to *Qtopia Desktop* on a Linux box, but you won't be able to restore from it after the upgrade, so this approach is not really to be recommended. The other thing to check is that the CF card you use to hold the ROM image for update is formatted with the FAT16 filesystem. This is the filesystem that the cards come with, so this won't be an issue unless you have reformatted the card.

GRAPHICS CINELERRA

Cinelerra is a complete audio and video production environment for Linux. This is an advanced compositing and editing system that is aimed at the professional or serious hobbyist videographer, rather than an easy way of editing and titling your home movies. This means it is a complex program that can take some effort to learn to get the best out of it. That's not a criticism of the program, more an observation. Power software often has a significant learning curve, it may be less friendly to start with, but pays dividends in the long run. What you

may find unfriendly is the installation process. Installing from one of the supplied RPMs should be fine, but the web site warns that compiling from source is rather difficult. I'm a coward so I took the easy way out and RPM installation worked first time.

Video processing probably puts more load on a system than any other productivity task (that's another way of saying "except games"). Serious processing of a long video can take a very long time, so *Cinelerra* has support for render farms. As it is open source, there's no extra licence fees to pay, no matter how many other computers you add to your system to help with the rendering. [LXF](http://www.linuxformat.co.uk)

Any requests?

We spend a lot of time trawling the Internet, looking for useful or fun programs to include on the coverdiscs. We hope that most of the time, you are happy with the results, but there may be something you think should be on the discs, something that we may not pick up on. In the forums section of the Linux Format website (www.linuxformat.co.uk) there is an area for CD/DVD requests. This is where you can ask for specific programs or distributions to be included on the cover discs.

It is helpful if you can provide a URL to the program's home page. Provided there are no obstacles to our including it, chances are we will do so. Some programs, while being free to download are not freely distributable, *Wolfenstein: Enemy Territory* is an example of this. To show that we care, *3d Desktop*, on the CD, and *ScummVM* on the DVD, were both requested by readers in recent weeks.

CoverdiscDVD

DVD CONTENTS AT A GLANCE

Desktop

Allin1	Fluxbox dockapp to monitor CPU, memory, network and power
BackgroundBuddy	An automatic wallpaper changer
Beaver	A lightweight, powerful programmer's editor
CRM114Discriminator	Controllable regex mutator and antispam filter
DirectoryAdministrator	LDAP management tool for users/groups and address book data
GenChemLab	A chemistry experiment simulator
Interceptor	KDE kicker applet for syslog monitoring and alerts
Jmol	Java/Swing-based molecular dynamics viewer
K-Track	A satellite tracking program
PartGUI	A Qt partition tool for Linux
Ratpoison	Window manager that lets you say good-bye to the rodent
Texmaker	A LaTeX development environment
WindowLab	Small and simple window manager of novel design
WMND	WM Dockapp for monitoring multiple network interfaces
X-CD-Roast	A program-package dedicated to powerful and easy CD creation

Development

ClusteredJDBC	A cluster of databases through JDBC
Graphics3D	High-performance 3D graphics codebase
J	A programmer's editor written in Java
Japple	Java Web application development environment
JLicense	Java library to create and validate software licenses
JRat	A Java runtime analysis toolkit
Kodos	Visually debug Python regular expressions
libieee1284	Parallel port library
LibTomCrypt	A comprehensive, modular, and portable cryptographic toolkit
Licq-osd	An On Screen Display plugin for licq
Taverna	Distributed compute workflow components in Java

Distros

GeeXboX	A standalone Linux media player system
LNx-BBC	Bootable rescue CD that fits on a business card CD
SlackwareLiveCD	A CD-based Linux distribution based on Slackware
TrustixSecureLinux	A secure Linux distribution for servers

Games

Cosmosmash	Clone of the 1981 Astromash video game
DustbowlClanTools	A community system for e-sport teams and gaming clans
Jools	Graphical puzzle game
Knights	Chess interface for KDE
KRconLinux	KDE tool for querying and managing game servers
PyTraffic	A Python implementation of the game Rush Hour
ScummVM	Classic gaming interpreter and emulator
Squicki	An arcade game

Graphics

AcidRip	A GTK Perl DVD encoding frontend
Avidemux	Graphical tool to edit video (filter/re-encode/split)
Cinelerra	Complete audio and video production environment for Linux
Coriander	A 1394 digital camera controller
Gtranscode	GUI frontend for transcode
JPatch	An amazing 3D Modeling/Animation tool
LiVES	The Linux Video Editing System
QuicktimeForLinux	Quicktime support for Unix
XMovie	High quality player for Quicktime, MPEG and DVD movies

Help

LDP	A complete mirror of the Linux Documentation Project
LinuxCookbook	Over 1,500 time-saving recipes and hints for Linux users
RUTE	The definitive guide for new Linux users

Internet

AdvogatoPoster	Post multiple diary entries to Advogato
Columba	A Java email client
Effusion	IRC-based BitTorrent management/sharing application
GeekBot	Personal information robot
Konversation	A user-friendly IRC client for KDE
MyPPP	PPP dialer for Red Hat Linux with a monthly usage logger
Scribe	A cross platform email client
SpeedtouchInstaller	Speed Touch ADSL configuration tool
SwingingMLDonkey	Java GUI for MLDonkey
Sylpheed	GTK+ based user-friendly email client

Mobile

AccessPointUtilities	Configure and monitor a Wireless Access Point
GPSManager	A GUI for downloading, organizing, and maintaining GPS data
HermesAP	BSS master mode for Hermes/Airport WaveLAN cards
Kismet	A 802.11b wireless network sniffer
SharpROMV3.10	Latest Sharp Zaurus 5500 ROM Version 3.10

Networking/Firewalls

FIAIF	An iptables firewall for Linux
Knetfilter	KDE frontend to iptables

Networking/RemoteAccess

LTSP	The Linux Terminal Server Project
PXESLinuxThinClient	Micro distro to convert PCs into thin clients

Networking/Routers

AstaroSecurityLinux	Firewall with packet filtering, proxies, filtering and VPN
ClarkConnect	Transforms an old PC into an Internet gateway

Networking/Utilities

AkeniLANMessenger	An instant messenger for LANs
--------------------------	-------------------------------

Office

BIE	A Java-based integration server
CyberTester	Create and conduct online tests and exams
GnuCash	A program to keep track of your finances
ManhattanVirtualClassroom	Web-based virtual classroom system

Server

Apache2	A high performance Unix-based HTTP server
Arachnophilia	Comprehensive and easy to use Web editor
Formmail-trap	Reports attempts to probe a formmail open relay
JBoss	An Enterprise JavaBeans application server
LDAPDNS	An LDAP to DNS gateway
MyClassifiedsSQL	Online implementation of a newspaper classifieds section
MySQL	A fast SQL database server
Porn.pl	Gets a list of domains from hosts linking to porn sites
Postfix	The Postfix MTA
Sympa	Multilingual List Manager with LDAP and SQL interfaces
WebMySQL	A web interface to a MySQL database server

Sound

BenMP3	An MP2, MP3, OGG, CD, WAV, and modules player
CSound	Powerful musical synthesis software package
MassRip	Convert and maintain Audio CDs in compressed audio formats
MP3record	A commandline MP3 recording utility
MusicPlayerDaemon	Control playback of music (MP3 and Ogg) over a network
SnackAmp	An MP3/Ogg/other format music player
Tuneroid	A tuner for musical instruments
Wavbreaker	A tool to split a wave file into multiple files
YaRET	Yet another Ripper-Encoder-Thingy

System

CommonUNIXPrintingSystem	An Internet printing system for Unix
Findutils	A set of utilities for searching a filesystem
G4u	Harddisk Image Cloning for PCs Using a NetBSD Floppy
PHPDumpMaker	A smart backup solution written with PHP
SureConnectUSBModemDriver	A USRobotics Sure Connect USB modem Linux driver

Essentials

GCC	The GNU Compiler Collection
jigdo	Utility to download/get small files and assemble into CD image



Sharp has released an updated ROM for the Zaurus SL-5000D and SL-5500. Make sure you read the installation guide very carefully before doing anything. Win your very own Sharp Zaurus on page 96!

User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally natter about all things Linux. You can find lots more information online at: www.lug.org.uk

1 HAMPSHIRE

URL www.hants.lug.org.uk
Contact Hugo Mills

2 BRISTOL & BATH

URL www.bristol.lug.org.uk

3 SCOTTISH

URL www.scottish.lug.org.uk

4 OXFORD

URL www.oxford.lug.org.uk
Contact Alasdair G Kergon

5 KENT

URL www.kent.lug.org.uk
Contact Kevin Groves

6 BRIGHTON

URL www.brighton.lug.org.uk
Contact Johnathan Swan

7 WORCESTERSHIRE

URL www.worcs.lug.org.uk

8 NORTHANTS

URL www.northants.lug.org.uk
Contact Kevin Taylor

9 ANGLIAN

URL www.anglian.lug.org.uk
Contact Martyn Drake

10 MILTON KEYNES

URL www.mk.lug.org.uk
Contact Denny De La Haye

11 DONCASTER

URL www.doncaster.lug.org.uk
Contact Andy Smith

12 MORAY

URL www.moray.lug.org.uk
Contact Stewart Watson

13 WEST WALES

URL www.westwales.lug.org.uk
Contact Dan Field

14 WOLVES

URL www.wolveslug.org.uk
Contact Jono Bacon

15 PETERBOROUGH

URL www.peterboro.lug.org.uk
Contact Steve Gallagher

16 EDINBURGH

URL www.edinburgh.lug.org.uk
Contact Alistair Murray

17 TYNESIDE

URL www.tyneside.lug.org.uk
Contact Brian Ronald

18 LEICESTER

URL www.leicester.lug.org.uk
Contact Clive Jones

19 GREATER LONDON

URL <http://glug.linux.co.uk/>
Contact John Southern

20 SURREY

URL www.surrey.lug.org.uk
Contact Jay Bennie

21 CAMBRIDGE

URL www.cam-lug.org.uk

22 DEVON & CORNWALL

URL www.dclug.org.uk
Contact Simon Waters

23 FALKIRK

URL www.falkirk.lug.org.uk

24 MANCHESTER

URL www.manlug.mcc.ac.uk
Contact John Heaton, Owen Le Blanc

25 HERTFORDSHIRE

URL www.herts.lug.org.uk
Contact Nicolas Pike

26 WEST YORKSHIRE

URL www.wylug.lug.org.uk
Contact Jim Jackson

27 SHEFFIELD

URL www.shefug.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.malvern.lug.org.uk
Contact Greg Wright

49 HUDDERSFIELD

URL www.hud.lug.org.uk
Contact Dave Naylor – knocker@caramboo.com

50 NOTTINGHAM

URL www.nottingham.lug.org.uk

51 ST ALBANS & LUTON

URL www.lust.lug.org.uk
Contact Michael Culverhouse – mike@easily.co.uk

52 WREXHAM

Contact Paul Kersey-Smith
Email paul@pkls.fsnet.co.uk

53 PRESTON & LANCS

URL www.preston.lug.org.uk
Contact Phil Robinson

54 DERRY

URL www.derry.lug.org.uk

55 ISLE OF WIGHT

URL www.iow.lug.org.uk
Contact David Groom – info@iow.lug.org.uk

56 SCARBOROUGH

URL www.scarborough.lug.org.uk

57 BLACKBURN

Email matt@consultmatt.co.uk

58 YORK

URL www.york.lug.org.uk

59 LINCS

URL www.lincs.lug.org.uk

**60 HULL**URL www.hull.lug.org.uk**61 WALTON-ON-THAMES**Contact William Mutch
Email rael@freeuk.com**62 GLOUCS & COTSWOLDS**URL www.gloucs.lug.org.uk**63 WEST OF SCOTLAND**URL www.wos.lug.org.uk**64 SOUTH STAFFORDSHIRE**URL www.staffs.lug.org.uk**65 MANSFIELD**URL www.mansfield.lug.org.uk**66 BORDERS**URL www.linux.bordernet.co.uk**67 BIRMINGHAM**URL www.sb.lug.org.uk**68 COVENTRY**Email info@coventry.lug.org.uk**69 NEWARK & LINCOLN**URL www.newlinc.lug.org.uk**70 BEDFORDSHIRE**URL www.beds.lug.org.uk**71 LINCOLN**URL www.lincoln.lug.org.uk**72 LOUGHBOROUGH**URL www.loughborough.lug.org.uk**73 EXETER UNIVERSITY**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@whatemail.com**76 CLEVELAND OPEN SOURCE GROUP**Email openlug@digitalmedia.co.uk**77 BEVERLEY**Email vladimir_lukyanov@hotmail.com**78 DUNDEE & TAYSIDE**URL www.dundee.lug.org.uk**79 SUSSEX**URL <http://sussex.lug.org.uk/>**80 WIGAN & ST HELENS**Email 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**REVISED
DETAILS****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

LUGS OF THE MONTH

GLUG – Goa Linux User Groups

Frederick Noronha writes:

More commonly known to Westerners as an island holiday paradise, it may surprise you to know that there are considerable numbers of Linux/Linuxers in Goa. Not only that, but we actually have THREE Linux User Groups – about the same density that exists in the UK's West Midlands! Collectively, we would like to invite all travellers, tourists and visitors to Goa to join us at any of our meetings, if they are at all interested in GNU/Linux and Free Software.

LUG-Goa (Panjim) meets at the Computer Society of India, near Panjim Market, on the 4th Saturday of each month. Contact: Frederick Noronha fred@bytesforall.org URLs: <http://groups.yahoo.com/group/ilug-go> or www.ilug-go.org
LUG-South Goa (Margao) meets at Adarsh Vidyalaya School on the 3rd Saturday of each month. Contact Arvind Clemente arvind@controlnet.co.in URL: www.ilug-margao.org

LUG-Goa Engineering College (Farmagudi)

Meets on an 'as & when' basis
 Contact: George Easaw geasaw@vsnl.com or Prof MS Rayadu msr@ele.gec.ac.in
 Of course, there are nearly 40 LUGs all over India. All users of Open Source, Free Software and Linux: if you're visiting India, why not contact the LUG local to where you are staying before you go, and pay them a visit? You're very welcome!

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.luv.asn.au
 Contact luv-committee@luv.asn.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
Alsund www.aalslug.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@cslug.org
HYDERABAD, SINDH, INDUS VALLEY
 URL www.geocities.com/slug_pk/
KASHMIR
 Coming soon!

China

BEIJING (GB encoding, but mostly written in Chinese)
 URL <http://mud.263.net.cn/~linux>
CHINESE LINUX USER GROUP
 URL www.linux.org.cn
NANJING
 URL <http://jllib.jlonline.com/njlug>

India

LINUX INDIA
 URL <http://linux-india.org>
ALIGARH LUG
 URL <http://linux.amupost.com>
BOMBAY
 URL www.ilug-bom.org.in
CHANDIGARH
 URL www.geocities.com/vipinb
CHENNAI AND MADRAS
 URL www.chennaiug.org/
CYBERABAD (CLUG)
 URL <http://seeknew.freesevers.com/clug/>
DELHI
 URL www.linux-delhi.org
KOLKATA
 URL www.ilug-cal.org
MADURI
 URL <http://linuxmadurai.tripod.com>
NORTHERN INDIA LINUX
 URL <http://groups.yahoo.com/group/lug-northindia>

Spreading the word

Linux use in schools isn't only a good idea in terms of cost, it's our most important advocacy asset, says Jono Bacon.

Schools essentially have two major areas for computing: Administration and Teaching. Administrative facilities are required for paperwork, grading and other organisational functions. In many ways this area is similar to other groupware and in some cases SOHO based software. PIM, Office and Internet apps are required here and Linux certainly has plenty of options.

The second area is in my opinion the more interesting of the two. To fully advocate Linux in schools, you must be specific about promoting the Open Source software in a way that compliments the direction of the teaching. It is important to stress here that the teaching content is rarely changeable by schools and is instead tied to the demands of government guidelines, so do not expect a school to change their database from *MS Access* to *MySQL* for example; but a free software solution is certainly advisable as an alternative to their current solution – there is a possibility that the guidelines will be applicable to free software.

What is important to remember is that free software in schools will get kids using it at an early age and hence promote early awareness of free software – a concept that has no doubt helped proprietary vendors such as Microsoft where most people of my generation (twenties) and younger have spent a large portion of their education using Microsoft products. Examples of possible software could include Open Source web browsers, *The GIMP*, Blender and other productivity software like *OpenOffice.org*. Remember that an entire shift to a Linux desktop will probably be a supplemental step, and a first step could be using a free software application on their existing systems – many Open Source and Free Software projects have Windows/Mac and Linux ports specifically to enable this. Next month we will identify the financial benefits of schools moving to free software. Keep on advocating, and let me know how you get on through linuxformat@futurenet.co.uk or www.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

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 Jono Bacon, Mike Saunders, Chris Brown, Hoyt Duff, David Cartwright, Richard Cobbett, Richard Drummond, David Coulson, Hans Huberland, Andy Chappelle, Frank Charlton, Michael J Hammel

COVER CD PRODUCTION

CD Editor Neil Bothwick

ART CONTRIBUTORS

Art assistance Martin Mulchinock
Photography Amanda Thomas, Smon Lees, James Wilson, Gavin Roberts, Getty, Powerstock, Eyewire, Photodisc, Digital Vision
Illustration Paul Bateman, Chris Winn, Shane Collinge

ADVERTISING SALES

Key account manager George Gill george.gill@futurenet.co.uk
Deputy portfolio Ad Manager Diane Clydesdale
Business development manager Damian Hughes
Senior Sales Executives Giles Crosthwaite-Scott, Harry Maltby
Classified Sales Executive Stephen Hall

MARKETING AND PROMOTIONS

Marketing Manager Mike Hawkins
Subscriptions Executive Marie Spicer

PRODUCTION

Production Co-ordinator Diane Ross
Group Production Manager Clare Tovey
Advertising Production Co-ordinator Jo Crosby

MANAGEMENT

Publisher Kelley Corten
Group Publisher Dave Taylor
Publishing Director John Weir
Managing Director Colin Morrison

DISTRIBUTION AND CIRCULATION

Circulation Manager Jamie Malley
jamie.malley@futurenet.co.uk

Distributed by Seymour Distribution, 86 Newman Street, London W1T 3EX **Tel** +44 (0)207 3968000
Overseas Distribution by Future Publishing Ltd. **Tel** 01225 442244.

Overseas Licences

International Licensing Director Simon Wear
simon.wear@futurenet.co.uk **Tel** +44(0)1225 822798
Fax +44 (0)1225 788105 **Cell phone** +44(0)777 5641493

Contact Details

Linux Format, 30 Monmouth Street, Bath BA1 2BW
Tel +44 (0)1225 442244 **Email** linuxformat@futurenet.co.uk

Subscriptions and Mail Order

Phone +44 (0)1458 271178. See page 100
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



Media with passion

NEXT MONTH

Issue 44 on sale Tuesday 12 August

The need for SOUND

Become a master of your sound channels with our expert guide to:

- Audio hardware ● OSS and Also ● Rippers and recorders
- Streaming servers and more...

ON TEST:

Ximian Desktop 2, Intel's VTune, JBuilder 9, Neoware thin client, Seapine Surround, SuSE Linux Desktop

DON'T MISS...

LXF's special Multimedia CD – it will be brim-full of the best software for getting the most out of your Linux box

PLUS

DON'T MISS NEXT MONTH'S LINUX PRO!

The legal pitfalls of domain names, Storage TCO 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 100

LINUX PRO

FROM THE MAKERS OF LINUX FORMAT

AUGUST 2003

Think training, think profit



Well-trained staff lead to better productivity. Find out what the options are for credible Linux certification **p4**

PLUS

Snort author speaks

Marty Roesch tells us how he sees the future of Intrusion Detection

Compiler writing


Flex and Bison – working with intermediate code

Ethernet Filtering

Filtering ethernet frames with ebttables

Serial ATA – the new SCSI?

New ATA devices that bury the old unreliable image



Welcome

Twenty-four pages of real-world Linux for IT professionals

At last year's Linux Expo in London, during both of the *Great Linux Debate* sessions, there were lots of questions about training and certification. Some seemed unaware of the LPI certification (even though they had a stand at the Expo), but many more were simply wondering where on Earth they could get training before they sat the exam. At that time, there wasn't an awful lot we could tell them, but the UK has finally caught up with countries like Japan and the USA with proper courses now available. This issue we have taken a look at the course for LPI Level 1 offered by the Training Camp.

There's more to the LPI than just the Level 1 tests though – more even than just Linux certification for admins too. We tracked down the Linux Professional Institute's President, Evan Leibovitch, for an exclusive chat about the challenges the LPI faces and how it will broaden its remit to include certification on the desktop.

If you've been following our series on *Flex and Bison*, and how compilers actually work, this issue's installment should be one of the best yet, as we get to grips with intermediate code generation for our embryonic language SKYLang. Skip on to page 10 for more.

Network Intrusion Detection Systems (NIDS) have become more and more important in the sysadmin armoury. On Linux, one tool stands out from the crowd when it comes to Intrusion Detection, so we were delighted to have the opportunity to speak to *Snort's* creator, Marty Roesch on page 20. As well as still providing direction on *Snort* development, he set up a company, Sourcefire, selling NID devices, and has plenty to say about what the future holds for software and intrusion countermeasures in general.

That's about all there is space for here, but there's plenty more in this issue of *Linux Pro*. Have a good read and email us your feedback!

Nick Veitch Editor
nick.veitch@futurenet.co.uk



“The UK has finally caught up with countries like the USA and Japan, with proper LPI certification courses now available.”

Contacts

EDITORIAL

Editor
Nick Veitch nick.veitch@futurenet.co.uk
Reviews editor
Paul Hudson paul.hudson@futurenet.co.uk
Art editor
Julian Jefferson julian.jefferson@futurenet.co.uk
Production editor
Matt Nailon matthewnailon@futurenet.co.uk

Group publisher
Dave Taylor
Publisher
Kelley Corten
Contributor
David Coulson

Photography
Corbis, Getty, Powerstock, Photodisc

ADVERTISING

Key Account Manager
George Gill george.gill@futurenet.co.uk
Business Development Manager
Damian Hughes
darnian.hughes@futurenet.co.uk
Senior Sales Executives
Giles Crosthwaite-Scott, Harry Maltby
Classified Sales Executive
Stephen Hall

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

AUGUST 2003

CONTENTS

TRAINING CAMP
Fast-track Linux certification **p4**


LPI
The future of certification **p8**

FLEX AND BISON
Intermediate code **p10**

STORAGE
Is SATA the new SCSI? **p15**

ETHERNET BRIDGING
...and bridging tables **p16**

SNORT
Marty Roesch speaks **p20**



4

20

Think training...

...think profit!

CNA, RHCE, MCSE, CompTIA, SCSA, CISSP – they're all professional IT qualifications of various sorts, but do they mean much in today's environment, and do we need another in the form of LPI? Professional certifications in non-IT disciplines have a long history of being trusted, however things are a little different in our field – many certifications, particularly those offered by Microsoft, are now often seen as paper qualifications that count for very little. In 1999, when I first heard that qualifications such as the Microsoft Certified Professional (MCP) were easy enough for a beginner to pass, I bought a book on Windows NT 4 Server – an operating system I had never used, and still have not used to this day – and a week later sat the exam. I was only a little surprised when I passed and received my “qualification”, and even a lithographed printout of Bill Gates' signature.

So, with some qualifications already being easy enough to attain without the least experience, why water the pool down with more? Before the Linux Professional Institute stepped in with LPI, the only meaningful attempt at widespread Linux certification was from Red Hat, with their Red Hat Certified Engineer (RHCE) programme. The key difference with LPI is that it is distro-neutral – it teaches Red Hat-like and Debian-like distro information, with no specific emphasis on either, which means that students are given much firmer grounding on all Linux distributions as opposed to on just one.

Training Camp is a successful worldwide company that offers intensive training courses in a variety of subjects, to which they have now added the LPI course. In just seven days they believe they can teach students how to install and use Linux, configure X-Windows, handle administration

Why did the penguin cross the road? To get to the LPI training camp, of course! PAUL HUDSON gets certifiable in Hampshire's leafy countryside...

INTENSE LPI TRAINING COURSE

TRAINING CAMP LPI CERTIFICATION COURSE
SUPPLIER Training Camp
WEB www.trainingcamp.co.uk
PRICE £2937 (inc VAT)

tasks such as backing up and adding users, and more, as well as passing two exams *en route*. Part of the reason they make this claim is because students are taught for most of the day (from 9am till 7pm) for the seven days, but perhaps the main reason is because Training Camp takes students away to a tranquil countryside environment in rural England, where they live on-site for the week – see the box, *The Environment* on the right for more information.

LPI Certification

Despite being restricted to one distribution, the RHCE qualification has attracted a reputation as being hard to get – the exams don't simply consist of multiple choice (multiple guess?) questions. Instead, you get a computer and you're asked to solve problems and provide solutions – things you can't really study for unless your “studying” is simply using Red Hat.

The LPI seeks to follow the same method, and so becoming LPI certified (LPIC) means you must be able to use Linux practically in addition to answering more general questions – it's no surprise, therefore, that over half of LPI students fail their exams. In order to provide structured levels of difficulty, the LPI is made up of three levels, known simply as LPI Level 1, LPI Level 2, and LPI Level 3, with each level split into two smaller exams – x01 and x02. Level 3 is currently under construction, but levels 1 and 2 are available right now. Roughly, LPI 101 and 102 are fairly straightforward courses covering basic Linux administration tasks, installation, and management, whereas LPI 201 focuses on more advanced administration topics, and LPI 202 is based around network administration – all fairly

general tasks in order to make sure students have a well-rounded understanding of the OS. LPI Level 3 is currently being developed, but is likely to cover much more specialised areas of administration than the previous two, more generic levels.

However, it's the process of LPI level development that is of particular interest, because the LPI is unique in that it attempts to be open source in its development process. That is, extensive comment is taken from the Linux community as to what should be covered and to what depth, and, because the LPI is entirely independent from distribution vendors they don't have to bias their questions to favour any particular product or solution. The end result is that each LPI course is very representative of what it means to be qualified to use Linux. To reflect this state of constant evolution, the course offered by the Training Camp is constantly evolving to take into account the emergence of relevant new ideas and information.

The LPI's lack of bias towards any particular distribution means that other companies can take and use the LPI certification scheme for their own purposes, adding their own specific content as needed, and indeed both UnitedLinux and Novell have both professed interest in customising the LPI courses for their own purposes. What this means is that LPI is rapidly approaching the point of critical mass – where the qualification is finally mainstream, and a large chunk of Linux users want to become certified through the programme.

“If you want to be an LPIC, chances are you're already an IT pro working as a sysadmin or other systems management.”

LPI Level 1, the course offered by Training Camp, is wholly encapsulated in the O'Reilly book, *LPI Linux Certification In a Nutshell* in that you can technically learn everything you need to know from reading that one book as long as you're a particularly capable self-learner. However, teaching yourself has never been easy, particularly for a topic as tricky as Linux can be, which is why the Training Camp method of learning is often preferred by learners.

At the Training Camp

If you want to become an LPIC, the chances are you're already an IT professional working in a system administration or other systems management role, and you're looking to either attain and qualify in core Linux skills, or you're looking to have your existing skills certified in a more concrete form.

Surrounded by fields and forests a short drive from Basingstoke lies Tadley Court, a quaint-looking arrangement of buildings and greenery which belies its high-tech interior. This is the UK headquarters of Training Camp, one of the world's leading providers of IT training solutions in a variety of topics, and also a company who

THE ENVIRONMENT

Nurturing the Linux admins of the future

FAR FROM THE CONCRETE OCEANS MOST OF us are used to in this day of age, the Training Camp facilities are located in Tadley, about mid-way between Basingstoke and Reading. Training Camp organise taxis from either of those rail stations, but also from Heathrow, Gatwick, and Southampton airport for students coming from abroad – you're met at the station/airport, and whisked away into the countryside to begin your learning. I must admit I was a little apprehensive as I approached the centre, because the taxi driver informed me that many people leave the centre absolutely exhausted after their intensive training was complete!

The building itself is air-conditioned throughout, and is split into four separate areas: accommodation, dining and relaxation, teaching, and testing. The quality of the accommodation is somewhere between student dorms from universities and a Holiday Inn – you get a double bed with ensuite bathroom, but the rooms are packed fairly close together. As the rooms are really only there for sleeping in and for private, after-hours study, they're more than comfortable enough. Add to that the fact that there's a fairly complete gym next to the rooms to make sure people can exercise all they



need while away, and it's fair to say that all the students at the training camp are well looked after in their living space.

On the dining front, there are three rooms open to students – the main dining area, where waitresses serve hot lunch and dinner replete with desserts, and a drawing room-style sitting area with several groups of comfortable seats to give students a place to get together and chat. A small bar area is available on request.

With regards to teaching, each class has its own teaching area, although there is the availability of moving to a larger room if the numbers are there. The teaching areas are simply downsized traditional classrooms – there's no decoration to distract, just comfortable seats (admittedly far more comfortable than most schools!) and rows of desks. Testing is kept quite separate, though –

there are several small testing rooms, each with their own CCTV cameras in to make sure that students don't attempt to cheat. Each student in the testing rooms has their own miniature cubicle and a laptop for them to use for the test, and, under the watchful eye of the video cameras, can get up and leave once they believe they have completed the 90 minute examinations.

To keep students relaxed and happy, there's a small snack room central to the complex, which has a coffee machine offering various types of coffee and hot chocolate, as well as a fridge containing over a dozen varieties of soft drinks, a water machine, and a selection of chocolate bars, crisps, and chewing gum/mints. Furthermore, there's a small pub only a short walk away from the estate, so students want for nothing during their stay.

COVER FEATURE **TRAINING**

◀ recently took onboard the LPI certification with much pomp and circumstance. Quite rightly, too – we visited Tadley Court for a day to experience their training methods, and were thoroughly impressed with the degree of professionalism and motivation we encountered.

We travelled down first thing on a Monday morning, ready for day one of the course. The other students, having arrived the night before, had already met over dinner and drinks that night, and also been given an introduction to the course. Each student, upon arrival, is issued with a laptop running Red Hat 7.3, a copy of O'Reilly's *LPI Linux Certification in a Nutshell*, a copy of the 150-page Training Camp LPIC course notes, plus a selection of high-quality stationery to last them for the week. They're also given a set of Red Hat 7.3 CDs and a Knoppix CD, both theirs to

Training Camp hopes to add to the breadth of respect that LPI certification is beginning to steadily garner in the US, Europe and other areas.

keep along with everything else except the laptop, as well as the option to have any other distro burnt to CD for them as desired.

The Tadley Court complex is both airy and relaxing, which means that students concentrate on learning, and not the air conditioning or external noise. The LPIC learning room is made up of four rows of desks, three of which seat three students each, with the front-most table being for the instructor. Each student has a network connection and

“Trained staff are more productive, more motivated and good for business, so...more and more sysadmins will be sent to certify.”

THE CERTIFICATION

LPI soon to be the new benchmark 'must-have' Linux qualification?

ONCE YOU PASS YOUR LPI LEVEL 1 course, you're qualified to administer – albeit on a somewhat basic level – Linux servers of all kinds. While some graduates might use “LPI” after their name, many consider it the equivalent of “Bronze Swimming Certificate”, with the thinking being that certification of your Linux skills isn't as important as your *actual* skill in the area. With many degree-level IT students declining to put BSc after their name despite having spent three or four years studying, it does indeed make sense that putting letters after your name after a one week course might come across as being somewhat pretentious.

It's yet to be proven that becoming an LPIC has any recordable impact on a graduate's job or salary prospects on completion of the course. Searching for LPI on Planet Recruit returned no matches, and searching Jobserve found just two – clearly LPI has a way to go with regards to convincing HR departments that the qualification is a good thing. However, this is likely to change over time. With so many students coming onto the courses now, it's likely to be only a matter of time before the qualification becomes widespread enough to be considered the *de facto* standard in Linux certifications.

laptop power supply waiting for them, and there's a large screen at the front for slides, backed up by white boards across two walls.

By about 8:55am, the last of the students had filed in from breakfast, ready for a 9 o'clock start. With the accommodation being only 100 metres away, it's not surprising that no one looks tired. Each of the students has had their Red Hat laptops overnight, and all look fairly comfortable as they boot up ready to start the class.

Back to School

The first hour of training discusses basic Linux information, answering questions such as “What is Linux?”, “What is the GPL?”, etc, and gives students a basic grasp on Linux concepts. However, the definition of “free” was skirted over very quickly, which could have left some students believing Linux was ‘free beer’ as opposed to ‘free speech’. Furthermore, terms such as “man pages” and “daemons” were used freely, despite many of the students not actually being aware of what either of them were. This problem exists because the recommended start level for LPI level 1 is somewhat vague –



I was informed by the lecturer that it was “recommended” that students had spent a year using Linux in industry. However only one person I spoke to in the class was in this situation, with most of the others either having dabbled a bit in Linux without much success, or being entirely new to it.

Regardless, the course proceeded on smoothly enough, and moved onto partitioning schemes. The students had been given required reading from their *Nutshell* books last night, covering the various popular partitioning schemes, so they were fairly clued in on what was good and why. In order to remain vendor-agnostic, *fdisk*, *cfdisk*, and *Disk Druid* were all covered in equal depth, with students being required to create their own partitioning layout based upon what they had learned.

With partitioning out of the way, we took our first of several breaks for the day, and the students went first to the snack room, and then off outside to bask in the sun for twenty minutes. Having only known each other since Sunday night, talk was fairly sparse and limited to course material and the weather. However, as the day wore on the group started to bond more, exchanging learning ideas, tips, and personal information more freely.

After the break we continued on with a discussion of hardware support, IRQs and interrupts, and finally rounded the morning off with an hour-long tutorial on how to use *Vi*. Many of the students were foxed by *Vi*, particularly when it came to regular expressions – these are covered later in the course, and only touched on briefly during the *Vi* tutorial.

We had another break, and moved onto installation. This is a particularly key issue in the course – the students need to have thorough knowledge of how to install at least one Linux operating system onto their laptops, as well as be able to partition the drives themselves, and view the administration information during installation in order to solve common problems with IRQs and interrupts. The students had to install Red Hat 7.3 several times each – once using the automatic partitioning wizard, once using their own, modified version of

SAMPLE QUESTIONS

Test yourself – could you benefit from a certification course?

HERE ARE SOME SAMPLE questions pitched at LPI Level 1 (the easiest level) – see how well you do without touching your computer for help!

1 Which command is used to change settings on IDE hard disk drives?
 a) hdparm
 b) hddparm
 c) diskparm
 d) hdparam
 e) ideconfig

2 Your logfile shows repeated connections to TCP port 143. Which named service is being accessed?
 a) imap
 b) smbd

c) nmbd
 d) pop2
 e) smtp

3 If you had a Linux system routing 3 different networks through 3 NICs and you were having trouble with your IP-Forwarding. Where would you look to ensure that IP-Forwarding is actually enabled?
 a) cat /proc/sys/net/ipv4/ip_forward
 b) netstat
 c) cat /proc/net/tcp
 d) iptraf -d eth0
 e) tail -f /var/log/messages

4 To cause a particular print job to be printed next, regardless of its

current position in the queue, what command would be used?

a) lpc topq
 b) lpc -t
 c) lpq -t
 d) lpq --next
 e) lpc move

5 You need to find out which package owns a file called /etc/paper.config. Which command will answer this question?

a) rpm -qf /etc/paper.config
 b) rpm -qalgrep /etc/paper.config
 c) rpm -Fq /etc/paper.config
 d) rpm -q /etc/paper.config
 e) rpm --requires /etc/paper.config

THE EXAMS

TO BECOME LEVEL 1 LPIC certified, you need to pass two exams: LPI 101 and LPI 102. Both are 90 minutes long and consist of a varying number of multiple choice questions. With Training Camp, students are given very realistic mock exams before their finals so they get the chance to spot any major holes in their knowledge, and the course instructor is always on-hand for after-hours tuition.

When you sign up at Training Camp, you get two exam coupons that you can either expend at their training suites or indeed at any VUE test centre. These coupons are valid for either LPI 101 or LPI 102, which means if you fail 101 the first time around, you can redeem your second coupon to resit the exam.

the partitioning wizard's recommendations, and once using their own partitioning layout. We broke for lunch at 1pm, with some students spending up to an hour (the maximum) at lunch, and others spending just fifteen minutes – some certainly appeared much more enthusiastic than others!

Back in the classroom, most people had managed to finish their Red Hat installations, and were moving onto Knoppix – students are taught both the Red Hat and Debian manner of installing software, although they only need to select one of the two during exams. By 6pm, the students had finished their installation run and were preparing for a roundup of the day's learning before dinner.

Each day of training ends with a similar roundup, and they rollover. That is, day 1 contains a round up of day 1's learning; day 2 has day 1 and day 2's learning; day 3 has day 1, day 2, and day 3's learning, etc. As such, the plan is that students get progressively more trained as the week goes by, whilst not being able to forget what they have already learned.

Conclusion

The LPI is a certification that's just waiting for the big time, and Training Camp is already there, waiting to get people certified in style. Successful managers know that trained staff are more productive, more motivated, and good for business, so it's almost certain that more and more system administrators will be sent to certify with LPI, particularly given how straightforward Training Camp make it.

What's necessary now is to widen the range of the courses to include desktop users, so that companies can ship off departments of people for Linux retraining. We've already seen that Training Camp can handle very large groups of people training at once – it's now just a matter of time. ■

CERTIFICATION

The LPI Vision

After just two years, the Linux Professional Institute has come from nowhere to be one of the most important organisations in the field of Linux training and certification. With testing partners across the globe and relatively low exam fees, the LPI manages to retain a reputation for thoroughness and credibility. *Linux Format* caught up with the LPI's President, Evan Leibovitch to find out the secret of their success so far, and where the LPI is headed.

LINUX FORMAT: How successful do you think the Linux Professional Institute has been so far?

EVAN LEIBOVITCH: As far as I'm concerned, it been incredibly successful. In our first few years of operation we have delivered more than 27,000 exams worldwide. We are told by the testing organisations that this is well ahead of the pace of some of the more established systems.

It takes a programme a good bit of time to get on its feet and get a critical mass. According to the people we have been speaking to, we are well ahead of that pace. We're extremely happy with the numbers we have done so far.

LXF: Do you think a lot of that is down to the fact that the LPI is 'vendor agnostic' as opposed to the more specific nature of many of the other programs?

EL: Well, we take it a step further and call ourselves 'vendor independent', in the sense that we are not just distribution-neutral, but that LPI is very much a product of the Linux community. There is a lot in LPI that shares its structure to the way that Linux in general works. We have a worldwide community of people helping us to make the exams, to promote the exams, to help improve on everything and to give us feedback on what people really need from Linux in education.

In terms of what makes LPI different – it's the community nature that not only allows us to cover multiple distributions but also work with different training vendors, different methods of preparation. Because we do just the certification we work with multiple training vendors as well as different training methods. We have people like IBM with a fairly substantial web-based tutorial system, you have companies doing CD-ROM-based learning, there are book companies like O'Reilly and so someone preparing for LPI really has many, many different choices.

We really do have the community, not vendors, dictating what LPI should be about. I'll give you an example – LPI does not expire certifications. This is something that is commonplace in IT, but the community told us they did not want a mandatory expiration. In a University program you



Evan Leibovitch – staying 'distribution-neutral' is the core idea behind the spread of Linux training and certification.

NICK VEITCH talks to Evan Leibovitch, President of the Linux Professional Institute.



don't have your computer science college coming back to you in ten years saying "the things that people are learning are different, so come back and do some more learning or hand back your bachelor degree." That just doesn't happen. The point we find relevant is that the knowledge you received back then never gets taken away. You can be encouraged to re-certify or refresh what you are doing, but the knowledge you had back then never gets taken away, and having a certification body say "OK, yesterday you were certified and today you are not, simply because we have done an upgrade of the software," we find to be driven by vendors, not by the people wanting to get certified. Since we don't exist to sell any particular product, whether it's a Linux distribution or a particular kind of training, we act and respond in the way the community wants us to.

LXF: Aside from the numbers of people taking the exams, do you have any feedback suggesting the industry as whole is taking note of LPI, insofar as asking for certification when hiring?

EL: Well the first step in this has been the support of other components in the Linux industry, where the LPI has been successful. We're working together with both UnitedLinux and Novell, both of whom have said they will base their own certification on the LPI base. They are going to trust

our certification for testing Linux core capability, and add on their own tests that will examine skills for their own specific additions. We have already received that sort of endorsement from people within the Linux industry. We are now starting to find examples of companies and consultants saying "hire us because we have certified people" or "we have a demonstrated level of skill because we have the certification." It takes a critical mass and acceptance from the community for this to happen, and it is to LPI's credit that this has started happening. We have the endorsement of the Linux community, our next step is to start talking more to the Human Resources industry and making sure that they are aware, not only that LPI exists, but of the high quality of the programme.

We are quite proud of the fact that 57% of the people who take the exams do not pass. We consider this significant. If people do not know the material when they come to take the exam they will not pass. This is because we have deliberately designed our programme to be extremely high quality, and it's something that the HR community is taking notice of – an LPI certification *does* mean something. They know that people who come in and think they are going to breeze through it just won't. We quite regularly get phone calls from people who say "I've been administering a Linux server box for so many years... how dare you fail me!" Those kind of things happen, but the fact is if they don't know the material well enough to pass the exam, they don't pass.

LXF: Which is the whole point of having an exam, isn't it?

EL: Precisely. Having an exam that is so easy to pass that anyone can do it, as far as we are concerned starts to have extremely little value both for those getting certified and those who would hire them.

LXF: I did have a look at some of the sample papers and there were some tricky questions there. I know the material that is covered and even the questions themselves are hotly debated within the LPI organisation. Has that been problematic – has having such a huge community involvement slowed down the introduction of the exams?

EL: Not at all. The main reason for that is while we do depend on the community to assist us, the community is driven by a very highly professional core of people making sure that the exams are edited and evaluated properly. On our staff we have a PhD in psychometrics (the study of making tests). The psychometrician is involved in evaluating the exams. There is a beta program for the exam, where we put out a test version and evaluate the response. Every question on an LPI exam is evaluated for relevance that doesn't necessarily mean it's evaluated for accuracy. The questions might all be accurate, but our criteria for what makes a question good is whether it separates those people who know their stuff from those who don't. A question that is accurate but which everybody gets right or everybody gets wrong, doesn't do what we need from it.

For editing and evaluation of the questions, we start with the community but we also have a very professional staff

"We have designed our programme to be extremely high quality, and it's something that the HR community is taking notice of"

that assist with the development process. We have a combination of a good enthusiastic community that is balanced off with a professional and skilled staff that put together the programme.

LXF: Considering the international scope of the LPI, the patterns of Linux usage seem to be slightly at odds with the traditional areas where IT is deployed. We know there has been some work on translating exams - is that something that is still a big priority for the LPI?

EL: Absolutely. Certainly we are stronger in some countries than we are in others. We have started to put together a highly decentralised network of groups in many different countries who are organising locally and assisting LPI in advancing Linux education in their countries, and at the same time making us aware of issues that may effect our program and making sure we are sensitive to doing this properly. At this time we have international groups formed or forming in countries such as Germany Japan, Australia, Brazil, USA and soon China. These are areas where members of the community are not simply assisting us, but going as far as putting together an incorporated non-profit body in their own country to promote LPI. At the moment the exam is available in English and Japanese, by the end of the year we fully expect to have the exam in German, Chinese and Portuguese.

LXF: We know the development of the exams is ongoing. Level 3 is in development, is that correct?

EL: Level 3 is supposed to progress very much in the same way as you would go through university – you would take generalised courses in the lower levels and then specialise in the upper levels. Now you may have heard about security because that looks to be the very first speciality that we will be working with. But there is also interest in Enterprise management, Linux clustering and possibly other things.

There is a high level of interest in that stuff for Level 3. However, before we finish level 3 there is a huge amount of interest in LPI getting involved in certification for Linux desktop applications. There is growing maturity on desktop software, and having a decent certification program is the best way of combating obstacles like the lack of targeted education. There has been a substantial amount of interest.

Our challenge is not to standardise on one desktop or set of applications, but making sure we can do a program that can accommodate multiple approaches

LXF: What other goals are there for the growth of the Linux Professional Institute?

EL: We do have an interest in trying to increase support. We have an interest in trying to set up an LPI group in the UK, much as we have been able to do in Germany, Japan and elsewhere. ■

FURTHER INFO

LPI (www.lpi.org) – the organisation's home page is a great resource for those wishing to study for an exam, or those wanting to find out more about what LPI certification means

Linux Certification in a Nutshell – This O'Reilly book covers the Level 101 and 102 exams, is available now, and costs £19.95 (ISBN 1-5659-2748-6)

Try an exam (www.linux-praxis.de/lpsim/lpi.html) – this site hosts an examination simulator, which should give you an idea of the type of questions in the full exam.

HELPING OUT If you are interested in helping the LPI set up a branch in the UK, please contact the LPI directly, or alternatively mail **Linux Format** at linuxformat@futurenet.co.uk and we'll pass on your details.

FLEX AND BISON

flex and bison. Compiler writing



Assembly: the final frontier. These are the voyages of the language *SKYLang*. Its continuing mission: to explore strange new functions; to seek out new loops and new compilation – to boldly code where no geek has coded before! This issue we’re going to hit warp speed and get into the real internals of compiler creation by exploring intermediate code and what it gives us, and also how it’s the last stop before we output some assembly code. But first, a recap to make sure we’re all clear where we are.

In the first part we discussed what compilers are, how they work, and what they can be used for. We started work with *flex* and *bison* to create a simple interpreter that spotted (but otherwise ignored) variable assignments. In the second part we took *SKYLang* forward to properly handle typeless variable assignment, operator precedence and association, we handled overflows and type conversion, and also outputting values of variables. In this instalment, the most important thing to do first is fix the bugs in the code to make sure it’s reliable using the bug fixing methods provided on page 12 – it’s best to get bugs fixed before moving on. Once that’s done, we’ll be looking at what intermediate code is, why it’s helpful, and then implementing it in our *SKYLang* compiler. I had hoped to be able to start covering functions in this issue, but there just isn’t enough room, which means it’s been pushed to next issue.

In part one of this series I mentioned you needed solid knowledge of C++, and that’s more important than ever now!

PART THREE
PAUL HUDSON
gets into the *really* complicated stuff - how to convert SKYLang scripts into intermediate code, what that means, and why it’s the way forward...

Intermediate Code

In a decent optimising compiler, a program goes through three distinct states: input code, intermediate code, and output code. Converting source code to intermediate code before outputting it is actually a big advantage: firstly, it forces a split between the parts of a compiler, which means that enabling a compiler to create assembly for more than one machine, or to handle other source languages, is simply a matter of changing half of it as opposed to all of it. As long as new languages compile down to the same intermediate code, or the same intermediate code can be made into various other assembly languages, the front-end/back-end plan is easy. Another advantage to using intermediate code is that it puts the source language into a form we can understand easily, meaning that we can recognise various code structures and can therefore perform optimisations on the code.

Intermediate code is a particularly tricky topic to cover, so we’ll be starting off relatively easily and building up. Switching to intermediate code means that we can add functions such as if conditionals, loops, and functions, so it’s very worthwhile.

Our intermediate code will look like assembly code, except it won’t be platform specific. By “assembly code” I mean that it will be quite basic, with each instruction having no more than two operands. For example:

```
$i = 40 + 10 / 2;
```

will become, in intermediate code

```
$temp = 10 / 2;
```

```
$i = 40 + $temp;
```

As you can see, our intermediate code is very close to x86 assembly, which will make assembly conversion a doddle later on if we go that route. Our intermediate code will have a variety of operations, which we will add to over time. To begin with, we’ll get setting variables to work, and move on to cover other code later. Adding the ability to set variables incurs a lot of other code too, so this will be the hardest part of intermediate code generation. Last issue, in order to fit into the space provided, we had to use a fairly simplistic way of passing variables around – this time it’s “the real thing”, so be ready to do a bit of rewriting.

CORRECTION

WHEN I SAT DOWN TO WRITE THIS article a few weeks after finishing last issue, I noticed a bit of a gaping error in the code from issue 42. `%left`, the *bison* operator associativity command, was used incorrectly last issue. The order in which `%left` is used is crucial, that much is true – however, the ordering must be defined by line numbers. So, on line 20 of last month’s code, you’ll see this:

```
%left T_DIVIDE T_MULTIPLY T_PLUS  
T_MINUS
```

That should read:

```
%left T_MINUS T_PLUS  
%left T_MULTIPLY T_DIVIDE
```

The reason for this is that the lower the token (on your screen), the higher the precedence it has. I must’ve been carried away trying to save lines of space, so the error lies between my chair and keyboard!

Defining operations

Our finished intermediate script will be a collection of operations, known as **op codes**. As with assembly code, this will be “flat” – there will be no tree structure involved, merely jumps to and from relevant parts of the code. Our script will be held in an **op code array**, and we'll be iterating through this array to execute actions, jumping around as necessary. At this point the intermediate code will not be optimised – we're interested in getting it to work right now; optimisation will come later.

First, here's what an **op array** looks like:

```
class SKYOpArray {
public:
    std::vector<SKYOp> opcodes;
    void execute();
};
```

An **op_array** is indeed just a vector of **SKYOps** – it's encapsulated inside a class so that, later on, we can add functions such as **optimize()**, or **output_to_c()**.

```
class SKYOp {
public:
    SKYOpType opcode;
    SKYVar *result;
    SKYVar *op1;
    SKYVar *op2;

    SKYOp();
    ~SKYOp();
};
```

A **SKYOp** contains a **SKYOpType** that stores what kind of operation this **SKYOp** is (does it add? Subtract? Test numbers?), plus three pointers to **SKYVars** to hold operand 1, operand 2, and the result of the operation. It's crucial that we use pointers here – as well as in other places – and we'll be looking at why that is shortly.

```
enum SKYOpType { SOP_ASSIGN, SOP_JMP, SOP_JNZ,
SOP_IS_EQUAL, SOP_ECHO, SOP_ADD, SOP_SUB,
SOP_MULT, SOP_DIV};
```

SKYOpType, as you can see, is simply an enumeration to handle the various operations we want to perform – this is a short list for now, and we'll look at adding more later. As the operations are used, I'll be explaining how they work – don't worry about the individual **SKYOpTypes** right now.

Those are the three key things we'll be using to handle operations. Now, let's take a look back at the intermediate code given above:

```
$temp = 10 / 2;
$i = 40 + $temp;
```

Now you know that each operation has two operands and a result, here's how that same code actually looks (you might need to re-read this a few times!):

```
operation1_result = operation1_operand1 /
operation1_operand2;
operation2_result = operation2_operand1 /
operation1_result;
```

As you can see, the result of operation 1 is also the second operand in operation 2. As the value of operation 1 (10/2) is computed and stored in operation 1's result, it should also update the value held in operand 2 of operation 2. Rather than trying to program some crazy

system of chain-updating variables, we simply switch to using pointers. Therefore, both operand 2 of operation 2 and the result of operation 1 both point to the same **SKYVar**. Note how the three **SKYVars** in **SKYOp** are all pointers to **SKYVars** – this is why.

The Big Switch

The problem with using pointers is that the rest of the system currently does not – it pushes and pops whole **SKYVars**, which is quite an inefficient thing to do. However, as you can imagine, teaching pointers is much harder – which is why it's been saved till now! To switch to pointers, a variety of changes need to be made. Firstly, we no longer want our lexer to pass back **SKYVars** – instead, it should pass back primitive values, which we convert into **SKYVars** inside the parser. Look for this line in both `skylang.y` and `skylang.l`, and delete it:

```
#define YYSTYPE SKYVar
```

That will revert the lexer back to returning integers as its default value. Naturally, we need to return multiple value types back from the lexer, so we use a new, built-in function of *bison* to handle this, **%union**.

```
%union {
    SKYVar* var;
    int intval;
    double floatval;
    char* charval;
}
```

Type that code into `skylang.y`, before the first tokens and after the **%}**. That informs our lexer and parser that the lexer can return any of the members of the union (a pointer to a **SKYVar** in **var**, an integer in **intval**, a floating-point number in **floatval**, etc). Note that there is no **Type** in there any more, which means you'll have to take out the two lines **yyval.type = ..** in `skylang.l` – we'll be assigning types in the parser from now on.

Now that a given token can be of several types, *bison* needs to be specifically told which tokens are of which type. This is done by specifying, in angle brackets, the name of the union value each token corresponds to. Here's how it looks:

```
%token <var> T_ASSIGNEQUALS T_SEMICOLON T_ECHO
T_OPENBRACKET T_CLOSEBRACKET T_FUNCTION
T_COMMA
%token <charval> T_VARIABLE T_STRING
%token <intval> T_INTEGER
%token <floatval> T_FLOAT
```

So, as you can see, **T_VARIABLES** and **T_STRINGS** are both **charvals** (that is, **char***s, according to the union), and **T_INTEGER** is an **intval**. Furthermore, whenever we assign to **\$\$** (the result of an operation), we're assigning to a non-terminal type. For example, last issue we had code like this:

```
if (dtemp > INT_MAX || dtemp < INT_MIN) {
    $$ .type = svFloat;
    $$ .floatval = dtemp;
} else {
    $$ .type = svInt;
    $$ .intval = $1.intval * $3.intval;
}
```

Now that **\$\$**, an *expression*, can be of various types, we need to clarify to *bison* that expression



FLEX AND BISON

◀◀ should be the **SKYVar*** part of our union. This is done with **%type**, and you should put it before the **%tokens**. Here's how it looks:

```
%type <var> expression statement
```

If you've made it this far, you should be able to recompile your application using pointers. If you receive errors, it will likely be fixable by changing **.s** to **->s**, or making objects pointers. For example:

```
sky_strtoint(&$3);
```

```
...
```

```
if ($1.type == svString) { .... }
```

The first line passes **\$3** as a pointer into the function, but, as it's already a pointer, you can drop the ampersand. The second line references **\$1** directly, which is incorrect as

we're using pointers now: **\$1->type** is correct. If you can't get it recompile, don't worry – there are lots of other changes to make, plus the entire code for this issue is on your cover disc this month, so you can either take that and use it as you want or modify your version to fix any errors.

Adding op codes

If you manage to get your code working entirely with pointers, you're doing very well indeed. If not, don't worry – **diff** against my version to see where you've gone wrong.

Now we're using pointers, we can get onto generating intermediate code. This involves another upheaval, but it's not as complicated as switching to pointers was! In essence, using intermediate code means that our compiler *compiles*

FIXING BUGS

Compiler corrections

THE ONLY THING WORSE THAN HAVING BUGS in your code is having bugs in your compiler, so we're going to nail down the problems we know about already before starting off on new work. Here are the problems that currently exist in our compiler:

- i) "123" as the charval of a **SKYVar** gets converted to the integer 3 rather than 123.
- ii) Assigning to the same variable name causes a memory leak
- iii) Echoing a non-existent variable causes a crash

I haven't ordered them by difficulty, so let's just jump in at i). Please note that before continuing you should have your SKYLang code fully up-to-date – that means filling in type conversion for division, adding, subtracting, etc, and having overflow checks in there too, as shown in the example multiplication code given in the last issue (with *LXF42*).

"123", a string containing a number, gets converted to 3 because our string to integer conversion is simply calling **strlen()**. Naturally it would be better if, rather than **strlen**, we attempted to properly convert to a number first, and if that failed (eg if the string was "cat") we could perform some other processing. The first step to get this done is to change this code:

```
if ($1.type == svString) { $1.type = svInt; $1.intval = strlen($1.charval); }
```

```
if ($3.type == svString) { $3.type = svInt; $3.intval = strlen($3.charval); }
```

into this code:

```
if ($1.type == svString) { sky_strtoint(&$1); }
```

```
if ($3.type == svString) { sky_strtoint(&$3); }
```

The reason for this is because we'll be calling a function, **sky_strtoint()**, to handle our string to integer conversions, as we'll be doing it quite a bit. For extra speed there's no harm in

inlining the function if you want to toy around – however, notice we put an ampersand before the variable in **sky_strtoint()**, which means we pass in a pointer to it rather than a full copy – inlining the function wouldn't make all *that* much difference.

Here's the function we'll use to convert:

```
1 void sky_strtoint(SKYVar *strvar) {
2   strvar->type = svInt;
3   char* errors = NULL;
4   int strconv = strtol(strvar->charval, &errors,
5   10);
6   if (errors == strvar->charval) { // failed
7     strvar->intval = strlen(strvar->charval);
8   } else { // succeeded
9     strvar->intval = strconv;
10  }
```

Note that we pass the value to be converted in as a pointer to the original so we can change it inside the function. The first thing we do is to convert the type to integer, because the result of this operation will be a converted integer regardless of its value. We then call **strtol()**, which is a standard C function that converts a string to an integer. **strtol()** takes as its parameters the string to convert, a **char*** where it can store errors, and the base to use for conversion. If **strtol()** encounters any errors, it will store the full value of the character string it tried to convert into the **char*** errors we defined. However, if it works, it return the converted integer value. So, on line 5 we check whether our **char*** errors has been set to the value we were trying to convert – if the two match, we know the function has failed, so we return the length of the string we tried to convert using **strlen()**, which is the behaviour we had previously. If the function succeeds, we return the converted integer value.

Make sure you replace all your string to integer conversions with the call to the new function, then recompile and try it out. Try executing these statements:

```
$foo = "wom" * "bar";
```

```
$bar = "12" + 12;
```

```
$baz = "wombat" / 2;
```

```
echo $foo;
```

```
echo $bar;
```

```
echo $baz;
```

All being well, you should get back the results 9, 24, and 3 – perfect! Note that there is also a function **strtof()**, which converts strings to floating-point numbers – why not try implementing this function yourself?

Onto error number two: assigning to the a variable more than once leaks memory. Take a piece of code like this:

```
$foo = 5;
```

```
$foo = 10;
```

The first line would create a new **SKYVar** in **UserVars["\$foo"]** and set its **intval** to 5. The second line would create another new **SKYVar** in **UserVars["\$foo"]** and set its **intval** to 10, and we lose our reference to the memory created in the first line – a memory leak!

The way to solve this is to only create a new **SKYVar** if there isn't one already. If the same map key is used more than once, we reuse the same **SKYVar** as we created previously. To do this, change the first line of our first statement rule from:

```
SKYVar `newvar = new SKYVar;
```

to this:

```
1 SKYVar *newvar;
```

```
2 if (UserVars[$1.charval] == NULL) {
```

```
3   newvar = new SKYVar;
```

```
4   UserVars[$1.charval] = newvar;
```

```
5 } else {
```

```
6   newvar = UserVars[$$.charval];
```

in the first pass and *executes* in the second pass – currently we execute lines as we find them. Just below the definition of **UserVars** in `skylang.y`, add this line:

```
SKYOpArray OpArray;
```

And at the end of our `main()` function, add this line:

```
OpArray.execute();
```

`main()` now runs `yyparse()` which will, once programmed to do so, create op codes for all our instructions, which are then executed when we call `OpArray.execute()`. Right now we have all our execution code inside our parsing code, which means we need to do a lot of moving around. Take all the code from inside your actions and put it somewhere temporary for now, leaving just the skeleton like this:

```
statements: |
```

```
7 }
```

Then delete this line from the end of the rule (about five lines below “`case svString`”):

```
UserVars[$1.charval] = newvar;
```

This new code declares a pointer to the **SKYVar** `newvar`, but doesn't fill it up with anything. Line 2 checks **UserVars** to see whether the variable is set already, and, thanks to the intrinsics of the `[]` operator, will create space for it if it doesn't, and return `NULL`. If the variable isn't set (it's `NULL`), we create a new **SKYVar** and put it into **UserVars**. If the variable is set, we set `newvar` to be the **SKYVar** that's already in there, thereby removing the memory leak.

The third and final problem is that echoing out a non-existent variable causes a segmentation fault. This happens because our `std::map` to hold **SKYVars** holds a *pointer* to a **SKYVar**, and not an actual **SKYVar**. Because we create a new **SKYVar** to hold the result of assignments, the only issue is when an unset **SKYVar** occurs in an expression. As such, the problem is very localised, and therefore very easy to fix.

In `skylang.y`, look for these lines:

```
| T_VARIABLE {
    $$ = 'UserVars[$$.charval];
}
```

That's the part that says “if an expression contains a variable, return the variable from **UserVars**”. We need to modify that to first check whether the variable exists in **UserVars**, and, if it doesn't, it needs to be created. Here's how that would look:

```
1 | T_VARIABLE {
2   if (UserVars[$$.charval] == NULL) {
   UserVars[$$.charval] = new SKYVar; }
3   $$ = 'UserVars[$$.charval];
4 }
```

Didn't I tell you that would be easy? That's fixed all the known bugs currently in the system, which means we're all set to go ahead and take a look at how intermediate code works.

```
statements statement;
statement: T_VARIABLE T_ASSIGNEQUALS expression
T_SEMICOLON { }
| T_ECHO expression T_SEMICOLON { };
expression:
  expression T_PLUS expression { }
| expression T_MINUS expression { }
| expression T_MULTIPLY expression { }
| expression T_DIVIDE expression { }
| T_VARIABLE { }
| T_INTEGER { }
| T_FLOAT { }
| T_STRING { }
```

Note that I've taken as much whitespace out as I can to save space. All the action code (the bits in between the braces that are now missing) should be stored somewhere safe for now. We'll be using it all later in the `execute()` function.

The actions for addition, subtraction, multiplication, and division are all similar because they are all binary operations – that is, they deal with two operands. As such, we're going to write a function that handles such binary operations, called `sky_do_binop()`. As binary operations return a result that we may want to use elsewhere, we need to return a value from this function. It also needs to accept the type of instruction, as it could be an addition, subtraction, etc. As well, it needs to accept the two operands to use for the operation.

Internally, the function needs to create a new **SKYOp**, add it to the our list of op codes, set it to use the two operands passed in, then return the result so that it can be assigned to be the operand of another operation. Based on that description, here's the function:

```
SKYVar* sky_do_binop(SKYOpType optype, SKYVar *op1,
SKYVar *op2) {
    SKYOp *new_op = new SKYOp;
    OpArray.opcodes.push_back(new_op);
    new_op->opcode = optype;
    new_op->op1 = op1;
    new_op->op2 = op2;

    return new_op->result;
}
```

It's all fairly self-explanatory, except perhaps the `push_back()` call. `push_back()` is a vector function that adds a value to the end of the vector, resizing it if necessary. In this case, we have our new operation being added to the list of existing operations – simple, really. In order to grab the result of the operation and pass it on for other expressions, we return a **SKYVar***, so we need to modify the actions that call `sky_do_binop` so that they make use of the return value. Modify your actions to this:

```
expression T_PLUS expression {
    $$ = sky_do_binop(SOP_ADD, $1, $3);
}
| expression T_MINUS expression {
    $$ = sky_do_binop(SOP_SUB, $1, $3);
}
| expression T_MULTIPLY expression {
    $$ = sky_do_binop(SOP_MULT, $1, $3);
}
| expression T_DIVIDE expression {
```



FLEX AND BISON

```

<<  $$ = sky_do_binop(SOP_DIV, $1, $3);
    }

```

So, now addition, subtraction, multiplication, and division all create operations. To handle assignments, believe it or not, it's almost just a matter of calling `sky_bin_op` again, this time with `SOP_ASSIGN` as the first parameter. The difference is that the first token in an assignment is a `T_VARIABLE`, which, if you remember, we defined with `%token <charval> T_VARIABLE T_STRING`, which means it's a `charval` – a `char*` – and *not* a `SKYVar`. To fix this, we first need to convert it to a `SKYVar`, then pass that to the `sky_do_binop()` function. Here's how that looks:

```

T_VARIABLE T_ASSIGNEQUALS expression
T_SEMICOLON {
    SKYVar *newvar = new SKYVar;
    newvar->charval = $1;
    $$ = sky_do_binop(SOP_ASSIGN, newvar, $3);
}

```

As you can see, switching to intermediate code can make your *bison* grammar more readable! But it does come at a cost – read the box *Clocking Speed* below for more info.

Before we get onto execution, there are two more changes that need to be made. As mentioned earlier, we now pass primitive variables back from *flex*, and these need to be converted into `SKYVars` when they get into the parser. Right now our expression non-terminal handles `T_VARIABLE`, `T_INTEGER`, `T_FLOAT`, and `T_STRING` – need to be converted. Here's how that works, although note that I've really cut out spacing in order to make this fit:

```

| T_VARIABLE { SKYVar *newvar = new SKYVar;
newvar->type = svString; newvar->charval = $1;
$$ = newvar; }
| T_INTEGER { SKYVar *newvar = new SKYVar;
newvar->type = svInt; newvar->intval = $1;
$$ = newvar; }
| T_FLOAT { SKYVar *newvar = new SKYVar;
newvar->type = svFloat; newvar->floatval = $1;
$$ = newvar; }
| T_STRING { SKYVar *newvar = new SKYVar;
newvar->type = svString; newvar->charval = $1;
$$ = newvar; }

```

As you can see, the conversion is simply a matter of creating a new `SKYVar`, making sure it's set to the right type, copying in the value we got from *flex*, then passing it back as the result of the expression. The second change

CLOCKING SPEED

LAST ISSUE WE RAN SKYLANG THROUGH a test – assigning a million integers. Using the code from last issue, it executed the script in just over seven seconds, however using the code from this issue we find the same script takes just under *eleven* seconds to execute – a substantial slowdown. You might think “What a swizz – who really needs intermediate code anyway?”, but consider this: in a loop, you need to execute the same instructions

multiple times. How would you manage that if you didn't have intermediate code? Put simply, you'd need to somehow re-read each line of source code every time you went through the loop, then execute it again, and that's incredibly slow. When you convert to intermediate code, you trade in a speed hit with very simple scripts for a massive speed boost when it comes to executing the same code again and again, as in loops and user functions.



you need to make is to add constructors for our `SKYOps` to `skylang.y`. Just after `main()`, add this code:

```

SKYOp::SKYOp() {
    result = new SKYVar;
}

SKYOp::~SKYOp() {
    delete op1;
    delete op2;
    delete result;
}

```

In the destructor we do a little cleaning up after ourselves to keep things tidy, but the constructor is the important part – if you don't create `result` here, your program will segfault because `result` will be null.

Execution at last

Now what you've been waiting for: executing our op codes. Execution of code should simply be a matter of iterating through all our op codes, checking what kind of op code it is (`SOP_ADD`, `SOP_MULT`, etc.), then running pretty much the same code we had before. The code you need to execute op codes is almost 200 lines long, so you'll need to bring up your copy from the cover disc and follow my explanation, as there isn't enough room to reprint the function here.

The variable `abs_op_count` is set to 0 and incremented by one for each instruction we execute. While this isn't used currently, it's a good way to track how many operations have been executed. The main *for* loop that follows it counts from 0 to the last operation in our `OpCodes` vector, and this is where the work happens. The lines “`SKYOp *op;`” and “`op = OpArray.opcodes[i];`” are just there to help make referencing the current operation easier.

Continuing on, we have “`switch(op->opcode)`”, which is where we select the correct actions to perform for the current operation. As you can see in the code, once inside a given case statement we have pretty much the same code that used to be directly inside our *bison* grammar. There are some differences – for example, `$1` is now `op->op1`, `$2` and `$3` are now `op->op2`, and `$$` is now `op->result`.

What else is there? Well, nothing, really – the hardest part was making sure we have everything as pointer so that we can chain operations together – handling execution is a doddle in comparison. Check through your program and compare it against the one included on your disc to make sure there aren't any obvious errors, then run `make` and try it out. This time you'll notice that code you enter isn't executed until you press Ctrl-D.

Conclusion

While some of the difficulty might have been saved this month if we had used pointers from the very beginning, you'll agree that it might only have confused less-experienced programmers. However, that's all been made up for this issue – SKYLang now supports intermediate code generation. This might not sound like a big thing, and it's really just a stepping stone for other features – next issue we'll look at the new-found freedom intermediate code gives. Implementing intermediate code was probably the hardest thing we'll have to look at for a while, so relax – it only gets easier from here! ■

Serial ATA The new SCSI?



Bizarrely it is the home user who is clamouring most for these devices at the moment – the slimmer cables are not only less cumbersome and presumably less fragile than the familiar IDE ribbon, but crucially allow better airflow in the case, highly prized amongst the overclocking and ‘modding’ brigade.

But it is in the server market that Serial ATA is likely to provide the most compelling benefits, at least for now. HBA manufacturers haven’t been quick to push their products to market due to drive manufacturers not rolling out devices in great variety or volume, but that’s about to change...

Altered images

The perception has always been that the only real selling point for ATA devices is the price. They are considered less reliable, challenged in terms of performance and cumbersome to deal with thanks to the dual channel, multiple device nature of the interface, never mind the ridiculously short cables.

The SATA spec deals with the last point. All SATA devices, and early controllers support 4 or even 8, use a simple point to point protocol. The drives no longer share the available bandwidth nor are there any more complicated modes and setting which must be applied to the level of the lowest common denominator on a shared channel.

Reliability is a factor, as is performance. It is certainly true that the MTBF for an ATA device is lower than a typical

The Serial ATA standard has been around a while, but, as with many technologies, has taken some time to emerge into the real world.

NATIVE DRIVES

INITIAL SATA DEVICES AVAILABLE WERE RATHER simply kludged, with a converter on the standard backplane of the drive.

A subsequent generation of devices has emerged with the conversion effectively embedded in the drive electronics. In most respects, these devices are identical to their ATA forebearers, with a slightly different backend.

Truly native drives will be the next step, probably arriving early in 2004.

NEXT MONTH

...we’ll bring you a detailed review of the Raptor, and take a broader look at the SerialATA options for Linux.

SCSI mechanism, and the fastest ATA drives come nowhere near the performance of SCSI. Largely, this isn’t anything to do with the specification of the ATA standard, but simply how the market expectation has driven the manufacturers to differentiate between the drives.

Raptor

Western Digital’s Raptor drives are their first attempt to promote SATA as an SCSI-beater. The idea is to manufacture drives to the same quality standards as SCSI, but at a much more attractive price, thanks to the lower cost overheads of producing SATA hardware. The Raptor is a single disc drive, which explains the rather limited 36.7GB capacity, running at 10000rpm – giving it an impressive seek time for a non-SCSI device of 5.2ms. The build quality is such that the manufacturers guarantee them for 5 years – much longer than the standard 1 year offered on commodity ATA devices.

With the Raptor retailing for around £125, it is certainly cheaper than comparable SCSI devices, but not by a huge margin. The real savings for heavy storage demands is probably going to be in cabling and controllers. The Raptor has narrowed the gap between SATA and SCSI though, and with the current SATA roadmap, the traditional view of “SCSI = better” may find itself challenged by the evolution of SATA over the coming years. ■

WHAT IS SERIAL ATA?

The subject was covered in *LXF37*, but briefly here are the key points

- Uses a serial interface, rather than the parallel one used by ATA. At high speeds synchronising the data from parallel streams becomes more problematic, so a simpler, serial connection can achieve higher speeds
- Current specification is 150Mb/s, hardly faster than the 133 currently offered by ATA. Later versions will push this to 300MB/s (probably in 2004), then 600MB/s within five years.
- Cabling is much, much simpler. The serial connection only needs a few wires, so the cables can be more flexible, better protected, longer and a lot cheaper – certainly a lot cheaper than SCSI cables.
- Point to point interface – each drive is connected directly to the drive interface, there are no other drives in the way hogging bandwidth. No chaining, no termination, no fuss.
- ‘Drop in’ replacement for ATA. SATA host bus adaptors are available, existing ATA devices can be attached via a converter if necessary.

Ethertables Ethernet bridging

Network capabilities of the Linux kernel have always been a particularly impressive reason for someone to use the GNU/Linux platform as a router, a web server, or simply a gateway onto the Internet from home. With *ipfwadm* in 2.0, *ipchains* in 2.2 and now *iptables* in 2.4, the IPv4 (IPv6 in the case of *ip6tables*) filtering possibilities when using a Linux kernel have made for a flexible and extensible system.


Of course, filtering IPv4 or IPv6 packets isn't the lowest level of the networking process, being Layer 3 in the OSI model. Below the IP layer is the Ethernet layer, layer 2, which operates with frames as the packet or payload of data being sent somewhere, along with 48bit MAC addresses being used to define each device on the link. However, unlike IPv4, MAC addresses are generally static and hard coded onto the device EPROM at the time of manufacturing, so one does not have the neat subnets which those who've used IPv4 or IPv6 will be familiar with. With Ethernet, everything operates on the link layer, such as between two NICs, or between hosts across a switch or hub. Generally Ethernet frames are not considered

If filtering IPv4 isn't enough, there's a kernel patch to allow filtering of Ethernet frames. DAVID COULSON explains.

A basic bridge includes two network interfaces and transparently forwards frames between them both.

routeable, but to all intents and purposes, an Ethernet switch routes Ethernet frames to the correct switch port through a system known as transparent bridging. Effectively, the switch, or bridge, will learn which ports particular devices with specific MAC addresses are connected to, allowing the frames to be delivered directly to that host, rather than having it broadcasted out of all ports, as with a hub or repeater.

While the switch is the most common device which performs transparent bridging, Linux supports functionality as a bridge device out of the box with 2.4 kernels. Creating a bridge on a Linux system with multiple Ethernet devices is a fairly trivial task (See Bridging Basics box) and it will happily forward frames back and forth without caring what they are for. Of course, on a large network where anyone can come along and plug their laptop into the local switch matrix, they can do a significant amount of damage, by sending spoofed frames to the switches to confuse them, or deliver faked ARP responses to hosts in order to capture packets. This can be solved by using an *ebtables* (<http://ebtables.sourceforge.net/>), which is the Ethernet version of *iptables*.



Using ebtables

The *ebtables* command line structure is almost identical to *iptables*. As with *iptables*, *ebtables* is split into multiple tables, each containing one or more chains containing our rules. Rather than having *filter*, *nat* and *mangle* which *iptables* has, *ebtables* has *filter*, *nat* and *brouting*. The first two are very similar to their *iptables* counterparts, in that they filter and rewrite source/destination addresses on packets, but the latter is specific to the *ebtables* system. We'll look at *brouting* later.

As with *iptables*, we can add, insert or delete rules in a chain. When we add a rule with the **-A** option, it will be added to the end of the chain, so it will be checked last. By inserting a rule with **-I** we can either insert a rule at a specific point in the chain, so **-I INPUT 2** would insert this rule as the second rule in the INPUT chain, or by specifying no value the rule will be inserted into the top of the chain. When we insert, it does not overwrite the previous rule, it will simply move them further down the chain. Deleting a rule can be done by either specifying a rule number, such as **-D INPUT 2** or by including the complete rule as part of the **-D INPUT** command.

Each *ebtables* command to insert, delete or add a rule must be followed by the name of the chain we want the rule to apply to. We may also specify a table with the **-t** switch, but if it is neglected then *ebtables* defaults to the 'filter' table.

To check our existing rules, we can simply do:

```
# ebtables -L INPUT
```

or alternatively this:

ebtables -t broute -L BROUTING

The **-L** option can be expanded to **-Ln** to include rule numbers for a future insert or delete command, or to **-Lc** to include byte counters for each rule.

Once the rules are set up, we can save our current tables and chains for reloading after a reboot:

```
# ebtables --atomic-dump -t filter --atomic-file
ebtables.filter.dump
```

The dump can be reloaded from this file with:

```
# ebtables --atomic-commit -t filter --atomic-file
ebtables.dump
```


Once we have a file specified by **--atomic-commit**, we can add the **--atomic** option to all of our commands to have the changes applied to our configuration file.

filter

The filter table contains three chains, which are used at various points of the bridging process. INPUT for packets destined for the local bridge, OUTPUT for locally generated frames and FORWARD which is for packets which are being forwarded between ports on the bridge.

Here it's worth noting the distinction between forwarding and routing in the *ebtables* documentation. When *ebtables* talks about forwarding, it means forwarding a frame between two Ethernet ports within the bridge. No IPv4 operations are made on the frames, and they simply come in one port and get thrown out the other. To 'route' in the *ebtables* world is to push the packet up to the IPv4 level and have it routed based on the local IPv4 routing table, or indeed the routing table for the specific layer 3 protocol for the data within the frame, on the host. While 'forwarding' is used by *iptables*, and indeed the Linux kernel, in the context of IPv4 and IPv6. As we're looking at *ebtables*, we're going to use the *ebtables* distinction, although once outside of the *ebtables* world, forwarding and routing generally mean the same thing.

With the 'filter' table we can block specific frames coming in or leaving the bridge, or simply stop specific frames being forwarded between ports. Rather than filtering based on IP source or destination, we instead have to filter based on



A router both forwards frames between interfaces and forwards IPv4 and other protocol packets.

source or destination MAC address, which makes life a little more interesting.

We can do a very basic filter rule, which will block frames coming from a specific device:

```
# ebtables -A INPUT -s 00:FF:5D:40:A7:72 -j DROP
```

Those who've used *iptables* will be familiar with the ACCEPT and DROP targets for rules, but *ebtables* adds a couple of other targets which change the functionality a little. The CONTINUE target will simply have the *ebtables* system continue down the chain if the rule is true. *ebtables* also has the RETURN target, which pushes back into the parent chain and continues through the rest of the rules in that chain. As one would expect, RETURN is not applicable in one of the defaults chains as they have no parent chain. While a MAC address can be specified, we can also use a broader definition of Unicast, Broadcast or Multicast. We can also filter based on destination too, using the **-d** flag:

```
# ebtables -A INPUT -d 00:FF:5D:40:A7:72 -j ACCEPT
```

Not particularly exciting at this point, but there are plenty of other things we can do with *ebtables* to make it useful.

We can, for example, block or accept packets coming in a particular interface. It is also possible

PATCHING THE KERNEL WITH EBTABLES

How to get them applied

THE FIRST STEP OF PATCHING THE LINUX

kernel is to obtain the current Linux source tree from ftp.kernel.org and appropriately configure it for the hardware. If the regular kernel has problems with specific devices, then adding patches to the kernel doesn't help track those down. Note that some distributions have customised kernels which contain numerous patches, so these may well already contain a version of the *ebtables* patch.

With the Linux kernel source in `/usr/src/linux`, we can grab the *ebtables* patch from http://ebtables.sourceforge.net/v2.0/v2.0/ebtables-v2.0.003_vs_2.4.20.diff and apply it:

```
# cd /usr/src/linux
```

```
# zcat ../ebtables-v2.0.003_vs_2.4.20.diff | patch
-p1
```

One may also want to apply the *bridge-nf* patch from <http://ebtables.sourceforge.net/br-nf/bridge-nf-0.0.10-against-2.4.20.diff>, allowing Netfilter to filter packets which are bridged, rather than routed. While this horribly breaks the OSI model, as one is filtering based on Layer 3 information on a Layer 2 device, it is quite useful.

```
# zcat ../bridge-nf-0.0.10-against-2.4.20.diff.gz |
patch -p1
```

Once the kernel has been patched, assuming there are no rejects, one can check for new configuration options and build the kernel:

```
# make oldconfig
```

```
# make dep clean bzImage
```

```
# make modules modules_install
```

For 2.5 users, the *ebtables* and *bridge-nf* code is already incorporated into the mainstream tree, although incremental patches can be found at <http://ebtables.sourceforge.net/sourcecode.html>.

One will also need to either install the package containing the '*ebtables*' tool, or download the tarball containing the source code from <http://prdownloads.sourceforge.net/ebtables/ebtables-v2.0.3.tar.gz>. Compiling the code is trivial, although one will want to check their distribution's FTP server to find out if there is a package available.

ETHERNET BRIDGING

to use the **!** character to denote a logical 'NOT' in the rule, so to block all packets coming in eth0 which were not from the NIC with the MAC 00:FF:73:86:2F:F4, we could do:

```
# ebtables -A INPUT -i eth0 -d ! 00:FF:73:86:2F:F4 -j DROP
```

We can also specify a output bridge device for the FORWARD and OUTPUT chain, using the **--logical-out** option:

```
# ebtables -A INPUT --logical-out br0 -d 00:FF:EA:D0:E8:FB -j DROP
```

Note that the **--logical-out** is for the bridge device, not the port device, which will be used for sending the packet.

One can also specify a protocol for the frames, so we can filter ARP or RARP packets and other Ethernet protocol specific frames. We can also filter IPv4-based packets, and so forth, as the protocol value is based on what generated the frame. A list of names and the associated protocol value is stored in `/etc/ethertypes`, although anything lower or equal to 0x0600 is defined as the length of the header, rather than the protocol type, so for these frames the protocol is called LENGTH.

Both the IPv4 and ARP protocols allow for further filtering, based on specific values, so if we wanted to filter a packet

ESSENTIAL LINK MORE EBTABLES INFO

WE HAVE LOOKED AT MANY features of *ebtables* in this article, but there's a great deal more documentation at <http://ebtables.sourceforge.net/documentation.html> along with many examples and practical uses of the *ebtables* system.

from the IP 10.1.1.2 which did not have the MAC 00:FF:73:86:2F:F4 we could simply do:

```
# ebtables -A FORWARD -i eth0 -p 0x0800 --ip-src 10.1.1.2 -s ! 00:FF:73:86:2F:F4 -j DROP
```

While the same can be done with *iptables*, this can be used as part of a bridging configuration, so the packets don't have to be involved in any IPv4 routing decision at all. We can also match the IPv4 protocol used, such as TCP or UDP, and by port number, as well as by Type of Service (ToS) value.

For networks which incorporate 802.1Q VLAN frame tagging we can match specific VLAN packets which are traversing our bridge. To block packets for VLAN 10 from being forwarded on the bridge, we can do:

```
# ebtables -A FORWARD -p vlan --vlan-id 10 -j DROP
```

We can also match the protocol of the encapsulated frame within the VLAN frame with the **--vlan-encap** option.

Unlike *iptables*, instead of a LOG target for our rule, we must use a separate switch of part of an existing rule. We can use the CONTINUE target if we wish it to transparently continue down the chain even if the rule matches. If we want to log the information from the Ethernet frame, we can simply use the **--log** option, although we must use another

BRIDGING BASICS

802.1d explained

ebtables isn't a great deal of use without a bridge, so here's the basics. For quite some time, the Linux kernel has supported 802.1d Ethernet bridging capabilities, and has grown to support more elaborate capabilities, such as the Spanning Tree Protocol. Creating an Ethernet bridge with Linux is fairly trivial, but it's always nice to understand exactly what is going on before disabling Ethernet devices left, right and centre.

Simply put, an Ethernet bridge is a Layer 2 device which forwards Ethernet frames from one device to another. While this sounds similar to what an IP router does, rather than having a routing table, the bridge will learn which LAN segment a particular node is on based on the frames sent from this node through the bridge. Obviously since MAC addresses can vary so greatly, there is no clear 'subnet' or definable collection of devices on each device, so it's impractical to use anything other than learning capabilities. Nearly everyone will have come across an Ethernet bridge, in the form of a Ethernet switch, although unless one has a more expensive managed device, some of the more complex features of the bridge will not be available.

A bridge will contain one or more Ethernet-capable network devices, although to actually function as a forwarding device it needs at least two. Each device which is part of a bridge is known as a 'port'. Linux has many

Ethernet capable devices, which include standard Ethernet NICs, some wireless NICs and TAP devices which are commonly used by user-space applications, including *OpenVPN*, *Vtund* and *User-mode Linux*. Each port on the bridge is set to function without an IP address in promiscuous mode, which enables the kernel to see all Ethernet frames entering a NIC, rather than those which have a destination MAC of the particular device.

While the individual ports on the bridge can't have IP addresses, the bridge device is also available as a network device, so can be *ifconfig*ed as usual. This allows the system running the bridge to have an IP address on the segments it is bridging for. Of course, the bridge does not need an IP address to actually do anything, as it operates in layer 2, so the system which is performing the bridging can be completely hidden from view.

Creating a bridge requires the *brctl* utility which is generally available as part of the *bridge-utils* package, or can be downloaded from <http://bridge.sf.net/>. If we wish to create a bridge containing eth0 and eth1, then we do:

```
# brctl addbr br0
# ifconfig eth0 0.0.0.0 up promisc
# ifconfig eth1 0.0.0.0 up promisc
# brctl stp br0 0
# brctl setfd br0 0
# brctl sethello br0 0
# brctl addif br0 eth0
```

```
# brctl addif br0 eth1
```

```
# ifconfig br0 up
```

While the **br0** device does not have an IP address, it won't do anything useful unless it is actually 'up'. A quick look at *dmesg* will show something similar to the following:

```
device eth0 entered promiscuous mode
br0: port 0(eth0) entering learning state
br0: port 0(eth0) entering forwarding state
br0: topology change detected, propagating
```

We have set the bridge to immediately start forwarding packets between ports, although if one is running STP then it's always wise to wait some time before forwarding so that the current bridge can be included in the overall topology of the network. Spanning Tree is a protocol which allows bridges to communicate with each other, so that forwarding loops are not created on the network. STP can also be used for basic layer 2 redundancy, as one can run two bridges in parallel and when one is disconnected from the main LAN, the other bridge will automatically start to forward the packets instead. Bridges communicate via Bridge Protocol Data Units, or BPDU packets and the **sethello** option defines the time period between the bridge sending BPDU packets out onto the network. The higher the value, the longer it will take for the network to notice a port outage, but if it's low then there will be more traffic on the network, lowering available bandwidth to other services.

option to log either IP or ARP information.

If we wanted to log the ARP information from packets coming in on eth0, we can do:

```
# ebtables -A FORWARD -i eth0 --log-arp --log-prefix
"ARP: " -p arp -j CONTINUE
```

nat

The nat table in *ebtables* allows us, as with the similarly named table in *iptables*, to rewrite the source or destination address of the packet. While NAT of Ethernet frames is not as widely used as NAT with IP packets, it can be particularly useful on bridges which are used on redundant networks, as rather than modify the MAC of the network device on the backup server, we can simply modify the frame and have it sent there instead. The nat table has three chains, PREROUTING, POSTROUTING and OUTPUT. While the chains are called PREROUTING and POSTROUTING, they are named so to be consistent with *iptables* and have no actual involvement in any IPv4 routing decision. Instead they are performed prior and following to the bridge decision, and OUTPUT is used for locally generated packets prior to the bridge forwarding it anywhere.

If we wanted to perform SNAT on packets coming in a specific device, we could do:


```
# ebtables -t nat -A PREROUTING -i eth0 -j SNAT --to-
src 00:FF:73:86:2F:F4
```

We can also specify a target following the SNAT, so further modifications can be made to the packet, or we can jump up to the parent chain, using the `--snat-target` switch. We could also jump to the DROP target, although it wouldn't make sense to drop a packet we just performed NAT on. If we neglect a `--snat-target` switch, *ebtables* will assume it is 'ACCEPT' and stop trying to match the frame to any other rule.

As well as SNAT and DNAT, we can also 'redirect' frames, which changes the MAC destination address to that of the interface it came in on. This can only be used in the PREROUTING chain, as it wouldn't be useful anywhere else.

broute

The *broute* table is the most obscure when comparing *ebtables* to *iptables*, although it is particularly useful. The



When using a router with *ebtables*, the 'broute' table is used to perform the decision between routing and bridging.

broute table is required to create what is known as a *brouter*, a device which can both transparently forward packets between Ethernet interfaces, or alternatively route packets based on the layer 3 protocol. This table allows us to decide if the packet should be bridged or routed, rather than actually permit or deny the frame from being moved from one system to another at all. A single chain is included in the *broute* table called BROUTING, which is checked at the very beginning of the bridging process.


If we 'ACCEPT' a frame in the *broute* table then it will be bridged, or if we 'DROP' it, the frame will be routed depending upon the protocol contained within the frame, such as IPv4, IPv6 and so forth:

```
# ebtables -t broute -A BROUTING -p IPv4 -j DROP
```

Conclusion

While *ebtables* has more limited applications than *iptables*, it is a powerful system, especially in its capabilities in creating a brouter from a Linux system. Of course, *ebtables* is a massively complex system which is changing frequently, so anyone looking at rolling out *ebtables* on a production system would be encouraged to check out the latest patches on the SourceForge site and to join the mailing lists to keep up with the latest bug fixes and patches. Happy bridging! ■

There are numerous hooks at various locations in the bridging process, allowing us to drop or modify frames at any stage.



Snort

and the future of IDS

Anyone who knows that IDS stands for Intrusion Detection System is probably aware of *Snort*. This GPLed software has not only become well respected in the area of IDS itself, but has also spawned many other projects. In addition, it has formed the key technology at the heart of Sourcefire, a company set up by *Snort*'s creator, Marty Roesch. Sourcefire builds on *Snort*'s IDS capabilities by creating standalone network sensor devices, backed up by training and support.

We managed to track down the hard-working author of what is believed to be the most used IDS in the world and asked him about his particular take on *Snort*, Open Source and the future of network security.

LXF: You have a lot of history in the computer security arena, and I guess that's why you embarked on the *Snort* project in the first place. Did you think there was something specific missing from the tools that did exist?

MARTY ROESCH: I started writing *Snort* because I got irritated with *tcpdump*. I wrote it for a few reasons, but one of the first things I wanted to do with it was use it as a debugger for another program I was writing. It was a piece of network software and I needed to see the payload of the packets, so I wrote *Snort* to do that.

It's really a classic 'scratch your own itch' Open Source story – I had a need for a better sniffer so I wrote it.

LXF: At what point did you realise there was a lot of interest in *Snort* in its own right?

MR: The first couple of months *Snort* was out, it was just a sniffer, and I started adding features to it to let it start behaving as a network intrusion detection system. Once I did that, people started to get interested in it. At the time there were no network IDS in the Open Source world, and intrusion detection was still somewhat an arcane, black art. Really *Snort* started hitting its stride about a year later, after I cemented in the architecture we are using today which actually makes *Snort* not just into an IDS but a flexible software framework for performing network traffic analysis that we just happen to use primarily as a NIDS.

You can make *Snort* do all sorts of tricks because it's extensible and has a plugin system.

LXF: I guess that's one of the reasons it's been so popular – because you can use it for all sorts of things.

MR: That's it. The research community liked it because

NICK VEITCH talks to *Snort*'s creator about Sourcefire and network security.

Snort provided an API and interface to a decoded packet stream that they could do all sorts of rapid prototyping of ideas – so we saw *Snort* getting picked up by universities and government environments really quickly.


It really gave people a solid platform for network traffic analysis with. Using it as an IDS is the primary reason it got popular, but once people understood all the interesting things you could do with it, it got popular for more than that – it became popular for being a neat piece of software that you can kinda use as a Swiss army knife.

LXF: It must have been gratifying for you that it was so popular. I don't suppose you imagined so many people would end up using it when you started?

MR: No, I was really surprised actually. It was a big shock to me. I put *Snort* out there and people started to use it, and that was as expected. But I didn't expect as many people to use it as are using it today, and also I didn't expect the organisations that are now using it to pick it up – we have large financial institutions, government and military institutions – really people who have some very specific and critical needs using the software. It was surprising to me that it was that useful. Or that I had done that good a job! The good thing is that it's so extensible that maybe they could trim out all the bad stuff I did and fix it to further their own goals!

LXF: I guess having people like that using the software, with very specific needs and aims, must

The [Snort.org](http://www.snort.org) website is hosted by Sourcefire.



have generated a lot of feedback on what you could actually do and how things might develop from there.

MR: It's interesting. We have a lot of different viewpoints on Snort in the world at large – we have the Open Source world view which is that Snort is neat and can do a lot of interesting things, and then we have the commercial competitors view of Snort. They try to box it in and minimise it. One of the things I enjoy is that some of the people who market products against *Snort*, they like to pick different versions when they do the comparison – they pick one version that didn't do something, another version that was missing something else, or a particular release that had a bug in it.

It is interesting how the whole political scene has evolved around it. It's definitely been a real education. Interestingly enough Open Source projects, once they get more than two people working on them – all of a sudden it's an exercise in politics and community management as much as it is in writing good software.

LXF: So the reason competitors do like to pick holes in it is because security is big business? Obviously with Sourcefire you are involved in that side as well. I would guess that you never dreamt about building a company off the back of *Snort* when you started it, but at what point did that become a feasible option?

MR: Yeah. If I'd thought where we wanted to be today and started writing *Snort*, I would be a lot smarter than I really am!

Basically in the fall of 2000 I came out of a startup that didn't do too well, where I was an engineer. I was looking around for what to do next, and because *Snort* was out there and I'd built a reputation around it. I didn't have a problem getting job offers, but I wanted to pick carefully because I hadn't had much fun at the startup.

I kicked around the idea of doing a company for a while, then another security company came along and made me a very nice offer to move to their company and bring *Snort* with me. That was the deciding point – I thought if it was worth that much to them then it must be really worth a lot more so I started Sourcefire. Capitalism at its finest!

LXF: These days, as shown by the current state of the market, it seems to be a very popular idea to build a company around an Open Source project, and there is more of a general roadmap of how that might work. At the time though, did you have difficulty in trying to base your business strategy around a piece of Open Source software?

MR: Well, I had a clear idea of what I wanted to do, but convincing other people that I was rational or sane was a difficult process. The investment community, when I was looking for funding in 2001, the investors were very down on Open Source. They didn't believe you could make money on it, they pointed to the collapse of VA and other hyped Open Source IPOs and said "Open Source is dead, don't you read the papers?" They basically said, "call us back when you make some money". So we did.

The idea with Sourcefire was that we would build proprietary systems around Open Source cores and we'd bring value to the Open Source system that would. It's kind



"I started developing *Snort* as an Open Source system just to see how the development methodology would work, and found it's a really great way to develop software."

MARTY ROESCH

of a basic idea. One of the problems corporations have with Open Source is that they are a little skittish about the notion of being supported by a community, so they like to have commercial support behind the products and we provide that as well as all the other things we provide here – all the technology we develop here to enhance snort to do Intrusion detection and what we call Intrusion Management better. So really it's a logical evolution. A lot of people have flirted with different business models and I think this is the best of both worlds. You have the good parts of Open Source and the good parts of commercial development and as long as you maintain the separation of church and state and you maintain dedication to your Open Source components, I think it's a very good way to do business.

LXF: Obviously, when you started Sourcefire, *Snort* was already quite popular in Open Source. Was that 'kudos' easy to transfer to Sourcefire itself - in the commercial space perhaps many weren't very aware of *Snort*. Was it easy to use the popularity of *Snort* to gain ground early on with Sourcefire?

MR: Yes, actually it was. Our original marketing campaign was my .sig file at the end of my emails. I was very active on mailing lists and that's how we got our first sale to Price Waterhouse, which was big win for us. The community is large and diverse enough, and I personally have a high enough profile out there that it wasn't too bad. Even if the business community didn't know about us, the technical community and the people doing security do know about us, so we were able to get their attention immediately, when we started operating.

That's how I ended up getting funded. These guys told me 'come back when you make some money', so I sold about \$300,000 worth of product from my living room, primarily to Fortune100s. Then all of a sudden we were very popular!

« **LXF:** It seems that the answer to my next question should be obvious then; but the continuing development of *Snort* is important to you and to Sourcefire?

MR: Oh yes. We continue to update and improve it. We had a new release a few weeks ago and there is another one coming up shortly. As we come up with new things for *Snort*, both here and community contributions, we add them into *Snort* and push it back out to the community. We spent a lot of money improving *Snort* and buying testing infrastructure and things like that to make sure it's a good system and hits performance marks. All the improvements and changes that we have made go back out to the Open Source community.

“If you work with the Open Source community, you can be successful in the face of large, entrenched competition.”


LXF: So there is a definite value to Sourcefire in *Snort* being Open Source?

MR: Sure, There's lots in it for us – we have a very large and diverse group of users who provide excellent QA for us. We have a huge installed user base, which also means a huge potential customer base. We have the ability to speak directly to our market via our Open Source presence. I have a big soapbox I can stand up on any time I want to and a lot of people will see it when I do. For us, it has some distinct advantages.

Also, with development, the quality of the feedback we get and the rapidity with which we get it when we turn out releases is much much faster than with proprietary commercial systems, which means we can evolve the centre software much faster than you would typically see in an organisation of our size. We were able to compete toe-to-toe with the biggest companies in this space when we were only 20 people. To use a military term, it's a pretty incredible force multiplier. If you have solid technology and you build smartly around it, and let the Open Source community continue to do the great things that communities can do, and you work with that community, then you can really be very successful in the face of large, entrenched competition.

LXF: With Sourcefire attracting ever-greater numbers of customers and *Snort* now getting a wider user-base and , is it easy to divide your time between the pure Open Source progression of *Snort* and the obligations and demands of your business?

MR: No, not really. It's actually very hard. In large part I've had to take a much more management-oriented position within Sourcefire – we've hired many of the core *Snort* developers here, so I kind of provide direction for them, and a leadership role, but day-to-day I'm not writing tons of code in *Snort*; but I am writing code on a product we are building in-house here. I'm busy with that, and also supporting our sales team and other efforts that involve the success of both *Snort* and Sourcefire.



There are several good books available on using *Snort*.

LXF: Do you have regrets about not being so hands-on with *Snort* as you used to be, or had you reached your limit with it anyway?

MR: I miss being as involved with it as I used to be, but I have enough things to keep me busy right now. I'm working on some pretty cool technology right now. It's another idea that I had that we're turning into a product here. I miss it, but I don't miss the politics and a lot of the other stuff I had to deal with on a day-to-day basis. Yes, and No.

LXF: If we talk about Intrusion Detection in general for a moment – a few years ago there probably weren't too many people investing much time or thought into it, apart from the obvious candidates who had to take it seriously. Nowadays it seems to be much more on the agenda for any business with a network presence. Do you think that enough thought is being given to the whole principle of IDS?

MR: I think a lot of thought is being given to it, the problem that I suspect is really plaguing people is that, you can think about IDS a lot, but unless you can implement things that are reasonable from a usability standpoint – from the notion of just improving the state of security on a network. A lot of people seem to fall down on that – they don't seem to understand it thoroughly. I think a lot of people who design systems aren't really practitioners, and maybe they never have been. Certainly their marketing departments haven't. One of the things that is really neat about *Snort* is that a lot of guys involved with the development have been practitioners in the past or are currently. Myself, I used to work for the government doing network analysis – we have other guys here who have worked for government or universities, where they have handled very large IDS deployments, and I think it's very important that you have some of this background to be able to put together solutions that are really useful to people.

LXF: I suspect that some of the problem is that these days IDS is one of those things you ought to have, and people go out and buy any old solution just to say that they have it, irrespective of whether it actually fits with what they really need to do. In order to get people to understand the problem, I guess there's a lot more education needs to be done for the customers first

MR: Certainly. One of the big problems is that, when the user goes to analyse the data, typically the data they are getting is so technical that the number of people who know how to take what they are looking at and turn it into any reasonable response is very small. It was often said of *Snort* that it's a great system, but when you go to look at the data you have to be Marty to understand it. To an extent that's true. When I wrote *Snort* I was doing forensic analysis for large government, military organisations. I really wrote *Snort* for me, so the expectation was when I wrote it that I would be the one doing analysis. I can look at raw packet dumps and tell you what's going on, but the number of people out there that can is maybe measured in the tens of thousands or maybe hundreds of thousands, but definitely not in the millions. That's a problem – if you want to roll out IDS to

people's homes for example, they're going to have a tough time getting anything useful out of it because the skill in interpreting the data is pretty severe.

LXF: Would you say that was the major challenge to IDS? Solving intrusion problems is about not merely reporting the data or even identifying attacks, but actually trying to make some sort of sense of the data, acquiring some intelligence and working out what is going on on behalf of the user?

MR: Yes definitely. This is where the whole correlation space comes from so people are trying to do correlation so they can take micro-events and turn them into macro events. That's interesting and there is a lot of interesting work being done there. I think that ultimately we get to having better IDS, and this is what my next project is based on, by having more context about the network environment we are operating in.

In other words, if I see an attack on my network, the way I figure out whether that attack was even capable of succeeding in the first place is that I go an either look it up on my network diagram or I look at the box and see what services are running and test if it's been hacked. So, IDS needs to get smarter and have the context about the network – the classic example is you get a 'Code Red' notification out of *Snort*, and you look at the box that got attacked and see that it's a Linux box which can't even be vulnerable to that attack. The IDS doesn't have that context, that one piece of information about the network – this IP address is running Linux which means that it can't be vulnerable to the following 500 things or whatever. I'm developing a system that will generate this contextual information and give it to the IDS. Then the IDS will do its job in term of the targets on the network instead of doing it purely in terms of the traffic.

LXF: that sounds really interesting. When is it going to be ready?

MR: It will be ready this fall. It is a product, not an Open Source system. Pieces of the idea will show up in *Snort* though, because *Snort* has to get smarter to understand the contextual data. So we'll have an automated system for producing all this data and *Snort* will be able to understand it, but if *Snort* users want to manually input the data they can do that as well.

LXF: We recently got a press release about a product, Real-time Network Awareness?

MR: That's what I'm talking about. The idea is if we are truly going to make intrusion detection better, then IDS has to start defining the targets on the network, not just defending the network traffic. So the only way to do that is if I know what's on my network. There are all these ambiguities that IDS has to deal with now – it doesn't know what OS you are running, it doesn't know how many hops it is from a host, it doesn't know the path MTU. If I don't know those three things there are a ton of things that can be used to evade me or make me report false data. If I can generate all that information, even that basic information, I can reduce the evadeability of the IDS and

```

04 21 06 37 73 89 4F 04 10 F8 9F 71 83 23 88 B3 11 78 01 00 00 00
44 46 98 F3 95 06 FD 5F 36 C3 04 16 B3 4F 15 5E 0E 00 00 00 00 00
8C 87 29 25 64 38 5C BC 93 22 64 10 41 01 74 13 00 00 00 00 00
7C 36 87 FC 44 8F 8B 86 CF 84 83 40 85 9C 4B F3 11 00 00 00 00 00
C0 71 0C 5E 73 50 97 26 13 65 C9 E5 91 00 8E 9F 00 00 00 00 00
E7 88 9F 27 5C 40 DD 30 CF F9 9E A6 1E 40 77 12 00 00 00 00 00
6F 23 D0 39 7D A9 0B 40 4E 34 00 D1 CC 12 69 26 00 00 00 00 00
28 8F 7E 96 9E E1 29 52 65 92 F1 E3 31 9C 96 5A 00 00 00 00 00
71 F6 57 4F E6 C3 94 E7 78 D0 78 72 86 8E 2F A8 00 00 00 00 00
08 EF 87 17 1F 82 D0 BC E3 41 39 2F 40 03 12 61 00 00 00 00 00
26 46 84 0E D0 C8 CC 9C FA 4E 9D 9C 32 E8 A0 46 88 00 00 00 00
57 70 2D 14 8E 77 03 80 0C 84 48 2A 9A 19 39 10 00 00 00 00 00
00 07 69 97 C5 67 90 75 E2 95 98 87 DE 20 98 39 00 00 00 00 00
65 CF 02 89 3F A4 56 C8 8F E2 E2 D0 0E F0 86 D1 00 00 00 00
72 BC 82 00 D3 28 44 2F 3E 58 F8 12 33 54 C2 FD 00 00 00 00
69 C1 19 EF DF 02 1B F4 6B 58 3C A6 28 47 7B 78 11 00 00 00 00
41 75 78 90 B9 03 8F 6C 1F 92 9A 5F 67 6F A5 A8 00 00 00 00
C1 80 48 AF E3 39 D0 73 7A A0 70 A4 C2 39 FD 37 00 00 00 00
8C 88 78 46 56 07 F8 0C EF 01 74 C2 76 75 8A 00 00 00 00 00
57 6A F8 01 51 05 D0 B3 31 10 0F 34 78 81 1C 38 00 00 00 00
0E DE 38 A0 8D 6F 4C 1F B8 11 40 12 2E 63 1C 00 00 00 00 00
5C 93 97 99 F1 70 12 75 51 66 8B AD 11 45 39 A8 00 00 00 00
88 8C E8 12 92 C2 4B 5E 60 77 3E CC 41 72 4C 67 00 00 00 00
40 9E 38 77 AF 38 62 00 00 00 00 00 00 00 00 00 00 00 00 00

=====
06/16 14:12:18.862619 192.168.130.91:33839 -> 212.23.32.7:443
TCP TTL:64 TOS:0x0 ID:3206 Iplen:20 DgLen:52 DF
***** Seq: 0x490E8566 Ack: 0x4968055C Win: 0x87C0 Iplen: 32
TCP Options (3) => NOP NOP TS: 2010076 322689382

=====
06/16 14:12:18.872621 192.168.130.91:33843 -> 212.23.32.7:443
TCP TTL:64 TOS:0x0 ID:5796 Iplen:20 DgLen:52 DF
***** Seq: 0x49C2C1EF Ack: 0x48DC7A6D Win: 0x1050 Iplen: 32
TCP Options (3) => NOP NOP TS: 2010077 322689384

```

If you know your network packets, you might be able to work out what's actually going on here...

reduce false positive. If I can take the next step and say I know these vulnerabilities are available at specific IP addresses, if I see those vulnerabilities exploited then I'm very interested – if I see anything else I'm not so interested – so I can prioritise appropriately.

LXF: Do you think that false positives are a real problem currently? IDS generally report so much that users take them with a pinch of salt...

MR: Well certainly. That's one of the primary problems of IDS. Without the contextual information we give the users a lot of work to do – we force them to turn into detectives and try to figure out if what the IDS is telling them is true, and if it is true, then if it's relevant or not. We are seeking to eliminate that and say 'these are the things that are truly interesting – here are attacks that try to exploit vulnerabilities that you are actually vulnerable to'. Plus we

"To make intrusion detection better, IDS has to start defining the targets, not just defend the network traffic."

can do things like detect change so I can say "your webserver is running *Apache*, I saw the *apache* chunked encoding buffer overflow and then 30 minutes later your webserver started offering IRC as a service as well, is that interesting to you?"

LXF: I think that covers everything we wanted to ask – is there any other advice for our readers?

MR: I think from the perspective of an Open Source developer, I started developing *Snort* as an Open Source system just to see how the development methodology would work, and I found it's a really great way to develop software. To all of your readers I guess I would say if they are thinking about starting an Open Source project, I would highly recommend it – it's very interesting and very fun. ■