

LINUX

www.linuxformat.co.uk

FORMAT



Customise it!



PERSONALISING THE DESKTOP

- GNOME 2.2 ● KDE 3.x
- Kernel 2.x ● OpenOffice.org
- Mozilla and more **PAGE 48**

EXCLUSIVE!

LINUX OFF THE SHELF!

UK users get pre-installed desktop machine. We review it first **p20**



OPEN PUBLISHING

Discover the wonderful world of Wiki, where your web pages write themselves **p60**

KEEP IN SYNC: PIMS & PDAS TALKING TO EACH OTHER **p82**

NETWORKING WITHOUT ETHO

Serial alternatives to Ethernet – cheap and effective **p56**

DVD issue also available

LXF39 APRIL 2003

Printed in the UK

£5.99



04

9 771470 423019

£5.99 (outside UK and Republic of Ireland)

THE UK'S BEST-SELLING LINUX MAGAZINE

Have it your way


One of the many joys of using a Free operating system is that a wealth of choice is opened up to you. Not only do you have the choice of many different tools to do the same job, but these themselves can usually be configured to work in many different ways. Whether you just want something to look different or fundamentally change the way the software works, it's usually possible with Linux. A complete guide to customising everything would probably take up many magazines-worth of space, so we have concentrated on the areas that may make the most difference to you in our main feature this month. Please feel free to send us your own tips on customising other applications, which I'm sure we'll be able to feature somewhere in future issues.

If you need any convincing of the wealth and flexibility of Linux software, look no further than the roundup of filemanagers we have this month. Often your most-used file manager will be the one you started using first, but you may find that

others are more suitable for the way you like to work. At the very least, you can check out the latest versions of all your favourites (and they're on the coverdiscs too).

If you are more a 'straight out of the box' person, then it will definitely be worth your while checking out the Evesham PC review on page 20. This pre-installed Lindows box could signal the first of a new breed of casual Linux user, and at the very least give us some idea how popular Linux is for the desktop. My guess is that if it's successful we'll see other offerings from the usual systems builders.

Don't forget to check out *Linux Pro* either. As well as some excellent practical pieces on firewalling and *Apache*, you'll also find a story on the record-breaking new SGI Altix 3000 series. While your IT budget may fall a little short of buying a 64-processor supercluster, the story of how they have added and extended the linux Kernel to provide real scaling performance should make interesting reading. Enjoy both mags!



Nick Veitch EDITOR

Linux the way you want it to be – tips and tricks for customisation **p48**

Could this be the start of something? Evesham's Linux box on test **p20**

Files – find them, open them, copy them, use them and manage them **p34**



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 SOME OF LINUX FORMAT'S TEAM OF WRITERS...



Andrew Channelle
With more time on his hands, Andy has been trying to get to grip with his finances – using Linux of course!



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



Richard Drummond
Now based in Flatville, Indianapolis, USA, Rich is *still* trying to find some decent beer.



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



And introducing... Paul Hudson
Formerly our PHP guru, Paul has taken the King's shilling and joined us as our new Reviews Editor.

Neil Bothwick
When not playing with his Zaurus, Nelz is the guy tracking down Gigabytes of the best software for your discs.

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

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

Mike Saunders
Currently in the hot seat for Hot Picks, Mike is keen to get your feedback on the best new software around.

Patrick O'Brien
Published by O'Reilly and IBM; and author of *PyCrust* has an enthusiasm for Python that is highly infectious!

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

LXF39 April 2003

Contents

LINUX
 FORMAT

 Welcome to another crammed-full issue of *Linux Format*, your guide to all things Linux!

COVER HIGHLIGHTS

Customise it!

48

 From minor browser tweaks to hacking the kernel, we show how you can make your OS and apps do *exactly* what you want

20 Linux off the shelf

Evesham leads the field, other retailers will follow...

56 Serial networking

Can SLIP and PLIP can do everything Ethernet does?

60 Open publishing with Wiki

Let others edit your work on the Web in real-time

82 Keep in Sync

Linux on your PIM and PDA is easier than you think!



“First dates are awkward and uncomfortable. An engaging conversation is critical to establishing common ground.”


78 Patrick K O'Brien loves Python. Really...

“Serious vulnerabilities in the Unreal 3D engine regardless of the OS give crackers a chance to launch attacks...”

11 NEWS – PivX asks Epic for patch

LINUXPRO



Record breaker SGI's Altix 3000 **Security** Managing threats from inside your company **Storage** External Maxtor 5000XT **Apache** Rewrite tricks **Tomcat** Access, logging and architecture **iptables** DNAT, SNAT





20 The price is right, but will the likes of Dixons and PC World stock them in time for Christmas?

<p>6 News Analysis of the latest Linux stories</p> <p>12 Mailserver Your points of view in print!</p> <p>20 Evesham E-scape review £250 Linux box with no self-build!</p> <p>22 Partition Expert review Every Linux user needs partitioning</p> <p>24 GnuCash review Appease creditors and the taxman</p> <p>28 SME Server review Administrate from your browser!</p> <p>31 Book reviews Perl, Mason, HTTP, firewalls & more</p> <p>34 File Managers roundup <i>Midnight Commander</i> et al</p> <p>42 Hot Picks The latest and best free software</p> <p>56 Serial networking SLIP and PLIP explained</p>	<p>60 What on Earth is Wiki? Hawaiian for "Quick", apparently</p> <p>66 Mozilla tutorial Tweaking the toolbars</p> <p>70 GIMP tutorial Textures bring your images to life</p> <p>74 PHP tutorial PHP as a CLI scripiter</p> <p>78 Python tutorial Installing and using <i>PyCrust</i></p> <p>82 PDA tutorial WinCE & PocketPC show the bloat!</p> <p>88 Answers Your question could win a PDA!</p> <p>98 Disc pages All the important information</p> <p>110 User Groups Asia and India LUG list, and UK too</p> <p>114 Next month SCO vs Linux, the <i>LXF</i> Awards...</p>
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------



COVERDISCS

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



CDS A AND B

GNOME 2.2 Full-featured desktop; **KDE 3.1 for Red Hat** Better late than never; **TightVNC** Platform-independent remote desktop operation; **BlueLava** control X10 devices from your mobile phone



CD C

SME Server We review it and include it on the disc – make sure you read the vital info on page 98!



DVD

GKRELLM watch everything that's really happening on your system
CENSORNET Internet access controller



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

SAVE MONEY!

SUBSCRIBE TO LINUX FORMAT

See page 96 or phone 0870 4448645



Newsdesk

Infosecurity Expo dates; MSN sites in 'Bork'; Mandrake 9.1 release cycle; Swiss tax returns submitted with Linux; Motorola pick Linux; Adobe to use Qt; Ogg-compatible Neuros Player; Unreal engine vulnerabilities

LINDOWS MOBILE PC

Desktop summit

Windows.com CEO Michael Robertson took a leaf from the "Steve Jobs Big Book of Computer Marketing" and used his keynote address at the first annual Desktop Linux Summit to announce a brace of new hardware products. The device that stole the show was the Lindows Mobile PC, a laptop that weighs in at under three pounds and has a price tag of US\$799. Robertson intends the device to bridge the gap between the traditional laptop computer and high-end PDA market.

"Until now, laptop computers have been either too expensive or too heavy to be useful for most consumers," he said. "PDAs are also a problem because they don't offer many features. The Lindows Mobile PC is a powerful machine that can go anywhere you want to take it"

So far, Robertson has shown an unerring ability to hit the psychological price points in his hardware dealings with the likes of Wal-Mart in the US and Evesham in the UK; bringing products to market which undercut rivals by serious margins and pushing technology into the realms of an impulse purchase for many buyers.

The Lindows Mobile PC also demonstrates the increasingly close relationship between motherboard/processor maker Via and Lindows. The machine is based on a 933MHz Via C3 processor and features a 12.5 TFT monitor, 256MB RAM, 20GB hard disk and ports for USB 2.0, Firewire and Ethernet. There is no



The Lindows Mobile PC is intended to sway those looking for a low-end laptop or high-end PDA.

modem but a PCMCIA slot provides an opportunity for expansion if users need to access the Internet over a phone line.

The other significant feature on the PC is removable storage. There isn't any. Users will be forced into acquiring a USB/Firewire CD-Drive, Zip or floppy. This makes direct comparisons with the twice-the-price iBook a little disingenuous, but it may also be regarded as a plus with corporate buyers used to their workforce misusing company resources and installing productivity-unfriendly apps like *Doom* on their laptops.

As well as using one of the Mobile PCs for his presentation, Robertson also offered a first look at Lindows 4 which improves support for plug 'n' play, especially for devices like Key-Chain disks, MP3 players and digital cameras. This wasn't a big deal if you're used to these features in Windows XP, "but its a big deal if you're from Linuxland," he said.

Robertson also announced the Lindows Family PC which is similar to the MediaCentre announced in *LXF* news last month but omits the TV tuner and features a subscription-based service to protect users from

the more unsavoury elements on the Internet. The PCs will cost \$349 with an annual US\$50 subscription to the filtering service.

Other summit news

After a controversial split which saw many of the original exhibitors pull out due to what was seen as an attempt by Lindows to hijack the event, the Summit finally went ahead with an eclectic range of speakers including Shawn Gordon from TheKompany.com and Gavriel State from TransGaming Technologies. After the departure of Hewlett Packard and Xandros last

GEOHERMAL ENERGY

Penguins headed north!

Linux NetWorx are taking their clustering technology from Utah to Iceland to scout for new locations suitable for geothermal power plants. The company will build a new system for running three-dimensional simulations

month, their places were taken by Epson, who promised Linux drivers for their latest range of multi-function devices, and ATI.

The conference was divided up into logical areas, covering general use, home, business and education. Via Technologies demonstrated their reference design Media Centre PC which should be available mid-March in the US\$700 range. The PC runs Lindows (there's that Robertson connection again), is the size of a hardback book and features a sub-woofer and satellite speakers, though strangely the standard version has no television or personal video capabilities, which appears to make little sense.

SuSE were also in attendance showing off their latest desktop distribution that comes complete with CodeWeavers' *Crossover Office* and the ability to run *MS Office*, as well as a complete copy of *StarOffice* from Sun.

On the conference side, there were talks on the current state of Linux in the business world, a discussion on the need – or not – for games support to make the platform more visible, and Robin Rowe from *MovieEditor.com* gave a detailed summary of the one area where Linux is the number one desktop operating system: film and special effects.

The final presentation looked at some real-world examples of schools that have migrated away from proprietary systems, and discussed the hurdles they faced and how these problems were tackled.

In all, the event was judged a success and positive coverage in much of the technology press has done some useful work in raising the profile of the soon-to-be-number-two desktop OS.

of geothermal reservoirs and Icelandic weather conditions. While the 48 processor cluster is a relatively small project for a company which in the past has built monster systems with over 2000 CPUs, it will have a real impact on Iceland's future energy use; almost 50 per cent of power consumed in Iceland comes from geothermal sources.

Olafur Rognvaldsson, CEO of Iceland's Institute for Meteorological Research said there was a real need for a powerful supercomputing solution that could produce high-quality simulations quickly. "Compared to more traditional supercomputers, price-to-performance ratio of Linux Networx clusters is unbeatable and the performance increase allows us to complete simulations and make decisions that were previously impossible," he said.

Linux NetWorx has also recently signed a deal with DigiGenomics (Taiwan) which should boost the Linux profile within the genetics sector.

INTENSIVE COMPUTING

Itanium 2 servers from SuperMicro



Server builder SuperMicro is aiming to become one of a select band to market Intel Itanium2-based servers with the launch later this year of a range of high-performance systems. New server boards will feature a pair of Itanium processors with 6.4GB/s data bus potential and support for 16GB of DDR memory.

Mike Graf, Product Manager for the new line, said the new servers would help to lower the cost of high-end intensive computing and expose new markets that have previously been out-of-reach for more traditional proprietary systems.

NEWSBYTES

■ Canadian software behemoth **Corel** have finally cut their last ties with Linux; closing down the linux.corel.com subdomain. The removal of the site draws a line under Corel's involvement in Linux, though their most ambitious release, *CorelLinux*, lives on in spirit at least, under the guise of Xandros.

■ There could be a significant drop in the cost of writeable DVD drives this year, as **LG** eye the market dominated by Sony. The first device off the LG production line is said to be compatible with all current 'standards' which should make buying a less worrying prospect.



■ Will your next **Red Hat** distro be Red Hat Advanced Workstation Desktop? In an attempt to boost sales of the profit making product line, Red Hat are to start aggressively marketing the *Advance* brand. The company claim a new lower price for Advanced Server will open new markets for the product and allow the company to compete more fully with Microsoft as well as traditional Unix sellers – where it has seen the most growth so far.

■ **Intel's** 64-bit flagship, Itanium2, has been getting a rough ride lately, with Linus Torvalds claiming the processor has 'lost all the good bits' of the x86 architecture while keeping the bad. IBM have also scaled back Linux support, pulling the core Linux/Itanium development team in to work on its own Power processor. The move, a spokesman said, was a response to the glacial rate of market adoption for the chip.

■ **Connectix**, maker of virtualisation software for Windows and Macintosh, has been purchased by Microsoft. VirtualPC, which allows Mac or Windows users to run other x86 operating systems, will apparently be integrated into the Windows core, but the real reason for the purchase is said to be the imminent release of a virtual server product, enabling 'hundreds' of operating systems to run concurrently on one machine. The news is particularly significant as just a week prior to the take over, Connectix announced a new edition of *VirtualPC* featuring Red Hat Linux, however, a Microsoft source said that there was no danger of future versions removing support for Linux.

Jono Bacon

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



COMMENT

Opening the box

“ You may have heard of the Xbox Linux project; aimed to get Linux booting on an unmodified Xbox. The project has a feverish development rate, due in part to the reward offered. At first I thought the project was another not particularly useful techy challenge, but there really is more to this than meets the eye.

It is quite clear that if Linux was to boot off an unmodified Xbox, Microsoft would be miffed; not because a stable, powerful competing OS is running on it, but that someone could buy an Xbox and theoretically have a cheap and powerful PC that runs Linux. If we take this a step further, Linux software could be developed for the Xbox hardware to make the unit into a Linux-powered games console, enabling free software developers to write games for a standard hardware format with all the 3D power that is needed. A company would only need to buy a busload of Xboxes, rebrand them and it could completely change the system into something else.

There are a number of legal ramifications of this move, and whether a company can license the acceptable use of its products; eg if I buy a can of Coke and use it for something else other than drinking, should the Coca Cola Company be able to take legal action against me?

My hunch is that “no” is the right answer, but then, I am a Linux person like you lot, and we like freedom generally. Microsoft have already asked the Xbox project to stop the distributed processing effort to crack the Xbox code that prevents non-native software running, but do they have a right to do this? ”

MDK 9.1 BEGINS RELEASE CYCLE

Mandrake not dead duck

Despite highly publicised continuing financial woes and bankruptcy proceedings, MandrakeSoft has begun the release cycle for the next iteration of Mandrake Linux.

Mandrake 9.1 features a number of interesting new elements for the first time, including the ability to resize NTFS partitions and ACPI support for laptops, though the latter is not enabled by default. The package also features new desktops in the form of GNOME 2.2 (RC) and KDE 3.1. On the visual front, there is a new theme called Galaxy which is currently GNOME only. The latest release candidate – number one as we go to press – should be available at www.mandrakelinux.com/en/ftp.php3. The release also boasts better support for hot-plugging USB devices and early reports suggests it is in a pretty workable state.


As has become traditional, the announcement is accompanied by a plea for more support from users and

downloaders via membership of the Mandrake Club. “Developing a Linux distribution is very costly, so it’s up to the community of users to ensure its health,” the announcement said. The Club also offers an extra tier of support for subscribers and access to additional services and commercial applications not otherwise available.

End of life policy change

Following the lead from Red Hat, Mandrake have introduced a raft of changes in its errata support for Linux distributions in a bid to cut support costs and/or improve upgrade revenues. Mandrake’s formula for end-of-life-ing a product is a little more complicated than Red Hat’s ‘12 months and it’s gone’ rule: desktop components will be supported for 12 months, while core components (base) including Apache and the kernel will be supported for 18 months.

For Mandrake 9.0 users, the first deadline is September 30 2003, while ‘base’ support will expire on



The Galaxy theme is currently GNOME only, but there is a KDE version drifting around the net.

March 31 2004 – important dates for all Mandrake users’ diaries.

While many users have expressed their understanding of why both companies needed to make the changes, it is feared it could pose a significant security threat in the short-term, as machines go unpatched. Jon Lasser, from Security Focus Online,

said that just because software is no longer supported, doesn’t mean it isn’t being widely used. “With network operating systems the consequences could be drastic. Those systems will be sitting ducks for vulnerability scanners, and the size of distributed denial-of-service networks may grow exponentially as a result”


Linux Web Watch/



Linux Half-life? Yes please!



CUBE will help FPS Linux devs.



LG Publishing has Linux vacancies...



Open Source Web Design.

Just a couple of things

Games at the forefront of Linux evangelism in gathering more Windows migrants

At the Desktop Linux Summit recently, a lot of time was spent discussing the role games can play in making Linux acceptable for home users. In this spirit, here’s a trio of sites useful for industry watchers, players and developers.

www.linuxgames.com is a useful site with plenty of info on both upcoming games and the technology that drives them. It also hosts a number of game

sites including *Linux Half-life* and the development effort for The Guild.

If world-building is your thing, you really ought to check out Cube (<http://wouter.fov120.com/cube/index.php4>), a first person shooter engine with both a high-spec and modest system requirements. The demo game looks outstanding and the engine is license free; let a million FPSs bloom!

And if you actually have a useful idea for a game, or even elements of a game, then Linux Game Publishing (www.linuxgamepublishing.com/) would like to hear from you. This fledgling publisher is seeking eight developers to create the killer Linux game. Good luck to them.

Web design is supposed to be easy, but actually putting together a site that

looks professional can be quite a task. One good way to save some time and effort is to visit Open Source Web Design (www.oswd.org/), a site which offers, and accepts, designs for free redistribution on the Net. Of course, if you make significant changes to a template, you should repost it to let others benefit from your labours – that’s what Open Source is all about.

NEWSBYTES

■ **Red Hat and Intel** have worked their way around some licensing issues which threatened to derail any possibility of collaboration on Open Source projects. The situation was resolved with Intel converting a proprietary license to a more BSD-like entity for its power management interpreter. Red Hat are now free to submit their range of patches which can be incorporated into the commercial offering.

■ A report in **India's Business Standard** claims that Linux is on course to achieve a 10 per cent market share of the country's desktop computing by the end of the year. Driven by cost-cutting, the government has made Linux the preferred choice in areas like educational IT purchasing.



■ **MandrakeSoft** have kicked off a new publishing venture with the publication of *The Definitive Guide to Using Mandrake Linux*. The 39 Euro book covers the basics such as installation and configuration as well as in-depth treatment of more advanced subjects such as file system structure and server maintenance.

■ **Taxpayers** in and around Geneva, Switzerland have been sent a CD-ROM containing *Mozilla*, *OpenOffice* and a tax application called *GeTax2002*. The CD is suitable for Windows, Linux and Mac and is designed to make filing online tax returns a little easier. Is Gordon Brown listening?



■ Despite what is euphemistically termed 'negative growth' throughout 2002, the **server sector** is looking up this year thanks to an increase in sales of Linux servers. Revenue from Linux servers rose by 63 per cent to \$2 billion. IBM lead the pack with sales of \$759 million. See *Linux Pro* for all the information on a record-breaking Linux server that's going to clean up in 2003.



While being thought by some as being very silly, the 'Borked' rendering of MSN.com does show off Opera's stylesheet skills rather well.

NETSCAPE 7.02 AND MORE New browsers

The major browser developers have all been busy recently, putting together a number of Linux releases and also, in the case of *Opera*, firing a satirical shot across the bows for Microsoft and its allegedly underhand marketing practices.

Netscape 7.02 has reinstated pop controls removed from previous versions. This appears to be a part of a wider AOL/Time Warner focus away from popup ads which are annoy by most users. You can download it from www.netscape.com.

OperaSoft has released a preview of *Opera 7*, following on just a few weeks after the Windows version. On the surface, version 7 ships with a very blue 3D-style interface and one-click install of new themes - including a quite convincing Crystal theme. Under the hood there is more robust DOM2 and CSS2 support as well as better ECMAScript and HTML 4.01 support. The new browser can also use multiple user style sheets, including a simulated text browser and, curiously, a setting which makes everything look like it was rendered on a C64. One very welcome

addition is a new mail client, M2, which features clever-ish spam filters.

For Windows, OperaSoft responded to a dispute with MSN.com (which allegedly saw the site serving up slow, broken pages to Opera users) by releasing a version that, when users navigate to www.msn.com, translates the page into Bork, the language spoken by Swedish Chef in *The Muppet Show*. So you could end up reading "Interteening, idoocteeeng und in tuooch! Feend duzens ooff games, ecteeefities, schuul tuuls, und noos fur keeds!"

Meanwhile, back in the world of Free Software, former Galeon mainstay Marco Presenti Gritti has gone it alone to create a new browser project called *Epiphany* (<http://epiphany.mozdev.org/>). Like *Galeon* before it, *Epiphany* has simplicity, elegance, GNOME integration and standards compliance in mind. The manifesto states: "Epiphany addresses simplicity with a small browser designed for the web - not mail, newsgroups, file management, instant messaging or coffee making. The UNIX philosophy is to design small tools that do one thing, and do it well!"



Another spare, svelte browser emerges from the Mozilla tree - *Epiphany*.

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

The last mile

“In the comms business, the ‘last mile’ is a metaphor for the difficulty in wrapping up a roll-out of the network to the far edges of the user base.

Linux is in a last mile situation of sorts. It is ‘almost there’ in terms of device driver support, usability, font rendering, packaging etc.

The ugly truth is that in a development model driven by volunteers, that last mile is not sexy, not interesting and not likely to get the attention it deserves. Linux may never make that last mile unless this is addressed. Admittedly, manufacturers are beginning to offer Linux support, but often on obscure web pages with versions far behind Windows and Mac iterations.

I was at an electronics retailer. There were many low-end scanners for sale, but none were supported in Linux. A number of older scanners are supported, but none of these old models are offered for sale. I suppose I could shop for an older model on eBay, but why should I be forced to buy used or older equipment?

Why can't manufacturers of low-end hardware offer current, SANE-compatible Linux drivers from the beginning? Don't they realize that they could easily sell this hardware to Linux users? A lack of drivers for low-end hardware hinders the adoption of Linux by Windows users, who purchase this stuff because it's so affordable.

Even installing closed drivers doesn't have to be hard. I recently installed pre-compiled (for a kernel version and distro I no longer use) HSF modem drivers and they automatically stepped me through a recompilation and installation. With some script tweaking they'd even be suitable for my grandpa's use. Manufacturers, what's the holdup?

MOBILE LINUX

Motorola migrating most phones to Linux

The architecture race in the mobile phone sector looks increasingly like a four-player affair with Motorola finally and decisively choosing Linux and Java to

power its high-end personal communicators over similar offerings from Microsoft, Palm and Symbian. The first Linux-based phone – the stunning looking A760 – should be available later this year.

As an advertisement for the pairing of Linux and Java, the A760 is hard to beat. As well as being a phone, the tiny flip open device also houses a fully-featured PDA, MP3 player, digital camera, video player and a high quality colour screen. It is also Bluetooth-enabled for wireless communication.

Motorola's Rob Shaddock said that though cost was of course a small factor in their choice of OS, the combination of Linux and Java had much more to offer.

"By supporting the Open Source Linux OS and Java technology, Motorola is creating the most open and flexible environment possible to help drive the development of compelling applications for rich, customised mobile experiences," he said. Committed early adopters of new technology will be important allies in advocating wider appreciation of Linux.



Motorola have committed to basing their next generation of high-end phones on Linux and Java.



Embedded Linux News

- **Sharp** has inked a deal with **Socket** to bring a range of CompactFlash accessories, including the company's Bluetooth Connection Kit, WLAN Card, Low Power Ethernet Card, 10/100 Ethernet Card and Single Port Serial Card to the Zaurus (see page 92 to win one!). Socket will also be supplying native drivers for the products. A spokesman was equivocal that CF-Card based peripherals would make the device more attractive in areas such as sales, service and healthcare where ready access to remote data is vital.

- **Cypress Semiconductor** has created a new range of embedded USB host controllers boasting a 16bit RISC processor, large memory and low power needs. The development kit includes a full complement of Linux drivers, framework firmware and functional examples.



- **Intrinsyc Software** announced a new iteration of its CerfCube reference design built on IBM's PowerPC 405EP. Designed for 'edge of network' apps like routers, cable modems and wireless hubs, the device should be out by the end of April.

LINUX PHOTOSHOP?

Adobe choose Qt


Graphics giant Adobe has opted to utilise the Qt cross-platform development environment for its latest consumer-focused editing and image management application, prompting speculation that future versions may make their way to OS-X and Linux.

The £39.00 application is the first Adobe product to bring the *Photoshop* brand within the range of general buyers and has already won the Best Of Show award at CES Las Vegas where it was unveiled.

"Trolltech has provided us with an intuitive, powerful tool. Qt simplified our

task of developing *Photoshop Album* by providing high-level tools that we could customise to meet our

needs," said Adobe's Mike DePaoli. Eirik Eng said Adobe's decision to adopt Qt to develop *Photoshop Album* was validation that Qt is a world-class development framework. "Qt is designed to take care of the horrible low-level details so the developer can focus on adding functionality to the application," he said.



HARDWARE AUDIO

Who let the Ogg's out?

The Xiph Foundation has

confirmed a deal which will see Ogg Vorbis audio compression technology incorporated into a hardware player. The Neuros Player is slated for a March release by maker Digital Innovations and is said to be as happy connected to a Linux box as any other PC. Xiph's Emmett Plant said the Neuros Player could mark a change in the way Linux is perceived by the industry, bringing "Linux interoperability for a portable player that's supported by the manufacturer, not an after-market hack supported by some guy in Johannesburg with a dial-up connection and a copy of Emacs."

The Neuros itself also looks pretty special with a unique feature its makers call MyFi. This, essentially is a low powered FM transmitter, meaning songs from the player could be streamed to a car radio or domestic stereo wirelessly. The device will ship with either 128MB of memory or a 20GB hard drive.



Neuros Player brings Linux interoperability and open source music encoding to a pocket near you.

infosecurity
EUROPE
29 April - 1 May 2003
OLYMPIA, LONDON
www.infosec.co.uk

Interested in security? Then put the above date and time in your diary.

Infosecurity is Europe's number one IT Security Exhibition. It brings together professionals interested in IT Security from around the globe with suppliers of security hardware,

software and consultancy services. It is dedicated to providing education that heightens awareness of commercial importance of secure and reliable access to corporate information. Watch out for the security survey coming soon on www.linuxformat.co.uk

PIVX REVEALS OLD BUGS

Unreal players vulnerable

A California security company

claims to have discovered a number of serious vulnerabilities in games built on the Unreal 3D engine. They say the hole can affect players *regardless of the OS they're playing on*, giving crackers the chance to launch a Distributed Denial of Service attack, crash game servers or run arbitrary code on a user's machine.

Luigi Auriemma at PivX said the holes had been there for five years. "They could be used by malicious attackers in worms or large-scale attacks that rival those of Nimda and Slammer," he said. PivX says it has known about the bugs for some time, but were waiting for a response from Epic, creators of the

Unreal game engine, before going public. Epic claim a server patch dealing with the issue is imminent. The following games may be affected:
Star Trek: Klingon Honor Guard
Unreal, Unreal Tournament, UT2003
The Wheel of Time
Deus Ex
Mobile Forces
Rune
Hired Guns
Navy Seals
TNN Outdoor Pro Hunter
Werewolf
X-Com: Alliance
Adventure Pinball
America's Army

KDESTUDIO GOLD

Corrections!

Following the C++ IDE roundup in

LXF35 we were contacted by representatives of theKompany, whose product *KDEStudio Gold* was covered in the roundup. They highlighted some inaccuracies and ambiguities in the review which we'd like to clear up.

- The review may have inferred that the source was not available for *KDE Studio Gold* when mentioning the trial version. This is not the case – the full commercial copy of *KDE Studio Gold* comes with the source code.

- The documentation for *KDE Studio Gold* wasn't rated. The full product does come with comprehensive docs.

- The DocBrowser context sensitive

library documentation wasn't covered in the review, which was based on the trial version. This is available with the Pro version of *KDE Studio Gold*

- The review erroneously stated that *Quanta+* was also a product owned by theKompany. This is an error – *Quanta Gold* is the software that the author intended to refer to.

- The review didn't mention that *KDE Studio Gold* also supports Java projects. As the roundup was of C++ IDEs, this wasn't considered as part of the roundup, but is obviously an extra feature potential users might like.

In light of some of these facts, we would also like to change the final score we gave *KDE Studio Gold* to **8/10** and hope that everyone now has a less ambiguous view of the product.

Paul Hudson

Having been our freelance PHP guru for a long time, Paul has taken the plunge of becoming Reviews Editor at *Linux Format*.



COMMENT

Not so hard?

“ In my inaugural piece, I've two matters to address.

Firstly, I have recently been experimenting with the Red Hat Linux distro. To be blunt, I'm a die-hard Debian-ite, so this (albeit temporary!) move may come as a bit of a shock to those who know my usual choice of distro. Quite frankly, it came as a bit of a shock to me, too – it seems that, as a Red Hat user, I am pretty much stuck with the same version of each piece of software that came with my release (7.3).

How on Earth do Red Hat users live with this? A user running Debian 2.2 can simply type 'apt-get dist-upgrade' and APT will download updates and upgrade the system so that the system becomes Debian 3.0. It's hardly rocket science, but it's something I dearly miss on a Red Hat box. Red Hat users: rise up and demand more from your distro!

Secondly, my colleague Hoyt Duff mentioned last issue that he believes that too many new versions of software are released without old versions being fixed first. Perhaps this is another case of users not realising that they have a *choice* with Linux? Debian 2.2 (Potato) released on August 14th 2000 after a long feature freeze and bug fix period. Potato r7, the final version of Potato, was released on July 13th 2002, almost two years later, and had precious few features added in that time. That's right: nearly every upgrade released in two years was a bug fix of some sort.

In the same way the joke goes: Man: "My arm hurts when I do this!" Doctor: "Well don't do that, then!" The answer to Hoyt's problem is simplicity itself: use software that concentrates on bug fixes; there's a huge amount of it about!

Mailserv

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

Floppy quirk

Not really a question for your *Answers* section, more a caveat for Mandrake users!

I was reminded forcibly when reinstalling Mandrake 8 on my second machine, of the following glitch which might well trip up newbies. If a USB mass storage device is plugged in and detected during install (in my case a 'Disk on key'). Mandrake assigns this to **floppy** in the /mnt directory, while **fstab** correctly nominates it as /dev/sda. /dev/fd0 appears correctly but is assigned to "floppy2" in the

/mnt directory. This can cause confusion... It certainly did for me when I stuck a 3.5" floppy in and mounted what I thought was the floppy drive (*Disk on key* not plugged in) Man came up helpfully saying access denied, even to root... lots of CH,ing but still no access! Finally checked **fstab** and behold the answer.

This applies to Man 8.0 and 9 and I imagine the in-between distributions as well.

James Baldwin, *via email*

Thanks for the tip, we hope they fix it for 9.1!

Snobs?

As a Software Engineer and Linux advocate surrounded by Windows developers and users alike, I can't help but wonder if I've at times been a little guilty of the 'OS snobbery' described by Simon Moore in the *LXF37* letters page.

However, I learned fairly quickly that any kind of extremism in relation to operating systems can spark amazing amounts of adoration and loyalty towards Microsoft and its prized operating system! Despite trying a much gentler approach to advocating the

joys of Linux, I seem unable to convince any of my fellow developers to switch. I've not used Windows on my box at work for more than a year now, but no one else in sight has even considered installing a Unix OS.

It strikes me that some people simply *don't want* to switch operating systems. They're happy with what they've got, and are willing to argue until they're blue in the face that it doesn't crash. They're also often unwilling to learn new skills, especially in the field of development.

★ Letter of the month

This month's winner receives a copy of **Tuning and Customizing: a Linux System book**



Mandrake tips

Let me begin by saying that I am happy with Win98SE. My problem is that if I stick to Microsoft I will eventually be forced to upgrade and I don't like XP, I like to be in control of my computer. Neither do I like to be treated as a suspected thief, when I have legally bought and paid for my operating system. I concluded that it is time to prepare for a future switch to Linux, and that learning would be fun.

Frustrated with the problems I had with my Mandrake 9.0 Internet connection, I deleted it and installed Debian (having downloaded 7 image discs!). I made a few errors but eventually installed with the latest Kernel, enabled the correct USB support for my motherboard, and my printer and modem (CDCether). However, at the end of the day, no USB devices would work. I visited the kde website where it boasted that there are system tools built into the new KDE3.1, but having downloaded the files (in Windows) I

found that Debian had set itself up with a number of unsatisfied dependencies, and **dselect** became a nightmare as a result. Basically my attempt to install KDE3.1 merely resulted in KDE2 and other progs being uninstalled, leaving me with GNOME 1.4 (Yeuk!)

After a few days of re-installing and trying various things, I decided to give up and re-install Mandrake 9.0. During my Debian experience, I had figured a few things out, so this time I re-installed Mandrake without selecting the Server software, so no NFS, no SQUID, the list of programs loading up and closing down is much shorter, and, to date, I have not lost my internet connection, even when the computer was left on overnight!

I have also had a few unpleasant experiences trying to reload Mandrake 9.0, and here are a few hints to help others:

1 Do not install your printer until you have got Mandrake up and running, and **cups/foomatic** updated. You'll only have to start over if it isn't.

2 DO NOT install server software unless you *definitely* need it. You can always add it later!

3 DO go into the Full Install list in expert mode, select CUPS but deselect *OpenOffice.org*. The Help file does not work. Better to download *OOo* from its site and install it.

4 If you have a PC Chips motherboard with SiS 630 graphics, and a 1280x1024 LCD display, selecting the **safe** version of Xfree4, results in you having the option of different colour depths, but the 24-colour version does not work, and causes the display to change to 1024x768. The **unsafe** version with 3D drivers fixes the colour depth at 16 but works, and works well (so far).

5 Deselect all the MySQL from the software upgrade. The files are faulty at present and after your machine has spent hours downloading them, the install fails and the whole download is wiped!

To install the latest *OpenOffice.org* 1.0.2 you have to

untar the download, open a terminal and, as root, go to the install directory and type **./install** When you have done that, close the terminal and, as the user, open *konqueror* or *nautilus* etc, go to the /usr/share/OpenOffice1.0.2 folder, and run **setup**. This should be done for every user. If you just run **setup** from the install folder, it loads in your home folders only! There are no instructions provided with this version, and the installation is different to 1.0.1.

Dave Spagnol, *via email*

Thanks for your tips on Mandrake installation. I have to say that we didn't have any trouble with installing a local printer during the install, but it is probably best to leave this step until you are sure the system is working anyway.

For your efforts, please accept our Letter Of The Month prize, which this issue is a copy of *Tuning and Customizing A Linux System* published by Apress (www.apress.com, ISBN 1-893115-27-5).

At times, who can blame them? Transitioning from a Windows to Linux programming environment will seem odd at best for someone used to the "joys" of Visual C++ 6.0. Yes, I know about K-Develop, and I love Anjuta, but the levels of documentation and resources available to developers just don't seem to approach the same scale as the Microsoft offerings. And we all know that people out-and-out dislike change. Which leaves a lone Linux voice in a Windows-loving company far too easy to shoot down.

On an unrelated note, I've only recently started buying your magazine – it has a refreshing perspective I assumed died with the Amiga! One suggestion though – could you consider shipping a version of the magazine with no discs? Your discs certainly deserve praise, but with a broadband connection I'm happy to download my own distros, (required only rarely). Software on a CD also goes out-of-date so quickly I'd rather do



Optimising would be easier with standard benchmarking tools.

Benchmarks

I'd like to suggest a follow-on to the excellent introduction to optimising Linux in LXF36. How about benchmarking performance? If I

“If anti-Linux developers were paying for their own copy of Visual Studio.NET we'd wager that they'd be much more interested in Open Source alternatives.”

the 2 minute download. The £6 price tag often makes me think twice about buying LXF when I don't need the discs anyway...

Lee Mallabone, *via email*

We'll wager if the other developers you mention were paying for their own copy of *Visual Studio.NET* they might be more interested in other options. Microsoft does have really good developer support, but it often comes at a hefty cost. Perhaps you ought to be given some sort of bonus for using less company resources!

I'm glad you like the magazine. I'm afraid at the moment we have no plans to do a disc-less version. As the shops already have two versions of the magazine to deal with, adding a third would be problematic – not least because most newsagents wouldn't want to stock all three, and the one they would choose to drop would almost certainly be the cheapest.

don't have time to recompile sources, how about methods for comparison of different distributions/versions on the same box? The optimising article provided *hdparm* tool as a good indication of disk performance but what about X (movie, 3D, standard desktop), kernel, memory etc.. etc?

I installed RH 7.1 and Slackware 8.0 on an old Pentium which was lying around, and Slackware was definitely faster, but by how much? It may not always be visually apparent and what's causing the slow-down, and it would be nice to have a test suite for benchmarking all aspects of performance to see the results of optimisation and pick the best distribution for a fast installation, especially if you run an office with desktops having the same specs.

Kevin Brennan, *Eire*

There are few reliable benchmarks for system performance, although some work is currently being done on a suite which should give reliable results across a range of Linux flavours.

Though of course, with a combination of top, which will show you CPU usage for tasks as well as plenty of other info, and time, which will record the time taken for a process to complete, you can easily cobble together some real-world tests and scenarios.

It may be the difference you notice between RH and Slackware is due to more services being enabled on RH. Also, for work in X, slight differences in versions can make a huge difference to overall performance.

Are you ready?

In LXF37 you posed the question "Is Linux ready for the desktop?"

As usual for Linux Format, the feature was of a very high standard: informative, balanced and promoting the platform without starry-eyed unreality. However, perhaps we should be asking the question the other way round: "Is the desktop ready for Linux?"

It seems to me that those who use Linux on 'personal' PCs probably fall into a small number of groups:

- IT professionals who want a Unix-like OS 'at home'
- Students exposed to Unix at college, ditto
- The anti-Microsoft brigade
- People like me, who built a box from mail-order bits and needed an operating system. >>




Ready for the desktop? We've had letters supporting both sides.

UNIHOST

Specialists in hosting dynamic web sites as well as e-commerce stores with secure SSL facilities.

- Linux Hosting Specialists
- CGI / SSI
- PHP 4
- Perl / DBI
- MySQL
- PostgreSQL
- JSP / Java Servlets
- FTP / SSH Access
- Full Logs & Web Stats
- Huge Network Backbone
- No Data Transfer Limits
- 100 MB Storage
- Full Shell Access
- Monitored 24 x 7 x 365
- UPS
- POP3 E-mail Boxes
- Unlimited Aliases
- SSL Option Available
- Reliable Service



£250 per year (ex VAT)

SPECIAL OFFER:

20% off 1st year's hosting if booked before April 6th!
Quote ref: **LF403M**

sales@unihost.net

0870 744 0721
UNIHOST.NET

Mailserver

« To be honest I can not see any of us in the above list making much of a dent in Microsoft's market share. We need to bring another group on board; the "What's an operating system?" user who just boots up and logs on to the Net or writes a letter or two. People in this category buy a PC with Windows preloaded and so why should they bother with something different? Having already paid up front they are just going to say, as one of your contributors recently reported, "It's good but why bother changing?" These folk need to be able to buy a PC with Linux preloaded. That is

To answer my own question: The desktop user is not ready for Linux because he/she does not know it exists... yet.

Ken Wilson, *via email*

I think that cheap systems like the Evesham one you mention could make a difference, as perhaps could the creeping popularity of Linux in Internet Cafés. In a desktop sense though, the biggest market, and probably the easiest one to sell the benefits of Linux on, is the 'corporate' desktop. As much office work centres around a few office apps and web-based applications, rather than cutting-edge video work or support

"In a desktop sense, the biggest market and the easiest one to sell the benefits of Linux on is the 'corporate' one as much work centres around a few apps."

why I was so pleased to see the Evesham website strongly promoting Lindows as a real alternative.

To reach the bulk of the home desktop market we need to promote Linux with PC sellers rather than individual users.

for exotic peripherals, it would probably be the easiest environment to support too.

Roundups

Thank you for the roundup article on Window Managers in *Linux Format* 37; I was looking forward to



Roundups: some readers think that we should approach these differently to the way they are done now: what's *your* opinion?

reading it as soon as I saw the front cover on the newsagent's shelf. I feel at this point I should say how much I value the style and content of your magazine – it is an informative, well-humoured publication that I am just about to subscribe to (crawl!).

However, the aforementioned article is an exception to this history, and I feel rather lost as to the real benefit of such a non-entirety: I appreciate the constraints that doing such a round-up involves, though to be honest you could have said what you did about

each WM in far less space (can anyone say "bullet points"?), and therefore leave room to discuss the other **unique** elements to each WM. The following sorts of information are vital, in my opinion, if one is to make an educated decision on their choice of WM:

- 1 Compatibility with other apps; *WindowMaker* advertises support for KDE, *NEXTstep* and *Gnome*; but what does this *mean* for people who may want to use the *Konqueror* browser; *Evolution* and, say, *XMMMS* (off the top of my head)? *ie* is it just a case of certain libraries not being

Helpdex

shane_collinge@yahoo.com



Mailserv

automatically installed, or is there something more fundamental about a Window Manager's support for other specific apps?

2 Just how much can you alter the eye-candy and behaviour of each? For example, in Mac OS X you have an animated/auto-hide application launcher which pops up when you place the mouse at the bottom of the screen. Such functionality may increase resource requirements; but it would be nice to know what sorts of things can be changed about a given Window Manager.

3 erm... Keybindings, DockApps, /whether you can use any app-launcher with something like FluxBox... and so on. Generally I think your article may well just have pointed people to the respective websites and let us find everything for ourselves (if indeed the website had all the info...). While the information on the history and toolkit is a good insight, it doesn't really make for a hugely relevant read when it comes to deciding what to use (unless you're actually a GTK developer, or something, I guess).

As an example, I can see no functional difference between FluxBox and IceWM. Neither are said to be able to use any kind of Graphical pager, as per Enlightenment, and I have no idea whether either have adopted (or can be made to use - this is important; just how modular are they?) some kind of DockApps. After reading your article I am none-the-wiser. Sorry folks.


On a personal note, what I would really like is to have a lightweight WM that supports a graphical pager and dockapps, but with the ability to auto-hide them when I maximise an application. I will also need KDE and Gnome app support, as per WindowMaker. :) - any ideas or comments?

On a different note, do you know where I can get the VideoLAN program from that you apparently included in the DVD edition of LXF35? (I get the CD edition.) James Thompson, via email Thanks for your comments. There isn't really much of a compatibility question when it comes to applications running on Window managers. If the KDE libraries are present, a KDE application will run. Similarly for Gnome. The only things that won't work are KDE components such as Kicker applets, which obviously need the Kicker.

I'm sorry you didn't get much out of the roundup, but there are constraints of size, within which we tried to give you some ideas of the features and useability of the different WMs. You can get the impressive VideoLAN from www.videolan.org - it's written by a bunch of French Open Source fanatics who are students at Ecole Centrale Paris.

Gringotts

Your article on the Gringotts was very interesting - this looks like an excellent solution to the problem of storing sensitive data. Unfortunately, it appears that the program is no longer available from www.prosa.com, and I



Opera weren't too happy at MSN singling their new browser out to send 'broken' pages to. Read the full story on www.opera.com

root: /#

£15 per month

Linux Virtual Machines:
simpler, cheaper, more reliable
than co-lo or managed hosting.

- Install any software you like, configure it however you like.
- Total control over your kernel: use iptables, IPv6 or IP tunnelling to build your own network and permanent online presence.
- Preconfigured software for superior hosting: Apache, MySQL, PostgreSQL, PHP, Perl, Python, Ruby & full gcc compiler suite (among others). Choice of Linux distribution.

**Included in our service
at no extra charge:**

- Watchdog with SMS and email alerts, or automatic reboots. Know about server problems as soon as they happen.
- DNS hosting: upload your domain data to our four DNS servers for resilient service.
- Mirrored (RAID 1) storage and nightly backups of your machine prevent data loss and down time.

Quality network & support

- 100Mb multi-peered, multi-homed network in Telehouse, Docklands.
- Friendly telephone or email support, whatever your experience of Linux system administration.

Prices per month ex. VAT

RAM	Storage	Traffic	Price
64M	3GB	5GB	£15
128M	6GB	10GB	£27
256M	12GB	20GB	£50

Bytemark
Hosting

☎ 08707 455026
or order online at

www.bytemark-hosting.co.uk

Mailserver

“The problem with peripheral support is the reluctance of manufacturers to create drivers or even publish the info to allow third parties to create their own.”

« was wondering if you happened to have a copy stored away in your computers. If so, please give me an address for downloading. Many thanks for your attention.

Terje V.Sørensen, *via email*

Not easy to find in a web search, among all those Harry Potter fan pages! There are stable and development versions available for you to download at: <http://devel.pluto.linux.it/projects/Gringotts/index.php?page=downloads>

Fonts and HTML

Andrew Joseph (Letters, LXF37) asks always to receive email in plain text. Whereas I would also support that in principle, there is a problem when receiving mail from far-and-wide. Foreign languages have no easy solution to this apart from HTML or UTF-8 encoding, but even then there is no guarantee that the recipient will have the necessary fonts, despite the age of the standard ISO-10646.

I wonder whether W3C could be persuaded to issue a reference font suitable for any system. It would not have to be of high quality, merely complete, and installed everywhere free of charge. Commercial organisations can then provide their own versions to improve the quality of the result, but at least the information would be accessible to all.

Then, WWW can really mean World Wide Web and Unicode and UTF-8 can start to replace ASCII as the text standard.

In addition, of course, it would be nice to have all e-mail programs able to create any character from any keyboard in some simple manner. Unfortunately, Microsoft is ahead of Linux in this area, but still a long way from ideal, and Apple is a very long way behind.

Andy Pepperdine, *via email*
Another very good point in this long-running saga. The idea of a W3C

reference font is an interesting one, but then there is the problem of what font format it should actually use, or managing the creation of such a font in multiple formats which may not always display the same way (eg TrueType and Type 1 versions of the same font can display differently).

Most modern Linux mail clients do actually support UTF-8 though if you care to use it.

MSN problems

I am sure someone already sent you this link, but I am making sure you have it, because it is worth publishing. I can't believe I haven't heard about it before.

Essentially, the article is about how Microsoft are serving deliberate incorrectly displayed pages on their MSN homepage to Opera users only, making the illusion that Opera can not render MSN correctly. Although when the same content provided to an IE browser is saved and viewed from the local drive in Opera it looks OK. www.theregister.co.uk/content/6/29219.html

Miguel de Melo, *via email*

Opera itself had plenty to say on this matter and even created a special version of their browser (see page 9 in our News section) The real question is whether MSN deliberately and cynically tried to undermine the use of non-IE browsers, or whether the team involved just don't know much about authoring web pages!

Get mapping

I work in UK Local Government. An essential application that your recent magazine article overlooks is Geographic Information Systems (GIS). I have used ArcInfo on a Sun workstation in the mid 1990s, but Unix is not now widely used – if it was at all then. MapInfo started out on Unix – I think, and I have worked with a crude early system on a BBC. More than 80% of data



GUI driven GIS systems for Linux are available. The excellent GRASS is one example, which you'll see more of in a future issue.

has a spatial element (address) – retailers are also interested in this. I have seen info on a system called GRASS but this really does need a GUI. Windows applications do have these and familiarity. Therefore they are “easy” to use.

I use / am getting more familiar with SuSE at home – dual boot with WinME. The cost of Microsoft licenses is going to force many organisations to reconsider their commitment, and look to migrate to Linux – the path will not be straightforward.

Linux is improving as a ‘desktop’ but there needs to be some improvements on simple things like program installation, and support for hardware such as my scanner (Visioneer) and USB MiniCDRW. JimK, *via email*

There are a number of GIS tools for Linux/Unix systems, but most of them are CLI-based. GRASS however, which you mention, does have a useful GUI – you could do worse than download the latest 5.1 release and give it a go. You'll find it at <http://grass.itc.it/>

I hope you don't give up on Linux. The overall problem experienced by Linux users with peripheral support is the reluctance of the manufacturers to create drivers, or even publish the information which would allow third parties to create their own, which is the case with your scanner.

Thankfully more manufacturers are beginning to realise that Linux support is a good thing due to the amount of quality feedback they

receive, and a lot of major names do cooperate with Open Source authors to enable use of their hardware. LXF

Posting to the forum

The LXF online community

Got a burning Linux issue you want to discuss with other LXF readers? Not only do our popular forums at www.linuxformat.co.uk have sections dedicated to your technical queries, hardware, programming languages and general help; but also there's always a lively discussion going on covering many different subjects.

Recent topics include:

- Retrieve filenames from MySQL
- Nvidia Kernel not loading
- apt problem!
- Wireless LAN card drivers
- iptables and passwords

Give a hoot – contribute!

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: lxfl.letters@futurenet.co.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?

Documentation / Value for money: Whichever score is most appropriate!

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



10 The close-to-perfect product.



8-9 Good, but has a few niggles.



6-7 Does the job, but needs work.



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...

Evesham e-Scape

A bold move by system-builders brings you off the shelf Linux. Find out if Lindows and commodity hardware cut the mustard **p20**

Acronis Partition Expert

Managing, resizing and formatting as many types of partitions as you can imagine with this Red Hat-based commercial tool **p22**

GnuCash

Manage your bank account and keep the revenue men off your back with this Open Source finance solution **p24**

Mitel SME Server

The GPL developer version of this 'fire and forget' office server solution gets the once over **p28**

Books

Perl, Mason, HTTP, firewalls, system administration and security – they're all covered in these latest tomes **p31**



COMING UP SOON...

IP* Works

We haven't often covered programming libraries and utilities in the past, but you're in for a treat as we put this commercial library for IP and networking support to the test – is it really better than Open Source alternatives?

Opera 7

MSN notwithstanding, the new client for Linux will be released any day now!

SuSE Enterprise Server

The first real test of United Linux, this is SuSE's new enterprise

offering, their first built on the collaborative technology – is it up to the task?

VariCAD 9

After a glowing review for the previous version, lots more features and fixes have been added to this CAD classic?

LINUX DESKTOP PC

Evesham e-scape Li

Evesham has become the first large UK PC seller to offer a real desktop alternative to general users. Andy Channelle finds out if there is more to their first offering beyond novelty value and a very, very keen price.

Via-based Linux PC. Linux boxes from www.sdit.co.uk have better spec, monitor and a higher price for around £399.

- **DEVELOPER** Evesham
- **PRICE** £249.99 and up
- **WEB** www.evesham.com

For some, the arrival of pre-installed Linux of the high-streets of Britain has been the subject of much anticipation. Proof – finally – that there is an alternative to the Microsoft Tax or wallet-busting Macs, and that at least one respected retailer is convinced of the viability of Linux for home users.

The first thing that hits you about the actual machine is its size and spartan frontage. The former is a consequence of basing the PC on one of Via's mini-ITX motherboards which integrate CPU, graphics, sound, network interface and monitor/TV output into a very compact space. The disadvantage of the integrated route is a lack of internal expansion potential, but Evesham would doubtless argue that the e-scape line is intended for the casual user who

has no interest in adding the latest GeForce graphics processor to their PC. And they'd probably be right. Integration is also key in bringing the price down to such an affordable level.

More problematic is the lack of a floppy disk. Okay, almost no one uses floppies these days, but its omission from the e-scape – and also the lack of space to fit one later – means there is no writable media available for the machine, and that makes backing up your data a real pain. It would certainly be advisable to consider opting for the

CD/RW drive which adds another £25(+ VAT) to the price, or a solid state USB 'key' disk. Of course if you're building or augmenting a network this won't be a problem as you can cover back-ups with a server.

The back plate of the PC features the usual array of ports (monitor, mouse, keyboard, audio, serial and parallel) as well as a pair of USB 1.1 ports, Network interface, S-Video port and a combined RCA Video and S/P DIF port. One of the expansion slots on the MoBo is fitted with a Linux-friendly modem which provides easy 'net access via BT Click by default, although it should be child's play to set up your preferred ISP. Around the front there is a further pair of USB ports obscured by a welcome child-proof cover.

Additionally the box contains a standard two-button ergonomic (right-handed) mouse, mouse mat,

lightweight but usable keyboard, leads and a pack containing four CDs.

Setting up

Setting up proved to be incredibly simple. Everything went where it should and the machine powered up smoothly and very quietly. Now I have to confess that last time I checked out Lindows I found getting access to network information quite difficult, and as I would need to note down the MAC address for the Ethernet adapter to add to my accepted connections list, I was prepared for a struggle. It was a pleasant surprise then to see that this magical number was displayed for a good 15 seconds on boot up as the machine attempted, and failed, to connect to the Blueyonder service. I then simply updated the MAC list from another machine, rebooted and it worked. Lindows updated itself while I made a cup of tea and within a few minutes I had a fully usable machine.

After adding the essentials that Lindows omits (calculator, *RealPlayer*, *Tetris* etc.) from the Click-'N'-Run Warehouse, I did:

```
# apt-get update
# apt-get install synaptic
# synaptic
```

You'll have to get a monitor to go with this setup – but it's still good value.



to get the nice front end to Debian's *apt* package management system, thus avoid the cost of Click-'N'-Run, and added a few extras to the basics.


The system itself boasts an 800MHz Via *Ezra* processor with 64k cache, 256MB RAM (8MB of which appears to be used for the AGP graphics system) and a 40GB hard disk. Though some way off the current 'entry level' it provides adequate performance for things like web browsing, basic photo retouching, office tasks, and media playback. Performance begins to degrade slightly – and I think this is one of the side effects of integration and loading so much onto the CPU – when you try to do too much at once. For instance, MP3 playback through *Xmms* tends to get very choppy when navigating around a large image in *GIMP* or even moving windows around on the screen, suggesting the graphics subsystem is stealing processor time from the audio.

There also appears to be an occasional glitch (which is not even consistently reproducible) that causes the active mouse area to be about 50 pixels to the left of the actual pointer. This happened twice over the period of a few days and wasn't resolved by logging out and back in again. I had to reboot. This can be an annoyance if you're in the middle of some work and I feel it would really throw many non-technical users.

For this non-technical user, Evesham offer a mixture of email and premium-line support at £1.00 per minute. The support area on their website is, unfortunately, a Linux-free zone and the spartan manual doesn't even have a troubleshooting page, which all adds up to a bit of a disappointment.

I also attempted to set up a few peripherals. Printer configuration is exceptionally smooth thanks to the brilliant KDE Control module which, for my Epson Stylus Colour 400, is the very essence of comprehensive.

Next I plugged in a digital stills camera. This is a basic 'USB Mass Storage' device which really should be automatically recognised. Many distros claim to do this, but I've yet to meet one that does it reliably. Fortunately I know that creating a new folder called 'camera' in /mnt, opening a console and typing `mount -t vfat /dev/sda1 /mnt/camera/` fixes the problem. This sort of stuff needs to be in the manual or on the web.



Xmms (running with OpenGL visualisation), *GIMP*, *OpenOffice* and connection to digital camera.

Software

We've covered Lindows recently but just to recap. It's a very tightly integrated operating system built around KDE 3.0.1 with a tweaked Ceramic/Crystal theme and software updates via the Click-'N'-Run warehouse. Click-'N'-Run costs \$99 per year and purchasing the Evesham PC gives you membership of the junior aisle. Click-'N'-Run itself boasts some 1,700 applications including the likes of StarOffice, Quake II and Photogenics. Currently Lindows is pushing a pair of family friendly applications – Virus Safe (a virus checker) and Surf Safe (a sort of net nanny) – and removable icons for these appear in the Kpanel. Clicking on them takes you to the exclusive Click-'N'-Run aisle, disappointingly they carry an additional charge.

Also installed (and housed on a separate CD) is a copy of *OpenOffice.org*, which is very handy if you're running on a modem and want a decent office suite; Lindows only features a very basic word processor.

Curiously, the Lindows CD in the box only contains an installer for the Click-'N'-Run Express service and is not bootable; I discovered this after taking over the machine with Mandrake (which installed and ran

okay) and then deciding to go back to the default. I'm hoping this is a case specific oversight by Evesham. Not everyone is going to have a Lindows CD lying around!

The package also contains a Via CD and manual but these relate only to Windows installations.

Conclusion

If the old adage – you get what you pay for – is true, What on Earth can you get these days for £250? The answer is a pretty good machine if you are staying away from kernel compiles, 3D and mad *GIMP* processing. Throw in a low-end 15-inch TFT monitor and you're still looking at a system price closer to £400 than £500, which is half the price of a traditional entry level PC or a third of the cost of an iMac. Of course the Evesham product isn't going to compete on raw power or expandability with these machines, but for day-to-day tasks, you could ask yourself whether you really need a supercomputer sapping power from the national grid while you play *Tetris* or surf through the daily news and your bookmarks?

This is a basic machine designed for basic tasks, and thanks to Lindows (with a little help from *apt*), Netscape and *OpenOffice.org* it does it all very

well. Hardware performance is unexceptional, but in the real world this barely matters if your needs are unexceptional. Even for power users this would make a useful second system, perhaps serving music and mail to the bedroom while providing a back up machine if the downstairs supercomputer throws a wobbly. A little extra support would be welcome to help iron out the few bugs in the system but this would probably kick up the price a little, and nothing proved to be insurmountable.

So if your funds and needs are not too extravagant this would be a worthwhile purchase. I've started saving already. – it won't take too long to get the cash together! **LXF**

VERDICT

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

A cheap and cheerful PC ideally suited to word processing, browsing, editing your digital snap and listening to music – and all for the price of a decent stereo system. Don't push it though.

LINUX FORMAT RATING
8/10

PARTITIONING UTILITIES

Acronis Partition Expert 2003

Hoyt Duff looks at a commercial partitioning offering based on a modified Red Hat kernel that installs under Windows or boots from CD in Linux.

Partitioning tool with support for various partition types. Also consider *Partition Magic*, or *diskdrake*.

- **PUBLISHER** Acronis (SWsoft)
- **WEB** www.acronis.com/
- **PRICE** US\$40

Many new computers arrive with WindowsXP installed, with XP becoming more prolific as consumers upgrade to new equipment. Linux and BSD-loving dual-booting users now encounter a new problem: How does one resize an NTFS partition? Previous Microsoft home-use operating systems were installed on VFAT-formatted filesystems, easily resized with modern commercial and free software solutions. NTFS is another beast entirely.

Exacerbating this dilemma is the bleak reality that many manufacturers do not include a proper WindowsXP install disk, but rather some annoying partition-image-restore-CD-ROM nonsense that's hard-coded to re-install Windows to the entire hard disk.

The Xandros distro includes a proprietary application to resize NTFS

partitions and Mandrake 9.1 includes a GPL re-sizer. Among a few others until now, there has been *Partition Magic* (www.powerquest.com/partitionmagic/, US\$70). The biggest drawback with *Partition Magic* is that it requires two floppies be built for use with Linux; it's just awkward. Worse, the floppy disks can't be used with the new breed of computers that have no floppy drives.

Enter *Partition Expert 2003*. You may have already heard the Acronis name before – it supplies their OS Selector Special Edition to SuSE for use in the new SuSE *Linux Office Desktop*. That app contains a boot manager as well as a minimal implementation of *Partition Expert*. For those that eschew commercial products for Linux, there are two significant features of *Partition Expert* that may intrigue them: it can be booted from the supplied CD-ROM and it is based on Linux (a modified Red Hat kernel, in fact).

While there exists the possibility of installing *Partition Expert* natively under Windows, sadly no such option exists for Linux. That's the only significant missing feature to the product that we can see. Still, booting it from the CD-ROM suffices for our needs.

GPL Compliant?

Community encouragement needed


If you examine the *Partition Expert 2003* CD-ROM, you'll see a directory listing of nothing. Even using the strings command yields little information. The CD-ROM uses Acronis' proprietary bootloader that can render all the partitions invisible. The included binaries are compressed using a GPL utility that can be found at <http://upx.sourceforge.net/>. Such obfuscation is common and to be expected with commercial products.

Acronis Partition Expert 2003 contains – by its own admission – GPL software, but are they compliant with the terms of the GPL? Certainly not at first they weren't. No notice appears in the printed documentation and it wasn't until February 2003 that <http://www.acronis.com/support/licensing/> appeared on their site. It's refreshing to see that a commercial enterprise adopts

the GPL and makes an effort to comply, as some appear to resort to outright skullduggery to avoid any disclosures. Enforcement of the GPL is, of course, up to the copyright owner, but reasoned encouragement from the community helps as well and provides a model for more widespread appropriate adoption of GPL software by commercial interests in every field of computing.

For Acronis, there seems to still be work left undone. We can't easily find a link to that GPL page on their site and a site search, using their own search engine, doesn't disclose it. Acronis might still be in need of a little appropriate encouragement from the community.

Read more about the GPL at www.gnu.org/copyleft/gpl.html and the LGPL at www.gnu.org/copyleft/lesser.html



The GUI is intuitive and easy to use. Wizards can be accessed or settings entered manually and manipulated graphically. You may create, delete, resize and copy FAT, NTFS, Ext2/3 and ReiserFS partitions, as well as “wipe” them with multiple overwrites to add extra protection to deleted data.

There is a downloadable (but crippled) demo as well as a full version at the website. A PDF copy of the User Guide is also available for download. The boxed version (reviewed here) comes with a bootable installation CD-ROM, a User's Guide (well-written and illustrated, but pertaining primarily to the Windows-installed software) and a registration certificate with the serial number for your copy. Registration entitles you to free downloadable upgrades (one is currently available, but appears to address only Windows bug-fixes).

Booting from the CD-ROM leads you to the main window after first displaying a choice of automatic or manual use of the application. The “automatic” choice will enable “wizards” suitable for easy repartitioning, but hides the manual options. Choosing “manual” permits all these options as well as allowing you access to the wizards; it's easy to switch back and forth.

Partition Expert also performs some partition repairs when it encounters errors it knows how to fix; it will not automatically repair damaged extended partition tables, nor re-align partitions that don't accommodate the 63-sector offset encountered when having used non-LBA (Logical Block Addressing) partitioning utilities like older versions of

fdisk or *ctfdisk* on a drive requiring LBA (large drives); *PE* accommodates drives larger than 180GB. *PE* promises to move data with no loss and can convert Ext2 filesystems to Ext3, as well as hiding and un-hiding partitions, setting the active partition, and performing FAT filesystem manipulations if required.

What could be improved? A native Linux installation would be nice but is not necessary, although a Linux-only upgrade path would be welcomed, especially when the next version of *ReiserFS* appears later this year. Support for JFS and XFS would be handy as well; their inclusion would make *PE* a must-have Enterprise-level utility. Compliance with the GPL is of more important concern in the long-run. **LXF**


VERDICT

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


Robust GPL compliance is sorely missed, otherwise it'd rate 9 for Features and score a 10 overall. A very useful utility.

LINUX FORMAT RATING
7/10

ReviewsGnuCash



The new account druid makes setting up accounts simple.



Scheduled transactions make managing your money a lot less hassle.

FINANCE MANAGER

GnuCash

After a long wait, *GnuCash* gets a host of much-needed features. **Andy Chappelle** straps on his parachute to test an application with a reputation for being "a little difficult to install".

GPL-wise, *GnuCash* is in a field of its own. The only real competition comes in the form of Capital (from TheKompany.com) or packages from Microsoft and Intuit on Windows.

- **DEVELOPER** GnuCash.org
- **PRICE** Free Software
- **WEB** www.gnucash.org

Over the past couple of years I've evaluated *GnuCash* on a number of occasions hoping to find the app that'd allow me to migrate my enormous investment of time and effort in *Quicken* on Windows to Linux, and thus remove another need to dual-boot.

It was with great trepidation that I downloaded the binaries for version 1.8 from the project's site. This is, in the words of the developers, an application that is "extremely difficult" to install. They recommend waiting to acquire a full distribution with it included or using an installation tool like *apt*, but even then, they say, you may still have to do some manual labour. Never content to wait, *LXF* began its descent into dependency Hades in search of the ultimate personal finance tool for Linux.

I started trying to install on *Desktop/LX* and soon discovered that it was indeed a laborious task akin to fighting the Hydra: solving one dependency throws up another 10. The website claims that a standard install of GNOME 1.4 will satisfy most of the dependencies with the exception of *Guppi*, *Guile*, *Slib* and *G-Wrap*. You'll probably also want to get *Finance:Quote* and *Finance:QuoteHist*, which can grab current and historical share and exchange rate quotes respectively from the Internet, and *LibOFX* which allows you to use Open Financial Exchange information which is used by many US banks.

After installing all of these – nothing. *KPackage* was still noting unsatisfied dependencies and I resolved to take another approach on a different machine, running Mandrake 9.0.

After confirming that Mandrake's *Cooker* – a repository for the latest RPMs planned for inclusion in future distros – contained the most recent version of *GnuCash*, I just updated the installation sources section in *RPMDrake* to include the *Cooker* and then selected the application. After noting the number of packages needed for the install and its size (55MB), *gpm* downloaded and installed everything and added an entry



An integrated web browser is useful for surfing financial data, or the *GnuCash* homepage for updates.

to the KPanel under Applications >Finances. Via the magic of Mandrake everything worked. Curiously it also killed access to Mandrake's package tools (on two boxes I installed on) from the Application Panel; you can still get access to them via the control panel.

Accounting for taste

Launching the app for the first time brings up a standard druid interface which provides options for starting a new account from scratch, importing existing data in QIF format or doing everything manually. After defining your chosen currency, (which defaults to the relevant currency for the system setup), you choose the type of account you wish to open. These range from a simple cheque account to mortgage and retirement options and, for the first time, there is a full selection of business account options making *GnuCash* a suitable choice for SoHo users. Each 'file' can hold a number of

accounts, so you could cover both personal and business finances in a single easy-to-backup document. It would be nice to have the option of password protecting a file or even discrete accounts from within *GnuCash*, though. Data are stored using an XML backend which, unfortunately, is readable using something like *Kate* or *KWord*. For most uses this won't be a problem, but there really needs to be a way to encrypt and password protect files from within to make it a viable proposition for most businesses.

GnuCash uses a system of accounts that divides each area of expenditure or income into payable and receivable accounts (which is the underlying concept of double entry book-keeping). While this pays dividends when it comes to compiling reports, it can be initially confusing.

When importing a QIF file from *Quicken* or *MS Money*, the druid transforms categories into accounts, so

you may end up with 'accounts' called *Amazon*, *Local Petrol Station* or any other random name you might have rashly given a payment in the past. There is an option to change these in the second stage of the import druid. Import is also fast: it pulled in and converted an *MS Money* file containing 1,500 transactions in about 12 secs. You can manually set opening balances for most options if needed, and once the druid has finished its work you're left with an account tree. In the case of most standard forms of account; the first level covers Assets, Liabilities, Income, Expenses and Equity and provides a quick overview of your financial situation. It can look a little overwhelming, but it soon becomes familiar. In fact, you can use either the main heading or subheadings to make your entries, which means the overall structure can be as simple or as complex as you want; though sticking within well defined categories make creating reports a lot easier later on. Navigating through the tree provides extra detail and double clicking on one of the branches opens up the Register which is where most *GnuCash*'s work takes place.

The register follows the conventions of a bank statement with each transaction confined to a single (or double) line. Each entry should contain a date, number, description and the details of the actual transaction, which can be either a transfer from another

section, income or expense. It all works as expected, and those familiar with other systems should have no trouble adapting to *GnuCash*.

In on time

One feature that makes *GnuCash* a serious proposition for day-to-day use (for me at least) is the advent of scheduled transactions. Finally, my Linux box can remember what day the mortgage payment goes out every month so I don't have to input it each time. This seems a small thing, but we're in an age where home owners could have 30 or 40 different direct debits coming out of their bank accounts each month. Adding these manually to the register is a pain and is the single reason previous versions of *GnuCash* haven't cut the mustard. The fact that they are here is a cause for celebration.

And setting them up couldn't be easier. You just need to specify the amount to be credited or debited from your account, the date on which the transaction takes place, its frequency and how you'd like to be reminded. Unfortunately *GnuCash* won't take your scheduled transactions from an imported file (they are, after all recorded as regular entries), so you'll be starting from scratch there. But it's not exactly time consuming and a little work at the start saves a lot of time.

Occasionally, especially if you're using the application to run a business,

you'll need some sort of output beyond the figures. *GnuCash* offers a range of charting tools to visualise income, expenses, profit and loss etc and, in the main, they work fine. There appears to be a small glitch in the calendar dialog that causes it to ignore February 2002 in the 'report from' box, which is odd, and there was an occasional issue with selecting years with the arrow widgets. This is not an insurmountable problem and is something I would expect to be cleaned up in a maintenance release. Despite these small problems, the charts are very configurable and offer a broad overview of where your money is going that is just not apparent from examining reams of figures.

Conclusion

There are elements of *GnuCash* that may seem a little confusing, especially if you're migrating from Quicken or Money, but if you take the time to read and digest the excellent help and tutorial section, you'll be using the application like a pro in no time. Support files are often the last thing that programmers think of so the *GnuCash* developers should be lauded for compiling a collection of fluidly written and well designed help pages that hold your hand from setting up your accounts to day-to-day use. It also includes a primer on good accounting principles too which is a bonus. There are some gaps, notably in


the area of online banking, but what is there provides a good template for future upgrades.

The only glitches I've encountered so far have been in the occasional failure of the UI to draw the calendar box within the Register – clicking on the icon again usually solves things – and the report difficulties mentioned already. Stability has been impressive, with no crashes or significant slow-downs even when browsing an absurdly large file.

The new way of working takes a little getting used to but, you could argue, it makes more sense to follow 'standard practice' rather than simply trying to ape existing applications. Anyone with even a basic knowledge of double-entry accounting will acclimatise quickly, and those without this experience will soon find themselves familiar with it through the thoughtful help system.

The addition of new account types and scheduled transactions makes *GnuCash* more suitable than ever for both home and small business accounts and it's free (surely the accountants favourite word). Integrated security or encryption will be a very welcome addition in future releases, the ability to access the contents of files from *Konqueror* would be a little worrying if, for instance, I was inputting sensitive information.

Overall though this is a great package, it does what you would expect and it does it in an elegant and intuitive way. Moreover, it does it without having to lay out any money. With an installation as smooth as that facilitated by the packagers at Mandrake, *GnuCash* would be a useful addition to any home or small business budgeteer. As it stands at present, the installation is a pain, but if you can't wait for a general distribution upgrade, it is worth it. [LXF](#)



The charts on offer are simple but effective, providing a coherent way of visualising cashflow (or lack of it!)

VERDICT

Installation	5/10
Features	8/10
Ease of use	7/10
Documentation	8/10

Not as feature-rich as the competition, but then not as bloated either. Some work on the installation routine would really make this a really top-class bit of software. As it is, Mandrake users will be laughing all the way to the bank.

LINUX FORMAT RATING
7/10

SERVER/GATEWAY DISTRO

SME Server 5.6 (Developer release)




Richard Drummond looks at a DIY server appliance distro for the SOHO networks. Does withholding free-version features adversely affect marketability?

A low-maintenance Linux server for the small business world. Try also Smoothwall GPL, Astaro Security Linux

- **DEVELOPER** Mitel Networks
- **PRICE** Free (unsupported)
- **WEB** www.mitel.com

Your average general-purpose Linux distro can more than adequately serve the needs of a SOHO network – and do so both efficiently and cost-effectively. The problem is that most small businesses won't have staff with sufficient skills to undertake the configuration and administration required to effectively use Linux, or afford to hire in the expertise to do so. One solution is to use a special-purpose distro such as SME Server, which turns a redundant PC into an easy-to-manage server appliance. SME Server can act as an Internet gateway (that is, a masquerading router and firewall) for your network and provide web, mail, directory, print and file-sharing services. It is designed to be easy to install, configure and administer.

SME Server is an Open Source server software product based on Red Hat 7.3. It was formerly developed by e-smith under the name *e-smith Server and Gateway*, but this name was changed when the company was acquired by Mitel Networks in July 2001. We're looking here at the freely-distributable, unsupported edition of SME Server. Mitel and its partners sell various products which use SME Server, including some all-in-one hardware and software solutions; they also provide various add-on and support services for SME Server (see the box, ServiceLink). Mitel calls the free version 'the developer release,'



SME Server can act as a dedicated email server or can fetch and deliver mail from a remote mail server such as an ISP's.

but don't be put off by this uninspiring title. It is a stable release. Moreover, a thriving community has grown up to support the product, providing forums, documentation and a number of open-source add-ons – all accessible for free at www.e-smith.org/.

All systems go

Installing SME Server is straightforward. The installer interface uses the ever-familiar text-based dialogs and prompts for a minimum of information. Boot the CD, pick a language, keyboard layout, and time zone and the software gets installed. Since it's a single-purpose distro, it's much smaller and quicker to install than many users might think. You should note that SME Server takes over the entire disk or disks of your system, so make sure there's no valuable data before you install. SME Server supports SCSI and IDE disks and RAID-1 in software (if you have two drives, the second drive is a redundant mirror of the first).

Once installed, your system is rebooted and you can perform the initial network configuration, and, again, this is via a text-dialog interface. First, you must enter a root password, and a domain and host name for your system. Then, SME Server will probe for an Ethernet adapter, and you must supply its network details. Then you pick an operation mode: SME Server can operate either as a server and gateway or as a server alone. (Actually, there are two 'server and gateway' options: public and private. With the former your web and mail server are accessible from the Internet; with the

latter they are not and the SME Server uses stricter firewall rules.)

If you go for the server and gateway option, you must then set-up your Internet connection. Dial-up modem and ISDN is supported, and Ethernet for ADSL and cable connections. DHCP and PPPoE are supported, as are various dynamic DNS services – such as that at www.dyndns.org – which you can use to give a server on a dynamic IP address a world-visible domain name. A neat feature for those on dial-up access – particularly if you're sharing a phone line – is that you can configure different connection policies for different times of the day: office hours, evenings and week-ends. You can choose whether the connection is always on, whether it times out after a short, medium or long period.

The last main task is to configure SME Server's DHCP server, if you wish to supply dynamic network addresses to your local network.

Now your SME Server installation should be ready to use. Any further configuration or administration can be performed remotely via a web-based interface, called the 'Server Manger'. A limited amount of local configuration is also possible with a menu-based 'Server Console'. This gives options to re-do the network configuration, to reboot the configuration settings, to reboot the system and so on. Interesting, you can also launch lynx to access the

ServiceLink

Buy add-on features for your SME Server

Commercial support is available from Mitel or one of its resellers for SME Server as a ServiceLink subscription. This can include optional extra features such as 24/7/365 server monitoring, virus protection (both for

the mail server and the files on it), URL blocking for the web proxy, point-and-click IPsec VPN functions and guaranteed email and DNS services. Prices start at around US\$795 for a year's basic support.

New in V5.6


Updates and add-ons

Improvements in this release include a 2.4 series kernel for the first time (and hence an *iptables*-based firewall is present) and a journalled filesystem (it uses ext3). The major system components have been updated to newer (stable rather than cutting-


edge) releases, including *Apache* 1.3.23, *BIND* 9.2.1, *PHP* 1.5.2, *Samba* 2.2.5. The SME Server installer supports the upgrading of earlier V5 series installations of SME Server, including upgrading the filesystem to use ext3.

Server Manager via the text console. Naturally, shell access is supported if you switch to another virtual terminal. You can also enable SSH via the Server Manager for remote shell access, and PPTP is supported to give secure remote access to the Server Manager from the Internet. (PPTP is a Virtual Private Network or VPN technology which creates a secure tunnel from a remote machine over the Internet to your local network; the Server Manager is only accessible from the local network.)

The Server Manager provides a simple and easy-to-navigate interface for administering the various services SME offers. It goes for functionality rather aesthetics, but everything is clearly laid-out and labelled. No online manual is supplied with the distro, but you can access a user manual at Mitel's web site. This has yet be updated for V5.6, but it is well-written. It clearly explains not only how to use and configure the system, but gives helpful background info too and includes sections on setting up client machines.



SME Server's i-bays provide a flexible and easy-to-use mechanism for file-sharing and group collaboration.



The local Server Console gives you text-based access to the web administration interface.

From the Server Manager you can complete the configuration of your system and perform the majority of routine maintenance. Thus you can control which forms of remote access are allowed, add sub-networks to the routing table, add host names and IP addresses for service via DNS, add virtual domains, configure the root LDAP directory entry, add printers, set the Windows workgroup, and so on.

Managing groups

You can also manage user accounts, groups, quotas and aliases. Each user automatically gets their own home directory, mailbox – accessible via POP3, IMAP4, or a web-based interface – and entry in the LDAP directory. SME Server can optionally act as a Windows domain controller and will export your user accounts to Windows NT/2K/XP clients if you need it to.

One of SME Server's top features is its support for file sharing. User's home directories are optionally accessible via HTTP or FTP and exported for Windows clients (with *Samba*) and Apple clients (with *Netatalk*). Besides regular user accounts, you can also create additional named collaboration spaces called information bays or i-bays. This is a quite unique and flexible system. Basically, each i-bay is a separate directory on the server, which, again, can be optionally set up for access via HTTP or FTP or exported for file sharing. You can control whether each i-bay is visible to the Internet or just to the local network, whether it is password protected and whether it is writeable. Each i-bay has three sub-directories: files, which contains the files exported for sharing and accessible via FTP; html, which contains web pages served up when that i-bay is accessed via HTTP; and cgi-bin, which can optionally contain CGI scripts for the i-bay's web pages. Creating an i-bay is a simple matter of clicking a link in the Manager interface, choosing a name for the i-bay, and setting the access controls.

The final aspect of the Server Manager is the administration section which provides monitoring and back-up functions. SME Server supports nightly back-ups to tape, or you can download a snapshot of your server configuration and all the server's data (including mail, user files and i-bays) to the desktop from which you are

accessing the Manager. As a nice touch, an additional option is the ability to create a re-installation disk, which is a bootable floppy which can be used with the installation CD-ROM to completely and speedily re-install SME Server along with your settings.

While back-up support is good, which is always a great attraction for the small business user, in contrast, the monitoring features are weak. SME Server simply provides functions to view the standard system and mail logs. There are no logs of server load, memory consumption, or network throughput. You cannot archive or email logs, remote logging isn't supported, and there are no alarm features. In the light of these facts, it could be said that management functions are weak altogether. The Open Source version has no support for checking for software updates and downloading and installing them, for example. You have to manually install RPMs as well.

These shortcomings highlight the main problem with SME Server. While it is easy to use, and it does what it claims to do competently, in this developer form, it is lacking in features. If you then ask why these particular features are missing, the reply would probably be that these are the features that Mitel sells in the commercial version – such as web filtering, virus protection and support for automatic updates. Nevertheless, the core Open Source SME Server could stand beefing up – with better monitoring features, improved printer support (for example, using CUPS rather than plain *lpd* queues) and an interface for tweaking the firewall settings. At the moment, while SME Server shines in the file-sharing department, as an Internet gateway it compares unfavourably with rival firewall distros such as Smoothwall GPL or Astaro Security Linux. **LXF**

VERDICT

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

Competent solution for small networks, with great file-sharing abilities, but otherwise a little light on features.

LINUX FORMAT RATING
7/10

Red Hat Linux Firewalls

Jon Kent has a squint at Red Hat's official guide.

- **PUBLISHER** Red Hat press
- **AUTHOR** Bill McCarty
- **ISBN** 0-7645-2463-1
- **PRICE** £37.50

Having reviewed Red Hat's rather poor administration guide, I did not hold much hope for this. Though the title suggests it is only of use to Red Hat users, this is not the case as nearly all of the examples featured hold true for any version of Linux.

It is split up into 3 logical sections, a quick overview of networks, firewall design/implementation and lastly firewall operation. Each section builds upon knowledge covered previously and is presented in a easy to follow format. The first section is designed to be a quick refresher but nevertheless covers firewalls and TCP/IP is enough detail to

be useful even to those who are not up to speed on the topics.

The firewall design section is where things start to become interesting. The discussion on the differences between *ipchains* and *iptables* will be of use, but is not presented until the difference between non-stateful and stateful firewall is explained. In addition to this, the various approaches to network topology are discussed, with the pros and cons for each outlined. The sections on *ipchains*, *iptables* and advanced *iptables* will provide you with more than enough knowledge to confidently implement a firewall solution. As you would expect NAT and masquerading are also covered in reasonable depth.

The last section on firewall operation spends some time on installing and configuring a Red Hat server for act as a firewall, with is probably the only Red hat specific section of the book. Something



that is usually missed is how you test and troubleshoot you newly configured firewall so it is good to see a chapter dedicated to this. Lastly no firewall installation is ever completed, so the last chapter covers some of the tasks required on a ongoing admin basis, which is, again, very useful.

This book is very well written, logical and will be of use to anyone, even if you only want to setup a firewall at home. Although there are some little problems which could have been corrected at the

proofreading stage, the contents of this book outweigh these slight oversights. If you want to know about Linux firewalls in general, this is the book to buy.

VERDICT

A great firewall book from Red Hat Press, useful for any distribution.

LINUX FORMAT RATING

9/10

Embedding Perl In HTML With Mason

After learning the intricacies of the funny handshake, Dave Cross rolls up one trouser leg...

- **AUTHORS** Dave Rolsky and Ken Williams
- **PUBLISHER** O'Reilly (2002)
- **ISBN** 0-596-00225-4
- **PRICE** £24.95


When building dynamic web sites, your life will be much easier if you use a templating system. This prevents you from having the same pieces of HTML appearing in multiple places on your website as each page is generated from a template which defines the parts of the page which are the same (site navigation, headers and footers etc) leaving you free to worry about the actual content which changes from page to page.

Mason is one such templating system that is available for use in Perl.

There are a number of such systems but *Mason* is one of the most popular and this book, written by two of the core *Mason* development team, goes a long way towards explaining why.

The book is available on the Web at www.masonbook.com/book/ for free, or you can buy the paper version. It starts with a look at the features that *Mason* gives you and a brief look at the major competitors to *Mason*. While this is useful, I'm not sure that it gave me enough information to decide to use *Mason* instead of an alternative system. The authors assume you don't need any coaching in Perl, so they don't explain anything other than the *Mason* features that they are using.

While it's possible to use *Mason* on a site that is created with CGI scripts, this can be a bit slow so most of the book is focussed on using *Mason*



under *mod_perl* – the Apache module which embeds a Perl compiler in your web server.

One particularly useful chapter is the one which gives an extended case study of how the authors used *Mason* to build a new website from scratch, with valuable examples of the stages you can go through when building a site using *Mason*. Another standout chapter is on scalable design.

Although the book starts right from the basics of *Mason* usage, it won't stop being useful once you've mastered that. There are sections on very advanced *Mason* features like

overriding the parser that extracts *Mason* commands from your source files and subclassing the various objects that make up *Mason*. *Mason* is Open Source, and can also be downloaded from www.masonhq.com

VERDICT

Good for both instructing beginners and as a reference for more experienced web designers.

LINUX FORMAT RATING

8/10

Linux System Administration

Jon Kent asks whether this is suitable for novices.


- **PUBLISHER** Sybex
- **AUTHOR** Vicki Stanfield and Roderick Smith
- **ISBN** 0-7821-4138-2
- **PRICE** £37.99

Linux System Admin books usually come in two guises covering a single distro – reference books or beginners guides. Though this book is very much a beginners guide, it still has info that seasoned users would find handy.

Although there is a slight Red Hat bias, the book covers other distros, such as Debian, and oddly even Corel gets a mention, though this died off quite a while ago. Oddly though, the authors recommend Webmin to administer, which is a good tool, but not something I would recommend to someone who needs to get up-to-speed with Linux. Once they know their way around a Linux system,

Webmin may prove useful, but I would not want a junior admin using Webmin until this knowledge is in place. Luckily enough, the book does cover the more conventional command line approach in detail as well.

Pretty much everything a novice would need to know about system admin is covered in enough detail for you to have a good knowledge base to build upon. If there's any major complaint about the book, it would be its structure – some sections feel like they're in the wrong place and sometimes unnecessary. The section on installing Red Hat is very well written, but the Red Hat documents supplied with the distro are as good – it might have been better to have covered a Debian installation which is not nearly so well documented. As this book goes come across as a beginners guide, it does seem a little odd to cover hardware configuration choices as it is highly unlikely that a junior



administrator would be involved in these choices; this section would perhaps of more use in the appendix.

But these are minor complaints when you look at the rest of the contents. This is possibly the first admin book that doesn't unnecessarily cover *Apache*, which is refreshing to see. There is everything you would need to know, from simple tasks such as setting up user accounts to performance analysis and tuning of your system. The most useful sections are on Kernel compilation, filesystem management and security, providing everything a newbie would need and could refresh the mind of an old timer as well. Even

the section on performance tuning, though not in great depth, still provides enough information to gain some benefit. Lastly there are some useful hints on troubleshooting some common system problems which are handy for reference as well as learning.

VERDICT

Novices and even old-timers will find this useful to have around. As a intro to Linux admin, it is one of the best, despite a lack of structure in places.

LINUX FORMAT RATING
 **9/10**

HTTP The Definitive Guide

It's not Linux-specific, but Maurice R Kelly reckons that all flavours of web users will appreciate this.


- **PUBLISHER** O'Reilly
- **AUTHOR** David Gourley and Brian Totty
- **ISBN** 1-5659-2509-2
- **PRICE** £31.95

Anyone who has ever read the HTTP specifications (RFC 2616) will know that they can be exceptionally dry in places. Designing a new Web application by referring to the RFC can be very time-consuming due to the sheer quantity of condensed information. Thankfully this book takes apart the specs and expands them into a more easily digested form.

It begins with a brief introduction to the protocol before getting stuck

into the meat of the HTTP messages. Rather than simply flesh out the RFC there follows a detailed section on the various applications of the HTTP protocol, such as Web servers, proxies, caches and robots/spiders, which is useful if only as an insight to the inner workings of the Web. HTTP security considerations are given a thorough treatment, as does the subject of transcodings and internationalisation.

These sections are all very detailed with plenty of excellent diagrams that really ease the understanding of the transactions being described. The book is well-written – it is technical enough without being over-confusing, although I sometimes felt that points were overly laboured and occasionally somewhat repetitive.



The final section covers "Content Publishing And Distribution" and included such topics such as *Frontpage* and *WebDAV* HTTP extensions, hosting and logging. While these are considerations for Web server administrators, I felt they really didn't warrant inclusion in this book.

Overall this book is great for people trying to understand HTTP for the Web or other protocols over HTTP. It is rounded off with 100 pages of appendices which form a decent HTTP reference. The programming examples

are not Linux specific, but it is aimed more at explaining the protocol than actually implementing it.

VERDICT

Good companion to the official specs but not the in-depth HyperText tutorial that many readers may actually desire.

LINUX FORMAT RATING
 **8/10**

Building Secure Servers With Linux

Want a secure network? **Chris Howells** examines a book that aims to help us all achieve this.

- **PUBLISHER** O'Reilly
- **AUTHOR** Michael D. Bauer
- **ISBN** 0-596-00217-3
- **PRICE** £31.95


With the exponential increase in computer crime and security incidents in recent years, there is a large need to ensure that computer systems are properly hardened against attack. Building Secure Servers with Linux has recently been released by the *LinuxJournal* columnist Mick Bauer in an attempt to provide a comprehensive security resource for system administrators.

While the name of the book suggests a lot of Linux-specific material, in actual fact virtually all of

the material would apply to any modern Unix-like system such as one of the BSDs. Most of the Linux-specific material of interest is confined to installing the software.

The book begins by taking a look at threat modelling and risk management. This sounds quite obvious, but in actual fact the process is often rather more complicated than it initially appears due to the large number of scenarios. Having a decent written guide is therefore invaluable.

From chapter two onwards, the book begins to get more technical. First, the physical designs of secure networks are considered; there is no point in having a network design that will aid attackers. Later, most popular network services are examined, with sample configuration options that can be applied directly to production



systems. Featured software includes iptables firewalls, the (in)famous *BIND*, as well as *Postfix*, *Sendmail*, *Apache*, and much more.

Of course, attempting to cover the whole of Linux security in detail in a single book this size is virtually impossible, and I would consider it more of a general overview. For example, with the growing popularity of wireless networks, it would have been quite nice to take a look at *IPsec* encryption since, since the built-in WEP encryption is totally inadequate. Additionally, if for example you have little idea about firewalls or DNS, then

you will almost certainly want to add titles such as

O'Reilly's *Building Internet Firewalls* and *DNS And BIND* (reviewed in *LXF* issue 38) to your reading list. **LXF**


VERDICT

An ambitious book which quite successfully covers a large number of topics.

LINUX FORMAT RATING
 **9/10**

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

DOMAIN NAMES
 WEB HOSTING
 E-COMMERCE
 VIRTUAL DEDICATED SERVERS


log on

and be seen

Pick up a new domain name from £3.50 or a web hosting package from £30 per year. Backed with robust support and technology, Easyspace is one of the largest ICANN accredited registrars in the world, and the UK's favourite domain name registrar.

log on today at: www.easyspace.com

Established in 1997 with over 1/2 million customers. Accredited ICANN registrar & nominet member. Prices exclude VAT.

Easyspace 
 your perfect partner for the web

Roundup

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



File managers

Too many files to deal with? **Marco Fioretti** tries to restore control with twelve (very different) file managers for Linux.

OUR SELECTION AT A GLANCE

- Midnight Commander
- UnixTree
- sfm
- Kc
- DFM
- FileRunner
- Emelfm
- Tkdesk
- Gentoo
- Rox
- Konqueror
- Nautilus

Theoretically, the ultimate File Manager on a computer is the user themselves. Classifying, moving or deleting our most important digital assets cannot be delegated to any piece of software, can it? When there is more than a handful of them, however, anything that can speed those tasks up, and possibly make them more pleasant, is a real blessing – that's why the File Manager is one of the first programs that every user sees.

Linux File Managers are *very* different from each other. First of all, "file management" means very different things to different people. System administrators handle configurations and binary packages;

programmers write libraries and (version-controlled) source code; end users fiddle with everything from word processing documents, spreadsheets to video, pictures, and web pages.

Many Linux users, however, wear all the above hats every day. Nobody could pretend that they can get away with only one tool 100 per cent of the time. For example, mouse addicts will have to quickly browse their file

system in text mode as soon as they have to debug (or repair!) their graphic server or window manager. Similarly, those diehard console devotees who claim that "*ls* and *find* are the only File Managers I need" will go down the graphics route the first time they have to sort out the snaps that they've downloaded from their camera.

The applications presented in this roundup will help in just about any

situation, have been tested (with one exception, see below) on Red Hat 8.0, and have been selected with the following criteria:

- Be usable even on not so powerful systems
- Be actively maintained
- If stable for several years, be easily compilable with the current gcc, libraries and toolkits
- Have an active home page with valid links to the latest source code

With quite a lot of candidates that initially seemed to be begging for inclusion in this roundup, these criteria turned out to be anything but granted, regardless of Google listings...

'Even diehard console devotees who claim "*ls* and *find* are the only FMs I need" will go down the graphics route sometimes.'


Midnight Commander

The most popular File Manager, and therefore a reference point for all console-based management.

■ **VERSION** 4.5.55 ■ **WEB** www.ibiblio.org/mc/

The *Midnight Commander* (binary name *mc*) is the most popular text-only file managers around. It works well on practically every kind of Unix under the sun (no pun intended), not to mention Win95/NT and OS/2. It is fully usable both locally and remotely,

and all you need to use it is the keyboard, even if some mouse support is available. This, and several other features that we will describe shortly, make it a popular choice among system administrators, both also home users should at least know that it



Midnight Commander sees no difference between binary packages and (remote) file systems.

exists, for the reasons mentioned in the introduction.

The default look of *Midnight Commander* is the one shown in the screenshot. The basic commands needed to get started are F1 for the online help, F2 for a context-sensitive menu, F9 to access the five pull-down menus on the top, and F10 to quit. File name completion and "help hints", as well as a shell prompt, are always available.

Looking more closely at the screenshot, you will immediately notice how extended the concept of 'file system' can be. In this case, the left panel is the content of an RPM file, not a standard folder. In general, *Midnight Commander* (at least in its Unix versions) knows how to deal with Virtual File Systems, which include FTP servers, and archives in tar or zip format. FTP can be both anonymous and (with commands in the form: **cd ftp://username@hostname**) non-anonymous: FTP-only proxies (not those for web browsing) are also supported.

All kind of 'file systems' can be displayed either as a flat listing, or with a tree view (this can be set separately on each panel). When needed, an

internal editor, which also has syntax highlighting for some languages, allows you to change configuration files or scripts without leaving the program. Something was deleted by mistake? No need to worry (at least on *ext2fs* filesystems): the *Commander* can undelete files. Other almost unique features are the capabilities to strip headers from news articles, and to call the info browser. Another reason why some consider the *Commander* more a System Manager than just a File Manager is the pop up window (in the Command menu) which lists and can kill background jobs.

With one exception (selecting 'Info' from the Left or Right menu simply froze the application) the *Commander* proved to be a stable, fast and quite powerful tool for file and system administration.

VERDICT

Features	9/10
Ease of use	5/10
Documentation	8/10
Performance	7/10

The tool of choice for system administrators and advanced users interested in trying new packages, but also a precious help in emergencies

LINUX FORMAT RATING

7/10

UnixTree

A *Midnight Commander* for DOS fans, with a GUI that boasts an altogether nicer look and feel.

■ **VERSION** 3.0.0 ■ **WEB** www.unixtree.org

UnixTree (the executable name is *xt*) is another console-based File Manager, inspired from the DOS program *XTreeGold*. It offers a nice threaded view of the file system. When it is needed to remove whole portions of the directory tree (**prune** in *UnixTree*) the portions that will be pruned are shown in reverse colour before actually deleting them.

The default view shows both the parents and subfolders of the currently selected directory: the actual files that it contains are displayed below, in a separate (obviously scrollable) area. The statistics of the whole file system and of

the selected node, are automatically refreshed. Several commands, like **find**, **tar**, and **grep**, are directly connected to hotkeys, and mouse support can be switched on or off. The program can also be run in read-only mode, to avoid accidental file deletions.

It is possible to split the screen to display two directories at the same time. Files can be printed, edited or viewed directly inside *UnixTree*, in ASCII or hexadecimal mode. If the file to be managed doesn't appear on the screen, pressing **G** (Goto) and then typing the first letters of the file name will jump directly to it.

Pressing F1 opens a sub window containing explanations of all the commands available in the current context. The F9 key, instead, gives an interesting Applications menu. In its default configuration, it allows users to select and start an editor, control the print queue, read news or email, or open any shell. Of course, this menu can be customised, ie the user can add any other application or special command (for example to mount or unmount devices).

Files can be selected (tagged) in several ways: when looking for files newer or older than a specific date, this can even be entered as 'yesterday', 'last sunday' and similar.

Almost all of the above features are also available for remote use: *UnixTree* can connect to an FTP server and let you manage the remote files just as the local ones. *UnixTree* can be entirely configured while running it: pressing F4 and then F10 goes to the main configuration screen. The

configuration section has eight pages of options, the display configuration four. All the settings can be saved to, or read from, any file: the last option makes it possible to try whole sets of different configurations without wasting time by having to restart the program every time.

Apart from the official website, a valuable resource for *UnixTree* fans is the XTree official fan page, at www.xtreefanpage.org/

VERDICT

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

A good alternative to *Midnight Commander*, with maybe a slightly smaller number of features, but a more intuitive and refined interface.

LINUX FORMAT RATING


8/10

RoundupFileManagers

sfm

GUI based File Management on a diet.

■ **VERSION** 1.7.2 ■ **WEB** www.chez.com/prigaux/sfm.html



No sub windows: *Sfm* always tries to occupy the smallest screen space.

sfm stands for **Simple File Manager**, and it is indeed as simple and minimalist a file manager as there can be today, using GTK, without going to the command line. Practically every command is bound to a keyboard

shortcut (one explicit design goal was to not be as mouse-dependent as the most popular competitors), and up to four *sfm* windows (identified as 'view 1.4' in the menu) are also tied to hotkeys, so jumping from one to

another is really fast. The menu is still there, so there is no need to study several pages looking like assembly code before you start using *sfm*: while you use the menu, however, do pay attention at the hot key listed aside *every* entry (more prized programs may really learn from *sfm* here!) until you memorise them, and then watch how your speed increases by a factor of two or three times. The command prompt for running generic actions on a selected file is there, of course.

sfm has no online help. The documentation is available on the website, or as a README file in the source distribution. Almost all the configuration is defined and changed by editing the \$HOME/.sfm file. This doesn't mean that that file must be typed in from scratch, however: First of all, there is a complete .sfm-example file included in the distribution, and another one is also available on the web page. In the second place, there are things that *sfm* 'learns' by doing, through the GUI, even if it doesn't give a dedicated window where the user can define commands and settings. When the user selects some file of a type never seen before by *sfm*, they are prompted with a dialog box, which asks

what action should be performed on the file. Once the action is typed in, *sfm* writes it in the configuration file for future use, and won't ask again.

In spite of this, the otherwise elegant interface is quite limited, and so are its capabilities, when compared to other programs of the same weight. The main, if not the only reason, is simply that the program has not been developed anymore in the latest years. This doesn't prevent it from still being a good choice for users who prefer a simple, but GTK themeable interface, and want to stand away from the mouse (or maybe just start GTK programming with a simple, but already working and useful project).

VERDICT

Features	4/10
Ease of use	6/10
Documentation	6/10
Performance	7/10

Nice idea, unfortunately not followed through. Ease of use would benefit with provision of some online help.

LINUX FORMAT RATING

 6/10

Kc

A lean and mean file manager and FTP client in one window.

■ **VERSION** 2.2.7 ■ **WEB** www.maik-heese.de/


Kc calls itself "A program for local and remote file management". This simple Tcl app is less featured than others, but still useful. Basic operations are all supported: disk usage, disk free, compression, changing name, position and permissions of files. One limitation of the compression interface (which seems to be common to other 'stable' programs here) is that it only supports zip and tar/gzip, not bzip2. Being Tcl, adding this support should be easy.

The graphical layout has a basic two-pane structure. The "info mode" of each pane (what kind of info is displayed) can be toggled separately: the screenshot shows modification time on one side, and file permissions on the other.

Simple features (and that *Kc* knows how to convert text files between Unix and DOS) makes *Kc* ideal for more

limited systems. What makes *Kc* a nice addition to any lean desktop is its built in support for mounting floppies, lomega ZIP disks, and other devices. This helps to get rid of another whole category of hand-typed commands, regardless of which window manager is used.

Another field of app (visible in the screenshot) is FTP. It is fast, simple and configurable: clicking the FTP button on top of the listing box gives a list of bookmarks, obviously editable from the *Kc* GUI. Recursive copy, moving and deletion of whole directories is possible both on local and remote files. Both the basic FTP statistics (speed, transferred bytes and files) and the complete log file of the FTP session are visible. Sadly, this last area is read-only: it is not possible to go inside it and type particular FTP commands. Something like this would



Kc provides just the essential info, but moves files back and forth from your PC to others easily.

have been a nice addition, and reduced the need to modify *Kc* when dealing with some FTP servers.

The author of *Kc* also makes available *mount-tool* and *tkcps*, two Tcl/Tk interfaces to mount local and remote disks and, respectively, print. Integrating these two tools in *Kc* should be quite easy, and is able to extend a lot the capabilities of this little nice tool.

VERDICT

Features	6/10
Ease of use	8/10
Documentation	6/10
Performance	8/10

A useful little program, with only the essential features, but clearly laid out and easily extensible


LINUX FORMAT RATING

 7/10

DFM

A simple manager with an OS/2 inspiration, and an essential desktop extension.

VERSION 0.99.9 WEB www.kaisersite.de/DFM/



DFM provides shortcut icons for both directories and applications.

Inspired by OS/2, its purpose is to make launching programs and managing files associations easier. It is probably the simplest one in this roundup which also attempts to create and manage a whole desktop.

Starting *DFM* without any options or source patch produces what is shown in the screenshot: the file manager window itself is really minimalist, but the user also gets, on the root window, shortcut icons to the (presumably) most

used directories and applications. These icons can be managed from the 'desktop context menu', the first window from top in the picture. Icons to launch new apps can be added by right-clicking on one of the existing icons, and selecting the 'Create>New program icon' entry. The icons can have any size.

Double-clicking on one folder opens another window showing its content. This can quickly fill the screen with *DFM* windows, but many users prefer this kind of display anyway. Furthermore, *DFM* offers a simple and nice way to keep control, ie to not confuse one window with another. Each window can be given a different background colour, which will be saved and reused when that window is opened again. In addition to this, the view of each window can be selected separately. In the screenshot, the one used for the desktop icons is called 'detail view', the one in the bottom window is the 'normal' (default) one. A hierarchical tree view is also possible.

DFM also places on the root window its own root menu, reachable with one right-click. With *Blackbox* 0.65 this only caused a moment of surprise, but it might interfere with

other Window Managers. Apart from that, however, *DFM* is declared to be conforming to the X Drag and Drop protocol (XDND), so there should be any problems interacting with other XDND-enabled X clients (basically all those that are designed properly...)

The online help for *DFM* is a shortcut to open an xterm which displays a man page. This one is quite synthetic (aren't they all?) but does give a general idea of the *DFM* way of doing things. All in all, *DFM* is a good choice for users who have limited resources, simple file management needs (or the know-how to perform the more esoteric operations at the command line) and still want shortcut icons.

VERDICT

Features	7/10
Ease of use	6/10
Documentation	8/10
Performance	7/10

Good for those with limited resources, and definable window colours simplifies working, but where's the 'wow' factor?

LINUX FORMAT RATING
 7/10


FileRunner

Makes repetitive operations on files a snap.

VERSION 2.5.1 WEB www.cd.chalmers.se/~hch/filerunner.html

With an executable name of *fr* it is written mainly in Tcl/Tk, with some C parts, and aims to be as fast and efficient to use as the command line, even for experts. To accomplish this, it features the common two pane

design, with some interesting additions. Each pane shows the content of one directory, with a lot of information (not all visible in the screenshot) accessible through the horizontal scroll bar right below the file listing.



FileRunner's mix of both an entry box and buttons is innovative.

Navigation is fast: any directory in the file system is reached immediately through the cascaded menus which are opened by pushing 'graph tree' buttons above each pane. FTP is supported natively, either in anonymous mode or with password, and, if the link goes down, it can be resumed from where it stopped. *FileRunner* can also download files from the web, directly or through a proxy. The most common operations (move, copy etc) are all directly linked to the buttons in the scrollable middle vertical row. Archives can be created or unpacked in the same way, with one or two mouse clicks, but perhaps the most interesting button is the one labelled 'ForEach'. As explained in the pop up menu, whenever one needs to perform some command line operation on a group of files, eg compressing, checking if they contain a particular string, or converting all GIF images to PNG selecting those files and pushing that button will give you a prompt where the desired command can be typed.

All commands that are not bound to buttons, and not iterative can be typed in the entry box at the very bottom of the window: their output will be displayed in the horizontal area below

the list boxes. The Help menu includes Quick Start, User's Guide, Installation, FAQ and Tips.

The configuration of *FileRunner* is not so rich, but based on an interesting mix of mouse clicks and manual editing. Edit Configuration opens another two pane window: the left side shows all options, grouped by type. Clicking on one option makes the right pane view jump to the point in the text configuration file where that option is defined, and the user must type the new value in the right place. Since each option is properly commented, there is very little need to know Tcl well, but advanced users can still tweak things.

VERDICT

Features	6/10
Ease of use	5/10
Documentation	8/10
Performance	7/10

Simple but with all the basic pieces available. Recommended for desktops on a budget, and to Tcl/Tk hackers

LINUX FORMAT RATING
 6/10

RoundupFileManagers

Emelfm

GUI-only configuration and easily added custom commands in one package.

■ **VERSION** 0.9.2 ■ **WEB** <http://emelfm.sourceforge.net/>


According to its author, *emelfm* is essentially a compilation of features he likes most in other FMs in this roundup, namely *FileRunner*, *sfm*, *Gentoo* and *Midnight Commander*. The default appearance is very similar to *FileRunner*, even if here the toolkit is GTK: two side-by-side panes (each one can expand to fill the whole window), separated by a column of buttons, and a lower area shell prompt, plus a sub window which can be hidden) to show the output of all the commands typed at that prompt, or bound to any *emelfm* button.

Even in *emelfm* the info available about each file is more than it fits in the listing box, and must be reached with the horizontal scrollbars. Columns can be added or removed, but their ordering cannot be changed: clicking on any column header sorts by that field. Files can be filtered by name, date, max or min size, from the corresponding menu.

There is bookmark support for the most frequently visited directories, and history depth is configurable. Navigation is via the history scrollable list on top of each pane or by clicking each directory. To reach the parent directory or your \$HOME instead – left or, respectively, right click on the folder icon.

There are several features useful to the system administrator, or just to the sophisticated user: First of all, the default User Commands menu includes entries like 'make a patch', 'find which RPM this came from' and 'Find SUID and SGID files'. In the second place, the row of small buttons in the lower right corner includes by default 'free' (to show memory usage) and 'su' (to execute commands as root, obviously after entering the password).

The configuration of *emelfm* is completely GUI-based: there is not much to play with as far as look and feel



emelfm starts with all the right buttons, and adding new ones with the particular functionality that you're needing is just one button away.

are concerned, but the actual functionality of the file manager can easily be configured or extended. Adding custom user commands to the already mentioned menu, for example, is very simple, and well documented. The same is true for the buttons in the central column, or for the small ones (Toolbar buttons) close to the shell prompt. Teaching *emelfm* how to recognise (from the extension) new file types, and to open them with the proper application, is also easy and fast.

VERDICT

Features	9/10
Ease of use	6/10
Documentation	5/10
Performance	8/10


A light prog perfect for users which like complex shell one-liners...but don't want to type them out in full.

LINUX FORMAT RATING
6/10

TkDesk

One of the first integrated file and desktop managers.

■ **VERSION** 1.2 ■ **WEB** <http://tkdesk.sourceforge.net>



TkDesk had a panel (partly shown left) years before KDE and GNOME.

TkDesk depends from the *incrTCL* library (*itcl*). At the time of writing, the current stable version of *TkDesk* (1.2) doesn't compile out of the box with the

latest version (3.2) of *itcl*. Therefore, for this review we used a Red Hat 7.2 system. The problem, however, has been brought to the attention of the *TkDesk*

developers, and it should hopefully be solved by the time you read this.

By default, this file manager comes with a File Browser Window, and a panel (Application Bar). The first one shows simultaneously three levels of directories opened. When more levels are active, the horizontal scrollbar can centre the view on the ones currently used. The user can quickly travel back and forth multiple levels of hierarchy in one window, and, when working inside a directory, look at the same time inside it, its parent, and one subfolder. If needed, more Browser Windows can be opened. A third separate List Window is basically equal to the listings shown in the figure, but covers only one directory. Right-clicking a listed folder makes available sub menus of all its subfolders or files, plus entries to archive the whole folder as a *tgz* or *zip* file with one click.

The Commands top menu manages, besides files, cron jobs and environment variables, not to mention (as packaged in Red Hat, at least) an RPM sub menu which lists installed packages.

The Help menu points to searchable copies of the User Guide and FAQ, and

a man pages viewer. The panel has several more buttons than shown here, and can be closed, like the KDE one, or dragged to any position on the desktop. Each panel button provides a pop-up menu with shortcuts to several actions. As the 'info' sub-windows show, files can be dropped directly into several buttons to be opened or printed. Style, orientation, and button combination of the panel can be changed, of course: just right-click on the first button, which accesses history, configuration, and other common actions. *TkDesk* can also place icons on the desktop.

VERDICT

Features	7/10
Ease of use	7/10
Documentation	6/10
Performance	7/10

If it were easily compatible, and had *bzip2* and mounting of external devices, it would be almost perfect.

LINUX FORMAT RATING
7/10


Gentoo

Even simple programs can make plain file listings look nice.

■ **VERSION** 0.11.34 ■ **WEB** www.obsession.se/gentoo/

Gentoo is written in C and uses the GTK toolkit. It works on several Linux platforms, Solaris, and on BSD systems. At the look and feel level, its declared goal is to borrow from the Amiga file manager *Directory OPUS*. Consequently, its appearance and file icons are more refined than those offered by other light

FMs. Below the surface, however, there are some common points, eg *FileRunner*. The two-pane layout is still there, even if, inside Gentoo, the two panes can have different widths. Each pane shows a lot of info on each entry: a horizontal scroll bar is still used to move back and forth the several columns. One



The graphical appearance of *Gentoo* is highly configurable, and all commands are easy to spot.

nice plus with respect to *FileRunner* is that the content of a folder can be sorted by any of the displayed parameters (size, owner, modification time and so on) just clicking on the column name. It is also easy to add or remove fields from the listing boxes, or change their order. One entry that can be valuable in some situations is the number of links that a file has.

There is no FTP support, or Help button (even if the man page is updated and quite complete). *Gentoo* offers three 'bookmarks' buttons, which can be pointed to any set of directories. Navigation may be slower than in other managers because there are no hierarchical menus. In order to descend inside lower levels of hierarchy, you must proceed one step at a time, clicking inside the pane. In some situations, however, for example with smaller screens, this may be preferable to multiple level of menus.

Gentoo natively recognises many different file types organised hierarchically: packages (Debian, RPM, tarballs...) all kinds of multimedia (images, audio, video), source code of several languages, devices (FIFOs, Sockets) and more. Each file type can be identified in various programmable ways (using extension, regular expressions, and/or the 'file' command)

and can then be associated with a style. In *Gentoo*, a style defines in detail, for each file type, what icon should be associated with it, how to list it in the panes, and how it should be edited or displayed. The *Gentoo* package includes about 100 original icons.

Two interesting *Gentoo* buttons are *RenameSeq* and *Split*. The first renames all selected files: prefix, suffix; numbering criteria (start number, decimal, octal or hex base) can be defined by the user. The second one splits a file in smaller parts (a lifesaver when the only 'network connection' is a box of floppy disks...), but the number of parts, their max size and naming are still defined by the user. All the predefined actions listed in the configuration panel can be performed on a selected file via the Run menu.

VERDICT

Features	8/10
Ease of use	7/10
Documentation	6/10
Performance	8/10

Easy to use and good looking without being heavy. The hierarchy navigation would benefit from tree buttons like the ones offered by *FileRunner*.

LINUX FORMAT RATING
 7/10

ROX

An alternative way to build a whole desktop.

■ **VERSION** 1.3.5 ■ **WEB** <http://rox.sourceforge.net>

Strictly speaking, the ROX File Manager (ROX-Filer) is only a part of the *ROX* desktop, which also includes a panel and other goodies. Users are not forced to invite the whole family though: the *Filer* is still fully useable (and quite a nice program) all by itself. Visually *ROX* is very clean and good looking at the same time. The design slogan was "no clutter": menus and prompts are just one or two clicks away, but show up only when explicitly called. Out of the box, the *Filer* already has a nice collection of file icons which make very easy to spot all the files of the same type, and extra collections are reachable from the website. Several pieces of documentation are directly reachable from the 'lifebuoy' icon.

ROX is also the first (but of course not the only one) of the programs reviewed in this article to offer thumbnails: since it also supports drag and drop, it is an excellent solution, especially if you cannot or don't want to use a full blown GNOME or KDE desktop, to manage image galleries. Sorting five years worth of holiday pictures is a snap with *ROX*. The speed is very satisfactory: the only task which requires a noticeable amount of time, (several seconds) at least on relatively old machines, is generating the thumbnails of a very large image directory.

This doesn't make the *Filer* unusable: when the directory is opened, all the images are immediately shown with the default

image icon and their name; the generation of thumbnails starts only after that phase (and only the first time a directory is opened), but doesn't freeze the program. Finally, since it can be turned off, it impacts performance only when you allow it.

The *ROX* filer, and its companion applications, also offer a couple of features of the kind which are defined "mission critical" or "totally irrelevant" in equal measure, depending on the user's need and their style of computing.

The first one is drag-and-drop saving, which is introduced, with its own protocol (XDS), at this address: <http://rox.sourceforge.net/xds.php3> DND saving is defined as either the natural complement, or the next step, of DND loading. The latter is what we do when we drag a file from a File Manager window and drop it on an icon representing an app, or some device (usually a printer). An XDS-enabled application, when the user wants to save a file, will not provide some file system view to reach the directory where the file should be saved. It will

just offer an icon that can be directly dropped in the *ROX Filer* window.

The second feature is native is support for Application Directories. Such directories contain an entire application: documentation, binaries, code and so on. To run that app, the user just opens its application directory in the filer; installing or uninstalling it is the same as copying or deleting a directory, and different versions can coexist without any special action. To know more, visit: <http://rox.sourceforge.net/appdirs.php3>

VERDICT

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

The *ROX Filer* may need its whole desktop to give its best, but it is nevertheless complete, fast and excellent for managing picture galleries.

LINUX FORMAT RATING
 8/10

Konqueror

The most used piece of KDE gets better and better.

■ **VERSION** 3.03.13 ■ **WEB** www.konqueror.org

Konqueror is probably the most used application of the KDE desktop. It is simultaneously an advanced web browser, an ftp client, a documentation viewer, a couple of more things which escape us at the moment, and a file manager in the more traditional sense.

From this last point of view, Konqueror has a lot of features, a nice, very themeable look and together with the rest of KDE, some spiffy audio and

visual effects. Don't worry, however: apart from the bell and whistles, it does a lot of concrete things too. For example, a very important *Konqueror* feature, obviously common to any recent KDE application, is support for bidirectional languages.


Back to managing files now: looking inside many places of the file system is easy. Any *Konqueror* window can be split in many different ways,

and each part can display in a different way the content of a whole directory, or of a selected files. We are not talking of thumbnails now, even if, of course, *Konqueror* does have them: another 'official' definition of this program is 'universal file viewer'. This may or may not be what an user expect from a file manager. In practice, it means that it can call for help other applications, or, much more frequently, *KOffice* modules, to show a spreadsheet, a web page, or many other things, inside its own window.

One last note about thumbnails: please have a close look at the 'LXFstyleguideROUNDUP.txt' icon below in the screenshot. Had you realised that what is inside it is the

actual text of that file? It adds a whole new dimension to "file preview", doesn't it? Imagine all those repositories of documents like company memos, Internet RFCs, and such: meaningless semi numeric names, and no clue about the actual content. Something like this makes much easier to find what you need without grepping your way through the whole folder.

In general, each *Konqueror* sub-window can show an icon view (right half of the screen) or a tree view (left half). This makes reordering pictures across multiple folders even faster (and with less screen cluttering) than with *ROX*. Drag-and-drop is supported even between different hosts, ie via FTP. Another thing that will probably make many users happy is the updating of directory views is automatic: no need for a Refresh/Reload button to check which files where created or deleted.



Konqueror offers a panoramic vista of the File System and a detailed view of a single folder in the same window.

VERDICT

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

Excellent file manager, and much more: consider that, even without the rest of KDE, and though heavier than some, it can make other programs redundant.

LINUX FORMAT RATING
 9/10

Nautilus

A file manager that understands *your* criteria.

■ **VERSION** 2.0.6 ■ **WEB** www.gnome.org/projects/nautilus/

Nautilus is defined as "The GNOME file manager, desktop, and graphical shell". Its goal is to provide one common interface to access both local files, applications or devices, and all kinds of Internet resources.

One cool feature is the 'emblem' concept. Emblems are additional flags that the user can attach on each file or folder, to classify them on any number of criteria. As the screenshots show, emblems can tell very different things, from where a file came from (Camera) to its status (Draft) or usage (Personal): the nicest part is that emblems can be also given to *Nautilus* as search criteria!

Folders and files can also be associated with custom text notes, and a history of the most recently visited folders is available.

There are two features especially helpful for users dealing with large image galleries. *Nautilus* obviously supports thumbnails views of images, but, on top of that, every thumbnail can be separately enlarged (not opened in the whole window!), to give a better idea of its subject. When it comes to image galleries, every folder can be displayed as an icon of one of the images it contains, anticipating content better than any text string. Even *Nautilus* extends

the thumbnail concept also to non graphical files, ie shows the first lines of HTML and text files inside their icons.

Nautilus supports the binding of scripts to custom menu entries. This makes it possible to extend the functionality of this already powerful tool with any kind of tasks, from archiving folders to making web galleries. Once the script has been created, placed in the proper directory (usually `~/nautilus/scripts`) and made executable, you can select any number of files and send them to the script itself by right-clicking on them, and reaching the script through the corresponding menu. The central repository of *Nautilus* scripts is <http://g-scripts.sourceforge.net/>

Another interesting *Nautilus* extension, very handy for programmers, is Apotheke (<http://apotheker.berlios.de/>), which gives to *Nautilus* a graphical interface to the most common operations on CVS repositories (CVS is

the Concurrent Version System normally used to keep large code bases under control. There is also an alt.binaries.* viewer (the NNTP module, found at <http://www.gnome.org/projects/nautilus/nntp.html>) which allows previews of the images and music files posted to those newsgroups, and, to some extent, can group related files together in directories.

VERDICT

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

New version is faster than before. The custom folder icon, emblem system, and support for external scripts are great.

LINUX FORMAT RATING
 9/10

FILE MANAGERS THE VERDICT




Konqueror is certainly the most popular application in this roundup.

Declaring any absolute winner for all scenarios would be impossible. Many of the candidates are optimised to manage different kinds of files in different contexts, eg admin versus leisure use. Apart from that, some programs are completely independent products, while others, namely *ROX*, *Konqueror*, *Nautilus* (and, partly, *TkDesk*), were designed from scratch, and give their best only as parts of more or less complex environments.

For these reasons the products have been grouped in three categories and ratings given separately inside each one. The first one includes console products focused on (possibly remote) system administration: in this space, we would pick *Midnight Commander* for the features, and *UnixTree* for the interface.

At the opposite end of the spectrum there are those (relative) heavyweights like *Konqueror* and *Nautilus*, perfect for end users with more or less recent hardware, who not only deal almost exclusively with a lot of images, sound, web pages and office documents, but also want to look at them right away. In this category, and obviously limiting the judgement to file management, we prefer *Nautilus*, because the emblem system is a pleasure to use.

All the other programs are meant to simply manage files in a window-based interface, (and, in some cases, launch programs) without looking at their content directly, and without slowing the system down: the exception here, as far as images are concerned, is *ROX*, which is however in the middle between this and the top category. We choose *ROX* if



UnixTree sub windows can be arranged in many different ways.



The default look, and the Documentation sub window, of the ROX Filer.

image galleries must be managed, otherwise *emelfM* for its easy way to add custom commands.

All the programs have been tested, on purpose, on two low-end systems

with 96MB and, respectively, 128MB of RAM, *not* running KDE or GNOME, with nothing else going on in the system but X and *Blackbox*. In all cases, the biggest speed difference was in the start up time, from almost nothing for *UnixTree* to more than 10 seconds for *Nautilus* and *Konqueror* (but these two times would be considerably smaller with GNOME or KDE already running).

One final note: should we ask for more from our File Managers? In our not-so humble opinion, yes. It would be great to see the *Nautilus* 'scripts' approach to be generalised and, above all, *shared* by all File Managers, through some common interface specification. "Make one tgz archive of this directory" can be accomplished by a short shell script. What if there where a standard syntax and location for File Manager Plugins, so that every new or existing file manager could just point to them? [LXF](#)

Table of features

Note: packages sizes and memory stats below are approximate. Memory usage varies from system to system, and is affected by the program's configuration and other factors. All programs are available under GPL. Konqueror package size varies depending from it being bundled with KDE or not.

Name	Toolkit	Pkg size	Avg mem use	Configurability	Alternatives
Midnight Commander	N/A	3.8MB	540KB	high	UnixTree
UnixTree	N/A	1.6MB	50K	medium	Midnight Commander
sfm	GTK	110KB	700K	very low	FileRunner
Kc	Tk	200KB	3.5MB	low	FileRunner
DFM	GTK	770MB	3MB	low	Kc
FileRunner	Tk	400KB	4MB	medium	Kc, emelfM
emelfM	GTK	250KB	4.5MB	high	FileRunner, Gentoo
TkDesk	Tk	1.1MB	6MB	high	ROX
Gentoo	GTK	1.6MB	5.5MB	medium/high	emelfM
ROX	GTK	2MB	7MB	medium	TkDesk
Konqueror	Qt/KDE	12/40MB	8MB	very high	Nautilus
Nautilus	GTK	17MB	15MB	very high	Konqueror

HotPicks

The best new open source software on the planet!



Mike Saunders

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

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

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

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

HotPicks at a glance

Kopete 0.6.0	42
Linphone 0.10.0	43
Netmap 0.1.2b	43
LSHW T.00.05	44
Evolvotron 0.0.3	44
CGI-shell 0.18a	45
Transcode 0.6.3	45
Starfighter 1.0	46

HOTPICKS AWARD


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



KDE INSTANT MESSENGER

Kopete

■ VERSION 0.6.0 ■ WEB <http://kopete.kde.org/>



The contacts list, user management menu and prefs box in action.

Back in issue 36 we took a close look at Linux instant messaging software in our roundup. At the time, the GNOME-based *GAIM* app came up trumps in providing a slick interface and competent multiple-protocol support – there was little to really challenge it, and the KDE desktop was very much in need of a rival. Four months since that roundup was written, though, *Kopete* has emerged from its early development state to become a fully-fledged IM client with superb multi-protocol support and a spiffy interface to boot.

As it's built around KDE 3, you'll need the appropriate development libraries and Qt 3 to compile – fortunately, binary packages are available for most major distros as well. *Kopete's* developers claim that the app will be integrated into the upcoming KDE 3.2 release (in the *kdenetwork* package), further strengthening the desktop suite's array of Internet tools.

Boasting support for *MSN*, *AIM*, *ICQ*, *Jabber* and *Gadu-Gadu* (although

no *Yahoo!* at present), *Kopete's* range of protocols is up there with *GAIM's* and makes it an ideal choice for those with accounts on multiple networks. In addition, support for IRC, SMS text messaging and the *WinPopup* system found on Windows networks is included; many of these are still under heavy development though. Each protocol can be loaded on demand using plugins – a wise design choice, as it allows new networks to be supported without recompiling the whole application.

Kopete's main window presents a familiar layout, with a tree list of contacts and a configurable toolbar (which can show icons, text or both). At the bottom right is a taskbar-like tray for the different networks you're using, with icons for each – you can check email, become "Away" and sign off with a click. It's an attractive interface (especially with the translucency option for chat windows), and the ability to customise toolbars and bindings heavily will be welcomed by power users.

Support for the various protocols is managed through the Plugin section of the configuration dialog – as each is loaded, a new option is made available for setting up user accounts and so forth. Similarly, some other useful plugins can be found in the main package, including a translator, GPG encryption, and connection status monitor (by *LXF* writer Chris Howells no less!). There's even an somewhat bizarre "motion detector" system which marks you as Away if it sees no action in a video device.

Other options

Other configuration options include contact list sorting, away messages and fonts. A nice touch is the choice of pre-set colour and typeface schemes for the chat window – it offers IRC, *MSN Messenger* or *Kopete's* own colours and layout options. Just about every aspect of *Kopete* can be fine-tuned, but fortunately the default settings are clean and usable, and those with only Windows client experience should have no problems adapting.

Kopete can use tabbed chat windows, either grouped by protocol or all in the same box – good for saving screen space and making it easier to find a conversation. Emoticons and file transfers are supported too, and sound alerts can be altered for a few events (WAVs, MP3s and even Oggcs can be used!). Right-clicking contact names on the list brings up options to manipulate groups, add a comment for the user or view a history of his/her messages.

In the works for the next release is *Yahoo!* support, a Perl script plugin, further development on the existing protocols and some cleaning up in the program's back-end. On the whole, we're hugely impressed with the progress that *Kopete* has made, and it's good to see a new app which can confidently compete [groan] with the mighty *GAIM*. With its superb feature set, great multi-protocol support, fast and stable code and plenty of configuration options, *Kopete* wins our Hottest Pick award this month. If you're a KDE user and IM regular, install it post-haste, and get chatting!

INTERNET PHONE CLIENT

Linphone


■ VERSION 0.10.0 ■ WEB www.linphone.org

With even most Internet dial-up connections being unmetered nowadays, it's very useful to extend the versatility of the Net beyond mere email and Instant Messengers. The concept of making voice calls over the connection is becoming more popular – this is a particular boon for those on single-line modem dialups, as it enables them to talk in real-time with friends and colleagues without having to hang-up. Linphone is one such client for online speech calls.

Developed in France and built around the GNOME 2 desktop (although a simple text-based version can be compiled from the source for console users), *Linphone's* RPM packages are converted from Debian .debs with *Alien*; this led to all sorts of problems on our test machines, so we

built it from the tarball. It uses "SIP" (Session Initiation Protocol) from the IETF to communicate with other clients – these include *eStera* and *Hotsip* on the Windows side, making it useful for communicating with non-Linux users.

With a properly configured sound card and microphone to hand, *Linphone* can be started as a typical GNOME program, an applet in the Panel, or as a background daemon to wait for an outside connection. To test the setup without bothering anyone, a "sipomatic" tool is provided which emulates another client over the loopback device. From there, making an Internet call is a simple matter of passing the correct SIP URL (ie sip:user@hostname) and waiting for a response – a connection is made to port 5060 on the remote machine,



Here's hoping that those tenacious double-glazing telesales people don't discover *Linphone*, or none of us will get any peace ever again...

and the other user can accept with a click and begin talking.

To provide adequate sound quality over slow links, a number of codecs are present for compressing the stream (Speex, ALAW, GSM and others). Along with the address book, another sweet feature is the ability to register with a SIP server, making it

easier for others to find you if you don't have a static IP – especially useful for broadband users. Ultimately, while *Linphone* proved to be usable in our testing the varying speeds and latency problems that plague slow links didn't help. Definitely worth experimenting with if you have a good Net connection though.

GRAPHICAL NETWORK MAPPER


NetMap

■ VERSION 0.1.2b ■ WEB <http://netmap.sourceforge.net>

Generally, most network administrators prefer to stay away from depending on graphical tools and instead get their hands dirty in text files. A point-'n'-click interface can provide a good helping hand at first, but later on the ability to grep, sed, diff and copy around config files and scripts is priceless. Still, occasionally a tool pops up where graphics are a necessity: enter *NetMap*.

Using traceroute, grep, awk and an external utility called *dotty*, *NetMap* creates a fully graphical representation of the layout of your network. Although the main tarball itself is small (the core being a C++ util) you'll need to have *GraphViz* installed – it contains *dotty* which is responsible for the final display. Compiling is a piece of cake, with no unlikely dependencies.

To create a map, the first required step is to edit *makelist/hosts.txt* and enter the hosts (one per line) you wish



Some routes out from a dialup – with more, it can become wildly complicated – and the authors want to map the whole Internet...

to pass to traceroute. The example file contains sites from all around the world, producing some interesting output. After that, running *netmap/netmap* generates the resulting image, which can be viewed with *dotty netmap.dot* (it's a structured plain text file format).

The coloured output is easy enough to understand, as you can see from the screenshot below. A yellow

parallelogram represents the initial node, with increasing intensities of red to signify the number of links from a host (ie light yellow for 2, strong red for 6). Standalone hosts are represented by grey ovals, with unknowns having a bright red triangle. Within each shape is the host's name, IP address and average ping time. *NetMap's* coders hope to write a system where users submit their own results, eventually creating a map of the entire Internet – a mammoth undertaking!

Although the dependency on the larger *GraphViz* package isn't ideal and colours and shapes can't be modified by users without directly hacking the source, a few command line options allow for some remedial tweaking here and there – input and output filenames can be changed, and you can specify the maximum number of hops and prevent it from trying to resolve hostnames. In all, *NetMap* could be a helpful little tool in getting a visual overview of a network and identifying slow nodes and trouble spots. We'll watch future development with interest.

HARDWARE INFO LISTER

LSHW

■ **VERSION** T.00.05

■ **WEB** <http://ezix.sourceforge.net/software/lshw.html>

Detecting and configuring hardware under Linux has always been a difficult task – not only do some x86 devices react badly to being poked and probed, but manufacturers that don't understand the new profit-generating advantages that are to be had from following an Open Source methodology aren't always overly keen to reveal the workings of their products either, making accurate peripheral device detection even harder than it has to be. Most mainstream distribution vendors have written code to try to make this better (vitally important in the installer), and now *ezIX* is another attempt at an ultra-easy distribution.

It's in the very early stages of development, but the coders behind it have already managed to create an excellent little detection tool – *LSHW* (*List HardWare*). Needing nothing but

the G++ compiler to build, the program works on both x86 and Power PC hardware, requiring the 2.4 kernel for full functionality. It aims to provide a comprehensive listing of all hardware and devices in a box – from the CPU and memory through to vendor names for individual devices.


Most of this functionality comes from running the program as root, which is necessary to probe the BIOS and extract the detailed DMI (Desktop Management Interface) information. This is presented in a tree format at the prompt; the indentation makes it sufficiently readable, with copious amounts of info for each device, bus and component.

On our test machines, it did a superb job of plucking out the vendor names and chipsets on each device plugged in, and could pick out fine detail such as the CPU's serial number

and cache memory levels. With the `-html` flag, *LSHW* outputs a plain but easier-on-the-eyes Web page for viewing in a browser or sending on to someone for help.

For such a small tool, *LSHW* could be a life-saver if you've picked up a second-hand box for a bargain and aren't sure of its contents, or are perhaps recycling donated equipment

for a school or community project – after running this, you'll know exactly what's present and can tune your kernel to suit. The hands-on method of poking around in `/proc/pci` and the like isn't as elegant and is more limited in the info it gives, and if you're sending in a bug-report which is hardware-specific, *LSHW's* format will be satisfyingly thorough to any coder.




Pastel-coloured HTML output, and the basic terminal equivalent, from one of our test boxes. Everything you need to find out what's what!

RANDOM ART GENERATOR

Evolvotron

■ **VERSION** 0.0.3

■ **WEB** www.bottlenose.demon.co.uk/share/evolvotron/



My eyes! They're burning! Won't somebody please think of the children?

Hieronymus Bosch liked to spend his Sunday afternoons painting outrageously disturbed visions of hell. Sinister creatures eating and mutilating one another against a burning backdrop is all very well, but those of us lacking in talent at the palette have to find other ways to make interesting imagery. Mandelbrot picture-generators have been around since the first colour displays, and now Evolvotron attempts to go a step further with "an iterative process of random mutation and user-selection driven evolution". Cripes.

You'll need the Qt 3 toolkit and the *qmake* program to build *Evolvotron* from source – the small application is written entirely in C++ and is multi-threaded for best performance on multi-processor boxes. As some of the larger images can put massive drain on the CPU when being generated, this is a good thing. You can also set the window size, number of tiles and threads to be used on the command line, so cutting it down to make it run on older boxes is possible.

When initially started, *Evolvotron* presents a grid of small gradient

images; nothing spectacular at first, but a few clicks in the more interesting tiles creates a redraw of the grid with new variations on it. It's best to choose those images with more detail and colour range to develop the rest – within a few clicks, interesting shapes, curves and colour blends start to appear in the cells, while "heating" and "irradiating" them exaggerates the pace of mutation. Equally, should the variations become too wild for your tastes, "cooling" and "shielding" will slow the process down.

Once a particular image has taken your fancy, a right-click allows you to expand the cell into a full-size picture (up to 4096x4096) and then save as PNG or PPM. Patience is required though – it can take an age to render for some of the more complex images. Similarly, individual cells can be "locked" to prevent them from changing, and others can be recoloured without changing their shape. The docs explain in detail how the mutations work, and anyone looking for some unique and pseudo-hand-crafted backdrops for their desktop should play around with it.

SHELL FUNCTIONALITY

CGI-Shell

■ **VERSION** 0.18a

■ **WEB** www.binaervarianz.de/projekte/programmieren/cgi-shell/


CGI-Shell, as the name suggests, intends to provide users with shell account functionality on a web server via CGI (in this case Perl). Although like many progs we feature, it's still in the early stages of development, it's already a fully usable and intriguing utility – the author states that “it intends to help webpage-owners to maintain their page comfortably” and provide an incentive for the host to enable SSH as well.

To run CGI-Shell, you'll need the LWP and Term Perl modules installed; they'll either be in the base Perl package supplied with your distribution or available as add-ons, but Red Hat users should note that they'll also need the perl-TermReadKey RPM installed as well

for it to work. After extracting the archive, you'll need to run `./cgi-shell -c` and enter a username and password to generate `.htaccess` and `.htusers` files for authentication, and then copy those files and the `cgi-shell-server.pl` file to the `cgi-bin/` directory.

With the server side of things now prepared, the user runs `cgi-shell http://www.myhost.com/cgi-bin/`, enters the username and password, and is presented with a familiar prompt. Impressive stuff you'll agree, and the shell program itself sports basic tab-completion and stored history – so it's pleasant and comfortable enough to work with.

Unfortunately, though, as only a dumb terminal environment is available, any programs which require user-interaction after the initial



Running interactive apps is generally a no-no, as seen here with VIM.

command are unlikely to work properly. This naturally prohibits the use of text-mode mail clients, editors like Emacs and Vi and other tools regularly used on shell accounts you'll be familiar with, but can also cause trouble with more critical commands like `passwd/chsh` etc.

On the whole, CGI-Shell serves its intended purpose of providing a simple system for performing the most basic

operations, and is an adequate alternative if you haven't got SSH access to a box. Of course, the lack of any encryption brings it down to Telnet's level of security (the author maintains that he's working on this though), and there's the aforementioned problems with running any ncurses-type programs, but it still does a decent job within its constraints.

VIDEO STREAM PROCESSOR

Transcode

■ **VERSION** 0.6.3

■ **WEB** www.theorie.physik.uni-goettingen.de/~ostreich/transcode/

As the number of video formats continues to grow – some open and documented, some proprietary and some too obscure to find decent players – there's been an immense effort by the free software community to make sure that they're all supported as much as possible. MPlayer is notably good in this respect as a viewer, and Transcode aims to be a flexible solution for altering video files and converting between a multitude of formats.

When building from source, the configure script detects which libraries you have available on your system and tailors the module set for it. The result is a bunch of command-line tools and modules for modifying video streams – de-interlacing, resizing (using assembler rather than C for pure

speed), flipping, mirroring, audio resampling, gamma correction and more are present.


For converting, support for DVD, MPEG, uncompressed AVI, the various DivX flavours and plenty of others through the supplied (or external) modules is available, with the `avifile` and `quicktime4linux` libraries giving a helping hand where needed (they're not distributed with Transcode itself in case of potential legal issues).

In a nutshell, Transcode pulls in the initial video data through its import modules, passing on the resulting stream to the main program for alteration, and then writing it back out through the encoding export modules. As such, it's ideally suited to converting DVD movies into smaller formats for CD-Rs, pulling the sound from a video

file, polishing up those old home movies and converting awkward formats into something with better Linux support. It performs very well and proved to be stable and robust in our extensive testing.

Without a doubt, the sheer number of options and tunable settings in Transcode will make it extremely daunting to the newcomer; consequently, we hope to see some Qt

or GTK+ front-ends developed soon, either as additional packages or included with the source. Still, there's no denying its versatility, the documentation is plentiful and precise, and the range of supported formats and effects that can be applied is astounding. Pop this month's LXF coverdisc into your box's drive and check it out – just be prepared to consult the docs quite a lot.



Transcode's pictorial representation of possible format conversions.



"What happen? Somebody set up us the bomb..." etc.

SPACE SHOOTER

Starfighter Demo

■ **VERSION** 1.00 ■ **WEB** www.parallelrealities.co.uk/starfighter.php

The blurb reads "after decades of war, one company who gained power supplying both sides with weaponry, steps forwards and crushes both factions in one movement. Using far superior weaponry and AI craft, the company was unstoppable. Thousands perished under its iron fist. The people cried out for a saviour..."

Sounds like Nostradamus' prediction for Microsoft's future? *Starfighter's* plot sets the scene for some fast-action space-shooting shenanigans. In the late 80s and early 90s, scrolling blasters were the staple diet of gamers, reaching their zenith in the face-smashingly hard but tremendously enjoyable *Thunderforce 4*. The cathartic effect of pumping ammo into beastly enemies, the unrealistically enormous end-of-game bosses – great shoot-em-ups are always remembered.

Amiga fans will may recall Parallel Realities' *TANX Squadron*; now the two coders behind it are bringing talent to the Linux gaming scene with this demo of *Starfighter*. Using the SDL libraries

for graphics and sound output, the game is fairly portable and easy to compile from source – versions for Windows and BeOS are downloadable.

As a young pilot going by the name of Chris Bainfield (makes a change from "Zlax O'Quasitron"-style character names), your goal is to battle through various missions before facing a showdown with a boss. Each task takes place above a certain planet – once finished, it becomes your temporary base and an opportunity to stock up on weapons or improve your Firefly ship.

Being only a demo version, though, which much still to be done, you're limited to four main missions and an encounter with a hefty enemy at the end. Most tasks involve two main objectives: a specific kill-and-collect run at the start, occasionally within a time limit, and then an all-out massacre on the remaining enemies. Pleasingly, effort has been made to ensure variety within the missions – the cargo-collecting job, for instance, requires careful control and precision

accuracy rather than mindless CTRL key hammering, and the wandering mines in another task call for extreme dexterity on the cursor keys. This is always the hardest gameplay aspect to get right in the "one-quick-go" genre – always presenting the player with something new and different, and *Starfighter* does a commendable job.

The Firefly handles with appropriate inertia, responding well to the controls and never falling into the dreaded "I didn't move there!" trap. Additionally, collision detection is spot-on so you never feel cheated, and the enemy craft dart around with a decent amount of unpredictability – things can get hair-rippingly hard when you're being

Recommended specs

No minimum system requirements are given for the game in the docs, but we'd say that a 300MHz or better chip with 32MB RAM will be more than sufficient to keep things smooth. You'll also need an display with 800x600 resolution or upwards – *Starfighter* can run in a window or in full-screen mode (settable in the options menu). A good sound card and 'phones or speakers will improve the sonic experience too, especially if you like pumping techno...

attacked from all sides. Still, the difficulty level and learning curve is well balanced for a short game.

Starfighter doesn't aim to be a graphical showcase but the sprites are well defined and clearly drawn (though the over-scaled final boss in the demo could do with some more attention we reckon), and the Anime-style artwork for the other characters is especially cute and fits in well with the overall presentation. The techno soundtrack isn't totally inspired but works fine, as do the crisp, aggressive sound effects.

In all, Parallel's stab at the Linux games market is to be applauded – it's polished, playable and great fun all-round, and we're very much looking forward to the final version. Those essential ingredients from classic shooters of the past are all present: frantic firing, swarms of enemies, weapon boosters and adrenaline-pumping background music.

Naturally, if you prefer something more cerebral than *Starfighter* won't float your boat, and it's still too small to last you more than a few hours continuous play at most, but as a well-designed quick-action blast fest it really deserves a look or three. [LXF](#)



After a hard day's shooting, it's a good idea to buy ammunition or enhancements for your ship if you want the following day to go as well.

CustomiseLinux



Create the Linux environment you want by taking advantage of the flexibility of Linux applications and environments. With expert tips from *LXF* regulars, it's time to give your desktop a makeover.

cover feature



Authors

This feature is based on the work of:

Jono Bacon, Richard Drummond, Paul Hudson, Mike Saunders, Nick Veitch

Linux the

One of the reasons for Linux's very existence is the freedom to choose. This philosophy has been very much at the heart of free software development since before Linux was even a few lines of code. Freedom in this sense means not only the freedom to look at source code, but the freedom to choose which applications to use and how they should work. The reason there are more than a dozen different email clients is that there are more than a

dozen groups of people who like to use their email in a different or particular way, or have different needs, whatever they may be.

This explosion of choice, and the general levels of interoperability between Linux software means that your own particular Linux installation is probably pretty customised already, merely by dint of the packages you have chosen to install.

But the average Linux developer has taken the concept of freedom of choice a bit further. Not only can you

KUSTOMISING KDE

From the *Kicker* to *Konqueror*, change almost any aspect of your desktop.

The KDE project has been around over five years with the express aim of creating an easy-to-use, flexible desktop environment. The project has not only achieved a great degree of functionality, but thanks to its army of dedicated developers, KDE is now one of the most popular pieces of Open Source software available.

With the raft of developers constantly working tooth and nail on KDE, the environment has become that incredibly configurable, and as part of this feature on customising Linux, we look at the important aspects of customising KDE. In this feature we focus on KDE 3.1, although much of what is discussed here can be applied to earlier versions of KDE.

Before we begin, we want to make clear that we won't be looking at all aspects of KDE customisation as space does not permit this; we will however examine at the major features that you will need to customise in your daily use of KDE, and how some changes can improve the usability of your system and let you get things done quicker.

In KDE, most of the action when it comes to configuration happens in the

KDE Control Center. **Figure 1** shows how the different categories are listed down one side of the Control Center, and when you click on a topic it is loaded on the right. The Control Center is a main point for configuration, although many of the modules are access in other areas (such as when you configure the desktop by right-clicking it). Now we are familiar with where most of our configuration lies, we will now look at the major areas that should be tweaked to get the most out of KDE.

The Panel – Kicker

The KDE panel, affectionately known as *Kicker*, is one of the most important components of the KDE environment. It is an area in which you can access running applications, start programs and see information. The default kicker layout is one that is good for general use but you may wish to reconfigure it.

If you right-click *Kicker* and select **Configure Panel**, we can bring up the configuration settings. The first thing to play with is the size – you may want to change the size so you can see

! way you want it!

choose between a multitude of mail clients, but you can fine tune them to work in the way you want to – whether that's in terms of the fonts they use or the way they handle mail. Security options, filters, folders – there's a huge range of options you can change on the average Linux system. But of course a lot of people don't bother. They install the software and use it, and rarely if ever visit any of the configuration pages or think about how they could make their favourite apps even more fun to work with.

For the ultimate in Customisation, you can even tinker with the kernel itself. When it comes down to it, it's just another piece of software that can be tweaked and prodded into doing what you want. Even simple changes to the kernel can make a huge difference in terms of speed and functionality.

Our aim in this feature is to show you just a small fraction of the things you can do with desktop environments like KDE and GNOME, as well as some popular applications

“FREEDOM means not only freedom to look at source code, but freedom to choose which apps to use and how they work.”

and the kernel itself. This article is intended to help you find out how to create a Linux environment that works for you, and be inspired to explore all the options that free software has brought you.




Figure 1: **The KDE Control Center**

applications more fully (an example is *Blender* which needs a smaller *Kicker* to see all of the buttons). You can select some predefined sizes, and also drag the panel border to resize it.

Other settings that are important include the location of *Kicker*, which can be selected on the same page, and the auto hide function by clicking on the **Hiding** tab to configure it. The **Menus** tab also allows you to select how the K Menu is displayed to you, and you can also edit the entries if

that is what you require.

Key Bindings

Key Bindings are important to some and useless to others, but KDE tries to account for both by having everything mouse and keyboard controlled. You can configure key bindings in the Control Center in Regional & Accessibility>Keyboard Shortcuts.

The configuration screen shows a list of actions that can be linked to a keyboard shortcut, and you will see

that many of them already have shortcuts assigned. To create a new one, click on an action and then click on the rounded square button below. You will then be presented with a box to create the shortcut. Simply perform the shortcut, and if it has not been assigned to another action, it will be assigned to the action you chose.

The configuration settings for shortcuts are very powerful and you can save shortcuts under schemes that are shown at the top of this window. There many options available in this screen and it is recommended that you investigate them as they can increase productivity dramatically.

Themes and styles

KDE by default comes with the beautiful Keramik style, but variety is the spice of life, and you can change the style of your KDE to suit your tastes. The **Appearance & Themes** category in the Control Center is largely dedicated to this.

The first item to change is in the **Style** part of the Control Center. This is where you can choose how the visual elements of your KDE desktop

are displayed. There are a number of options within the box, and more styles are available at KDE Look website www.kde-look.org. You can select the theme from the box, and the buttons below it will show a preview of how the style will look. To choose a style, click the **Apply** button.

On the **Style** page there is also an **Effects** tab. This page enables you to apply special visual effects and animations to your desktop experience. Many of these are visually very appealing and professional, but unfortunately can often slow down lower-performance computers.

Another element you may wish to change is the window decoration and behaviour. This is in the **Window Decorations** section. You can select the Window theme, and there is also a tab in which particular window decoration themes can be configured (**Note:** not all available themes have configurable options).

The Desktop

The desktop is an important tool in the use of the KDE environment; not only do you have icons on it, lay



CustomiseLinux

KUSTOMISING KDE



applications over it, but you also want to have some snazzy wallpaper spread all over it. All of these settings and more are accessible by right-clicking the desktop and selecting Configure Desktop. The settings for the desktop are in a number of different categories, and they are all worth investigating. It is a good idea to first take a look at **Multiple Desktops**. In this page you can set the number of desktops shown on the panel; a feature that can be used to isolate applications in different desktops to make running lots of applications easier. Next, you may wish to take a look at the **Background** page. This page allows you to select either



Configuring the KDE desktop.

Wallpaper or patterns as your background. It is also possible to mix the two to have gradients displayed over your wallpaper.

A final important part in this configuration box is the **Screensaver**

page. In this page you can select your screensaver from the box; you can then use the **Setup** button to configure the screensaver that you have selected. To set the amount of time before the screensaver is

Broken KDE?

Every so often any application can get a little confused and become erratic in its behaviour. This usually happens when the configuration files have not been updated properly due to a bug or problem. If this happens to you, you

can remove all of your configuration for your user account by removing the **.kde** directory in your home directory.

When you remove **.kde** and restart KDE, the environment will be set to the defaults.

MAKING GNOME YOUR HOME

Now more configurable than ever, make GNOME work the way you want it to.

In the new GNOME Control Center, modules may be accessed with **Nautilus**.



With the recent release of GNOME 2.2, the GNOME desktop has made leaps and bounds in catching up with its arch-rival, KDE. GNOME 2.2 has had some significant changes over its predecessors, not least in the area of configurability. In this article, we'll look at some of the ways you can configure

GNOME 2.2 to look and feel the way you want it to. This article assumes you are using the *Metacity* window manager – which is now the default. You are free to use another window manager should you wish to, but note that support for the previous favourite, *Sawfish*, is still somewhat incomplete in GNOME 2.2

From the user's point of view, the GNOME Control Center is the focus for all GNOME configuration. But with 2nd-gen GNOME, the Control Center is not a stand-alone application: it's a plug-in for the browser *Nautilus*. You can navigate the tree of installed Control Center modules via the URL **preferences:///** and click-on a module's icon to launch it. If entering a URL seems like too much trouble, you can access the Control Centre quickly from the **Start Here** page, which is opened by clicking the desktop **Start Here** icon. Or, you can launch individual Control Center modules from the GNOME menu (or even by invoking them from the shell).

Everybody wants their desktop to look different, and visual dressing is the fastest way to make your mark when personalising a system. Luckily, GNOME 2.2 has a new theme manager that makes tweaking your desktop's appearance a doddle. Just click on the Theme icon in the GNOME Control Center to open it.

Theming

There are three main aspects to theming in GNOME – the GTK+ (or widget) theme, the icon theme, and the window manager theme. The theme manager lets you bind all these together (plus backgrounds and splash screens) as a single desktop theme, and so it makes distributing, installing and managing complete themes really easy. You can still set the theme components individually, of course. Click on the **Details** button to access and manage separate lists of widget, window manager and icons themes.

started, you can specify the time in the box to the right.

Icons are often used on the desktop as a quicker method of starting applications and accessing files, and as such you need to be able to add and remove them easily. This can be done by either dragging a program from the K Menu to the desktop, or by right-clicking the desktop and selecting **Create New** and selecting the type of icon to add.

Konqueror

Although *Konqueror* is not part of the KDE desktop and is an application, it is handy to configure a few important aspects of it as many KDE users use *Konqueror* for browsing. You can access these settings via the Settings>Configure Konqueror option.

The first thing to configure is Java/JavaScript. You can set these up in the **Java & JavaScript** category. To enable Java select the **Enable Java globally** checkbox and then use the **Path to Java executable** text box to


find the program (usually in /usr/bin). If you don't have the *java* program, you can obtain the official Sun JDK from <http://java.sun.com/>. To enable JavaScript, click on the JavaScript tab and select the **Enable JavaScript globally** checkbox.

The other major element to set up in *Konqueror* is *Netscape* plugin support. To find your Plugin, select the Plugins category and click on the **Scan for New plugins** button. The new plugins will then be shown in the **Plugins** tab.

Summary

KDE is a huge environment in terms of configurability. The developers have always wanted to ensure that new features can be turned on or off as required, so there is a vast array of settings that I could not cover in this article. I have instead tried to focus on the practical changes that I hope you find useful in your daily use of KDE.

KDE is a system that is worth exploring, and It is a wise idea to load



up the KDE Control Center and just go through all the pages of settings and perform a few trial-and-error tests with them to see what you can do with the environment. You may very well be surprised with just how much you can configure to your exact requirements. Good luck!

Setting up *Konqueror*.

GNOME 2.2 comes bundled with several default themes, but, if you want more, go to <http://art.gnome.org/>. Themes are distributed as tarballs, and you should be able to drag-and-drop themes from your browser to the theme manager to automatically download and install them. This can still be a bit buggy, however. Icon themes should now be installed in the folder **.icons** in your home directory, but the theme manager sometimes puts them in **.themes** along with the other theme components. If that happens, try unpacking the icon theme tarball manually into the correct folder.


Your desktop wallpaper can be set by right-clicking on your desktop and selecting the **Change Desktop Background** option. This lets you pick an image to use as your wallpaper (with options to centre, scale or tile that image) or to cover the desktop with a solid or gradient fill. *Nautilus* also comes with a range

of pre-set patterns and colours that you can use, but these are accessed differently. Select the **Backgrounds and Emblems** option from *Nautilus*'s **Edit** menu, and you can drag-and-drop a pattern or colour from the dialog that pops up onto your desktop to use it as a background fill or onto *Nautilus* itself to use it as a background in the file manager.

Keys and clicks

Now your GNOME desktop is looking pretty, how about getting it to respond to controls the way you want? We start with the window manager, since this is an area that I find doesn't quite work the way I like by default.

Assuming you're running *Metacity*, then Control Center gives two dialogs to manage window behaviour: **Windows preferences** and the *Metacity Setup* module. This split in functionality leads to a lot of unnecessary duplication – for instance you need to use both to



properly set window focus and raising policy; and the latter also lets you set the window theme (which you can also do in the theme manager). ➤➤

GNOME 2.2's new theme manager makes installing and setting widget, window manager and icon themes simple.

CustomiseLinux

MAKING GNOME YOUR HOME



Configurability is one of *Metacity's* weaknesses, but the Window prefs dialog lets you choose between shading or maximizing actions when double-clicking a window's title bar and also lets you pick a hot-key for window dragging, while the *Metacity* module lets you pick a title-bar font and set the number of workspaces to use.

More window manipulation controls are set via the Keyboard Shortcuts control centre module. You can assign keys to various desktop operations, such as popping up the **Run command** dialog or opening the GNOME menu, and window management functions – such as activating, shading or maximizing windows and switching workspaces.

Simply click an operation and press the desired key combination to bind it to that operation (or press backspace to unbind the key). Here you can also choose the text-editing shortcut mode that GNOME applications which expect text entry will employ. You can choose between the GNOME default (which works like Windows and the Mac) or an *Emacs*-like mode.

Talking of keyboard shortcuts, one of the neat features of GNOME 2.2 is the new multimedia key daemon, known as *Acme*. This program lets you bind key-presses to various multimedia actions such as muting or un-muting your sound card's volume, ejecting a CD, or opening a web browser (see *Acme keys*

for a complete list of actions). *Acme* will even work with keys that haven't been assigned by your X keyboard mapping, so is good at putting all those extra keys on your Multimedia keyboard to use. The *Acme* settings dialog can be launched via Multimedia Keys option in Control Centre, and works like the main Keyboard Shortcuts preferences tool.

Files and types

Making file manager *Nautilus* use your preferred applications for opening and viewing files is quite straightforward. GNOME manages a database of file types – or MIME types – and, for each type of file, it associates such info as typical filename suffixes and a default icon to use when listing that a file of that type in *Nautilus*. In addition, for each file type, GNOME stores a list of applications that can view or open files of that type. Your GNOME installation will be set up with useful defaults for these – eg *GIMP* will be set up to open most types of image, but chances are you want to modify the default settings.

If you select a file in *Nautilus* and right-click, then the **Open With** sub-menu will present a list of apps that GNOME knows can be used to open that file. Choosing **Other application** in this sub-menu pops up a dialog where you can configure settings for the **Open with** menu. At the top of the dialog is a

ACME keys

Actions that ACME can map to key-presses

Mute	mute/un-mute volume	Search	launch GNOME search tool
Volume down	increase volume	Email	launch Evolution or Balsa.
Volume up	decrease volume	Sleep	enable APM or monitor power saving
Power	log out of GNOME session	Screensaver	blank screen with screen saver
Eject	eject CD-ROM (with eject command)	Finance	launch GnuCash
Media	launch default MP3 player (XMMS)	Help	launch GNOME Help (Yelp)
Play	mapped to X keysym XF86AudioPlay	WWW	launch default browser (Mozilla)
Pause	mapped to X keysym XF86AudioPause	Groups	opens http://www.gnomedesktop.org/
Stop (Audio)	mapped to X keysym XF86AudioStop	Calculator	launch GNOME Calculator
Previous (Audio)	mapped to X keysym XF86AudioPrev	Record	launch GNOME sound recorder
Next (Audio)	mapped to X keysym XF86AudioNext	Close window	close current window
My Home	Opens your home folder in Nautilus	Shade window	shade current window
Refresh	mapped to X keysym XF86Refresh		

Customising XMMS

Don't stop at the desktop, your software applications can be customised too!

The popular multimedia player *XMMS* has a large following of users creating and using skins, and also fully supports *WinAmp 2.0* skins to boot. Shown in the screenshot below is the ever-popular *Ultrafina SE* skin, although *XMMS* comes with a healthy supply of others packaged up with it. To see what you've got installed, right-click on a blank part of the *XMMS* interface, select **Options**, then **Skin Browser**. For the more adventurous out there, select the box that says **Select**

random skin on play to have *XMMS* change skin after each song.

Another regularly customised feature is the ability to play visualizations in the background while you're listening. These eye-candy gems display all sorts of graphics synced to the tune currently being played, and range from surreal shapes and colour wipes to dancing penguins. Right is how the popular plugin *Infinity* looks mid-action – look at it for too long, and you'll get cross-eyed!



XMMS is a great all-round media player, with lots of room for customisation.



Can you see the hidden dancing penguin?

list of programs that can open files of this type, and the list will probably include programs that weren't shown in the menu. For each app, you can hit the **Modify** button to configure whether that app is shown in the menu or shown only for that particular file.

To add an additional app that can open files of that type, click on the **Go There** option in the **Open with Other Application** dialog or launch the **File types and programs** preferences from the **Advanced** section of the Control Center and select the desired file type in the tree shown. To add another tool, look at the bottom of the **Action** section. In the **Default Action** chooser, pick **Custom** and enter the command to launch the new tool. Click **OK**.

You may notice that the Action section of the File type dialog has two controls – Viewer component and Default action (which we used above). The first of these specifies a GNOME browser component that can be embedded in *Nautilus* to view that file in line, typically used for viewing images. If you have a viewer component specified, then double-clicking on a file of that type in *Nautilus* will open that file with the viewer rather than any tool specified by the **Default action** setting. You need to set the **Viewer component** control to **None** for the **Default action** to be used on double-clicking a file.

GConf editor

How to manipulate hidden settings in GNOME

GNOME 2.2 employs a system called *Gconf2* for storing and managing application settings. *GConf* uses a central repository or registry for storing settings, and applications access settings in this repository as named keys and associated values. Each application's settings are stored by *GConf* as an XML file, and the names and data types of keys that an application can understand are specified

by that application's *GConf* schema, also specified as an XML file. Often an application's *GConf* schema will specify settings – or *GConf* keys – that it doesn't provide a GUI for modifying. This may be because the author hasn't yet implemented it, or because it's an advanced option and doesn't need to be easily changed by the user. In either case, the *GConf* Editor is just the tool you need. It provides a GUI for directly

manipulating *GConf* keys and values. Since it reads and parses *GConf* schemata, it can be used to manipulate all the keys that an application employs.

Using *GConf* may not be easy, however, and may even be dangerous. While some applications document the schemata that they use, other's don't; you'll just have to go by whatever, possibly cryptic, key names the developer chose.

GNOME of the future?

For more info, the GNOME user manual and the GNOME home page are great. While GNOME 2.2 is a big improvement on earlier releases in the GNOME 2 series, the truth is that it still has a way to go to match the configurability of KDE. The focus so far in the GNOME 2 series has largely been in the opposite direction. The developers have concentrated on usability and ensuring consistency and integration throughout the GNOME suite. All worthwhile goals. But now that we have a solid and stable foundation, work can continue in making GNOME much more tweakable. Hopefully, as GNOME 2.4, approaches, as more applications are ported to GNOME 2, and as more users get on board, the platform will develop into



Drag-and-drop a pattern or colour from this dialog to use it as your desktop wallpaper or as a background in *Nautilus*.

something that is as just as flexible and configurable as KDE.



Customising Mozilla

Follow the continuing *LXF* series on page 66

Built upon the flexible *Netscape Portable Runtime*, *Mozilla* was designed to be customisable to the extent where it could look and feel at home in several different environments. *Mozilla's* powerful skinning engine, powered by XUL (XML-based User Interface Language), allows skin designers to completely customise colours, icons, and dialogs inside *Mozilla*, even to the extent where people may not even realise they are using *Mozilla* at all!

However, the real power of *Mozilla* comes with some of the plugins available for it, made using JavaScript and XUL. As mentioned, one of the core design concepts of *Mozilla* was that it should be as flexible as possible, and this is demonstrated perfectly by the *Optimoz* project. The goals of the *Optimoz* project are two-fold: firstly,

they have created a gestures module to *Mozilla* in order that users can perform various commands by moving the mouse in a particular manner, and secondly they have implemented *pie menus*, which are a series of circular pop-up menus that replace *Mozilla's* normal popup menu and have their options selectable merely by moving the mouse in a given direction.

If you ever played the game *Black & White* under *WineX*, you'll be familiar with using mouse gestures as a means to control your computer. For example, the default gesture selection when the gestures plugin is first installed includes basic movements like dragging left to navigate backwards in your history and dragging right to navigate forwards – although it does require you to let it know you're performing a gesture by



One of *Mozilla's* more useful plugins, *pie menus* are as addictive as they are original.

holding down the right mouse button. If you're like me and you sometimes find it difficult to gesture competently, never fear: the gesture recognition can be configured to accept shaky gestures, or you can even force it to require more than just the right mouse button being held before it executes an action.

The *pie menus*, pictured above wholly replace *Mozilla's* default context menu, and the idea is that you select an item simply by moving in its direction and

blindly clicking the mouse. Because the *pie menu* will follow your mouse if you move too far, you're always assured that moving down and clicking will, for example, refresh the page.

There are also *sub-pie-menus*, which are slightly more complicated. For instance, to close the tab you're currently working with, the movement is "diagonally up, down, click". It takes a little getting used to at first, but eventually you will find yourself instinctively knowing which movement does what, and that's when *pie menus* have got you hooked, line, and, er, crust. Of course, until you've reached that point, you can just pause your cursor for a second on any given menu, and a description of what each option does will appear to remind you.

CUSTOMISING THE LINUX KERNEL

For the ultimate in customisation, you need to tinker with the kernel itself.



Recompiling the Linux kernel can seem like a daunting job initially – after all, it is the heart of the operating system and the most crucial software component of a working machine. However, once you're familiar with the terms and steps involved, it's quite simple; there are only a few commands to enter and the easy-going configuration tools help immensely.

This guide will explain how to get hold of the kernel source, methods of tailoring the configuration for exactly what you need, how to compile and install it, and finally give a few tips on applying patches. If you have a basic familiarity with the command line, you should have no problems. Before we start, though, consider the kernel versioning scheme: the middle number is even or odd and indicates a 'stable' or 'development' series respectively.

So 2.4.x is the current stable tree, with 2.5.x being where all the heavy development goes on (eventually to become a stable 2.6.0). The final number reflects the revision level – 2.4.20 was most recent at the time of writing, with 2.4.21 due shortly (with minor bugfixes, driver updates etc.).

Need to recompile?

If your system is running smoothly and all of your hardware is well supported, there's no pressing need to rebuild right now. But if you want to improve performance a little, upgrade for a

security fix, add support for a new driver or try out a patch from the Net, you'll need to follow these steps. Check out the **ChangeLog** of new kernel revisions. If there's a security hole fixed or new driver for your hardware, it's definitely worth patching and rebuilding.

Get kernel source

All mainstream distros include a package containing the kernel source; in Red Hat, Mandrake, SuSE and Debian it's called **kernel-source**. You'll find the package on your distro CDs or the Net, but check out `/usr/src` first to see if it's already installed.

Bear in mind that distro vendors tend to make heavy modifications to their supplied kernels – consequently, it's usually too messy to apply revision update patches or others which are created against the plain kernel releases. So, if you'll be adding patches or simply want the latest and greatest version, go to: <http://www.uk.kernel.org/pub/linux/kernel/>

The latest 2.4 revision is there, the 2.5 development series and incremental patches to boost your kernel from one version to the next (see the boxout below left). Find whatever you're looking for – like **linux-2.4.21.tar.gz** – and download it. To extract, get a terminal prompt and **su** to the root user, then:

```
# tar xfvz linux-2.4.21.tar.gz
```

or whatever file you've downloaded. If it's a **.tar.bz2**, you'll need to use **tar xfvj**. This expands the compressed archive, creating a new directory with 130MB+ of files – make sure you have enough disk space! Move into the directory with **cd linux-2.4.21/** (or different version number) and you're ready. If you're using the distro-supplied kernel source package, you'll need to **cd /usr/src/** and then wherever your vendor placed it (ie **linux-2.4/**).

Configuring kernel

With the kernel source in place, you're ready to begin the configuration process. As a first check, read the **Documentation/Changes** file and make sure the tools listed are up to date – you'll need at least **GCC**, **binutils**, **make** and some other

development programs installed. Note that all of the following commands should be run as **root**. If your source tree has already been used for a compilation, and you want to start completely clean, enter:

```
# make mrproper
```

This removes any old cruft, allowing you to start completely from scratch (*Mr Proper* being a Finnish cleaning product which inspired Linus Torvalds). Now, if you want to base your configuration on the kernel currently running, you'll need to get hold of its **.config** text file and copy it into this directory. Red Hat, for example, stores theirs in **configs/**, but you can avoid this step and just construct your own. To begin:

```
# make xconfig
```

If you're running the X Window System (ie KDE, GNOME or another window manager), a box will pop up offering a point-and-click system for choosing options. Meanwhile, those at the plain text console should use:

```
# make menuconfig
```

which is equally easy to navigate, but will work anywhere. Don't be terrified by the amount of options; the setup system is intelligent and will reveal or hide other options depending on what you choose. Work from top-left to bottom-right (or just top to bottom in *menuconfig*), choosing the features and drivers you need.

Most kernel features can be compiled as modules (**m** instead of **y** or *****) – these are individual drivers and features loaded when needed, rather than being part of the single kernel file. It's a good idea to keep things modular, as it improves flexibility and keeps the main kernel small and simple.

Note: You should never try to select critical boot devices and drivers as modules! If you boot off an IDE drive with the ext2 filesystem, for example, you must compile these into the kernel – otherwise it can't understand the drive to load the modules in the first place. When in doubt, compile straight in.

For each option, clicking on **Help** will bring up a useful box with

Applying patches

Update your kernel

With new kernel versions and features being released so frequently, downloading each large source archive would be a horrendous chore. As a result, developers produce patches: these are essentially text files describing the differences between one version and the next (being generated by the *diff* program). As most of the kernel remains the same between releases, these patches are always far smaller than the full source bundles.

Equally, developers working on flashy new features don't intend them to be part of the main tree, so they're distributed as patches too – you can add them to

your own source archive quickly and easily. Most patches are distributed as gzipped text files, so for a revision update patch you'll need to **cd** into the main kernel source directory and enter:

```
# zcat patchfile.gz | patch -p1
```

If it's a **.bz2** file, use **bzcat**, while plain uncompressed files just need the normal **'cat'**. This will work with the vast majority of kernel patches around on the Internet, but always check out the supplied instructions if it's anything more esoteric. Adding **--dry-run** to the end of the above command is a useful diagnostic tool for seeing what changes could be made, without it actually doing them.

guidance; the text is mostly very thorough and well-written, and should give you a firm idea whether or not you need something. Most items are self-explanatory, but always consult the **Help** box when you're unsure.

Having good knowledge of your hardware can help considerably – you'll know exactly what to choose and what to leave out. Certain options like the CPU type in **Processor type and features** can provide a slight performance boost, and you should select SMP support if you've got more than one processor. In all, it's best to cut out anything you definitely don't need (e.g. telephony support, etc), resulting in a smaller, leaner and more efficient kernel. Once you're done, you're ready to start compiling.

Compiling

Exit the config program (saving your changes), and type in:

```
# make dep && make clean &&
make bzImage && make modules
&& make modules_install
```

This is a series of commands joined together with **&&**, which ensures that the following operation will only begin if the previous one succeeded. Now, grab some tea/coffee/beer/absinthe and wait – depending on your hardware, the compilation can take from 15 minutes

(around 1 GHz, plenty of RAM) to a few hours (old 486).

Hopefully, you'll eventually be back at the prompt with no error messages. If not, read the lines indicating what went wrong – in most cases you'll need to go back into the config program and change some options.

Now, copy the new kernel image and supporting files into place:


```
# cp arch/i386/boot/bzImage
/boot/vmlinuz-new
# cp System.map /boot/
# cp .config /boot/dotconfig-backup
```

Note that the modules are already in place under `/lib/modules/<version>`, and the arch bit will be different on non-PC systems. The final step is to configure your bootloader: if this is GRUB, then you need to edit `/boot/grub/grub.conf` and add a new title entry for `/boot/vmlinuz-new` (copy and paste the original, changing the kernel image file name to the new one).

LILO users should open `/etc/lilo.conf`, add a new image entry (again, easiest to copy and paste an original) and run `/sbin/lilo` to update it. If you have trouble, see **'info grub'** and **man lilo.conf**.

Reboot!

The fun part. Shut down and restart the machine, choosing your shiny new kernel in the bootloader, and watch it



fire up. Keep a close eye on the boot messages – or enter **dmesg** after logging in – and look out for problems. If there are module errors, entering `/sbin/depmod -a` as root should cure them (most distros do this on startup).

And that's it! Play around with your machine, make sure everything works as expected and enjoy. If you've missed anything, just go back and rebuild the kernel using the above process, or if something more serious is wrong, you can restart the machine and choose your old kernel from the bootloader. Any other queries, why not try the forums at www.linuxformat.co.uk for advice from other LXF readers?

kernel.org contains just about every kernel and patch that you need to customise and update.

Customising OpenOffice.org

It's easier than you think...


One of the problems with *OpenOffice.org* is that, while the default button icons might be easy to recognise, they aren't pleasant to look at. *OoO* allows you to customise all the icons used in the interface.

OoO will only load BMP bitmap files. To translate existing icons you can save them out as BMP files from *GIMP* or use the *ImageMagick* tools to convert the format. You'll also want to stick to the size of the current icons to avoid problems. Put the results in the default *OoO* icon folder: `/usr/lib/openoffice/share/config/symbol/` Now you can assign these button icons to actions with `Tools>Configure` item. Click **Customize**, select a function and click on button to get an index of the icons.

Also worth tinkering with are the keyboard shortcuts. Fairly straightforward

this – assign a useful key combination to the often used functions and save yourself time and effort. There are plenty of key combinations unused by the default settings, and plenty of functions to attach them to. Go to `Tools>Configure` and click on the keyboard tag. Select the key combination from the top scroll box view, and the action to perform from the lower group of scroll boxes. Then just click on the **Modify** button.

Finally, a few tweaks that can make your *OoO* experience a little easier. For example, *OoO* can preview the fonts in the fonts list. This renders each font name in the relevant font, which can make selection a lot easier if you aren't familiar with all the fonts installed. This option isn't enabled by default on many



A million and one configuration options await. Approximately.

versions, but you can turn it on by visiting the `Tools>Options` menu item. Click on `OpenOffice.org>view` in the tree on the right, then click on the checkbox for

Preview in fonts lists. While you're here, you can also change the generic look-and-feel to be more like *MacOS* or *OS/2* for instance. Worth a try! [LXF](http://www.linuxformat.co.uk)

SLIPandPLIP

Cheap and easy networking with

SLIP and PLIP

Getting his wires all crossed once again, **Mike Saunders** looks at simple and inexpensive networks using serial and parallel cables.

Back in the early days of personal computing, the expense and complexity of home machines usually meant that few households had more than one computer. Today, though, cheapness of commodity hardware has led to the typical home having several systems in use – often to provide family members with their own Internet stations. Equally, there's a greater need for ways of linking boxes together, as shifting files on floppy or CD-R is inefficient and tedious.

As a result, home networking is hugely popular now and being able to share files and printers (or even a Net connection) is immensely useful. In most cases, this is achieved through the common Ethernet approach; however, it's not always ideal and other methods can be cheaper, easier or simply more appropriate for the task in hand, and that's what we're looking at here. This guide will show you how it works, what cables to buy, how to configure the machines, and ways to make full use of your newly-networked boxes.

So what are SLIP and PLIP? Both acronyms stand for “<type> Line Internet Protocol”, with the first letters meaning Serial and Parallel respectively.

These are ways of networking machines via the stated ports – most machines have at least one of these available, so the only additional hardware required is the correct cable. You simply attach a cable to both machines' serial or parallel ports, and after a bit of tweaking you have them networked.

Why use it over Ethernet? If you've got a couple of boxes and you don't need the high speeds offered, it's overkill. Good quality PCI NICs are fairly cheap these days but a simple serial or parallel connection is even more cost-effective. Similarly, it's very quick and easy to set up, is supported out-the-box by many major operating systems, and is great for older hardware.

SLIP and PLIP are godsend for users networking aging or esoteric systems – PCMCIA NICs for old laptops can be pricey (and ISA ones becoming rare), and finding a network card for that Acorn, Amiga or Atari you're running Linux or BSD on can be a nightmare. The methods in this tutorial are available on almost all Linux systems with the appropriate ports, making it low price and easier than a very simplified pie which comes with a comprehensive user manual and 24/7 telephone support.

SLIP or PLIP?

It's your choice...

Which networking system you choose will depend mostly on the types of ports you have available. If your lucky box has both serial and parallel ports free, and you're not planning to link to a Windows box (which as we see later can only use SLIP), check out the following comparison. Systems where speed isn't so crucial are better off with SLIP, but if you'll be putting larger demands on the box (like NFS for example), PLIP is more ideally suited.

SLIP (SERIAL):

- Slightly easier to set up

- Cables tend to be cheaper
- Slower (approx. 10 kilobytes/sec is about the usual speed)
- Longer maximum cable length
- More of an “industry standard”, so there's better OS support

PLIP (PARALLEL):

- Can require a little more work
- Faster (up to 40 kilobytes/sec)
- Short maximum cable length (approximately 15 metres)
- Can be heavy on the CPU
- Could block off your printer
- Cables tend to be more expensive

1 GETTING THE RIGHT CABLES

After deciding which method you want to use (see boxout above), you'll need the right types of cable. Luckily, this is easy once you know exactly what you're looking for. Small local stores can usually equip you with the right

“Simply attach a cable to both machines' serial or parallel ports, and after a bit of tweaking, you have them networked.”

The screenshot shows a Mozilla browser window displaying the product information for a Belkin parallel cable. At the top, there's a banner for 'get broadband now for just £27.99'. Below that is a navigation menu with 'Main', 'Features', 'Specification', and 'Reviews'. The product title is '25 Way DB Male to 25 Way DB Male PC-PC Parallel'. A table lists product details:

Mfr Code	F3D508-10-GLD	ex VAT	£10.80	buy now
Quicklink	117DWS	inc VAT	£12.69	
dabspoints	12			

Below the table, it says 'available now' with a '5' and 'avail 1-3 days' with '3790'. There are icons for 'print window', 'e-mail to a friend', and 'wish list'. An image of the cable is shown, with the caption 'Image for illustrative purposes only.' Below the image, there's a paragraph of text: 'Belkin's Direct Cable Connection lets you unleash the full potential of Windows 95. This high-quality cable allows two computers to share files, folders and peripherals such as e-mail, modems and CD-ROM drives. By connecting this cable to two parallel ports, you will enjoy the full power and flexibility that only Belkin and Microsoft can offer!'

Belkin's solidly-made parallel cable on show at www.dabs.com.


kit, but buying online can provide the most choice. Our favourite online retailers are www.dabs.com and www.maplin.co.uk.

For SLIP, the type you need is commonly known as a “null-modem cable”; Amiga buffs will remember linking their systems in such a way for multiplayer games like *Stunt Car Racer*. If you're linking PCs, you'll need a 9-pin female-to-female cable (ie both ends have holes instead of pins), but some older non-x86 boxes may need the alternative 25-pin type. Maplin sell a good 2 metre 9-pin lead for £5 (as well as longer ones), but if you'd rather shop locally then asking for a null-modem serial cable at any decent store will get you sorted.

For PLIP, you have to be slightly more careful in choosing. Typical names for this type of cable are “direct connection” or “LapLink” – the latter being familiar to Windows users. This is a male-to-male 25-pin cable, and among Belkin's range is a superb 3 metre lead for £12.69 (see www.dabs.com for the full range).



SLIP and PLIP



Kernel configuration looks much like this – here selecting SLIP as a module.



Again, your local computer shop should have something similar available, but make sure you explicitly state that it's for data transfer between two systems. If you're feeling adventurous, you can even fashion together a cable yourself – see `drivers/net/plip.c` in the kernel source.

With the right cable in hand, plug the ends into both machines and you're ready to set up the network. If you're using PLIP, it's best to drop into the BIOS (hit DEL, ESC or F2 after powering up) and make sure that the parallel port is set to ECP or EPP mode. Failing that, choose bi-directional – these modes allow for faster data transfer. Also, take a note of the IO port and IRQ settings, as you might need this info later.

2 CONFIGURING THE KERNEL

Most popular distributions include SLIP and PLIP support as modules in their supplied kernel packages. To confirm yours has this, take a look inside this directory:

```
/lib/modules/<kernel_version>/kernel/drivers/net
```

and see if `slip.o` and `plip.o` are

present. If not, enter `dmesg` and scan through the output to see if the support has been compiled-in, and if you're still not sure, just skip ahead to the setting up stage and try the commands – the error messages if they fail will let you know whether or not it's supported.

Still, you may have recompiled your kernel and disabled support (or your distro's kernel doesn't have it), in which case you'll need to roll your own. Recompiling the kernel can be a complex task and is beyond the scope of this tutorial; there are plenty of guides on the Net, or dig out a venerable copy of *Linux Format* issue 3 for a full walkthrough. We'll just skim over the basics here as a refresher.

Install the kernel source package (or download and extract it yourself), get a terminal with a root shell prompt, and input:

```
cd /usr/src/linux
```

(This might be `Linux-<version>` depending on your distro, or wherever you extracted it). Then issue `make menuconfig` or `make xconfig` – the latter being easier if you're on a graphical desktop.

Jump into the "Network device support" section, and scroll down to find the PLIP option. Select as a module. Go further down to SLIP, and again select as a module – this may enable three other options, so select them as well. Exit the config program (saving your changes) and issue:

```
make dep && make clean && make bzImage && make modules && make modules_install
```

After building, copy `System.map` and `arch/i386/boot/bzImage` to `/boot`, then edit `/etc/lilo.conf` for the new `/boot/bzImage` file and run `/sbin/lilo` if

you're using that bootloader (Grub should be fine), and reboot. You should be all set for using SLIP and PLIP now.

3A SETTING UP THE SLIP CONNECTION

First of all, boot the machines and get a shell prompt – you'll need to log in or `su` to the root user for this to work. Under Linux, the serial ports are named `/dev/ttySn` (where `n` is 0 for COM1, 1 for COM2 etc.) and on older kernels these may be `/dev/cua0`, `cua1` etc. Now load the SLIP module as follows:

```
modprobe slip
```

Then type `dmesg` and read the last few lines of output to confirm that it has been loaded correctly. (If you have trouble and you've compiled serial support as a module, try `modprobe serial` first, and also look at `man setserial`.) Then set up the serial port as a network device:

```
slattach -p cslip -s 115200 /dev/ttyS0 &
```

Of course, change the serial device name to whatever is appropriate. The `-s` flag sets the speed – on most PC systems that is the highest you can get. The `-p` flag sets the SLIP mode: `cslip` is slightly faster (but only compresses the packet headers). Then:

```
ifconfig s10 192.168.92.1 pointopoint 192.168.92.2 up
```

This creates the network link between the two IP addresses given to each machine, the first being this box, and the second being the other. Note the single `t` that is in the middle of `pointopoint` – it's not a typographic error, just the way it is. Finally:

```
route add 192.168.92.2 dev s10
```

which sets up the kernel's routing tables for the other box. Now you need to do the exact same process on the other box, but reversing the IP addresses; *ie* the `ifconfig` line will be `.92.2` and `.92.1`, and the route line will use the first box's IP address. Once done, try `ping 192.168.92.2` from the first machine, and vice-versa, to confirm that they can both actually reach each other.

3B SETTING UP THE PLIP CONNECTION

As with SLIP, you'll need a root prompt on both boxes. Make sure LPD

Talking to Windows and BSD

Cross-platform communication

Helpfully, Windows machines can also be linked to Linux boxes with serial cables (sadly there's only a very old MS-DOS driver for PLIP). This is particularly useful when your main box is Windows-only, or for sharing an Internet connection with other family members – you can enjoy using *Konqueror* and *Mutt* as normal while others browse with *MS Internet Explorer*. All you require is the same type of null modem serial cable described in the main text. You'll need to delve into the depths of PPP

configuration if you want to make the most of connecting to non-Unix boxes, and Microsoft's PPP implementation is characteristically quirky in places. Still, with a bit of work you can establish a link – the intricacies of getting it configured deserves a whole tutorial to itself, so check out the following HOWTO for a well-written guide: www.tldp.org/HOWTO/Serial-Laplink-HOWTO/

Meanwhile, as we saw earlier, SLIP and PLIP are excellent networking

methods for old or non-PC systems where Ethernet can be rare or expensive. Few operating systems in the history of computing support as many architectures as NetBSD, so being able to link with it is massively useful on the stranger platforms that OS caters for. See the man pages for `slip` and `s10` – the process is very similar to Linux. NetBSD's PLIP driver appears to have suffered from severe bit-rot though, so if you want a BSD and a parallel port connection, check out FreeBSD.

is not running (so it doesn't hog the parallel port), and do `/sbin/lsmmod` to see if the parallel port modules (**parport** and **parport_pc** on PCs) are loaded. If not, enter:

```
insmod parport && insmod
parport_pc
```

Check with **dmesg** that it has been configured correctly, and then issue:

```
# modprobe plip
```

Again, check with **dmesg** to confirm – you should get something like this:

```
plip0: Parallel port at 0x378, using
IRQ 7.
```

If this isn't working, edit

`/etc/modules.conf` to provide the modules with the correct parameters – you'll need the equivalent of this:

```
alias parport_lowlevel parport_pc
options parport_pc io=0x378 irq=7
```

Change the IO and IRQ settings to whatever is correct, **rmmod** the modules and start again. Once you have the configuration sorted and modules loaded, enter:

```
# ifconfig plip0 192.168.92.1
pointopoint 192.168.92.2 up
```

(Note that you might need to change **plip0** to **plip1**, depending on the **dmesg** output, and also be wary of that single **t** in **pointopoint** that we mentioned earlier.) Again, we've assigned those IP addresses to each box. Finally, provide the kernel with the necessary routing info with:

```
# route add 192.168.92.2 dev plip0
```

Now, the same as with SLIP, you need to do the same operation on the other box, reversing the IP addresses to match (switch them in the `ifconfig` line, and change to `.92.1` in the `route` line). Then enter:

```
ping 192.168.92.2
```

from the first box to check that the machines are communicating, and do the same on the second box (using the first's IP address this time).

4 MAKING USE OF THE NEW LINK

With the machines now networked and talking to each other, you'll want to do more than just ping them for fun. Firstly (unless you already have some kind of DNS set up) it's wise to edit `/etc/hosts` on each machine, and provide them with memorable names – something like **192.168.92.1** **firstbox** etc. is fine.

If one of the boxes is going to be a headless server or mostly inaccessible, you'll need to configure a way of

logging in remotely – the two best ways of achieving this being Telnet and SSH. Generally, the former is frowned upon for its poor security, as anything you type is sent in clear text, and it's lacking in features. Install *OpenSSH* on it (included with most distributions), and:

```
# ssh -C user@hostname
```

on the other box, replacing the **user** and **hostname** with whatever user accounts you've set up and whatever you put in `/etc/hosts` (or just use the IP address). The **-C** flag is for compression, which is a good idea on slower links.

From there, you have a fully accessible box to log into from your main system, and you can now go about installing *Apache* for web-serving, configuring NFS to share directories, or anything else you fancy doing. PLIP is just about fast enough to do remote X: try **xhost**

```
+192.168.92.1
```

in an X session on the `.92.2` machine, and then **xterm -display 192.168.92.2:0** on the `.92.1` box (or use the hostnames).

5 INTERNET CONNECTION SHARING

Being able to access the Internet with both boxes is enormously useful, and quite simple to set up. Most home and small office connections have a single outgoing IP address; consequently, you need to alter the kernel's packet filtering tables to pretend that the data from one box is actually coming from another.

So, in our case, say box A (192.168.92.1) is connected to the Internet – you want any packets from box B (192.168.92.2) to go out onto the big bad Net as if they're from box A. This allows you to have multiple machines all hiding behind one IP address. For the setup we've been discussing in this article, you would use this on the first box:

```
# iptables -t nat -A POSTROUTING
-s 192.168.92.2 -j MASQUERADE
```

Or if you're using a pre-2.4 kernel, use this:

```
# ipchains -A forward -s
192.168.92.2 -j MASQ
```

Either way, you'll then need to:

```
# echo 1 > /proc/sys/net/ipv4/
ip_forward
```

Then, on the second box, enter:

```
# route add default gw 192.168.92.1
```

That makes sure all non-local traffic goes through the first box. Also, unless you fancy using plain IPs all the time, you'll need to edit

`/etc/resolv.conf` on the second box and put **nameserver X.X.X.X** (usually the IP address of your ISP's name server).

And finally...

While we've looked at proper networking with serial cables here, it's worth bearing in mind that they can be useful in other circumstances too. Machines without a graphics card or monitor (headless) can have the OS




“Most of the popular Linux distributions include SLIP and PLIP support as modules in their supplied kernel package.”

installed and administrated with the “serial console” approach – the new box sends its output down the cable to another machine running `minicom`, and afterwards an entry can be put in `/etc/inittab` allowing logins using the same technique. Debian is ideal for installing on old headless boxes via the serial console.

Similarly, the *Serial Terminal Linux* single-floppy distro downloadable at www.eskimo.com/~johnnyb/computers/st/ can turn an ancient hard drive-less laptop into a handy repair station for servers, once again by booting up a simple `minicom` setup and logging in over the serial cable. From here you can leave a “top” display running to monitor the box, or move around a bunch of servers to admin (negating the need to have Telnet/SSH running on them).

And of course, if you get stuck at all with your networking, or have hints for others then pop into the forums at www.linuxformat.co.uk or find real-time help on IRC ([#linuxformat](http://irc.freenode.net) at irc.freenode.net)! [LXF](http://lxf)

The output of the **dmesg** command can be useful to troubleshoot when things go wrong.



What on Earth is... WIKI?

For many people that stumble across this cornerstone of the Open Source ideal, the idea of letting strangers edit their work is a bit scary. **Andy Channelle** allays our fears...

»» What on Earth is *Wiki*? It sounds like some sort of long-lost religion or art movement...

See, just one question in and already you've made a fundamental error. What we are covering here is the *WikiWikiWeb*, an alternative hypertext authoring tool created to allow collaborative development over a network. Some would argue it is actually the realisation of what the Internet promised in the first place; but then different people would probably say the exact same thing about blogging bringing the prospect of personal publishing to the masses. So now you've made a devastating *faux pas*, don't you just wish you could hit some sort of magical link and start all over again?

»» Ah ha! I know what you're doing. Isn't some sort of 'magical link' the Big Idea behind *WikiWikiWeb* sites?

It is; and I think we'll dispense with the full name and just stick to *Wiki* for the moment, it's easier. Anyway, *Wiki* sites are indeed editable. By anyone. Beyond the normal content, which is mostly text-based, *Wiki* pages tend to have a series of set links covering things like revision history, recent changes and, of course, the Edit button. In theory this means anyone accessing the site can hit the link and make their own additions or deletions to the page. For instance, you're reading away on a *Wiki* site and you notice that the author claims the great fire and the plague hit London in 1665. Of course this is wrong, so you hit the edit link and change it to 1666. The old version gets backed up – you've specified that this is a big change – and your new version is the one the next visitor will see. The next visitor, in turn, may think that it's important to note that the fire started on Sunday September 2nd so they click the edit link and add their details. The next visitor may add geographical detail as to where the fire started, and so on.

»» And the next visitor may like to add that Kylie Minogue started the fire by osmosis before single-handedly ruining pop music for the 21st Century. What stops people demolishing everyone else's work by defacing the page? It sounds like a recipe for chaos.

True enough, as the *Spider-Man* movies and comics say, power and responsibility go hand-in-hand. The answer is that there is nothing within the system to prevent a malicious user from going on the rampage. In the original *Wiki*, editing a page so that its only content reads 'delete' will actually delete the page.

This much control is regarded as a reason why the system works. In practice though, such damage will be picked up quite quickly – especially on a popular *Wikisite* – and can (sometimes) be rectified by rolling back to a previous version. As with any documents, it's wise to back up your *Wiki* on a regular basis to ensure you can recover from any disaster, whether it is machine- or man-created. An entry called *Why Wiki Works* claims that "to make an impact on *Wiki*, you have to generate real content... anyone can play, but only good players last".

If you're really paranoid, some of the available *Wikis* let you password protect content or even make it non-editable, but then you have to ask whether your site is a genuine *Wiki* or just simulacra.

There is a bigger issue though: when you contribute to a *Wiki* you give your permission for that contribution to be re-edited by strangers and the end result may not look anything like the original product. We don't want to get too deep, but this is the philosopher's axe in real life. For instance, two years ago, *LXF* began an entry in the *Wikipedia* (http://www.wikipedia.org/wiki/List_of_British_monarchs) on the kings and queens of England. The initial page consisted of a simple chronological list of monarchs, but it has now acquired a lot more information. The list has been divided into the various 'houses' and most names have spawned a

page (or series of pages) of their own. The question is: to what extent can *LXF* be said to have 'authored' this page? Does anything of the original work still exist?

Open Source developers may have wrestled with this new concept of authorship and ownership, but for many *Wiki* may be their first exposure to genuine collaborative development, and if they're at all precious about their prose, it might also be their last. Having someone else edit your work can be quite disturbing.

Many users who have stumbled across *Wiki* think it will be the perfect tool for blogging and this has inspired a long-running debate within the original *Wiki* at <http://c2.com/cgi/wiki> on the concept of what has been termed *Wiki Squatting*, using a *Wiki* as your homepage. One contributor wrote: "A *Wiki* squatter puts something on *Wiki* and calls it 'his'. For example, you feel very strange editing the diary page because it seems very personal". *Wiki*, the argument goes, is for everybody.

➤ It's an intriguing idea. Who came up with it?

Ward Cunningham is credited with both the idea and the first *Wiki* implementation in 1995 – though Cunningham himself traces its origins back to a *Hypercard* stack he wrote at the tail end of the Eighties – at the Portland Pattern Repository which publishes *Pattern Languages* and related information. Pages were (and still are in the original *Wiki*) served up by a CGI script written in Perl, and as time passed new features were added.

The name, if you were wondering, comes from the Hawaiian phrase for *Quick*, which was chosen by Cunningham to evoke the speed at which pages or sites could be created and modified. He hated the phrase *QuickWeb*. The first *Wiki* clone was written by Patrick Mueller in REXX for OS/2.

➤ Wiki clone? Is there is more than one Wiki then?

Er, yes and no. There is only one *Wiki*, but there is an awful lot of *Wiki* Engines created in everything from ASP to VisualBasic. The most popular languages for *Wiki* serving are Perl, Java and PHP, but there are also implementations in Awk, Python, Lisp and even Erlang (among others) for people that prefer those languages. One of the most popular engines is *UseModWiki* (<http://c2.com/cgi/wiki?UseModWiki>) which is an easy to install and configure script written in "100 per cent pure Perl". The script itself comes in at a hefty 2,300 lines, but has a comprehensive featureset that makes it suitable for even the largest projects.

And at the other end of the scale is *TinyWiki* (<http://c2.com/cgi/wiki?TinyWiki>) which ➤➤



WhatOnEarthWiki

◀ manages to do its work in a surprisingly compact 35 lines of Perl code. Both *TinyWiki* and *UseModWiki* are available under the GPL.

▶▶ Am I going to have to learn a whole new markup language to make any sense of Wiki pages?

In the main, writing *Wiki* is simple. You'll probably have a few hiccups at first, especially if you're trialling various Engines, but within a few minutes you get used to whichever system you're currently using. As it is primarily a text-based medium, design is limited which puts the emphasis on content rather than style. You could argue that this is one of the best reasons to use it, especially in something like a documentation project where clarity of content and presentation is more important than colour, branding or some other distracting visual noise. In a sense it is a little like the textual origins of the Internet, but with that all-important element of completely open access for all visitors.

While each *Wiki* Engine has its own style of tagging, most tend to follow Ward Cunningham's original scheme. For instance, creating a link in a document – perhaps the most fundamental formatting element – is most often accomplished by writing two capitalised words together without a space, i.e. `LinuxFormat`. This is called a *WikiName*. And here comes the clever part: if you type the *WikiName* of an existing document within your *Wiki*, the phrase will become an underlined link to the relevant page in the traditional fashion. However, if the page does not yet exist, the link becomes tentative and will have an underlined question mark at the end of it; clicking on the question mark creates the new page and opens up an edit screen to add some content.

Beyond the bounds of your particular *Wiki* there is also a scheme for building an *Interwiki* which is a system spanning a number of servers. It's a pretty complex operation, but for all its complexity though, it's a remarkable system to use, and – I might add – is surprisingly addictive.

Beyond the magical links, there are a number of formatting tags to cover styles such as bold, italic or bulleted text as well as provision for creating standard hyperlinks. These are designed to be simple and memorable. See the *Some basic formatting* table opposite for a quick introduction to these tags.

▶▶ Obviously I'll need some sort of web server, but what are the general requirements for building my own Wiki?

The main prerequisite is indeed a web server and most are extremely happy running on *Apache*, which is generally available as part of most common distributions. If you don't have it, you'll have to download from www.apache.org and install it manually. Next up is a database to house your data; most of the *Wiki* engines available from <http://c2.com/cgi/wiki?WikiEngines> tend to favour *MySQL*, but there are also implementations for *PostgreSQL*, *MS SQL Server* and *Access*. And finally you'll need to decide which type of *Wiki* you'll be installing as this may entail additional technology. Opting for the JSP *Wiki Very Quick Wiki*, for instance, will involve installing a J2EE web container such as *Tomcat* (for *Apache*). See this last month's and this month's *Linux Pro* for more on *Tomcat*.

The most common technology for handling the management of the *Wiki* is Perl, so if you've not got that installed (open a terminal and type **which perl** to find out if you have it and where it is) you'll have to get it from your distribution disk or download it from www.perl.com. You'll need to know its location to adapt the *Wiki* script you intend to use. The usual location is `/usr/bin/perl`. Similarly, if you haven't got PHP (type **which php**) and you're using a PHP implementation you'll need to download it from www.php.net/downloads/ and install before getting to work on configuring the actual *Wiki* script. A setup like this should (hardware notwithstanding) handle all you can throw at it.

Once the infrastructure is established you can download your chosen *Wiki* script and install it in *Apache's* `cgi-bin` directory.

Of course I'm simplifying here – there's never enough space – but most of the *Wiki* scripts available will have a fairly comprehensive set of installation instructions and, once you've actually dealt with the server basics, what's involved in adding a *Wiki* is child's play.

▶▶ Well, at least all these things are available in a standard distribution. Isn't there an easier way?

There's always an easier way, though if you envisage creating an enormous *Wiki*, you will definitely need a database backend and the security and stability under pressure that *Apache* provides. If, however, your needs are more rudimentary you can ditch most of this and set yourself up in less than ten minutes. Firstly instead of *Apache*, try *Abyss Webservice* from Aprelium (www.aprelium.com) this is a free (but not GPL) webservice that runs on Linux or Windows, features a very simple web interface for its various configuration options and is both PHP and Perl able. Using the Linux version is simply a matter of unpacking, cd-ing to the correct directory and typing `./abyssws`, though you'll need to enable support for your chosen CGI script (see the Aprelium website for more info).

In order to dispense with the need for a database, download the PHP *Wiki*

contribute

add

Some basic formatting

Simple and memorable

Here are a few of the main formatting elements that you'll come across using the original Wiki that's available at <http://c2.com/gci/wiki>. Be aware though, that other implementations of the Wiki scheme may have different ideas, some will even take standard HTML, which is useful when copying information across from other sources.

---	Horizontal line
<tab>*	Bulleted list
1. 2. etc	Numbered list
<space>(at the start of a paragraph)	Monospaced font
" "	Emphasis, usually italic
''' '''	Extra emphasis, usually bold
http:, ftp:, mailto: etc	Automatic link to external website, ftp server, email address
ISBN<10 digit number>	In Cunningham's Wiki and others, an ISBN number will link to the Amazon.com page for the book with that number

Processor (PWP) from <http://www.net-assistant.de/wiki/static/StartPage.html> and unzip it into the .htdocs (or the specified root directory) of the webserver and open localhost in a browser. PWP comes with a pre-made index.html, uses a flat file backend and each page you create is stored as a standard text file. Despite being simple to set up, PWP uses CSS to define formatting which means it's quite simple to integrate the visual elements into a regular website. The downside is that, as there is no database serving up pages, it is realistically limited to 300-500 pages. This is ideal if you're looking at a corporate project (say documentation development) but is not going to scale up too well.

Ah ha, that would be fine if I was in anyway technical, but I'm new to this. Anything easier?

So far, the simplest way I've found of building your own Wiki server is *EddiesWiki* (<http://www.tinyted.net/eddie/wiki/>) which is an 85K download complete with HTTP server and no other requirements, it is genuinely an 'unzip and go' application. The downside is that currently there is no Linux version; it was written in Visual C++ and currently compiles on Visual C++ 6.0 and [Windows] GCC 2.95-2. However the download includes source code, so it's possible that someone could sort out a Linux version. Anyone want to try?

Again *EddiesWiki*, which is Postcardware, uses a flat file system for storing pages so is only suitable for small projects. It makes a perfect personal *Wikiserver* though.

There's still a snag. I'm in no position to run my own accessible webserver and my ISP won't let me run Perl or PHP scripts on my personal web space unless I cough up even more money than I do already. Can I do it without having to manage a server?

The simplest (and it doesn't come any simpler than this) way of starting your own *Wiki* is to sign up with a *Wikifarm*. This has the benefit of off-loading all the

hassle of configuring and maintaining servers onto someone else and, in most cases, it is free for non-commercial uses. You give up some freedoms and you are at the mercy (sort of, potentially) a commercial entity, but really there aren't many cons to this route of WikiOwnership (that markup scheme is habitual you know).

Visit www.seedwiki.com, click 'Start Your Free Wiki', choose a name, a *Wiki* Type (so *Wikis* of similar types can be corralled into communities) and language, set a code word so you can manage the backend of the site and hit 'Save'. And that's it. You can now start editing and creating pages. The good thing about *Seedwiki* is that you can format pages using standard-ish HTML tags and widgets that can, for example, display the number of visitors a *Wiki* page (or entire site) has had, send out a mail whenever a page is updated or include the contents of another page.

Seedwiki hosts some 40,000 pages ranging from a serious discussion of *Wiki* as a means for publishing to a forum for participants in Open University degrees, to more light-hearted and esoteric subjects. It's a nice place to hang out and edit.

So I've set up my Wiki, clicked a few edit buttons, written a few swear words for fun. Now what?

There's plenty you can do. Wiki's text bias and incredibly easy collaborative features makes it perfect for the creation of documentation, and this is seen in the growing popularity of *Wiki*-based HOWTOs. Think about it: a developer creates an application and writes some elemental documentation. Users begin using the application (or other developers begin adding to it) and, in search of help, come across the HOWTO *Wiki*. Instead of just consuming the instruction they can add their own experiences, tips and questions which may then be answered by the developer or subsequent visitors.

As more people visit, the documentation becomes more comprehensive and the world becomes a better place... (Excuse the rampant idealism!)

It's also finding favour among commercial developers and testers who can code and document side-by-side without having to worry about the complexities of versioning in the most popular word processing packages.

Or how about really blurring the boundaries of author and reader by starting your own *Wiki*-novel or fanzine? If you can put it into text, it can be *Wiki*-fied, though of course, you not only have to submit to other people's opinions of your work, but also let them try to do better.

If it's so good, why hasn't it been used for any thing big?

Apart from an enormous number of HOWTOs and the original *WikiWikiWebsite*, the most exciting *Wiki* project (and, some would argue, the ultimate use of the Internet) is *Wikipedia* (www.wikipedia.com), an attempt to create a free and open encyclopaedia based squarely on the concepts of peer review and freedom of information in the truest sense.

Wikipedia was launched in January 2001 and in just two years has managed to accrue over 100,000 English language articles and some 37,000 articles in other languages. As its reputation as a haven of freedom of expression and knowledge has spread it has attracted the attention of academics, experts and researchers keen to expand its content. Users are no longer just consumers of information, they are creators, editors, owners.

Even better than an open encyclopaedia are the prospects of a *Wiki* news site (there is a nascent news community on *Wikipedia*) offering a diverse range views on current affairs or highlighting news that the mainstream media – for whatever – reason can't or won't cover and a *Wiktionary* which would be able to react to neologisms and changes in usage much faster than printed works.

And contrary to accepted wisdom, being subject to the whims and wishes of biased individuals often leads to some sort of equilibrium in articles if the common-sense guidelines are adhered to. With enough interested participants this is the very definition of seeing things from all sides.

Hold on, Mr Boundless-Enthusiasm, it can't all be flowers and romance. What are the downsides?

To be honest, the downsides are exactly the same as the benefits. Some people won't get the concept and will plough through your ideas and effort like an ill-mannered tank commander. If subjecting your words to the potential of such an onslaught scares you, then you will find more reward in plain old static websites and should probably steer clear of *Wiki*.

And if you're using a *Wiki* for the first time, just remember to tread lightly, don't rant, and don't 'over-write' without thinking about it first. The page's progenitor may have spent hours or days honing his/her prose to the knife-sharp point you're just about to hammer flat with a throw-away phrase. 

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 for a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact us, by post, by email (linuxformat@futurenet.co.uk) or log on to our website and post your suggestions in our forums? (www.linuxformat.co.uk). Hope to hear from you soon!

Nick Veitch EDITOR

How code is represented

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap
between lines:
begin
mniWordWrap.Checked := not
end;
```

Usually, you'll find the code on our CD/DVD too.

THIS MONTH TEACH YOURSELF...

Custom Mozilla

In our on-going rebuild of the browser, the next step is tweaking toolbars **p66**

The GIMP

Texture magic as we explore the creation of effects like rust and grunge **p70**

PHP

As well as being useful on the web, PHP can be a great CLI scripting language too! **p74**

Python

First steps with Python introduce us to objects, variables, *Pycrust* and some dreadful puns **p78**

In sync

Having a PDA is one thing, but what happens when you're double-booked on your desktop PIM? Synchronising is the answer **p82**



TIP OF THE MONTH!

Grokking grep

One of the most useful, and therefore popular tools for administrators is *grep* and its associates *egrep* and *fgrep*. Put simply, *grep* searches through specified files for a given pattern and outputs matching lines.

A basic scenario might be using `grep error /var/log/*` to print out all error messages in your system logs. Similarly, you can replace **error** with **warning** or **notice** to get different information. However, if no filenames are specified on the command line, *grep* searches standard input for matches, which is where it often comes in most useful.

Take, for example, the command `ps aux`. On most computers, this will output around thirty lines of text describing the processes running, of which you may only be interested in three. Using *grep*, we can filter the result of *ps* so that we only see results based on the criteria we pass to *grep*. So, to filter for programs being run by user **hudson**, use:

```
ps aux | grep hudson
```

This will technically match any programs being run that are called *hudson*, but that's pretty unlikely! This could be extended by using:

```
ps aux | grep -c hudson
```

which returns the total number of programs being run by **hudson** (`grep -c` counts the number of matches). To finish, the **hudson** search pattern can be substituted with a call to another command-line program, giving a command like this:

```
ps aux | grep -c `whoami` -v
```

Note that **whoami** is surrounded by back-ticks (next to the 1 key on UK keyboards), not single quotes. The **-v** parameter greps on the inverse of the search criteria, which, when combined with using the result of the command *whoami*, means that *grep* will pick up every program *not* being run by you.



BESPOKE BROWSERS

Mozilla: toolbar tweaks

Andy Channelle shows how to extend the browser's features with a selection of toolbar hacks. Plus, a look at the godfather of markup languages – XML.



Between *Mozilla's* menu bar and content area are the toolbars. By default, the standard browser suite ships with the main navigation toolbar, which we pulled apart in the first episode of our tutorial, and the 'Personal Toolbar' which, as the name suggests, is a space to add personal items such as links. The latter also has a Home button, which some designers suggest is one of the more eccentric decisions made by the UI builders. The rationale appears to be that the Home button is merely a link and would be out of place on the main navigation bar. If you disagree, you can install an XPI work-around which adds a traditional home button; though this widget isn't theme-aware, so you could end up with an ugly toolbar if your current theme doesn't support it. Visit <http://home.no.net/triand/mozilla/home/en/> for more info. That the two standard toolbars neither look nor work alike just demonstrates how much you can remake the interface. In fact, putting any XUL content into a toolbar opens up a whole new avenue of customisation.

The most basic change you can make to the Personal Toolbar would be to drag and drop a hyperlink onto it. This would give you one click access to that site from the toolbar as well as adding a link into the Personal Toolbar Folder within the bookmarks menu – manually adding a link to the folder would also add it to the toolbar.

However, you are not just restricted to static links. Go to Bookmarks>Manage Bookmarks and select File>New>Bookmark. Give your new link name (Dictionary would be appropriate) and then add the following to the URL box:

```
javascript:q=document.getSelection();for(i=0;i<frames.length;i++)q=frames[i].document.getSelection();if(q)break;if(!q)void(q=prompt("Enter a word?"));if(q)location.href='http://dictionary.reference.com/search?q='+escape(q)
```

Close the Bookmarks dialog and click on the link in the toolbar. You should be presented with a small Javascript Application requesting a word. Type in something and hit Enter to be taken to Dictionary.com's definition (You can also do a second link, substituting **dictionary** in the URL for **thesaurus** to find synonyms). This is actually the second part (from **q=prompt** onwards) of the URL we typed in at work. To see the first half, simply highlight a word on any web page and hit the link to get an instant definition. The great thing about this script is that, once you've found the correct syntax for the **location.href='URL'** element you can reuse the script for almost any search engine. For instance, you can add Google's news service to your toolbar with the following:

```
javascript:q=document.getSelection();for(i=0;i<frames.length;i++)q=frames[i].document.getSelection();if(q)break;if(!q)void(q=prompt("Google News search?"));if(q)location.href='http://news.google.com/news?1&q='+escape(q)
```

Simply nesting a number of these links – with appropriate URL changes – in a single menu is the first step towards creating, in the case of the latter, a Google Toolbar. There is, though, already a well-developed toolbar covering Google at <http://mozdev.org>, so we'll go no further down that road.

Getting complicated...

More interesting is the idea of adding an additional, installable toolbar which entails a headfirst dive into XML/XUL. We'll start with the most basic toolbar which offers a single link to another URL. This is the genesis of a 'writers' toolbar' so we'll link to www.dictionary.com.

The program begins with an XML definition followed by the 'window' element. Everything in the toolbar will be encapsulated within the **<window>** tag, which is the most fundamental XUL element. We also need to give the document an id, define the orientation of the window and declare the XUL namespace, which defines valid elements.

```
<?xml version="1.0"?>
<window
  id="DictionaryXULWindow"
  orient="horizontal"
  xmlns="http://www.mozilla.org/keymaster/gatekeeper/
there.is.only.xul">
```

A touch of Javascript takes care of 'event', which is a simple link:

```
<script>
function dictoolbarLoadPage(url){
  // jump to a new page
  window.content.document.location.href=url;
}
</script>
```

We then provide a toolbox name, which will change when the toolbar is finally packaged up, a toolbar id, define the look of the button – in this case a standard **toolbarbutton** element. We also need to give it a label, call the Javascript and provide a link for the script to jump to.

```
<toolbox flex="1">
  <toolbar id="DicToolbar">
    <toolbarbutton id="DicToolbarMenu" label="Dictionary" on
command="dictoolbarLoadPage('http://www.dictionary.com')"/>
  </toolbar>
</toolbox>
</window>
```

The beauty of *Mozilla* development is the ability to test your handiwork without leaving the confines of the browser. Simply save your text file with an .xul extension and open it as normal in *Mozilla*. The other advantage is that *Mozilla* has a pretty robust XML validation tool built in, so if your code won't run, you get a good idea, including line and column number of where problems may lie.

We can extend this button by turning it into a dropdown using the **<menupopup>** tag. This first involves changing the adding a 'type' definition to the toolbar id.

```
<toolbarbutton id="DicToolbarMenu" label="Dictionary"
type='menu'>
```

Next we add the **<menupopup>** tag and the menu items which, handily, appear inside a **<menuitem>** tag and follow the formula of the original link. A standard separator can also be inserted using **<menuseparator/>** and, on testing you'll see that the button automatically acquires a small down arrow to signify that it is a menu button.

```
<toolbarbutton id="DicToolbarMenu" label="Dictionary"
type='menu'>
<menupopup>
  <menuitem label='Dictionary' oncommand="dictoolbar
LoadPage('http://dictionary.reference.com')"/>
  <menuitem label='Thesaurus' oncommand="dictoolbar
LoadPage('http://thesaurus.reference.com')"/>
  <menuseperator/>
  <menuitem label='Word of the Day' oncommand="dictoolbar
LoadPage('http://dictionary.reference.com/wordoftheday')"/>
</menupopup>
</toolbarbutton>
```

Of course, you can add as many elements and toolbar separators as you fancy. To complicate things further though, we ought to add a search box – everything should have one!

We start by adding the various elements to the block of Javascript. After defining the function, we need to get the menuitem element (dictoolbarSearchText), extract its value and finally create the search terms:



The Personal Toolbar can be used to provide quick access to your favourite sites.



Here is our fledgling toolbar rendered within *Mozilla*'s content window.

```
function dictoolbarSearch(){
  var mySearchBox = document.getElementById
('dictoolbarSearchText');
  var mySearchTerm = mySearchBox.value;
  var mySearchString = "http://dictionary.reference.
com/search?q="
  + mySearchTerm;
  dictoolbarLoadPage(mySearchString);
}
```

Once the Javascript is done, we need to build the XUL and add the input box where the user can enter a search term. This is done with the **<textbox>** tag. We add an initial phrase to the box using the **value=** element and ensure these words disappear, or are reset to an empty string, when the user clicks the box with the **onfocus** event handler. Finally, as we're aiming at writers, we can assume they'll type a word rapidly and hit RETURN, so we use another event handler (**event.keyCode**) on that key to launch the search.

```
<textbox id="dictoolbarSearchText" value="Find a word..."
onfocus="this.value=""/>
onkeypress="if(event.keyCode == 13){dictoolbarSearch();}"/>
```

So far, so texty; but the web is a visual medium, and our toolbar should be too. Inserting a graphic at the far left of the toolbar is simply a matter of adding an **<image src='[path]'** tag with a local or remote location, just as you would a standard **** tag in HTML. *Mozilla* can handle PNG, JPG, GIF and BMP format, but if you stick with PNG you can include alpha blends to



Netscape's CC Kit

Client Customisation, not Carbon Copy!

If you're intending to build a browser for distribution with only a few alterations, Netscape – *Mozilla*'s corporate sponsor – has a *Client Customisation Kit* which, Netscape claims, will let you create a customised distro in 10 minutes. The kit requires a 39MB download, is backed up with a hefty online PDF customisation guide and is aimed primarily at ISPs and sysadmins intending to install across a large number of seats. To this end, the application makes it fairly easy to configure the Autorun screen, browser title bar, default homepage and a range of links through the sidebar, Bookmarks menu and Personal Toolbar.

While much of this can be done by even a casual programmer, the kit does provide access to tools which simplify the process of defining application selection and packaging a distribution. The documentation also has a lot of information on the layout of CDs for the various architectures supported and also additional information for installing on OS-X.

Be warned, though, to access the *CCK*, you have to sign up for Netscape's browser distribution program which, we fear, may entail receiving spam – sorry, 'carefully selected user-targeted advertainment' – from AOL.

« create translucent elements and realistic drop shadows (if that sort of thing concerns you). It also ensures the image is displayed as you intended. Insert the tag between the toolbar and toolbarbutton id tags.

```
<image src='wordbar.png'/>
```

This places the image before the interactive elements of the toolbar, but could also be a link itself, using either standard HTML or as an XLink, in which case you'll need to define the appropriate namespace at the head of the document:

```
xmlns:html='http://www.w3.org/1999/xhtml'
```

and:

```
xmlns:xlink='http://www.w3.org/1999/xlink'
```

Introduction to applications

What this toolbar has done, although still within the framework of a browser, is introduce the concept of application development. And the various methods employed in its creation can easily be adapted and extended to build real applications capable of being distributed to all the major platforms.

```
<?xml version='1.0'?>
```

```
<!-- The obligatory Hello Word app -->
```

```
<window xmlns='http://www.mozilla.org/keymaster/gatekeeper/there.is.only.xul'>
```

```
<box align='center'>
```

```
<button label='Click here' onclick='alert('Hello World');' />
```

```
</box>
```

```
</window>
```

This is a traditional Hello World application rendered in *Mozilla's* browser window. It is possible to view the app, before getting involved in the work that packaging entails, by opening a console and typing:

```
./mozilla -chrome://[path]/[filename].xul
```

Amazingly it's not an enormous leap from here to a web browser; communication technologies are so intrinsic to the *Mozilla* project that a simple web browser can be thrown together in a few lines.

```
<?xml version="1.0"?>
```

```
<?xmlstylesheet href="chrome://global/skin" type="text/css"?>
```

```
<window id="browser-window" title="My first Web Browser"
```

```
xmlns="http://www.mozilla.org/keymaster/gatekeeper/there.is.only.xul"
```

```
width="800" height="600">
```

```
<browser src="http://www.linuxformat.co.uk" flex="1"/>
```

```
</window>
```

We've added an xml-stylesheet reference, given the window a name and size and, with the use of the **<browser/>** tag, provided a URL for *Mozilla* to render. The next visual element in a browser window (we're going from the top) is the menubar, and creating this is done in a similar way to the toolbar we made earlier.

```
<menubar id="topbar">
```

Moz-tech 101: part three

A brief look at some of the underlying technologies that make *Mozilla* so flexible

XML, the eXtensible Markup Language, is infiltrating every aspect of computer use from data storage (*MS Office II* will have XML-based file formats) to graphics – thank to the magic of SVG – and *Mozilla's* own XML User interface Language. It is currently the buzz word that subsumes all other buzz words. But what exactly is it, and why all the fuss?

It's probably best to start with what it isn't. XML is not a replacement for HTML; while HTML was developed to display data and concentrates largely on the visual, XML has been devised to describe data and addresses what data actually is. In fact, these technologies make great partners, with XML defining the structure of data and HTML displaying it. The big difference in the way XML and HTML are written is that while HTML is restricted to a standard set of tags, XML tags are user defined which, when you think about the volume and types of data out there, is vital. Using XML also enables you to keep data and formatting information separate, which means designers can concentrate on designing without worrying about the specifics of the content.

There is a standard way to format a valid XML document, beginning with the version declaration, which we've seen in the various scripts we've done in this series:

```
<?xml version='1.0'?>
```

Next we have the actual data in our file:

```
<root>
```

```
<child>
```

```
<subchild>
```

```
</subchild>
```

```
</child>
```

```
</root>
```

It's worth noting that the tags above are user defined so if we were creating a data format for an online supermarket we would use tags to represent, for example, aisle, product, ingredients, nutritional information, price and special offers. And we might end

up with something like this:

```
<aisle>Tinned Vegetables
```

```
<product>
```

```
<name>Baked Beans</name>
```

```
<ingredient>Beans (49%)</ingredient>
```

```
<ingredient>Tomatoes (27%)</ingredient>
```

```
...
```

```
<nutritional_info>
```

```
<energy>158k cal</energy>
```

```
...
```

```
</nutritional_info>
```

```
<price>29p</price>
```

```
</product>
```

```
<product>Mushy Peas
```

```
...
```

```
</product>
```

```
</aisle>
```

We could take care of displaying the data by referencing a remote CSS file which would tell the browser how to render the various tags. It's also possible to use Javascript to format and display XML data.

Standards must be a good thing, there are so many of them? Into the already brimming cauldron of World Wide Web Consortium (W3C) standard is a new point at which XML and HTML meet: the recently ratified XHTML. This standard has been designed with the future in mind, offering the extensibility of XML and the compatibility of HTML, and encouraging web developers to get into the habit of creating well-formed documents. XHTML is a good starting point if you're already familiar with HTML and want to acquire XML skills.

Mozilla and XML

As well as being an XML compliant web-browser, *Mozilla* uses a number of applications of XML, the most obvious being XUL, which is used to define various elements of

the user interface. In our basic browser, for instance, we have used it to configure the content window, menubar and navigation buttons, but these elements are just pretty widgets without something to connect them. In most cases, the connector will be Javascript, but it could just as easily be any programming language.

There are a number of inherent benefits to basing everything on XML, including the ability to embed other XML technologies such as MathML and SVG within a well-formed XUL document. It also makes localisation a less problematic task and works in a uniform way on platforms ranging from PCs to mobile phones.

BOOKS

XML is a huge subject, so you're going to need a good reference if you intend to use extensively.

■ **XML in a Nutshell**, O'Reilly UK; ISBN: 0-5960-0292-0, £28.50 – Covers everything from basic syntax and Document Type Definitions to XHTML, though info on more peripheral elements like MathML and SVG is limited.

■ **Professional XML**: Programmer to Programmer, Wrox Press Ltd; ISBN: 1-8610-0505-9, £32.99 – For a more comprehensive view on everything XML is capable of, and weighing in at 1,200 pages, *Professional XML* is an invaluable resource, though it is getting a little long in the tooth and may soon be updated with the publication of XML 1.1.


WEBSITES

<http://www.mozilla.org/catalog/architecture/xul/> -

Mozilla.org's enormous collection of documents on XUL is a good place to start.

<http://www.xulplanet.com> - XULPlanet is almost as comprehensive as *Mozilla.org* but is far more readable.

<http://www.w3.org/XML/> - Home of the official XML specification and also a large selection of links to tutorials, primers and programmers references.



A toolbar can contain any elements that *Moz* can deal with including images, Javascript and other interactive elements.

```

<menu label="File">
  <menupopup>
    <menuitem label="New"/>
    <menuitem label="Open"/>
    <menuseparator/>
    <menuitem label="Close"/>
    <menuitem label="Exit"/>
  </menupopup>
</menu>
<menu label="Edit">
  <menupopup>
    <menuitem label="Copy"/>
    <menuitem label="Paste"/>
  </menupopup>
</menu>

```

There are various attributes we can assign to the menubar. 'disabled' which can be set to true or false, determines whether the menu will be available or greyed out when the application is launched – false is the default. **accesskey** enables you to define which key will activate the menu (or **menuitem**). When called, *Mozilla* scans the **menu label** and if the defined letter is included, it underlines it on the menubar. Furthermore, we can also set keybindings for speeding up various common tasks such as cutting and pasting.

Good news

The good news is that as we can nest XUL elements to our hearts' content, adding a submenu to, for instance, the **New** element simply involves inserting the following over **<menuitem="New"/>**:

```

<menu id="new-menu" label="New" accesskey="N">
  <menupopup id="new-popup">
    <menuitem label="Page"/>
    <menuitem label="Image"/>
  </menupopup>
</menu>

```

Using a similar layout you can add a popup menu which appears when the user right-clicks on an element in the window using a **<popupset>** **<popup>** container. This, the context popup is one of three available to *Mozilla* developers. And finally on this topic, you can build tooltips which will help your viewers navigate around the various elements of your application. You can opt for standard tooltips, which will be drawn in *Mozilla's* easily readable light yellow/black scheme or you can make your own style. The former involves simply adding a **tooltiptext='your text'** element to your menu, while for the latter, you'll first have to create a tooltip id and then add a reference to it from the relevant **menuitem**.

```

<tooltip id="open" orient="vertical" style="background-color:
#F47373;">
<description value="Click here to open a new document"/>
</tooltip>

```



A very simple Hello World app needs user input and provides a visual response.



The LXF Browser gets some – important – menubar items.



Tooltips default to pale yellow, but you can define any colour you wish.

BELOW: Menu buttons follow the look and feel of your main browsers installed theme.

```

...
<menuitem label="Open" tooltip='open'/>
  Traditionally below the menubar we would find the navigation toolbar, and creating this is a similar job, substituting buttons for the menubar entries above:
  <toolbar>
    <toolbarbutton id="back-button" class="toolbarbutton-1"
      tooltiptext="Back"
      oncommand="goBack();"
      observes="canGoBack">
    </toolbarbutton>
    <toolbarbutton id="forward-button" class="toolbarbutton-1"
      tooltiptext="Forward"
      oncommand="goForward();"
      observes="canGoForward">
    </toolbarbutton>
  </toolbar>

```

You'll notice that, as we have previously configured the `/global/skin/` Chrome at the head of the document, the browser has inherited the look from whatever skin/chrome you're currently using, which saves you some time.

The one thing missing from our menus and buttons, and the one thing that actually makes them worthwhile, is functionality, which we don't have the space to cover this month. [LXF](#)



NEXT MONTH

Now we've got our toolbar code together we need to package it all up, which will be the main task next month when we probe the mysteries of *Mozilla's* Cross-Platform Installer (XPI). We'll also attempt to shoehorn in a brief look at adding RSS sources to the sidebar and delve a little deeper into application development.

TutorialGIMP



ART AND IMAGE EDITING

GIMP texture magic

PART 2 Realistic surfaces make all the difference between your images looking merely good or being really *great*. Michael J Hammel takes us through every Noisifying detail.

An unmarked surface of metal skin wrapped around an alien ship adds little to the story behind the ship or its occupants. But add some patches of rust, a few scorch marks and impact creases in the metal and you find yourself face-to-face with battle-scarred travellers recently escaped from a nightmare on a watery planet.

Textures are one of the basic building blocks for both 2D

illustration and 3D animation. The skin of a human, the liquid feel of text, and the surface of metal are all common forms for computer generated textures. Creating realistic textures in the *GIMP* is less complex than it might seem. And the key lies where you might never have guessed: the generation of noise-filled clouds.

GIMP provides two plug-ins for generating clouds, both of which are called "Plasma". The first comes with the stock *GIMP* distribution while the other can be downloaded as source code from the *GIMP Registry* or retrieved from various online resources in ready-to-use form. Both will load themselves into the Filters->Render>Clouds submenu found in the Canvas menu (click and hold the right mouse button in a Canvas window to get the Canvas menu). The version from the *GIMP Registry* should show in the menu as "Plasma2". A "plasma" is just a type of gas, and clouds are roughly a gaseous form of water, thus the name of the filters.

In all our examples we'll be starting by rendering a desaturated cloud. A desaturated cloud is like a grayscale cloud except that a desaturated cloud is really a colour cloud with all the colours washed out. You can add colour back into a desaturated cloud. A true grayscale cloud cannot be coloured. After we create the desaturated cloud, we'll be adding different amount of desaturated noise. By "noise" we mean random dots distributed about the current layer. The noise will then be used as a bump map to generate height within the cloud, or blurred to provide streaks.

Printer resources

For your *GIMP* masterpieces

www.cups.org/

CUPS provides networked printing support and a spooling interface.

<http://gimp-print.sourceforge.net/>

GIMP-Print goes way beyond just a *GIMP* print plug-in. It's a full-featured set of drivers for many printers. Its drivers are so well done they are considered better than some of the original

manufacturers drivers and are being used by many MacOS X users now as well.

www.linuxprinting.org/foomatic2.9/


Foomatic is the glue between print spoolers and printer drivers.

www.linuxprinting.org


This is the place to start to find out if a driver is available for your particular printer.

CONCRETE


One of the easiest textures to create is a simple concrete or plaster background. Such textures can be used to turn a basic 3D box into a cement block or any flat surface into a pavement (US: sidewalk!) or road.



1 Start with white background in a new Canvas window. Add a new layer and fill it with grayscale clouds using the Plasma2 filter (Filters>Render>Clouds>Plasma2). Set the Noise distribution to Exponential and the Noise Amplitude to 0. Then set the Scaling function to Norris-like and the Scaling parameter to about 1.7. This should produce a slightly spotty cloud shape as opposed to soft wispy clouds. The amount of detail you'll get will depend on the Random Seed. You can either try clicking the Time button several times to find a reasonable preview or try using the value 1044227968 (which was used in this example) as the Random Seed.



2 Now apply some noise to this image using the Noisify filter (Filters>Noise>Noisify). Set the color channels to be uniform by deselecting the Independent button. Then move one of the channel (Red, Green, or Blue) sliders to about 0.30. Apply the plasma layer as a Bump Map (Filters>Distorts>Bump Map) to the white layer. The settings of the Bump Map filter depend on your cloud layer. Play with the Azimuth, Elevation and Depth and scan around the Preview a little to see what you will get from those settings. The Map Type options will change the image a bit, but are less important than these other three settings. Be sure to turn off the visibility of the cloud layer. After you have applied the Bump Map the white layer will hold your concrete texture.



3 Variations on this technique will produce many different textures. There are two plasma filters: Plasma and Plasma2 (the latter is available from the *GIMP Registry* but is not included in the stock *GIMP* distribution). Using the Plasma (Filters>Render>Clouds>Plasma) filter instead of Plasma2 will generate a more evenly distributed pattern of concrete. This example used the same settings for Noisify and Bump Map as the other example, but started with a desaturated (Image>Colors>Desaturate) cloud pattern generated with the Plasma filter instead. While it looks like concrete, adjustments to the Contrast and Brightness can soften the appearance and provide you the initial pattern for more complex textures like skin or cloth!

WATER


This next texture is again one of the more simple ones, due largely to the availability of two plug-ins in the stock *GIMP* distribution: Ripple and Waves. Another plug-in, Water Surface, is available from the *GIMP Registry* or its own Web site, but we'll look at making this texture from scratch, as it is a useful skill.



1 As usual, start with a cloud layer. Plasma2 is best for this. The usual Plasma will work, but generates a cloud that fills the layer more uniformly than we want. Apply a radial motion blur (Filters>Blur>Motion Blur) with a length of 3. After this completes, apply a Gaussian IIR (Filters>Blur>Gaussian IIR) blur using a radius of 35. We want this to be heavily blurred.



2 Next, add some Waves (Filters>Distorts>Waves). Make the amplitude 47, the phase 210, and the wavelength 43. Duplicate this layer. Flip the duplicate both horizontally and vertically using the Flip Tool in the Toolbox. Then set the layer Mode to Difference on the duplicate layer.




3 Merge the two layers. The result is good, but for a little variation in these too-perfect waves, we'll apply a ripple (Filters>Distorts>Ripple). Set the period to 128 and amplitude to 6. Make it a sine wave with wrapped edges. Finally, colorize the layers using the Color Mode of the Bucket Tool. A good color for an aqua appearance would be a hex triplet value of #009199.




TutorialGIMP

RUST


A more complex pattern to produce is that of metallic rust. The key ingredients for rust are both color and pattern. Be sure to experiment with your color choices to find one that meets your minimum standard.



1 Once again we start with a single white layer. Fill this layer with a dark, reddish-brown color. Try a hex triplet value of `#a46711` in the Color Selection window. Your tastes will vary here. Next, add noise using the Noisify filter (Filters>Noise>Noisify). Set the color channels to be uniform (not independent, just as in the concrete examples) and the channel settings between 30 and 45. Then duplicate this layer.




2 Make sure the duplicate layer is the active layer. Distort the layer with some wind (Filters>Distorts>Wind). Set the Threshold between 7 and 20 and the Strength between 2 and 5. Make the wind blow from the left across both edges using a blast-style effect. Add a transparent layer above the duplicate layer. This will make the transparent layer the top layer. Render some grayscale clouds into this layer. Like the earlier concrete example, the clouds here should be spotty. If you have the Plasma2 filter, use it with a Uniform noise distribution set to an amplitude of about 1.08 or a Cauchy distribution with an amplitude of about -1.12. Norris scaling for both should be between 1.15 to 1.25. Use Brightness/Contrast to reduce the amount of white so that there is about a 50/50 ratio, and preferably a little more black in the clouds layer than white.




3 Colorize the cloud layer using the same reddish-brown color you used on the original white background layer. To do this, make sure the cloud layer is active and type `Ctrl-A` to select the entire layer. Then double click on the paint bucket in the Toolbox – this is the Bucket Fill tool. In the Tool Options window that opens, change the Mode to Color. Then click in the Canvas window. The image should take on a reddish brown tint. Now set the layer Mode for the Cloud layer to Multiply. You'll see some of the texture show through the white parts of the cloud layer now.



4 Add a layer mask to the duplicate layer (right click on the layer name, then select Add Layer Mask). Select the cloud layer by clicking on its layer to make it active and typing `Ctrl-A` to select the entire layer. Copy this selection into the layer mask. Use `Ctrl-C` to copy the selection. Then click on the duplicate layer's mask in the Layers, Channels, and Paths window and type `Ctrl-V` in the Canvas window. With the layer mask still selected, double click on the Flip tool (the left/right arrow in the Toolbox). With the Tool Options window open, click once on the Canvas window, then switch directions in the Tool Options window and click again in the Canvas window. This has the effect of flipping the layer mask both horizontally and vertically. Change the layer mode for the duplicate layer to Divide (aka Dodge).




5 Turn off the visibility of the Cloud and duplicate layers for a moment. Click on the Background layer to make it active. Now apply some more wind to this layer using the same settings as you used on the duplicate layer. At this point you can add some waves to the layer, although it may or may not add any real distortion to the overall effect, depending on the amount of white in the cloud layers. Set the Waves (Filters>Distorts>Waves) Mode to Smear, disable Reflective, and set the Amplitude to about 7, the Phase to about 100 and the Wavelength to about 16.



6 Select by color the black regions in the top layer (the cloud layer). Increase the Fuzziness threshold to grab more of the black area. Add a transparent layer to the top of the layer stack and fill it with the original rust color (you should still have this set as your Foreground color). You can use the Bucket Tool to fill the selection in this transparent layer but be sure you reset the Mode in the Bucket Tool Options window to Normal. Finally, change the Layer mode to Addition for the transparent layer. Turn off the Selection by typing `Ctrl-Shift-A` in the Canvas window. Viola! Rust.

GRUNGE

Grunge is a popular style these days for print work. Generating grunge is a very personal thing – much like the music form – so this texture example may or may not produce what you might call grunge. But it does for us.




1 Start with clouds. The clouds should be relatively soft, little to no harsh edges if possible. Less cloud fill is probably better for a grunge look that doesn't overpower the rest of the image. If you can't get one yourself, try a Random Seed of 1044492096, a Uniform Noise distribution with an amplitude of -1.12 and Norris-like scaling with a Scaling parameter of 1.24. Add noise to this with Noisify using non-independent color channels set to between 0.05-0.10.




2 Use the bump map option of the Emboss filter (Filters>Distorts>Emboss) on this layer. Use a low azimuth and elevation (25-35), and a moderately high depth (60-70). We'll call this our "bump map" layer over which everything else will be blended. Duplicate this layer.



3 Now comes the trick. We need to set up the *GIMPressionist* plug-in (Filters>Artistic>*GIMPressionist*) – it's complex but can produce fantastic textures. We're going to apply this to the duplicate layer. Select the Cubism preset and apply it. This sets the initial parameters to something useful, but we're going to modify them a bit. First, change the Paper to canvas2.pgm. Change the brush to sponge02.pgm, you can scale it up using the Aspect Ratio if you like. Next change Orientation span angle to 360 with a start angle > 20 and using Random orientation. In the Size tab, leave the Sizes value at 1, with a minimum of 4 and a maximum of 13 using the Value size setting. Finally, set the Placement Stroke density to 15 and make them Evenly distributed. This might be a little slow to calculate. Give it time. When you click on OK the filter will begin processing. Don't be surprised if it takes a while to finish. You can Update the preview before you do so, but that will add to the time it takes to finish this step.



4 Use Color Select (Select>Select by Color) to select a region of the Bump Map layer. Create a transparent layer above this layer and fill the selection with black. Move the transparent layer to the top of the layer stack. Apply the Twist filter (Filters>Distorts>Twist) to the *GIMPressionist* layer. In the Twist filter window, select the Cosines function and Effect 4 in this filter. You shouldn't need to change any of the other parameters.



5 Use Select by Color on the *GIMPressionist* layer. Again, create a transparent layer above this and fill the selection with black in this new layer. Be sure the new transparent layer is moved to the top of the layer stack. The order of layers and blend modes should be:

1 Black filled selection from *GIMPressionist* layer with layer Mode set to Divide.

2 Black filled selection from Bump Map layer with layer Mode set to Multiply.

3 *GIMPressionist* layer with layer Mode set to Screen.

4 Bump Map layer with layer Mode set to Normal.



6 Variations of blend modes and possibly the order of layers may help produce the effect you want. Also, flipping the direction of the original or duplicate layer may help. Color can be added to any layer or a top layer of colour can be added with the layer Mode set to Color or perhaps some other blend mode.



7 A selection of grunge end results – it all depends on what values that you choose.

NEXT MONTH

Two issues digital artists face in their daily work are the need to undo changes made to previously saved work and the desire to modify colour content with simple brush strokes. For the former, they turn to

GIMP's layer masks to generate non-destructive changes. For the latter, they look to GIMP's blend modes. Next month, we'll aim our tutorials at both of these highly desirable yet easy-to-use features. [LXF](#)

USING SHELL SCRIPTS

Practical PHP Programming

This month, **Paul Hudson** takes a look at how you can make use of PHP as more than just a web scripting language.



As you probably know, PHP is a recursive acronym for *PHP: Hypertext Preprocessor*, so why should a product designed for the web ever stray from its original goal? Quite simply, because PHP is cool, and if I could use PHP to do *everything*, I'd certainly give it a try, if only to satisfy my programmers' love of a challenge!

A Command Line Interface Server Application Programming Interface (CLI SAPI) has existed for PHP since v4.2, but only with the release of 4.3 has it come to maturity – it's now built automatically, and can be configured to be installed automatically also. **This article assumes you have PHP 4.3 installed**; the CLI version of PHP 4.2 doesn't have some of the necessary functionality to run the scripts below.

Why use shell scripts?

If you're not already using shell scripts as part of your daily life, you're probably expending too much energy on little tasks! Shell scripts are a great way to automate repetitive tasks, receive and process user input, and create basic applications.

Using PHP to build your shell scripts allows you to not only take advantage of the power and flexibility inherent to the language, but it also brings you one more step closer to the magical day when you can type `rm /usr/bin/perl` and not worry about the repercussions!

PHP-CLI can also be used as part of a daisy-chain between other commands. For example: `php -f foo.php | sort -u > myfile.txt`. In my experience, slotting PHP into my existing toolkit in this way has been the most profitable use – once PHP is up there with `/sort/`, `/grep/`, and `/strings/` in terms of common use, you know you're on the right track!

Before you start

Owing to the fact that PHP was originally developed for web use, there are quite a few variations between normal PHP use and CLI use. Among other things, various `php.ini` settings are automatically overridden, some helpful constants are set for you, and no HTTP headers are output when the script is called.

Take a look at the box *PHP-CLI Differences* on page 76 for a list of the new constants and also how your `php.ini` will be affected. Unless you have already installed the CLI SAPI, you should also read the box entitled *Installing the CLI SAPI* on the opposite page before you run the scripts.

Also, it's important to note that there are two ways you can execute your PHP scripts. The first way is to use `php -f`

`/path/to/yourfile.php` – this simply calls the binary and passes your script as a parameter and is a little clumsy, if effective. My preferred manner is to `chmod +x` PHP files I wish to call from a shell, then add an appropriate shebang line at the top, something like this: `#!/usr/local/bin/php`.

Your first script

Writing CLI scripts is, for the most part, similar to writing PHP scripts for web use – you have all the same functionality available to you as you would when writing for the web. As such, the scripts we're going to look at will specifically highlight special functionality and clever possibilities for the CLI SAPI.

To begin with, we are going to look at `argc` and `argv` – two variables often overlooked in recent times. These two combined allow you to iterate through parameters passed to your script, with parameter `0` being the script itself.

```
#!/usr/local/bin/php
<?php
echo "$argc arguments were passed. In order: \n";
for ($i = 0; $i <= $argc -1; ++$i)
echo "$i: $argv[$i]\n";
?>
```

Save that in your public HTML directory as `cli1.php`. Note that you may need to amend the first line depending on where your CLI SAPI resides.

Line one is simply a shebang line (contraction of “sharp” (#) and “bang” – how else could you pronounce an exclamation mark?) that points to where the CLI SAPI is on your system.

Line three prints the value of `$argc` out to the system. The usage of this becomes clear in lines four and five; it correlates directly to “the number of parameters passed to the script minus 1”. It's minus one because `$argv`, the values of the parameters that were passed, is a zero-based array like all others in PHP.

In lines four and five, we use `$argc` to iterate through `$argv` and output each element as we go. So, with the script saved as `cli1.php`, run `chmod` like this:

```
chmod +x cli1.php
You should now be able to run the script by itself by typing:
./cli1.php
For output, you should see:
1 arguments were passed. In order:
0: ./cli1.php
```

Try running the script again – this time, pass in random arguments. For example:

Important notice

A security hole was discovered in the release of PHP 4.3.0 regarding the availability of the CGI binary on servers. Please download and use PHP 4.3.1 if you wish to upgrade to use the information in this article!

```
./cli1.php --foo --bar baz=wombat
```

This time, your output should be:

```
4 arguments were passed. In order:
```

```
0: ./cli1.php
```

```
1: --foo
```

```
2: --bar
```

```
3: baz=wombat
```

And so, as you can see, you have a building block upon which you can base your first shell application.

Getting down and dirty

Before we continue with reading and writing from standard streams, I would like to give you an example as to how you can integrate a PHP shell script with other applications to form a useful tool.

Back in LXF32, we talked about handling database information for a fictional golf club. If you don't have the article to hand (backissues are available – see page 97), use these SQL commands to recreate some of the data:

```
CREATE TABLE staff (Name VARCHAR(255), Age INT, Job VARCHAR(255), Pay INT);
```

```
INSERT INTO staff VALUES ('Joseph Smith', 29, 'Security Man', 29000);
```

```
INSERT INTO staff VALUES ('Harold Barnes', 32, 'Security Man', 29000);
```

```
INSERT INTO staff VALUES ('Carmen Hobbes', 30, 'Security Man', 30000);
```

Now, here's a PHP script that will demonstrate a little of the power of using PHP for shell scripting:

```
#!/usr/local/bin/php
```

```
<?php
```

```
if ($argc < 3) {
```

```
    echo "Usage: phpcli2.php [where clause] [match clause]\n\n";
```

```
    exit;
```

```
}
```

```
$whereclause = $argv[1];
```

```
$matchclause = $argv[2];
```

```
mysql_connect("localhost", "yourusernamehere", "yourpasswordhere");
```

```
mysql_select_db("lxfdb");
```

```
$result = mysql_query("SELECT Name, Age, Job, Pay FROM staff WHERE $whereclause = '$matchclause'");
```

```
$nummatch = mysql_num_rows($result);
```

```
if (!$nummatch) {
```

```
    echo "No rows matched!\n\n";
```

```
    exit;
```

```
}
```

```
echo "Found $nummatch rows:\n\n";
```

```
while ($r = mysql_fetch_array($result)) {
```

```
    extract($r);
```

```
    $Pay = '£' . number_format($Pay);
```

```
    echo " Name: $Name\n";
```

```
    echo " Age: $Age\n";
```

```
    echo " Job: $Job\n";
```

```
    echo " Pay: $Pay\n\n";
```

```
}
```

Accelerators

CLI SAPI + PHP caching solutions

Just as a small side note, both *Zend Performance Suite* and *ionCube PHP Accelerator* do not cache CLI scripts, so you should not accept any application speed-ups here.

Installing the CLI SAPI

As easy as 1-2-3...?

The CLI SAPI only really finalised as of PHP 4.3, so it's imperative that you install this particular version. Some Linux distributions (particularly Debian) are unlikely to have PHP >= 4.3 available for easy install, so you may need to compile from source. Please refer to LXF30 with regards to complete configuration instructions – this reference will cover only compiling and installing the CLI SAPI.

On a bare machine, the process is quite simple:

- i) tar xpvj php-4.3.0.tar.bz2
- ii) cd php-4.3.0
- iii) ./configure
- iv) make
- v) make install-cli

All but the last step is predictable enough. **make install-cli** takes the CLI SAPI from `./sapi/cli/php` and, on my Debian box using the commands above, copies it into `/usr/local/bin`

Depending on the version of PHP you download (see the note in the box at the bottom of the opposite page), you may need to alter the commands above ever so slightly – eg **cd php-4.3.1**, etc.

You may encounter problems when 'rolling your own' PHP due to the naming clash between the CGI and the CLI. Sadly, the situation isn't really very simple. Here's how it actually breaks down in its most basic terms:

When executing **make**, the default action is to build both the CGI and CLI and put them into `./sapi/cgi/php` and `./sapi/cli/php` respectively. Both the PHP CGI SAPI and the CLI SAPI are named "php" – the distinguishing feature here is the file structure. If you get mixed up later, running both of them with the **-v** argument will distinguish between the two: the CGI version reports **PHP 4.3.0 (cgi)** and the CLI version **PHP 4.3.0 (cli)**.

The result of **make install** depends on your **configure** line. If you have opted to use a specific SAPI in your configure line, for example `apxs`, the CLI is selected for use. If no alternative SAPI is entered when configuring, the CGI is selected. The selected option, CLI or CGI, is then copied to `{PREFIX}/bin/php` for use. Owing to the fact that both binaries are placed into the `./sapi` directory after **make**, you can do this copying yourself and rename them to **php-cgi** and **php-cli**, if you find it simplifies things.

The installation of the CGI binary can be overridden by using the special make target **make install-cli** after **make install** – **make install** will copy the CGI binary into `{PREFIX}/bin/php` and **make install-cli** will overwrite the CGI binary with the CLI binary. Alternatively you can use **--disable-cgi** in your configure line, which works just as well!



```
?>
```

Save that as `phpcli2.php`, and **chmod +x** it. You'll need to modify the `mysql_connect()` line to provide valid credentials for your server.

The script might look a little complicated, but all it does is query the database and extract people with the data specified in the arguments provided.

To start with, we make sure that the user has provided sufficient arguments for the script by checking whether `$argc` is under a certain amount. We need at least three arguments provided (although everything above the third is ignored), and these are: the script name (argument `0`, always passed), argument `1` (becomes assigned to `$whereclause`) and argument `2` (becomes assigned to `$matchclause`).

If you're using the staff data given above, you could run `phpcli2.php` in the following way. Our SQL query is built up by using the argument `2` and argument `3` to match specific records, and we then iterate through these records and print out each individual result. So, to pull out all records where a staff member is being paid £29,000, the script can be called like this:

```
./phpcli2.php Pay 29000
```

As the recordset is limited, only two results are returned. Getting the feel of how this can be useful? Try combining the script with other tools – here's an example to get you started:

```
./phpcli2.php Pay 29000 | grep Name | sort
```

Here we take the same script and pipe it through `/grep/` so that it will filter for lines containing **Name**, then finally pipe the result of `grep` through `sort`, so it will show the results in alphabetical order. The end result, as you will see when you run the script, is that the names of all people in the database earning £29,000 a year are outputted, sorted alphabetically.

Getting into the swing of things

Now that you're comfortable with the idea of integrating PHP directly with your other tools, I want to look at the CLI SAPI streams that are opened for use in all CLI scripts. Owing to the fact that it's very common to use `/stdin/`, `/stdout/`, and `/stderr/` inside shell scripts, the process of opening and closing streams to access these has now been automated through the use of the constants `STDIN`, `STDOUT`, and `STDERR` respectively.

These constants, when combined with basic I/O functions like `fgets()` and `fputs()`, allows you to easily read and write to the console. `fputs()` writes to a given stream, and takes two

parameters: the stream to write to and the string it should write. It also takes an optional third parameter that is the length of the string to write, but this can be omitted if you wish to write the entire string. Its opposite, `fgets()`, reads data in from a stream, and takes one parameter: the stream it should read data from. It also has an optional length parameter, to specify how much data should be read. If you don't specify this parameter, the length defaults to 1024 bytes. The return value of `fgets()` is the string that was read in from the stream.

The concept of streams, which is relatively new in PHP, is a unified approach to handling files, pipes, sockets, and other I/O resources. What this means for us as developers is that `fputs()` and `fgets()` can read from and write to files, HTTP connections, or the console – all using the same command. Here's an example script to demonstrate `fputs()` and `fgets()` in use for working with the console:

```
#!/usr/local/bin/php
<?php
fputs(STDOUT, "\nThe Amazing Favourite Colour Script\n");
fputs(STDOUT, "What is your favourite colour? ");
$someText = strtolower(trim(fgets(STDIN, 256)));
fputs(STDOUT, "Your favourite colour is $someText!\n\n");
?>
```

As you can see, the script makes extensive use of `fputs` to output data to `STDOUT`. This may seem merely like a trickier way of writing using `print()`, but it's much more flexible because one could use `STDERR` in place of `STDOUT` and redirect error output to a specific place.

The only part of the script which might confuse you at first is the call to `fgets()`. In the example, we read from `STDIN` in order to allow the user to enter up to 256 bytes of text. This input is then passed through `trim()` and then `strtolower()` before it's placed into the `$someText` variable. The `trim()` function is used here because if users do not wish to enter the full 256 bytes of information, they can terminate their input by hitting Return. If their input *isn't* sent through `trim()` first, their string is stored in `$someText` with the extra `\n` (new line) intact.

The last line of the script outputs our user's string back to the screen, to complete the circle. Reading and writing data inside the console is on the whole quite easy and straightforward, as the above example shows. Save the script as `cli3.php` and try editing it a little yourself in order to get a good grasp of how it works – be sure to try redirecting different parts

PHP CLI differences

Key things to remember

In many ways, the CLI SAPI works in the same way as the CGI SAPI and other SAPIs. However, it's important to keep track of where it differs and why, because it will directly affect how you write your scripts.

The three easiest to remember differences are that the CLI SAPI doesn't output headers, doesn't use HTML formatting for its error messages, and also doesn't change the working directory to that of the script. The latter is particularly important and should be noted – when using the CGI SAPI, calling `/home/paul/foo.php` would set the working directory (where it references files and other objects from) to `/home/paul`. This is not the same for the CLI SAPI – it uses the location of the CLI SAPI binary as its working directory.

Four `php.ini` directives are automatically overridden when using the CLI SAPI. These are: `html_errors`, `implicit_flush`, `max_execution_time`, and `register_argc_argv`. `html_errors` is set to **False** by default because, naturally, one is working inside a shell where HTML would only clutter things. `implicit_flush` is set to **True** by default. For those unfamiliar with how PHP works, it internally buffers output, sending it to output in chunks. When `implicit_flush` is set to **True**, PHP sends any output from `print`, `echo`, etc, straight to output without trying to buffer it. In PHP's normal operating environment, the web, this can slow things down tremendously, and is generally not advised. However, it wouldn't be very helpful for PHP to buffer things when the output is a

shell, because users want their feedback immediately. As such, `implicit_flush` defaults to **False** for the CLI SAPI.

Owing to the fact that shell scripts can run for much longer than web scripts, `max_execution_time` defaults to **0** for the CLI SAPI. This allows your shell scripts to have an infinite execution time. Finally, `register_argc_argv` is automatically set to **True** to allow you easy access to `argc` (the number of arguments passed to your script) and `argv` (an array of the arguments themselves).

The final key difference between the CLI SAPI and other SAPIs is that the CLI SAPI automatically sets the constants `STDIN`, `STDOUT`, and `STDERR` for you, simply because you are likely to work with at least one of these three fairly regularly.

of the output, as this is one of the best ways to take advantage of the STD* constants.

Sending code direct to PHP

For more advanced users, there is one more way to take advantage of PHP for console programming: passing code directly to the PHP CLI binary.

Through the use of the `-r` argument, it becomes possible to enter PHP code on the command line. The opening and closing PHP tags are not required, as the `-r` tag is designed to execute pure PHP code.

Here is a very basic example of `-r` usage:

```
php -r 'phpinfo();'
```

As you can see, you get the PHP module information printed straight to the console. However, it is a very basic example, and doesn't highlight the key problem of sending code direct to PHP with the `-r` argument.

The problem, which may cause many users headaches because of its elusive nature, is that many shells (I use *bash*, which is "affected"), have variables of their own, and if you use double quotes, will perform variable substitution even before PHP gets hold of the code. As a result, this seemingly innocuous code will fail:

```
php -r "$abc = 'def';"
```

Bash, and many other shells, will evaluate `$foo` and find it unset, and will pass a blank on to PHP. Therefore, what PHP will see is this:

```
= 'def';
```

The solution to the problem, as I mentioned already, is to use single quotes with `-r`. However, you will need to constantly be on guard when using backslashes for escaping, or, even harder, trying to get shell variables in your code.

The grand finale

To finish this article, I wanted to show you something to help widen your horizons with regards to PHP CLI development: *dialog*. This neat little program, probably already installed on your system (if it isn't, you will need it to continue), allows you to create basic dialogs for shell environments. Furthermore, we can call *dialog* direct from our PHP scripts and grab the response.

In order to demonstrate how *dialog* is called, some of the options that can be passed to it, and how to grab and make use of its return value using PHP, here's a big chunk of code:

```
#!/usr/local/bin/php
<?php
`dialog -- msgbox "This demonstrates PHP interacting smoothly with dialogs." 10 40 `;
$result = `dialog --menu "Please select an option" 0 0 0 "1" "File Selection" "2" "Calendar" "3" "Password Box" "4" "Checklist" 2>/dev/stdout `;
switch ($result) {
case "1)":
$result = `dialog --fselect /usr/bin 0 0 2>/dev/stdout `;
echo "The filename you selected was $result\n\n";
break;
case "2)":
$result = `dialog -- calendar "When is your birthday this year?" 0 0 22 12 2003 2>/dev/stdout `;
echo "The date you selected was $result\n\n";
```

CLI 1-2-3

Online resources to give you some ideas

At the time of writing, PHP 4.3 is still fairly new, and there are no articles concentrating on the new SAPI as yet available online. However, there are one or two that cover beta versions of the CLI SAPI and/or the CGI SAPI, which should be enough to get you started. If you're interested in learning more, your best

starting bets would be these sites:

www.zend.com/zend/spotlight/shellscrip1.php
<http://phpbeginner.com/columns/ray/cli>
www.devarticles.com/art/1/231/
www.phpbuilder.com/columns/darrell20000319.php3

```
break;
case "3)":
$result = `dialog --passwordbox "Please enter your MySQL password." 0 0 2>/dev/stdout `;
echo "The password you entered was $result\n\n";
break;
case "4)":
$result = `dialog -- checklist "Please select your Foo" 0 0 0 "1" "Some Option" "on" "2" "Another Option" "off" 2>/dev/stdout `;
echo "The options you chose were $result\n\n";
break;
}
?>
```

As you can see from the script, *dialog* does most of the hard work. Note that I'm using backticks (`), NOT single quotes. On UK keyboards the backtick key is to the left of the 1 key and above the Tab key. In PHP, the backtick is known as the *execution operator*: PHP tries to execute the contents of backticks as shell commands, returning any output.

So, in line 4, the command ``dialog -- menu "Please select...dev/stdout ``; is executed as it is contained within backticks, and any output from it is placed into `$result`. `$result` is then put through a switch/case to decide what menu option the user chose, and then *dialog* is re-run with a layout dependent on what the user chose in the first dialog.

This isn't a *dialog* tutorial, so I'm not going to go into any further depth regarding the usages above – type `man dialog` to learn more. The important thing is that you can hopefully see another way you can make good use of PHP – one could run an SQL SELECT query on a database and create a dialog with values in for the user to select, for example.

Save the script as `phpcli4.php` and try it yourself. Remember, you will need to have *dialog* accessible to try the above. Typing *dialog* by itself will bring up a list of all the parameters it can take – try amending the above example to include some new functionality of your own.

Conclusion

With the release of PHP 4.3, shell programming using PHP is likely to really start taking off. The ability to transfer your existing PHP skills direct to the CLI SAPI means that you can make the switch with the minimum of fuss: nearly everything you have learnt to far can be used immediately from the command line.

It is probable that only time and experience will allow you to fully take advantage of the CLI SAPI – certainly, the more usage that you get by playing around with it, the more you will think up new ideas for it. Given time, I can see the CLI SAPI slowly eating away at Perl's market share; make sure you're one of the first on the bandwagon! 

About Paul Hudson

Paul Hudson is a Bath-based web developer specialising in PHP and Perl. He can be emailed at HUDZILLA@PHP.NET

NEXT MONTH

Next month we'll be continuing our examination into alternative uses for PHP by looking at how PHP allows you to develop graphical applications using PHP-GTK. Advance warning: this is not for the faint-hearted! If you have any comments or suggestions about this series, please be sure to write in.

FIRST DATE

Out on the town with Python



Patrick K. O'Brien chaperones our first date with Python so we don't get up to anything until after the engagement.

Last issue I introduced you to my true love: the Python programming language. I described how we met, what Python means to me, and what I find charming about Python. I told you that I wanted you to share my love for Python, and that I would facilitate your familiarization with this fabulous programming language. So, are you ready for your first date?

First dates can be awkward and uncomfortable. It may take a while to relax and establish common ground. And what will you talk about? You can't possibly get through a first date without talking. Good communication is critical: an engaging conversation may be the only way to overcome the fact that your date didn't find the movie nearly as funny as you did. Likewise, your first date with Python should have its share of stimulating conversation. Follow this tutorial and you, too, will gain the gift of gab, develop a dexterity at discourse, and become a consummate conversationalist. Let's begin by finding out whether Python is already installed. You can determine that by typing **python** at the command line. If Python is installed you'll see something similar to the following:

```
$ python
Python 2.2.2 (#5, Feb 2 2003, 10:00:24)
[GCC 3.2 (Mandrake Linux 9.0 3.2-1mdk)] on linux2
Type "help", "copyright", "credits" or "license" for more info.
```

What you are seeing is Python's interactive mode, also known as the Python shell. The first line displays the version of Python; the second line gives some information about the operating system; the third line suggests some words you can type for more information; and the final line prompts you for input. (Type Ctrl-D when you are ready to leave the Python shell.)

The basic Python shell is a useful tool, but to really engage with Python tete-a-tete, you'll need the features provided by a graphical Python shell: a program that prompts you for input, sends your commands to the Python interpreter, displays the results, and helps you every step of the way. There are several alternative Python shells, and your humble chaperone happens to have written one himself, called *PyCrust*. (Alas, I did give in to the temptation to litter the technology landscape with yet another silly pun.)

PyCrust, the flakiest Python shell

If you have wxPython installed, running *PyCrust* is as simple as typing **pycrust** at the command line. When *PyCrust* starts it looks quite similar to the regular Python shell. But as you start using it you'll see how it leverages Python's introspection capabilities to bring you helpful features, such as syntax coloring, autocompletion, and calltips. The accompanying screen shots on these pages illustrate some of these *PyCrust* features.

With features like these, I happen to think *PyCrust* is one of the best tools for getting to know Python. But if you find that you aren't able to install wxPython, you can still follow along using the standard Python shell. Whichever you choose, I'm expecting you to hold up your end of the dialogue. Please see the box below for more details about installing.

I expect that, because a productive programming session is like a good conversation: you start with a few ideas, you never know the exact course the conversation will take, and you end up knowing your companion (and often yourself) better than when you started. Let's see how that applies to Python: launch *PyCrust* and get ready to type.

Python Objects

Let's start by talking about objects. An object is really just a way of managing information. The fundamental purpose of computer programs is to do interesting things with important pieces of information. Those pieces of information often take the form of numbers, and strings of text. Usually a variety of numbers and text strings are combined to describe important things like people, restaurant menus, movie listings and show times. Even digital images, perhaps forming an entire movie, can be described and manipulated using a programming language.

When we bundle together certain pieces of information, along with the operations that can be performed on that information, we've created what is called an 'object'. Object-oriented programming is centered on the idea that the information and relevant operations for something like a customer, restaurant receipt, or movie ticket, should be bundled together into a discrete package.

Installing Python and wxPython

GUI toolkits and such...

Both Python and wxPython are Open Source and free of charge. Many Linux distributions include Python, and some are starting to include wxPython as well. If you don't have Python installed, or you have an old version (anything prior to version 2.1.3), you can get the latest stable version (version 2.2.2 at the time of this writing) from the Python website: www.python.org/download/

RPMs are available, as is a source tarball that you can install using the typical **./configure; make; make install** routine. Check the README file for more details.

The wxPython GUI toolkit is available from the wxPython website: www.wxpython.org/

Most people find it easiest to install wxPython from the RPMs. There is a separate package containing a demo program and sample applications. Be sure to install that as well.

The wxPython package includes the *PyCrust* shell, an interactive, graphical Python shell that leverages Python's introspection capabilities to give you an enhanced environment in which to learn Python. If wxPython is installed you should be able to type **pycrust** at the command line to launch *PyCrust*.

In Python, you can define your own class of object, specifying the information and operations for instances of that class. We'll cover that capability in a later issue. For now, you should realise that every piece of information in Python is actually an object. Even things as simple as the number 2, or the string of text "I love you," are objects, and have a number of built-in capabilities. These capabilities are called 'methods,' and the following examples should help you understand the method to this object madness.

Numbers

Our first example will involve calculations using numbers. When you type in a number, Python recognizes it as such. (Don't laugh, most other languages need to be told something is a number.) Python also recognizes all the common numerical operators, like + (addition), - (subtraction), * (multiplication), and / (division). So you can use the Python shell as a desktop calculator. Try the following examples. Type in the characters you see after each >>> prompt, then hit Enter. Your Python shell session should look just like this:

```
>>> 2
2
>>> 2 + 2
4
>>> 27 / 9
3
>>> 3 * 5 - 1
14
>>> 3 * (5 - 1)
12
>>> 4 ** 3
64
>>> type(2)
<type 'int'>
```

The last example shows the result of using the type function on a numeric value. The type function is built into the Python language. It returns information about the type of any object. Some functions, such as the type function, require values to be given to it. You do so by enclosing the value in parentheses following the name of the function.

We'll discuss functions in detail in a later tutorial, but they are so handy it's hard not to introduce them early on. If you aren't familiar with functions from another programming language, you can think of them as sort of magical operations. It's enough to know that Python comes with several built-in functions, they sometimes take in values that you supply, and they always return a value.

Whenever something returns a value, such as the type function or the evaluation of a numerical calculation, the Python shell responds by displaying a text representation of that return value. That's what you see on the screen each time you hit Enter. Unless you happen to enter an erroneous statement, in which case you'll see an error message instead:

```
>>> 42 / 0
Traceback (most recent call last):
  File "<input>", line 1, in ?
ZeroDivisionError: integer division or modulo by zero
So far the numbers in our example calculations were all integers. Python is equally adept at using floating-point numbers:
>>> 2.5 + 8.4
10.9
>>> 27 / 9.8
2.7551020408163263
>>> 3.6 * (5 - 1.1)
```

Priceless Projects

Gadgets, contraptions, and treasures

BioPython – www.biopython.org/

The Biopython project provides Python libraries and applications that address the needs of the bioinformatics community. These are the folks who deal with biological sequence data, such as strands of DNA or the sequence of amino acids that make up a protein.

Numerical Python – <http://sourceforge.net/projects/numpy/>


For mathematical folks, numerical Python adds a fast, compact, multidimensional array language facility to Python.

SciPy – www.scipy.org/

SciPy supplements the Numeric module, bringing together a variety of high level science and engineering modules for such things as graphics and plotting, optimization, integration, special functions, signal and image processing, genetic algorithms, ODE solvers, and more.

Python Imaging Library – <http://www.pythonware.com/products/pil/>

The Python Imaging Library brings image processing capabilities to Python. The library supports many file formats (such as PNG, JPEG, GIF, TIFF, and EPS) for applications performing image archiving, display or manipulation.



```
14.039999999999999
```

```
>>> type(2.5)
```

```
<type 'float'>
```

Those of you with mathematical interests will be happy to know that Python recognizes other numeric types (long integers, complex numbers) and notations (hexadecimal, octal, scientific).

Strings

Text is represented by a sequence of characters, called a string. You create a string literal by enclosing the characters in single ('), double (") or triple ("'" or """) quotes:

```
>>> 'I love you'
```

```
'I love you'
```

```
>>> "I love you"
```

```
'I love you'
```

```
>>> """I love you"""
```

```
'I love you'
```



Autocompletion: whenever you type a dot after an object name, PyCrust displays a list of attributes and methods for that object.

```

<< >>> "I'm in love with you"
"I'm in love with you"
>>> "'I said, 'I'm in love with you.''"
'I said, "I'm in love with you."'
>>> """Roses are red,
... violets are blue."""
'Roses are red,\nviolets are blue.'

```

Notice how the choice of quotation marks is simply a convenience for you, allowing you to embed quotation marks or newlines (`\n`) within a text string, and has no impact on the type of string object created by Python. A string is a string is a string.

And a string is an object. That means it has innate abilities, or methods. Remember when I said I'd explain the method to this object madness? Let's do that now: let's talk about methods.

Methods look a lot like functions, except that they are associated with an object, or object class. The way to access one of those methods is to follow the object with a dot operator (`.`), the name of the method, and opening and closing parentheses. Just like functions, values can be passed inside those parentheses. Here are some examples of string methods in action:

```

>>> 'I love you'.upper()
'I LOVE YOU'
>>> 'I said I love you and I do'.count('I')
3
>>> 'Do you love me?'.split()
['Do', 'you', 'love', 'me?']

```

As you can see, a method such as the `count()` method operates on the string object with which it is associated. In our example, we counted how many times the string `'I'` appears in the string `'I said I love you and I do.'`

The last example split a string into pieces that were separated by whitespace, returning a list of strings. We'll cover lists and other container-like objects next month. For now simply recognize that a list is a sequence of comma-separated objects enclosed in brackets.

Speaking of lists, another useful built-in function that returns a list of strings is the `dir` function. When you pass it an object, it returns a list of all the attribute and method names for that object. Here is what the `dir` function has to say about a string literal:

```

>>> dir("something")
['_add_', '_class_', '_contains_', '_delattr_', '_doc_',
'_eq_', '_ge_', '_getattr_', '_getitem_',
'_getslice_', '_gt_', '_hash_', '_init_', '_le_',
'_len_', '_lt_', '_mul_', '_ne_', '_new_', '_reduce_',
'_repr_', '_rmul_', '_setattr_', '_str_', '_capitalize_',
'_center_', '_count_', '_decode_', '_encode_', '_endswith_', '_expandtabs_',
'_find_', '_index_', '_isalnum_', '_isalpha_', '_isdigit_', '_islower_',
'_isspace_', '_istitle_', '_isupper_', '_join_', '_just_', '_lower_',
'_lstrip_', '_replace_', '_rfind_', '_rindex_', '_rjust_', '_rstrip_',
'_split_', '_splitlines_', '_startswith_', '_strip_', '_swapcase_', '_title_',
'_translate_', '_upper_', '_zfill']

```

See how it identified the `upper`, `count`, and `split` methods that we used in our previous example? The `dir` function is but one of Python's many introspection capabilities. *PyCrust*, and other Python shells like it, make heavy use of Python introspection. If you've been following our examples using *PyCrust*, you should be getting a sense of why I'm such a fan of Python's interactive mode, and why programming in Python is very much like a conversation. Got a question about a Python object? Try the `dir` function. Better yet, navigate through the *PyCrust* namespace tree control. (Also known as *PyFilling*. Get it?)

Variables

If you happen to enter a sequence of characters without quotation marks, you'll probably see an error message, such as:

```

>>> love
Traceback (most recent call last):
  File "<input>", line 1, in ?
NameError: name 'love' is not defined

```

This `NameError` hints at the fact that `love` could have been the name of a defined value, in which case we would not have gotten an error. Earlier we mentioned one type of name value recognized by Python: built-in function names. Another type of name value recognized by Python is the variable.

So far we haven't used any variables. Instead, we've been working with literal values, by typing in literal representations of numeric and string objects. As you might imagine, literal values can be rather limiting. For example, you may want to refer to the same value without typing it in each time, you may want the value to reside in an external file, or you may want the value to be entered by the user of your program.

The way to overcome this limitation of literal values is to refer to an object using a name of your own choosing: a variable. One way to create a variable is by assigning a literal value to a name using the assignment operator (`=`), as we do in these examples:

```

>>> a = 5
>>> b = 20
>>> l = "Patrick"
>>> love = "couldn't possibly live without"
>>> you = "Python"

```

Once assigned, a variable is indistinguishable from the object that it references. So we can perform calculations with our variables, and invoke their methods, just like we did with our numeric and string literals:

```

>>> a * b
100
>>> b / a
4
>>> l.upper()
'PATRICK'
>>> love.split()
['couldn't', 'possibly', 'live', 'without']
>>> you.count('Perl')
0
>>> l + ' + love + ' + you
"Patrick couldn't possibly live without Python"

```

As you work with variables, methods, functions, and other elements of the Python language, pay attention to their case. Python is case sensitive, which means that Python treats items in uppercase, lowercase, and any mixture of the two as distinct items. That case sensitivity applies equally to variables you create:

```

>>> item = "something"
>>> item
'something'
>>> Item
Traceback (most recent call last):
  File "<input>", line 1, in ?
NameError: name 'Item' is not defined

```

Strong, Dynamic Types

Until now we've looked at numeric and string types separately. But it's time to see how well they mix. That requires an awareness of Python's strong, dynamic typing. To understand how objects of

Meet...

The Father

You didn't think you could go on a first date without meeting the father, did you? Known as the BDFL, or "Benevolent Dictator For Life", Guido van Rossum created Python around 1990, and continues to oversee its development.

van Rossum's home page: www.python.org/~guido/

Guido's Python tutorial: www.python.org/doc/current/tut/tut.html

Python Style, according to Guido: www.python.org/peps/pep-0008.html

Interviews and chat with Guido van Rossum: www.python.org/~guido/interviews.html

Guido's hangout (along with the other core Python developers): <http://mail.python.org/mailman/listinfo/python-dev>

different types interact, you need to understand two things:

- Python objects are strongly typed
- Python variables are dynamically typed

When I say that Python objects are strongly typed, I mean that objects have a definite type (whereas some other languages store everything as strings, for example), and that objects of one type (such as strings) won't be coerced into another type automatically, just to make an operation work properly. For example, you can't add a string to a number, even if the string is made up of numeric characters. Here is what you'll see if you try:

```
>>> 2 + '2'
```

```
Traceback (most recent call last):
```

```
File "<input>", line 1, in ?
```

```
TypeError: unsupported operand types for +: 'int' and 'str'
```

```
>>> '2' + 2
```

```
Traceback (most recent call last):
```

```
File "<input>", line 1, in ?
```

```
TypeError: cannot concatenate 'str' and 'int' objects
```

Notice how the error messages are different. The `+` operator following a number invokes a numerical addition, whereas the `+` operator following a string invokes a string concatenation. In both cases, Python detects that the type of value following the `+` operator is inappropriate for the type of operation to be performed. If you want these operations to succeed, Python allows you to explicitly convert values to different types (part of the Python philosophy is that explicit is usually better than implicit):

```
>>> 2 + int('2')
```

```
4
```

```
>>> '2' + str(2)
```

```
'22'
```

This example may seem a bit silly since we could have also solved the problem by typing in the correct literal value. But remember, most programs are made up of variables, not literal values. And many times those variables, perhaps entered into a GUI or captured in a web form, will begin life as strings, and will have to be converted to other types.

Speaking of variables brings us to our second point: dynamic typing. When I say that Python variables are dynamically typed, I mean that a variable can refer to any object, and the same variable can begin by referring to an object of one type, and later refer to an object of a completely different type. Let's look at some examples of dynamic typing in action:

```
>>> a = 2
```

```
>>> a
```

```
2
```

```
>>> type(a)
```

```
<type 'int'>
```

```
>>> a = 3.4 + 8.2
```

```
>>> a
```

```
11.6
```

```
>>> type(a)
```

```
<type 'float'>
```

```
>>> a = '2'
```

```
>>> a
```

```
'2'
```

```
>>> type(a)
```

```
<type 'str'>
```

```
>>> a = int(a)
```

```
>>> a
```

```
2
```

```
>>> type(a)
```

```

PyCrust
File Edit Options Help
1 PyCrust 0.0.2 - The Flakiest Python Shell
2 Sponsored by Orbtech - Your source for Python programming expertise.
3 Python 2.2.2 (#5, Feb 8 2003, 10:00:24)
4 [GCC 3.2 (Mandrake Linux 9.0 3.2-1mdk)] on linux2
5 Type "help", "copyright", "credits" or "license" for more information.
6 >>> 2 + 2
7 4
8 >>> x = 3
9 >>> y = 7
10 >>> x + y
11 10
12 >>> people = ['Susan', 'Bob', 'Mary', 'Steve']
13 >>> people.sort()
14 >>> dir()
dir([object]) -> list of strings
Return an alphabetized list of names comprising (some of) the attributes
of the given object, and of attributes reachable from it
No argument: the names in the current scope
Module object: the module attributes
Type or class object: its attributes, and recursively the attributes of
its bases
Otherwise: its attributes, its class's attributes, and recursively the
attributes of its class's base classes.

```

```
<type 'int'>
```

The flexibility reflected in dynamic typing is a most appealing aspect of Python. In fact, much of the productivity gained working with Python can be traced to the combination of strongly typed objects, dynamically typed variables and simple variable creation.

Reserved Names

There are some names that you can't use as variables. You see, every language has its own vocabulary. That's what makes it a language. For Python to be a language, and not just a tool for data manipulation, it must reserve a few identifiers for its own purposes. That means you cannot use these reserved words as variable names, otherwise you would take away Python's ability to hold up its end of the dialogue. Rather than simply list them for you, I'll also show you how to find them using Python itself:

```
>>> import keyword
```

```
>>> keyword.kwlist
```

```
['and', 'assert', 'break', 'class', 'continue', 'def', 'del',
'elif', 'else', 'except', 'exec', 'finally', 'for', 'from',
'global', 'if', 'import', 'in', 'is', 'lambda', 'not', 'or',
'pass', 'print', 'raise', 'return', 'try', 'while', 'yield']
```

These Python keywords make up the basic vocabulary of the Python language. As you can see, there aren't a great many of them. We'll cover all of them in the course of this tutorial series. The thing to realise now is that programming in Python is simply a matter of understanding a handful of Python keywords, a variety of operators (such as `+`, `-`, `=`, `==`, `.`), a few built-in functions and basic object types (numbers, strings, lists), and the rules for combining them in ways that suit your purpose.

Summary

I'd love to continue this conversation into the wee hours, but a good chaperone knows when to give his guests some time to themselves. So I will leave you alone with Python until the next issue. I trust the two of you will have plenty to talk about. There is so much more we could have said about numbers and strings, objects and methods, and other wonderful things. [LXF](#)

Calltips: whenever you type the opening parenthesis for a function or method, PyCrust displays information on what parameters can be passed along, and what value is returned.

NEXT MONTH

We continue conversation with Python, talking about containers: strings, lists, tuples, and dictionaries. If you have any feedback or opinions, please send them to the usual address.

POCKET TIME MANAGEMENT

Take control of your life!



Never enough time? Deadlines whoosh by, demands pile up and your partner fumes at your inability to remember anniversaries and other significant dates. Thankfully there are tools to manage your time. **Andy Channelle** plans his days with military precision...

Charles Baudelaire said: "We are weighed down, every moment, by the conception and the sensation of Time. And there are but two means of escaping and forgetting this nightmare: pleasure and work. Pleasure consumes us. Work strengthens us. Let us choose." In the UK we appear to have chosen the latter. British men and women are, it is said, the most hardworking people in Europe; we're office-bound for longer, take more work home with us and often forego the breaks that we're entitled to. In order to pack our lives with such "strengthening" work, time management has become very a important skill, and some are better at it than others.

Enter the Personal Information Manager (PIM) and its portable cousin the Personal Digital Assistant (PDA), a pair of technologies that could banish forever the need to hunt for a significant snippet of information amid a forest of irrelevance.

On its own, either of these innovations can simplify your life; but combine the two and you have access to the same data on the move as you have when sat at your desk, and you may never be lost again. In the past you may have had a difficult time getting

Ximian's Evolution offers a familiar user experience for those raised on **MS Outlook**.



If you want to win a Sharp Zaurus running Linux, turn to page 95.

the various tools needed to sync your local and portable data, but thanks to a lot of openness from 3COM/Palm, users of devices based on PalmOS should find their sync cradles work 'out of the box'. You may need to do a little manual configuration, but if you have Pilot-Link installed (and most distributions include this as a matter of course), nothing too taxing. Unfortunately the news isn't so good for users of other OSs.

Epoc, Symbian's old PDA system, has one useable – and we use that word advisedly – syncing tool that attempts to ape *PsiWin* with some degree of success. In fact it turns out to be easier, though not very convenient, to deal with your data manually on these devices.

PocketPC devices use Microsoft's ActiveSync technology, the fundamentals of which the company are reluctant to divulge.

Big boxes...

There are quite a few applications that promise to take care of your data at the PC end of the equation. Many though tend to concentrate on doing one job, and when it comes to choosing a personal information manager, 'joined-up thinking' is often a very important aspect; for instance adding an appointment to the calendar should link in with your contact's details within the address book and perhaps even send them a copy of a ToDo list for the meeting agenda. If you're part of a group, the suite could also email the other members of the group to let them know you won't be available. Fortunately Open Source developers understand this and there are three very good integrated solutions offering the full spectrum of PIM elements.

The basic tasks most Personal Information Managers will be called on to do are managing contacts, appointments and ToDo lists, though the applications highlighted below have abilities above and beyond this core featureset.

Evolution

For users reared on Microsoft's *Outlook* product, *Evolution* from Ximian makes a perfect replacement. New non-Linux users will warm to it quickly because using it is fairly instinctual as well as there being visual similarities. As well as providing similar functionality, *Evolution* follows many of the design conventions of *Outlook*, and also, via its *Connector* add on, gives desktop users access to shared resources on *MS Exchange 2000 Server*. *Evolution's* *raison d'être* is email and time management, but it goes further than that, taking in task and contact facilities and display of news and weather information.

Evolution has a range of filters for importing email and calendar settings which makes migrating from other platforms simple, in fact a wizard does all the hard work, including searching your machine for valid files. The formats *Evolution* is able to handle are:

vCard this is the standard format for contact data used by KDE, GNOME and many other applications and will have a .vcf or .gcrd extension. The standard, along with its appointments counterpart *vCalendar*, is managed by the Versit Consortium


www.imc.org/pdi and was designed to incorporate the information you'd expect to find on most business cards.

iCalendar the iCalendar format is shared by *Evolution*, *Outlook* and the standard PalmOS calendar application, which obviously makes data exchange a no-brainer.

.mbx Microsoft's *Outlook* has its own mail format which, as expected, is closed. The only way to get *Outlook* contacts information imported into *Evolution* is by emailing them to yourself, but there is a workaround which allows the importation of your email. In Windows, import *Outlook's* .mbx files in Netscape's mail client or *Eudora*. Move the resulting files, which will now be the more manageable mbox file to your Linux partition/machine and import them into *Evolution* the standard way. When importing mail from Netscape, make sure you select File>Compact All Folders if you don't want the Trash file to come back and haunt you.

Mbox finally, this is the mail format supported by *Mozilla*, Netscape, *Eudora* and many other mail clients. Synchronisation is handled by the useable (but development-stalled) GNOME Pilot, which itself relies upon Pilot-Link (www.pilot-link.org/) which has just had a maintenance release and apparently works reliably with *Evolution* 1.2.1. GNOME Pilot can sync most data between *Evolution* and Palm device but is a little unreliable updating contact details. This, I understand, is a problem with the daemon itself but 'official' details are sketchy and the project's home page is in serious danger of becoming a cobweb-site.

Still, synchronisation between Palm and Ximian's *Evolution* is very easy once the devices are actually communicating. Firstly launch the *GNOME Control Center* and select the Pilot Conduits section. The elements needed for *Evolution* synchronisation are *EAddress*, *ECalendar*, *ETodo* which deal with contacts, diary and ToDo lists respectively. You can configure each conduit in a number of ways:



- Synchronize – copies new data from the computer to the palmtop and vice versa. Items that were on both machines, but are now only on one will be deleted.
 - Copy from Pilot – if there is new data on the handheld, it gets copied to the PC.
 - Copy to Pilot – As above but the other way round.
 - Merge from Pilot – Copies new data onto your Linux box and deletes data that has been removed from the PalmPilot.
 - Merge to Pilot – Vice versa.
- The most obvious choice is 'Synchronize' but your needs may differ. Once this is set up, simply put your PalmPilot in its cradle and hit the Hotsync button.

KPIM

KPIM is a suite of applications covering the various personal information tasks. Though less integrated than *Evolution*, KPIM's discrete parts can interoperate fairly consistently. Again the main elements are the Calendar, which in this case is *KOrganizer* and the mail client *KMail*.

Though less integrated than *Evolution*, the KPIM suite reigns supreme when it comes to importing your vitals from other applications. It is able to handle both contacts and mail from Outlook 4/5, .cnm files from Pegasus Mail, MS Exchange's Personal Address Book (.pab) files, .LDIF from Netscape and Eudora's proprietary address book details. Additionally the address book application can import standard list files with a user defined separator value. The calendar is similarly well-rounded and can take information in *vCalendar* or *iCalendar* format.

The mechanics of the actual synchronisation is, again, handled by Pilot-Link but has a few more configuration options in the conduits including options for conflict resolution which can be set to 'Palm overrides PC', 'PC overrides Palm' or 'ask' on an item by item basis. Kpilot also has a file installer which can install applications from the desktop (where there's lots of space) to the Palm (where often there isn't).

The KPIM bundle covers just about every time, contact and project management task you would expect, and a few you wouldn't!

A Palm PDA like the Zire in conjunction with a mobile phone could even be used to control your desktop remotely using an app like BlueLava.



Tutorial PDA&PIM

Online PIMs

The time-management equivalent of webmail

PDA's, even low-end ones, are not cheap. So what if you need access to your calendar, contacts and other data from a variety of locations without having to shell out for a palmtop? Fortunately web portals such as Yahoo! and Netscape offer a range of time and contact management applets which are available from any Internet connected PC regardless of location or operating system. Moreover, these online solutions often


have features that their local counterparts don't such as remote sharing of whole calendars or specific events with selected colleagues and emailed alarms.

The most comprehensive free solution comes from Yahoo! which, in addition to the popular email element, has extensive calendar, contact and project management tools. Yahoo! also offers facilities for importing and exporting calendar data as .dba files for

PalmOS devices or comma separated value (.csv) files for other PIMs, and for the address book there is support for .csv, Palm's .aba and Netscape's .ldif formats.

If you have a Linux box at home, a Windows PC at work and whatever you can lay your hands on 'on the road' this may be a worthwhile option.

An online PIM is a useful alternative to a hardware palmtop organiser.



◀ And once you've defined your conduits, it is merely a case of hitting Hotsync and sitting back as your data is unified.

Psion Revo is a great handheld, but Psion don't make consumer PDA's any more.



JPilot

This is a serious, and successful, attempt to recreate the connectivity suite that Palm bundles with its hardware. Beyond

the contact, task, appointment and memo features, JPilot can also be used to install applications or files remotely onto the Pilot.

The only real problem with it is a lack of flexibility in the calendar views. There are no week, month or annual views so there's no chance of getting a broad overview of your commitments. That said the layout of the application is surprisingly intuitive and mirrors the feature-set of the Palm itself rather well.

To beef up the features a little, JPilot is able to use a number of plugins including a useful expenses tracker and a mailer applet which syncs incoming and outgoing mail on the two machines. Of course, it uses Pilot-Link so once the conduits are set up, synchronisation is actually very easy.

...and little boxes

PALM OS For the last couple of years Palm has held sway over the PDA industry, licensing its operating system to the likes of Handspring and Sony while building their own range of extraordinarily successful machines. Palm devices cover the whole range of prices and needs, from very basic diaries to the latest Tungsten multimedia monsters. As you'd expect Linux support for Palm devices is very good, easily rivalling the officially supported Windows software. In fact, Palm connectivity is included in most standard kernel versions and hackers are very adept at keeping up with Palm's release schedule; support for the recently released low-end Palm Zire, for instance, has had kernel support since 2.4.20. But if you have a kernel release previous to this and no desire to rebuild your kernel or wait for the next distribution upgrade, there is a quick hack which will allow you to HotSync your Linux box and Zire with any of the above PIM suites.

Firstly you'll need to make sure you have kernel sources installed and then find the files **visor.c** and **visor.h**. They should be within `/usr/src/linux<version number>/driver/usb/serial/`. As root, edit **visor.c** and add the line:

```
#define PALM_ZIRE_ID 0x0070
```

Next, open up 'visor.c'. You can extract the format of the entries needed to add Zire support from one of the other devices such as the Palm m505; first copy the line and then substitute **m505** with **Zire**. You'll also need to add a table variable for the device:

```
static __devinitdata struct usb_device_id palm_zire_id_table []
= { { USB_DEVICE(PALM_VENDOR_ID, PALM_ZIRE_ID) }, { } } *
Terminating entry * };
```

As root, you should then compile this module with something like:

```
gcc -D __KERNEL__ -I/usr/src/linux-<version number>/include -Wall -Wstrict-prototypes -Wno-trigraphs -O2 -fomit-frame-pointer -fno-strict-aliasing -fno-common -pipe -mpreferred-stack-boundary=2 -march=i586 -DMODULE -c -o visor.o visor.c
```

Finally replace the original **visor.o** module in `/lib/modules/<version number>/kernel/drivers/usb/serial/` with your new one and start syncing. If you have a post 2.4.20 kernel it's simply a case of HotSync and go!

Zaurus and Qtopia

It's good to talk...

Surprisingly Linux connectivity came late to Sharp's miniature marvel. This lack was satisfied with the release of Qtopia Desktop for Linux which provides a full complement of PIM applications, software and document transfer tools and conduits between desktop and Zaurus. As with the Palm software, setting up synchronisation is easy and merely involves selecting the connection method – USB, Serial or Network – and doing a little conduit configuration. The software also makes migrating from Palm devices a breeze with its import filters and has a similar range of conflict resolution options as the KPIM suite.

The latest version of the Qtopia desktop software makes quite a nice PIM in itself,



The latest version of Qtopia desktop could work as your main PIM if you only have basic needs.

with quick access to the most relevant functions and a nice overview of appointments. It would even be useable as your main PIM if the bigger suites

were too resource hungry for older hardware or you wanted a single system for a mixed OS environment. The only thing missing is mail integration.



Using *plptools*, you'll be able to sync Psion PDAs with your Linux box serially.

PSION The PDA market moves so fast and one system that was the bees' knees one year, falls foul of fashion the next. Psion's capable machines are just one example of this phenomenon; kicking off the whole PDA market and then fading into the background as Palm and Microsoft stole its thunder. Psion is still in the operating system game, but the company has spent the last year building it into mobile phone solutions and expensive Netbooks. Many corporate early adopters found the hand unit-based Psion system initially very useful for entering stocktaking information, often on combination with barcode readers, then started using them for other more office-based tasks.

Still, there are many Psion/Epoc devices out there, and many reasons for keeping them – not least the ability to type at a reasonable speed. They may also be carrying around important data that needs at least to be backed up and synchronised.

In contrast to the ubiquitous PalmOS-based machines, Psion devices require a little more convincing to work with Linux as the tools are less sophisticated and hardly ever get included in the default installation procedures of most distributions. But the software to help you access and convert your data is out there. Unfortunately you have to do more than just hit the sync button to make sure everything is up-to-date.

The first thing you'll need to do is actually get the files from PDA to PC, and there are a couple of ways you can do this.

The essential download is *plptools* (<http://plptools.sourceforge.net/>) a collection of libraries and utilities which enable the serial link between PC and Psion. *plptools* consists of:

- *plpbackup* – a command line backup utility, it can be configured to back up the entire contents of your machine or specific files and directories.
- *plpftp* – a file management utility.
- *plpnfsd* – allows the Psion storage to be accessed as a regular nfs-mounted directory.
- *plpprintd* – a printer daemon for printing via the PC.
- *sisinstall* – a command line utility for installing Psion's native .sis files to the machine.
- *ncpd* – this is the standard daemon that handles the serial connection.

plptools is available as a single binary which, once installed, runs at startup and provides consistent access to the device. For some reason the first time we ran it, it absolutely hogged system resources, but on subsequent uses its needs have toned down.

You can use *plpftp* to access drives which is not a problem if you prefer to use the shell to do things like copying or moving files. However, if like many users, you feel more at home using a GUI, your device should be available in `/mnt` and you can manipulate files to your heart's content.


By default the directory is owned by **root** so in order to access it, you'll need to use the superuser file manager in KDE. Whilst you are here, you can also change the permissions to provide user access; a process that turns out to be slightly more difficult than 'simply' changing permissions. To sort it out, you

Mozilla

Calendar and PIM functions

If you have downloaded a recent *Mozilla* build and grown accustomed to its very good mail client and address book, you may be interested in the calendar applet which was supposed to be integrated into the core last year, but seems to be forever delayed.

The calendar looks much like the KDE and Evolution efforts, but with that lovely grey/blue *Mozilla* sheen, and uses the standard iCal format for its data. It also has an integral tasklist manager which, while basic, does the job. Though it doesn't support *Outlook's* proprietary data format directly, *Mozilla* can import and use the .ics files *Outlook* is capable of exporting and also allows the publishing of public calendars on any WebDAV server using the Tools>Publish option.



An example of creeping 'featuritis'? *Mozilla* is also a pretty good PIM as well as mail manager and browser.

The lack of automatic PDA synchronisation is a problem, and if you're using a PalmOS-based machine, you're better off sticking with either *KPIM* or *Evolution* rather than trying to

persuade *Mozilla*. However, if you're using a Symbian-based PDA you'll be jumping through hoops to get the data there anyway; one more short hurdle won't kill you!

need to edit the text file `/mnt/psion/proc/unixowner`. Change the name from **root** to your user name and you should have access to all the Psion's documents through *Konqueror* and, so, through any application that uses the standard KDE file selector. Note that when you reboot the PC, this gets reset back to **root**. Not pretty, by anyone's standards!

Once you've got full access to your information, the next task is to make it readable by the applications you use on the desktop and there are two main methods to achieve this, either taking care of the conversion on your Epoc device or on the PC.

The former is best accomplished using a utility such as *nConvert* (www.neuon.com/apps/nconvert/index.php3). This is a shareware app that can convert the device's native files to a variety of formats among which are a number of .tsv and .csv options that can be read by all the applications above. You can install *nConvert* using the *plptools* **sisinstall** command. **cd** to the relevant directory then do:

```
sisinstall -w nConvert.sis
```

The **-w** argument provides a little extra feedback via the *Newt* UI. You can then convert your documents using the Psion and then import them into, for instance, *KOrganizer* using the standard KDE dialog.

The other option is to use *PsiLin* which is an attempt to recreate *PsiWin* (the standard sync tool for Windows) on Linux. The only version we could find was all in French and the GUI front-end refused to show a connection to the device. Fortunately the conversion tools, which can make html from native .wrd files, and turn Agenda, Data or Sheet files into tab separated text, are all available via the command line.

The real problem that's most obvious is that there is no standard sync option with any of these tools, so success is all predicated on the user remembering which elements have been updated where. Those in search of an easy life should probably upgrade to a new machine! [LXF](#)



Palm Tungsten has a 'proper' keyboard, but many users prefer to use a Graffiti pad.

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 – fire off an email to lxf.answers@futurenet.co.uk or send a letter by snail mail 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



Refusal to boot

Q I have a Compaq Deskpro 5200MMX on which I have successfully loaded

Smoothwall 0.98 from your earlier Cover CD. The load went faultlessly! The PC booted perfectly from the CD. And still does! I burned a CD from your recent DVD, copying the *Smoothwall 2.0* directory.

However, this does not seem to produce a bootable CD. As mentioned above, the earlier Cover CD containing v0.98 still boots. I deduce that the problem lies in what I should put on the CD to load v2.0, as opposed to any problem with the PC.

The files that I copied over are – piped under a DOS DIR.

```
1) DESCR-6      94
04/10/02 12:00 .descr
2) LINK-8       26
04/10/02 8:57  .link
3) SHORT-10    38
04/10/02 12:00 .short
4) 20B2--12 GZ 578,832
04/10/02 8:52  2.0b2-bullet-
fixes1.tar.gz
4) SMOOT-18 ISO 23,975,936
01/10/02 15:03 smoothwall-
2.0-bullet.iso
```

I feel that in consequence the Boot guidance you have published in an earlier issue is probably not relevant in this instance as the PC seems basically OK. Is my approach correct? Are the files from the appropriate distro directory 'normally' sufficient if burned on to a CD. Is anything else needed? Perhaps a fix would teach me something!

Mike Davies, via email

A When you build a bootable ISO, you need to install a boot program into that image, such as the Linux kernel or an installer. Simply including the directory in the ISO isn't going to do anything, as the

```

david@macha:~ (pts/24)
05:18:47 up 23 days, 10:35, 7 users, load average: 0.19, 0.12, 0.15
182 processes: 156 sleeping, 1 running, 4 zombie, 21 stopped
CPU states:  0.2% user,  2.2% system,  0.0% nice, 97.6% idle
Mem: 155152K total, 153692K used, 14672K free, 4552K buffers
Swap: 629148K total, 182783K used, 446357K free, 128835K cached

  PID USER      PRI  NI  SIZE  RSS SHARE STAT %CPU %MEM    TIME COMMAND
 3685 david      17   0 1200 1200  820 R   4.5  0.0   0:00 top
 1068 uml        9   0 6984 6820 6588 S   0.1  0.4   0:12 linux-2.4, 19-45
    5 root       9   0   0     0   0 SW   0.0  0.0   0:00 bdflush
 158 root       9   0   0     0   0 SW   0.0  0.0   0:00 lockd
 275 root       9   0  680  628  552 S   0.0  0.0   0:00 ypbind
 284 daemon    9   0  320  264  264 S   0.0  0.0   0:00 multilog
 277 root       9   0  680  628  552 S   0.0  0.0   0:00 ypbind
 338 amaild    9   0  456  392  392 S   0.0  0.0   0:00 tcpserver
 331 root       9   0  256  212  212 S   0.0  0.0   0:00 splogger
 374 root       9   0  432  372  372 S   0.0  0.0   0:00 getty
 341 root       9   0  960  540  540 S   0.0  0.0   0:00 snmptrapd
 375 root       9   0  432  372  372 S   0.0  0.0   0:00 getty
 376 root       9   0  432  372  372 S   0.0  0.0   0:00 getty
 377 root       9   0  432  372  372 S   0.0  0.0   0:00 getty
 378 root       9   0  432  372  372 S   0.0  0.0   0:00 getty
 379 root       9   0  432  372  372 S   0.0  0.0   0:00 getty

```

While it looks as if we're using all of our memory, most of it is actually used for caching files, rather than running processes.

CD isn't going to have anything to execute when you run from it. You may want to look in the Smoothwall directory, as there will be documentation for burning a boot CD there. You can also check out the smoothwall site at <http://smoothwall.org/>, which should describe the steps you have to go through in order to burn a bootable Smoothwall CD.

Drivers needed

Q Is anyone at *Linux Format* aware of any possible video card conflicts when running a Linux distro? Can't seem to get to the X window screen when running your November '02 distro of Debian 3.0. I've tried Red Hat 7.3, Knoppix, and FreeBSD and your product. The kernel installs, and boots clean, but will not move past the screen colour of total blackness.

My video card is a ATI Radeon 7500 which runs Win XP Pro. On a Soyo mobo in a non raid config. I have a dual boot OS on a western Digital 40 gig hd.

I've tried dropping my screen resolution down to 16-bit at 800x600, but to no avail. If its a new/different video card needed,

what would you recommend?

Doug Einarson, via email

A Looking at the information on XFree86 4.1.0 for ATI cards (www.xfree86.org/4.1.0/ati.html), which is the version of XFree86 you are likely using, there is no support for the ATI Radeon chipset. ATI have some XFree86 drivers available at www.ati.com/support/faq/linux.html, although again, only more recent (8500+) chipsets are supported. Unfortunately, the current release of XFree86 doesn't fare much better, so unless you're happy using 640x480 at 8-bit, another card is in order.

As to which card to pick – it purely depends on your budget and requirements. NVidia have supported the Linux driver market for their devices particularly well, and considering the rate at which they produce new chipsets, you can likely pick up an older GeForce2 card for your system for a mere pittance if you shop around a bit.

CDR where art thou?

Q Excellent mag – can't fault it at all. I am looking forward to being able to understand all of it though! Yes,

you have another Linux newbie! When I plug in my USB CDR it is recognised by Linux but where is it? How do I mount it and find it? And the same with my PCMCIA card. How do I mount and find it?

Richard Brown, via email

ATo use a USB CDR device, you will need to compile in support for the `usb-storage` module, or load it with `modprobe usb-storage` if it is distributed as part of your Linux installation. This will make your CDR available as a `/dev/sr0` device in Linux, which can be used by `cdrecord`, and in turn used by `gtoast` and other CD burning utilities. You can mount this `/dev/sr0` device using the `mount` command, as you would with a regular hard drive or floppy disk.

The PCMCIA card should work fairly similarly, although you've not mentioned exactly what the PCMCIA card does. You may wish to check <http://pcmcia-cs.sf.net/> to ensure that the card is supported and what you need to do in order to get it working happily. Depending which distribution you have installed, you may have to look for upgrades to

`pcmcia-cs` if it is only supported properly in more recent distributions of the PCMCIA utilities.

Veni, videocard, vici

QI have just upgraded my Mandrake 8.2 config to Mandrake 9.0 off your coverdiscs, but X (version 4.2.1-3) crashes immediately on startup with errors:

MGA(O) : Static buffer allocation failed, not initializing the DRI and

Need at least 9216kb video memory at this resolution, bit depth.

The video card is a Matrox Millenium II with 4 MB RAM.

Looking at XF86Config-4, the card & chip have been correctly detected. I have tried commenting out higher resolution modes. This got rid of the memory error, but the DRI error remains. I have never had a problem with X before on this machine.

I've searched for fixes on Mandrake and XFree86 websites, but I found nothing. I suspect the MGA driver has actually regressed between versions....?

```

david@tailtiu:/usr/local/apache/conf (pts/9)
listen 80
</IfDefine SSL>
Listen 443
</IfDefine>

LoadModule php4_module         libexec/libphp4.so
LoadModule gzip_module         libexec/mod_gzip.so
LoadModule auth_up_module      libexec/mod_auth_up.so

User nobody
Group nobody
ServerName tailtiu.dnz.davidcoulson.net
ServerAdmin david@sourceforge.net

<Directory />
Options None
AllowOverride None
Order allow,deny
Deny from all
</Directory>

```

Apache's httpd.conf file decides which user your web server will actually be running as.

I now have a useless Linux system, and I have wasted hours trying to fix it. This sort of issue when upgrading should be a thing of the past, as neither Linux nor the Open-Source development model will ever gain mainstream acceptance whilst they have a reputation for core software being fragile.

Any help would be appreciated, as if I can't fix this, I will have to restore 8.2 from Drivelmage disks, to save my work.

LXF is an excellent read, worth every penny. Pity some Linux software and distros aren't up to the same standards...

Martin Lawrence, via email

AThe error you are seeing is because the direct rendering system doesn't like your video card. This may not have been enabled for you under Mandrake 8.2, but fortunately it's quite easy to fix. You need to edit your XF86Config file, which will live in either `/etc/X11` or

A QUICK REFERENCE TO: PGP


Privacy is a major concern when using the Internet, particularly following the RIP Act. Of course, there are many other reasons why one would want to ensure privacy, such as within a business environment or for personal information being passed around the global Internet. Of course, it is possible to do end to end encrypting with tunneling technologies such as `IPSec`, `PPTP` and `SSH`, although these require that you've got a point-to-point connection to the other end. They also fail to address storage issues, as once the message ends up on the local mail server, it gets dumped into the mailbox in plain text.

Pretty Good Privacy, more commonly known as `PGP`, is an ideal fix for all of these problems. Using `PGP`, we can encrypt a message using the recipients public key, so that only they can unscramble it, or we can alternatively simply sign the

message using our own key, ensuring that the person receiving our message can be certain it came from us and has not been modified in any way.

`PGP` is a commercial package, although there is an Open Source implementation, known as `GPG`, *GNU Privacy Guard*. `GPG` uses exactly the same command line interface as the commercial `PGP` alternative, so it can replace `PGP` without having to modify any programs which utilise the `PGP` subsystem. Indeed, most programs which use `PGP` will offer `GPG` as an alternative, although most open source applications will use `GPG` as default and anyone using `PGP` will have to reconfigure it.

`PGP` is useful on its own for encrypting files and documents before sending them on, making it perfect if you store sensitive data on your system. However, the most common use is for signing and encrypting email, so the mail client needs support for `PGP` or `GPG`. Many



Mozilla's PGP plugin Enigmail allows users of the popular browser to send encrypted and signed emails

command line clients support `PGP` out of the box, including `mutt`, although `pine` requires a script to handle it, such as `pgp4pine`. For other clients, there are various scripts and utilities to integrate `PGP` into them, which can be found on <http://freshmeat.net>.


Mozilla has a popular project known as *Enigmail*, which provides `PGP` capabilities to this popular web browser/mail client. Anyone using nightly builds may have issues with it, but for users of milestone releases, it makes *Mozilla* just as useful as all other clients available.

FREQUENTLY ASKED QUESTIONS WIRELESS NETWORKING

FAQ HOW CAN I CONNECT MY LINUX BOX TO AN EXISTING WIRELESS NETWORK?

There are a wide variety of wireless networking products on the market. On the network adaptor front, we have PCMCIA, PCI or USB connected adaptors, as well as a few Ethernet connected devices. However, most of the PCI devices are simply a PCMCIA card with a PCMCIA bridge, so we use exactly the same method to get the card up and running.

PCMCIA devices are managed by the *pcmcia-cs* suite from the site at <http://pcmcia-cs.sf.net/> which includes the stable builds of the device modules for numerous chipsets. There is a great list of supported PCMCIA devices available from <http://pcmcia-cs.sourceforge.net/ftp/SUPPORTED.CARDS>. Should the specific card not be supported, there is the *wlan-ng* driver suite, which can be downloaded from <http://linux-wlan.org/>. This still requires the *pcmcia-cs* installation to actually manage the PCMCIA sub-system, but it provides more up to date modules for certain chipsets, as well as adding the ability to use some other chipsets.



If the *pcmcia-cs* suites doesn't support your wireless NIC, the *wlan-ng* modules support a wide range of chipsets.

FAQ DO I HAVE TO USE PCMCIA CARDS, OR DO USB WIRELESS DEVICES WORK TOO?

wlan-ng supports a number of chipsets used by USB wireless adaptors, although there are only a few devices which actually use these. For example, Linksys has a WUSB11 device which comes in a number of different versions. Only the v2.5 device will work happily with Linux, as the rest have chipsets which do not have their specification available for the creation of a Linux driver. A list of

supported devices can be found at www.linux-wlan.com/linux-wlan/index.html#Compatibility. There is no disadvantage to using a USB device, as the USB bus supports a data transfer rate up to 15Mbit, which is slightly higher than the 11Mbit maximum of 802.11b.

FAQ DOES LINUX SUPPORT WEP FOR SECURE NETWORKING?

The wireless modules from *pcmcia-cs* and *wlan-ng* support WEP encryption using multiple keys, as well as both

ad-hoc and infrastructure based networking. The command line utility *wlanconfig*, or in the case of *wlan-ng* *wlanctlng*, allows the keys and SSID for the network to be specified, although one would normally use one of the supplied scripts to enable WEP on the network.

One should note that it is particularly easy to crack WEP using tools freely available on the Internet. For a home network this may not be much of a problem, but for a business LAN it's a great way for an attacker to instantly get behind the expensive firewall sitting on the network border. WEP should be used if available, as it at least makes it a challenge for an attacker to penetrate the LAN.

FAQ SO, HOW CAN I MAKE MY WIRELESS NETWORK SECURE?

One of the best ways to do this is to put the wireless system outside of the company network, on the other side of the firewall, then require that wireless users tunnel through the firewall onto the network. This can be achieved with a technology such as *IPSec* or *PPTP*, as well as some proprietary tunneling technologies, including *Vtun* and *OpenVPN*.

« /etc/, and you need to comment out or delete the line which reads:

Load "dri"

This should avoid the nasty problems you are seeing. Unfortunately, most Linux distributions have issues

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 *lxadmin* (when there's no deadline!) and the ever-present 'anonymous', the forums are also frequented by Linux heroes like **Jeremy, Neiz, Fingers99, Rhakios, Erin** and many others brimming with knowledge and experience of using Linux in a wide variety of situations.

between versions, and it's worth noting that Mandrake 9.0 is the start of a new distribution build, so for anyone who's thinking of upgrading, it may be worth waiting for 9.1 or 9.2 to ensure that any obvious bugs have been worked out. This is one issue which United Linux aims to avoid, as it will ensure a clean upgrade between releases and consistent package base from all distributors.

Mandrake mirth

Q I recently reloaded Mandrake 9.0 having deleted it to try Debian 3.0.

I have these questions:

1 SOFTWARE UPDATES

During the final part of the installation, I got timed out during the Software Updates section. No matter, I would do it later. Which I did. The problem seemed to be that



There are so many Mandrake-related user groups and websites that the availability of too much information rather than too little can be a problem.

Answers



Chronic cron

Q I am running a script with cron to back up certain directories to a second hard drive. This is the script:

```
#!/bin/bash
SRCD="/home/"
TGTD="/mnt/backup/"
OF=home-$(date +%Y%m%d).tgz
tar -cvfP $TGTD$OF $SRCD
```

However, I get the following error:

```
tar: /mnt/backup/home-20030301.tgz: Cannot stat: No such file or directory.
```

It then runs through the whole directory and ends with the error:

```
tar: Error exit delayed from previous errors.
```

The tar file does not get created, thus no backup. If I manually create home-20030301.tgz using *touch* and then run the script again it works, but I do not want to create the file every time before it runs. This script is part of a cron job that runs once a week. Where am I going wrong?

A The problem in the script you sent is the manner in which the parameters are being supplied. The **-f** tells **tar** to write the output to a file. It also expects the next parameter to be the filename you want to write to. I would restructure the last line of your script to read:

```
tar -cvf $TGTD$OF -P $SRCD
```

Furthermore, you have specified the variable **OF** to contain today's date and the extension **.tgz**. It is generally accepted practice that **.tgz** indicates that the file is also gzipped. Either rename the target file to **.tar** or (preferably) **gzip** the output by using the **-z** switch:

```
tar -cPvzf $TGTD$OF $SRCD
```

Mail log messages

Q Hello, I've recently been finding a lot of messages like the following in `/var/log/maillog`

```
NOQUEUE: server.domain.com [192.168.1.39] (may be forged) did not issue MAIL/EXPN/VRFY/ETRN
```

during connection to MTA

Can you tell me if this message is trying to tell me something meaningful? And if there is anything I can do to get rid of it?

A Sadly I cannot tell much from this message alone. It basically means that someone or something has connected to the SMTP port but has not sent a message and then broken the connection (or been disconnected by the server). Maybe you have a spam blacklist configured and it will not allow this sender through, or it may be a probe to check what mail daemon software you are using. It could be as simple as a dropped connection during a mail send. You will probably find that there is another entry in your logs just before this one, which will tell you more as to why this is happening.

Lost passwords

Q I hope you can help me. I'm fairly new to Linux system administration and I'm trying to learn to do everything

from the shell. However, I admit that I do still use *Webmin* once in a while to help me when I'm stuck. I have a problem now that I've lost my *Webmin* username and password and I don't know where to reset it. Thanks for your help.

A There is no shame in using a graphical configuration utility like *Webmin* to make day-to-day tasks easier or to seeing how config files should be set up. To discover a *Webmin* username you need to cat out the following file:

```
cat /etc/webmin/miniserv.users
```

This file has two fields: the username and the encrypted password. To change the password you need to run the following script that comes included with *Webmin*:

```
/usr/libexec/webmin/changepass.pl /etc/webmin username password
```

The first parameter is the configuration directory for **webmin**, the second and third are the username and the new clear text password.

iptables timeout

Q Unfortunately, I have recently had my system security compromised and I am now running *iptables* to filter traffic into my SuSE 8 server. The server runs a corporate website and anonymous FTP for software downloads. I am allowing incoming

traffic to port 80 and port 20 and 21. However, I am unable to download files from the server via ftp when connecting using passive transfers – the connection just times out. Can you help?

A The reason passive ftp is not working is because in passive transfer mode the server tells the client to open a new connection to it on an arbitrary high port (above 1024). However, the firewall needs to be configured to accept traffic on this port. If you open up all the high ports then essentially your server will be wide open again. This was a classic problem with older software firewalls and packetfilters (such as *ipchains* and *ipfwadm*.) However, *iptables* can do stateful connection tracking. First you'll need to verify that you have *ip_conntrack* and *ip_conntrack_ftp* compiled into the kernel or compiled as a loadable module.

With this done you can add the following *iptables* rules

```
# The following 2 rules allow the inbound ftp connection
```

```
iptables -A INPUT -p tcp --sport 21 -m state --state ESTABLISHED -j ACCEPT
```

```
iptables -A OUTPUT -p tcp --dport 21 -m state --state NEW, ESTABLISHED -j ACCEPT
```

```
# The next 2 lines allow active ftp connections
```

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

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 application supports POP3 or IMAP email servers. It offers the ability to read, 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 will be transferred to your outbox the next time you synchronise. SD and Compact Flash card slots allow upgrades, addition of software or video files.

Email all your sysadmin questions and problems to Hans Huberland: sysadminqa@rackspace.co.uk



It's only 74mm wide by 138mm long and a svelte 18mm thick – so the Sharp Zaurus really is Linux in your pocket!

```
iptables -A INPUT -p tcp --sport 20 -m state --state ESTABLISHED,RELATED -j ACCEPT
```

```
iptables -A OUTPUT -p tcp --dport 20 -m state --state ESTABLISHED -j ACCEPT
```

These last two rules allow for passive transfers

```
iptables -A INPUT -p tcp --sport 1024: --dport 1024: -m state --state ESTABLISHED -j ACCEPT
```

```
iptables -A OUTPUT -p tcp --sport 1024: --dport 1024: -m state --state ESTABLISHED,RELATED -j ACCEPT
```

In the active ftp transfer rules above the client sends the server a high port to connect to and the server connects to this port from port 20 to initiate the transfer. However, if the client is also behind a firewall, which is not stateful, then this will not work and passive transfers will be required.

Old data transfer

Q I have an old Linux server that is being replaced by better hardware and a newer version of Linux. I would like to take the old server's SCSI data disk and put it into the new server. The new server will be running ext3 but the old disk is ext2 and the old server does not have ext3 support. Is there any way to change the drive to ext3 without formatting it?

A Yes, in essence ext3 is totally the same as ext2 with journaling added onto it. It is possible to add a journal to an ext2 filesystem to turn it into an ext3

filesystem, it is also possible to mount an ext3 filesystem as ext2 (and ignore the journal), this can be useful if the journal information becomes corrupted. The turn an ext2 filesystem into ext3 you will need to run the following tune2fs command (assuming the partition you want to use is **sd c1**):

```
tune2fs -j /dev/sdc1
```

Similarly, if you would like to mount an ext3 filesystem as ext2 you can either change the filesystem type to ext2 in `/etc/fstab` or you can specify the filesystem type with the **-t** switch on the mount command.

Return to sendmail

Q I have recently added a new domain for mail to my server and whenever I send mail to it it bounces with the following message:

```
553 mail.server.com. config error: mail loops back to me (MX problem?)
```

```
554 <user@server.com>... Local configuration error
```

I have checked and rechecked my MX records and they are exactly the same as my other domains (which work). I don't understand why it is still doing this? I'm running Red Hat 7 with sendmail.

A This is a pretty misleading but very common error message. Your mail server has not been configured to accept mail for the new domain. You need to add this domain and all its aliases to `/etc/sendmail.cw` and restart sendmail.

Sleepy server

Q I've been having some strange issues on my server. Connections from certain customers sites to both POP3 and ftp seem very slow to log on. Once they are connected the speed is fine it's just while the connection is being established. This only happens from some locations and I cannot replicate the problem locally but these locations can connect to other ftp sites without any problems. I'm using `wu-ftpd`, `ipop3` and `xinetd`, which shipped with Red Hat 7.3.

A I believe that the problem lies with the fact that the clients do not have reverse DNS set up for their IP addresses or have invalid reverse lookup information or that the clients are coming from behind a firewall or connection sharing device and are using Network Address Translation (NAT) to access the internet from non-routable IP addresses. These services are compiled against `libwrap`. To turn this reverse lookup off on every connection you need to remove the following entries from `/etc/xinetd.d/wu-ftpd`:

```
log_on_success += DURATION USERID
```

```
log_on_failure += USERID
```

The same can be done with `/etc/xinetd.d/ipop3`. Remember that you will need to restart `xinetd` for this change to take effect. This should

increase the speed of the initial connection for people coming from these locations.

Talk to database

Q I am using MySQL as a database on my server and I would like to add DBD/DBI support so that perl can communicate with my databases. How do I add these modules?

A The installation of DBD and DBI will want to connect to your MySQL database with the username root and no password. If you have set a password for the root MySQL user, you will need to start MySQL skipping the grant tables. To do this stop MySQL and then restart it:


```
/usr/bin/safe mysqld --skip-grant-tables &
```

When you're done don't forget to restart MySQL normally:

```
killall /usr/sbin/mysqld
```

```
/etc/rc.d/init.d/mysql restart
```

Type `perl -e shell -MCPAN` to connect to CPAN. It will ask some configuration questions. Choose the most suitable answers (the defaults work in most scenarios) until you get a `cpan>` prompt. At this prompt type: `install Bundle::DBD::mysql` to initiate the installation. It should ask which drivers you want to install. Press `1` for MySQL only. You should get acknowledgement that installation completed successfully and be back at a `cpan>` prompt. Type `quit` to exit CPAN. Don't forget to restart MySQL if you need to.



NTL woes? Visit here and www.petitiononline.com/NTL/petition.html

Permissive SuSE

Q I would like to know if there is a correct approach to setting directory/file permissions for php programs that reside in my webserver directories. I am using SuSE 8.1, and it created a directory structure of /srv/www/htdocs. For instance, I have the popular phpgroupware program at /srv/www/htdocs/phpgroupware.

The Install instructions are usually sparse for PHP programs; they usually tell users to untar or copy the directory to the webserver and set the proper permissions. I typically use *Konqueror* to extract the tar kit in /home/winter and copy it using File Manager – Super User Mode to /srv/www/htdocs. Next, I invoke *Mozilla* and type: `http://home7/phpname.php`. Sometimes I get file permission errors, so I simply use file Properties of *Konqueror* to set ownership to root/root and every access permission of every file to 666. But content management PHP programs that create files in their subdirectories do not like this, and they refuse to operate properly.

I am using: SuSE 8.1, PHP 4.2.2, MySQL 3.23.52, Apache 1.3.26. What should I be doing to

properly set the file ownership and access permissions?

Kevin Winter, via email

A When running something such as *PHPGroupWare* through a web server, all file accesses will be performed as the UID and GID which the web server is running as. This can easily be found out by checking your httpd.conf file and looking for the 'User' and 'Group' directives. With most distributions this is nobody/nogroup or www-data/www-data, so if *PHPGroupWare* wants to write to the filesystem, you will need to chown the directories it wants to write to as that particular user. It's worth noting that 666 is an insecure permission, as it means any user on the system can write to that file. You will also need to ensure that all directories are executable, otherwise you and the *Apache* processes will be unable to read or write to these locations.

NT-unwell?

Q I have had connections staying up for ages, and I'm sure that NTL changes my IP every couple of hours, so I cannot see why it should cut off when it is.

The other day, I went to a site: www.parkers.co.uk to do a price

check on a used car. I had already successfully carried this procedure out in Windows, but noticed that a grey box appeared before the four pull-downs, suggesting that a Java applet was at work.

When I tried to do this later from Linux using *Opera*, I only got two pull-down fields (and no grey box), and when I completed them they asked to be done again! When I did them the second time, it appeared to accept the input, but it was at that point I lost my connection.

Undaunted, I reset the computer, and tried the same thing. Exactly the same result. Next time I connected to Linux, I had saved the desktop, so the machine returned to exactly the same point and the cable connection stopped again. This time I typed another address into *Opera* (even though it could not load it) and then closed *Opera* before exiting. I reset the computer and on bootup, everything was fine. I ran *Opera* but there was no attempt to access the Parkers site and it went to the address I had entered earlier. The connection stayed up the rest of the evening, and was still there when I closed down before bed!

The question is: Is the Linux security system jumping in to deal with something it sees as a threat, and disconnecting me for my own good; or is NTL disconnecting me for the same reason? Would downloading and installing Java help? Alternatively are there two possible renderings of the Parkers site, one where Java is present, one where it is not, and the non-Java (probably javascript) code is faulty and this causes the disconnection.

All great theories, but I need an answer. I have tried ifup eth1. I have tried closing and restarting dhcp. What file carried out the bootup activities? Maybe I could take a few lines from this, and run it rather than restarting the computer when I get cut off.

Dave Spagnol, via email

A The first thing to note is that you need to be root to take down an interface. As you're hopefully running your browser as a non-root user, neither the browser nor the Java plugins can disconnect you from the Internet. To our knowledge, there is no security system for Linux

which would actively disable your Internet access, as this isn't a particularly effective way of handling security if you're remote to the physical machine, particularly if you're experiencing any false positives.

When it disconnects you should take a look at /var/log/messages and /var/log/syslog. These will log what your DHCP client is doing. It could be that it has been unable to obtain a lease from NTL, so set the interface up without an IP address, or it is having other issues. You've not mentioned which distribution you are using, but many have /etc/rc.d/init.d/networking or /etc/init.d/network, which you can 'restart' to try a new DHCP request: `# /etc/init.d/networking restart`

However, this isn't a nice way to solve the problem. If you can find some debugging information for DHCP and send it in to LXF, we should be able to provide you with a more suitable resolution to your problems. Of course, you could also check out the LXF forums and find out if any other NTL users have been having the same problems and have cooked up a fix of their own. [LXF](#)

Submission advice

We are happy to answer all sorts of Linux related questions. If we don't know the answer, we'll find out for you! But in order to give you the best service, it helps a lot if you read the following submission advice.

- Please be sure to include any relevant details of your system. 'I can't get X to work' doesn't really mean anything to us if we don't know things like what version of X you are trying to run, what hardware you are running on.
- Be specific about your problem. Things like 'it doesn't work' or 'I get an error' aren't all that helpful. In what way does something not work? What were you expecting to happen? What does the error message actually say?
- Please remember that the people who write this magazine are NOT the authors or developers of Linux, any particular package or distro. Sometimes the people responsible for software have more information available on websites etc. Try reading the documentation!


We will try and answer all questions. If we don't answer yours specifically, you'll probably find we've answered one just like it. We can't really give personal replies to all your questions.

WRITE TO US AT:
Linux Format, Future Publishing, 30
Monmouth Street, Bath BA1 2BW or
email: lxf.answers@futurenet.co.uk

missed one?

LINUX FORMAT BACK ISSUES

Every month *Linux Format* brings you the best tutorials, the essential reviews and the latest news. But if you've missed out on a must-read feature or a vital programme from our expertly compiled CDs and DVDs, order your back issue NOW! And remember, you need never miss an issue of your favourite Linux mag, if you subscribe to *Linux Format* (see overleaf for more details).



March 2002

Product code:
LXFB0038(cd)
LXFD0038(dvd)

DVD HIGHLIGHTS:
Python, The GIMP, Lycoris, Yellowdog, Knoppix, Cygwin, DansGuardian, Senken, AbiWord, X Northern Captain, Eclipse

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

CDs HIGHLIGHTS:
KDE 3.1, UNIX Tools For Windows, Konstruct, phpMyAdmin, X Net Strength, AutoUpdate



February 2003

Product code:
LXFB0037(cd)
LXFD0037(dvd)

DVD HIGHLIGHTS:
IPCop, Postnuke, Lilypond, Squid, Allegro, Cocoon, Fandango, gPhoto, Xrmap, Ted, AdvancedPHPDebugger, Flightgear updates

MAGAZINE FEATURING:
Lindows on test, Serial ATA, Window Managers roundup, Arkeia review, Xft2, build your own Linux backup box, Java compiling, Jool Kwartz review

CDs HIGHLIGHTS:
Yoper, KOffice 1.2.1, Parted, DOM Tooltip, Guarddog, GL-117




January 2003

Product code:
LXFB0036(cd)
LXFD0036(dvd)

DVD HIGHLIGHTS:
Evolution PIM/calendar/email client, Freevo personal video recorder, OpenZaurus, PilotLink

MAGAZINE FEATURING:
Optimisation tips and tricks, Netbox Cubit review, build your own desktop, Xandros, Cubit, CORBA, Realsoft 3D, Instant Messenger roundup

CDs HIGHLIGHTS:
Racer, Freeduc, Linux From Scratch, GNOME Meeting, Opera 6.1, Film GIMP, phpOpenTracker, Tuxpaint, Snort



Christmas 2002

Product code:
LXFB0035(cd)
LXFD0035(dvd)

DVD HIGHLIGHTS:
Mandrake 9.0 - 3 CD download edition, jEdit, BastilleLinux, Phoenix, VideoLAN, KDE 3.0.4

MAGAZINE FEATURING:
Huge C++ IDE RoundUp, LXF Awards nominations, gadgets for Xmas, User Mode Linux, compiling code, process accounting

CDs HIGHLIGHTS:
3CDs, including 2CD Mandrake 9.0, SpamX, Stellarium, Krename, ImagePress, Nogger, DigitalMusicCentre



December 2002

Product code:
LXFB0034(cd)
LXFD0034(dvd)

DVD HIGHLIGHTS:
KDE 3.1beta2, Scribus, Smoothwall2, Movix, GNUPrivacyGuard, GCompris, Drip

MAGAZINE FEATURING:
Distro reviews of SuSE 8.1, Red Hat 8.0, Lycoris and a mini-distro RoundUp, scanning, PHP acceleration, Linux Expo UK 2002

CDs HIGHLIGHTS:
Unreal Tournament 2K3 (Demo), DemoLinux, Cardfile, RUTE, OOoDictionaries, Mozilla binary packages



November 2002

Product code:
LXFB0033(cd)
LXFD0033(dvd)

DVD HIGHLIGHTS:
Debian 3.0, mFighter, OpenOffice.org (bugfix release), DVDrip, Cyrus IMAP Server

MAGAZINE FEATURING:
Sun's move into the Linux server market - with LX50 review, The Liberty Alliance, Systems programming, using OpenOffice.org, Homebase

CDs HIGHLIGHTS:
Kylix 3, GNUCash, BXPro, KDevelop, Opera, Vega Strike, Parted, AnjutaIIIDE, GTransferManager

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**

Benefits – why subscribe?

- **Subscribe by Direct Debit – it's easy! Fill in the form below or call 0870 444 8645* and have your bank details ready.**
- **Huge savings on a year's supply of *Linux Format*.**
- **Subscription copies are sent hot off the press – be one of the first to get yours!**
- **Guarantee your copy – never miss an issue.**
- **All you need to do is fill in the form. If you live in the UK, you don't even need a stamp!**

* or call the overseas number shown at the bottom of the coupon, below

RISK-FREE GUARANTEE!

You can cancel your subscription at any time and receive a full refund on all unmailed issues, no questions asked

Title _____ Initials _____ Surname _____
 Address _____
 _____ Postcode _____ Country _____
 Email address _____
 Daytime telephone number _____

Instruction to your Bank or Building Society to pay Direct Debits



Media with passion
 Future Publishing Ltd, Cary Court,
 Somerton, Somerset, TA11 6BR



Originator's Identification Number
7 6 8 1 9 5

Please send this form to address shown

1. Name and full postal address of your Bank or Building Society branch

To: The Manager _____ Bank/Building Society
 Address _____
 _____ Postcode _____

2. Name(s) of account holder(s) _____

3. Branch sort code _____
(from the top right-hand corner of your cheque)

4. Bank or Building Society account number _____

5. Instruction to your Bank or Building Society

Please pay Future Publishing Direct Debits from the account detailed on this Instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Future Publishing and if so, details will be passed electronically to my bank or building society.

Signature(s) _____ Date _____

Ref No (Office use only) _____

Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

You may photocopy this blank form and then complete it. You cannot fax or e-mail Direct Debit forms.

QUOTE ORDER No. LXFP39

LINUX FORMAT

Please choose your method of payment **1** or **2**

- 1.** Direct Debit (UK only – double CD Edition) – I wish to pay £27.45 every 6 months
 or Direct Debit (UK only – DVD Edition) – I wish to pay £29.95 every 6 months

Please complete the form on the left

I understand that I will receive 13 issues during the next 12 months.

2. Cheque/credit card 13 issues

	CD	DVD
UK	<input type="checkbox"/> £54.90	<input type="checkbox"/> £59.90
Europe (including Eire)	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
North America	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
Rest of the World	<input type="checkbox"/> £91.90	<input type="checkbox"/> £96.90

- Cheque (payable to Linux Format. Sterling cheques drawn on a UK account.)
 Visa MasterCard American Express
 Switch Issue No/Valid date _____

Card No _____ Expires _____

Signature _____ Date _____

UK READERS RETURN THIS COUPON TO:

Linux Format, Subscriptions, Future Publishing Ltd, Freepost BS4900, Somerton, Somerset, TA11 6BR.

OVERSEAS READERS RETURN (POSTAGE PAYABLE) TO:

Linux Format, Subscriptions, Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6TB, UK

WEB: www.futurenet.com/secure/subs1.asp?ID=92

Please tick this box if you would prefer not to receive information on other offers.

ORDER HOTLINE ☎ 0870 444 8645
OR FROM OVERSEAS ON +44 870 444 8645

OFFER ENDS: 24.04.03

EVERY ISSUE IS A SOFTWARE BONANZA!

Every issue *Linux Format* comes with either 2 CDs or a DVD crammed full of the latest Linux software. Get all this delivered direct to your door every month with a subscription costing from just £27.45 for six months!

In the last six months you would have received on DVD:

The latest **FULL LINUX DISTROS** including ● KDE 3.1 ● Lycoris ● DemoLinux ● YellowDog ● Sorcerer Linux ● Knoppix ● Yoper ● Red Hat 7.3 ● Slackware 8.1 ● Debian 3.0 ● Gentoo
PLUS Over **750** of the best Linux apps and utilities

Don't miss these great software packages – save yourself hours of download time with every issue of *Linux Format*!

« **Subscribe from just**
£27.45

every 6 months by Direct Debit
saving £22.97 a year (CD issue)

DVD issue from £29.95 for 6 months



Essential disc info

Read this important information before you use your *Linux Format* coverdisc – CD or DVD. We've collated some helpful info to help you get the most from these jewels of data!

Finding the essentials

Missing something?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases, the latest libraries and

other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

Package formats

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

Documentation

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

What are all these files?

If you are new to Linux, you may find the profusion of different files and extensions confusing. As we try to give as many packages as possible for compatibility, there will often be two or three files in a directory covering different types of Linux, different architectures and usually source and binary versions – so which do you install? They can be identified by their filenames, and usually just by the file extensions.

Someap-1.0.1.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.0.1.i386.deb – The same, but a debian package.

Someap-1.0.1.tar.gz – This is usually source code.

Someap-1.0.1.tgz – Same as the above, tgz is abbreviated form of tar.gz

Someap-1.0.1.tar.bz2 – Same, but uses bzip2 compression instead of zip

Someap-1.0.1.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.0.1.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.0.1.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7.x PPC Linux.

Someap-devel-1.0.1.i386.rpm – A development version.

Installing from tarballs

A tar ball is a two stage archive. First the files are archived into a single file with tar and then compressed with Gzip or Bzip2. To unpack, cd to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
```

```
tar xvf --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 CD/DVD being physically damaged we'll send you a new, working version within 28 days. Send your defective disc – complete with your name, address, and a description of the fault – to:

**Linux Format, Future Publishing Disc Department, 3B
Athena Avenue, Elgin Industrial Estate, Swindon, SN2 8HF.**

Creating install CDs with cdrecord

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord -scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v  
/path/to/image.iso
```

You can simplify the command by saving some default settings in /etc/default/cdrecord. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the "Burn" tab and the "ISO 9660 Image" gadget near the top of the window. Put the path to the image file in the gadget and press "Combust!". Now put on the kettle while the CD is created for you.

Other OS?

You do not have to use Linux to burn the ISO to a disc. All the Linux-specific bits are already built into the image file. Programs like *cdrecord* simply dump it to the disk. If you don't have a CD-writer, find someone who does have one, and a DVD drive, and use the CD burning software on their computer. It can be Windows, MacOS, AmigaOS whatever.

No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions it is also possible to mount the images and do a network install, or even a local install from another disk partition. The methods often vary between distributions, so check on the distro vendors website for more information. [LXF](#)

Coverdisc



If it's a happening program in the Linux world, we've tracked it down for you. Neil Bothwick is your trusty guide through this month's jam-packed *LXF* DVD.

Whatever your personal views on censorship of the media by laws, governments or vested corporate interests, you're probably concerned with the content of the material that is being viewed over your Internet connection. This month, you'll find the answer on the *LXF* DVD. If you're suspicious that family members or employees may be browsing to sites that contain material that you wouldn't want your ISP to associate with your account, *CensorNet* is a great solution an individual computer or over a network. It's not just about restriction of access though – *CensorNet* is more concerned with regulation, as shown by a key feature – the time option. If


you are one of the increasing numbers of employers that don't mind your staff surfing the Internet to non-work related sites or checking their webmail accounts out of working hours, you can use the program to protect your productivity while appearing to be non-restrictive at the same time. It's a no-lose situation.

The Desktop section of the DVD is a bit of a catch-all section this month. For making backups of your home system to store remotely, you'll find *Split2CD* ideal; and while you may recognise the functionality provided by the likes of *K3b* and *GKrellM*, there are a few others with more diverse interest. Computing can get a bit dry sometimes – if things are getting a bit too serious, give *TalkFilters* a try – it'll bring a smile to your face. Even if you've only a passing interest in Astronomy, there's some mileage to be had from programs like *QastroCam* or *Solar System Simulator*.

DISTROS CENSORNET

We regularly feature Linux distributions that use an older PC to provide Internet access to a local network, now here's one that is designed to deny, restrict or delimit access. *CensorNet* can be used to prevent a company's Internet connection being abused by staff, or stop people using an Internet cafe to indulge in unwelcome activities. It is capable of filtering content by URL, keyword/phrase, file extension, MIME type or PICS rating. *CensorNet* can also be used to place limits on bandwidth or times.

Filters can be applied on a per user basis, or to specific machines on a network. It contains a list of banned sites, *CensorNet* provide a subscription service to keep this automatically updated. Filtering can be applied at a user level or to individual computers. *CensorNet* also acts as an Internet



K3b makes creating audio and data CDs fast and easy.

gateway and caching proxy, you don't need a separate machine for these functions. *CensorNet* is supplied on the DVD as an ISO image, ready for burning to CD-R. See the *Essential disc info* page at the end of this section for information on how to do this. Once you have a CD-R, boot from it and follow the prompts. After installation, *CensorNet* is configured and administered via a web browser over the local network.

DESKTOP K3B

There are certain topics that are commonly asked about when people first try Linux after using another OS. These include scanning, printing and, very commonly, CD burning. While Linux has some extremely powerful command line tools for this job, they aren't exactly user-friendly, especially to the new user. So when the topic of CD burning is raised, certain GUI packages are often recommended for both novices and power users, and *K3b* is one of them.

K3b supports the creation of a wide range of CD types; data, audio, video and mixed mode. It also

supports ripping of audio CDs and DVDs, although the latter obviously requires a DVD-ROM drive. The DVD ripper pulls movies from DVD and encodes them to DivX. It uses transcode for this, which you'll find elsewhere on the DVD.

In common with almost all other point-and-click CD creation tools, *K3b* uses the faithful command line tools like *cdrecord* and *cdparanoia* to do the dirty work, and you'll find these in the dependencies directory on the DVD, apart from transcode which you'll find in the dependencies directory of Graphics/DVDrip.

DESKTOP GKRELLM

GKrellM is a set of system monitors, displaying the information in a series of vertically stacked meters. This layout makes it ideal for leaving at the side of the screen, so you can keep an eye on what's going on in your system. The system is totally flexible, each of the monitors can be individually configured or disabled, so you only see the information you need. *GKrellM* has many monitors built in, including CPU usage, network throughput, online



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.

time, disk space usage, system temperatures and voltages (read from `lm_sensors`), memory and swap usage, battery status, uptime and many more. If the built in monitors are not enough, there are plenty of third-party plugins, several of which are on the DVD. For example, I use *GKrellmBUPS* to keep track of the status of my Belkin UPS.

A new feature of version 2 is that it can operate in client/server mode, so you can have a stack of monitors on your desktop computer that show the status of a headless server elsewhere on the network, or even over the Internet. *GKrellm* is on the DVD as RPM and Debian packages as well as source. Many of the plugins are supplied as source only, so if you want to compile these you should also install the *gkrellm-devel* RPM. Incidentally, the name is a reference to the Krell meters in the Sci-Fi classic film, *Forbidden Planet*. A film that's so old that Leslie Nielsen plays a straight role with absolutely no 'beaver' jokes.

GRAPHICS VIPS

Mention graphics manipulation and most people immediately think of *The GIMP* or *Image Magick*. However, there are other graphics and image processing programs for Linux, and this is one of the latter. There are actually two packages here. *VIPS* (VASARI Image Processing Software) is the image-processing library that does the hard work, *nip* is the graphical interface to *VIPS*. Developers may be more interested in using *VIPS* from their own programs, and the API is documented in the included packages. The emphasis is on image processing rather than graphical manipulation, *VIPS/nip* are not the best combination for retouching your digital




image processing, specialising in large images, with *VIPS*.

photographs, but they excel at things like colour manipulation. They are also particularly suited to working with very large images that do not fit into available RAM. The packages are supplied as source and RPMs, with separate packages for *VIPS*, *nip* and the documentation. The *nip* documentation includes a tour of the program and tutorials on its use.

GAMES FROZENBUBBLE

It can be difficult to describe a game in a sentence or two. It is all but impossible to give an impression of what it is like to play in the same space. It may be easiest to say that *Frozen Bubble* is an arcade game that has received countless rave reviews, including having been a *Linux Format Hot Pick* in the past, and that you really should try it for yourself. The game's own description is *Frozen Bubble* is a game in which you throw colourful bubbles and build groups to destroy the bubbles". However, this seems to totally miss the key feature of this game, it is hopelessly addictive.

The web site at www.frozen-bubble.org/ has an interesting 'making of' section, charting the development of the program. Surprisingly, for such a popular and polished program, this only started in October 2001.

GRAPHICS LINUXOPENMEDIABOX

The question "is Linux ready for the desktop?" is asked so often that we're in danger of forgetting that there are other areas where it could be useful. The *Linux Open Media Box* project targets one such area, the home entertainment system (often referred to as a set top box, although we've yet to see one on



LinuxOpenMediaBox making a home entertainment system based on Debian.

top of any TV set). The aim of this project is to create a free and open source based multimedia console to plug to your TV. *LOMB* gives you the ability to play most multimedia streams, from DVD to DivX, mp3 to Ogg Vorbis, and it can play games as well.

It is designed for the EPIA mini-ITX motherboards, as reviewed in *Linux Format 30*, but that's mainly because of their compact size and quiet running, there's no technical reason for not running it on any sufficiently powerful X86 hardware with TV out on the graphics card. For the most silent operation, which is

important for a system used in the living room, *LOMB* is designed to boot from a flash memory device such as a Compact Flash card, with media files being read from CD or DVD. This uses Emdebian, an embedded version of Debian. However, the latest hard drives are so quiet, especially the 5400rpm ones, that hard disk booting would be feasible, as would network booting from your main PC sited elsewhere. This is an interesting project that is worth further investigation. Install *TuxPaint* on the system and you'll soon have the kids addicted to Linux too. [LXF](#)

Finding files on the CDs and DVDs

Using the Index

What happens when you want a particular program that you think was on a previous cover disc? You could try digging out all your discs, putting each one in the drive and trawling through the index files for what you need. Although you would eventually find it, it would probably have been quicker to download it (unless it's a 650MB ISO image). There is a much faster way, two in fact.

If you have a Java-enabled web browser, load the `index.html` file of the current DVD and you'll see a search box towards the top right of the page. Type in part of the program name and press return and it will search the recent discs. Select the program you want from the results window, making sure the appropriate disc is mounted, and it will load the program's directory into the browser window. There are two check boxes below the search gadget, the first restricts the search to the current disc only and is selected by default. The other causes the applet to search the descriptions of each program as well as the names. This is useful when you aren't sure of the name of the program and want to search on what it does.

The searching is fairly simple at the moment, a straight substring search. If you enter more than one word, it will search for that phrase, rather than for descriptions containing any combination of those words. This is a fairly new addition to the cover discs, so the index files for these only go back a few months. The alternative, which has been available since issue 20, is a set of CSV files in the `Essentials/CSV` directory of the disc. These are ready for import into any database program, with each line of the file containing the following data:

```
Disc, Name, Path, Description, Home page, Version
```

Or, if you're really lazy, there's another option. You can simply search the CSV files directly with `grep`, it works for me.

CoverdiscDVD

DVD CONTENTS AT A GLANCE

Desktop

Coreutils	Collection of core GNU utilities
FBpanel	Lightweight X11 desktop panel
Fpted	Small, fast, and feature-rich multiple file text editor
GKrellM	System monitor package
GKrellMBUPS	GKrellM Belkin UPS plugin
GKrellMailwatch	GKrellM plugin which monitors mailboxes
GKrellMMS	Plugin for controlling XMMS from within GKrellM
GKrellMTimers	GKrellM plugin to define timers and stopwatches
GKrellMVolume	GKrellM plugin that controls your mixer device
HamFax	Send and receive amateur radio facsimiles
K3b	CD burning software built using KDE
Katalog	A set of scripts to create and search an index of CDs
KPopUp	Sends and receives Winpopup messages
Nano	Pico editor clone with enhancements
PresentingXML	A framework for XML/XSLT applications
Qastrocam	A Webcam and telescope program for astronomy purposes
Setedit	Editor for C/C++ programmers with a nice text interface
SolarSystemSimulator	An OpenGL/Glut-based solar system simulator
Split2cds	Splits large directories to chunks of a predefined size
TalkFilters	Collection of humorous text translators
X10-wishGUI	Remote access to a Linux box with X10-wish drivers
XAutomationTools	A command line interface to XTest for automating X Windows

Development

AnthillBuildManager	Promotes knowledge sharing within an organization
CrystalSpace	A free and portable 3D engine
DoubleChocoLatte	Software Configuration Management/Bug/Enhancement Tracking
EiffelWrapperGenerator	An Eiffel wrapper generator for C
FOX	C++-based library for graphical user interface development
Kcurses	A simple widget set for ncurses
LinCVS	Graphical frontend for the CVS-client
NCC	A C flow analysis compiler
Pyxine	Python bindings for xine
QuantLib	C++ library for financial quantitative analysts
SFTPTeamProviderForEclipse	File/directory synchronization for the Eclipse workspace

Distros

AstaroSecurityLinux	A powerful and flexible firewall
CensorNet	Internet management tool
IDMSLinux	A server distribution
SME-Server	Convert a PC into a Linux Internet server

Games

Enigma	A free Oxyd clone
FrozenBubble	An arcade game
IterVehemensAdNecem	A graphical rogue-like game
Mupen64	MUPEN64 is a n64 emulator developed on/for Linux
SearchAndRescue	Air rescue flight simulator
TuxPaint	Simplified, entertaining drawing program for young children
WineXGameManager	Graphical interface to install and play Windows games

Graphics

Cinelerra	Complete audio and video production environment for Linux
DVDrip	Full featured DVD Ripper GUI
EffecTV	A real-time video effect tool
EXIFprobe	Show content and structure of TIFF/JPEG/EXIF images
KPlayer	KDE media player based on mplayer
LinuxOpenMediaBox	Linux multimedia box with TV output
Lumiere	GNOME 2 frontend to mplayer
MartinsPictureViewer	JPEG viewer
VIPS	Image processing system
WhiteDune	Graphical VRML97 editor and animation tool

Internet

ccIRCRelay	IRC relay tool
CoffeeshopAccessMgr	An Internet access control front-end for NoCatAuth
Galeon	GNOME Web browser
GlobespanADSLModemsDriver	Globespan-based USB ADSL modems driver
ICEMail	Java email client based on JavaMail API with SMIME support
MosfetKonquerorImageLoader	Replacement image loader for Konqueror
PhpDig	HTTP spider/indexing search engine using PHP and MySQL
Qtella	Gnutella client for Linux
StatFreak	Generator of skinnable, customizable IRC statistics
TinyMonitor	URL content monitoring program
WMget	Background downloader dockapp

Mobile

ipkg	Install Opie packages from OpenZaurus with a Sharp-ROM
PortaBase	Database program to create, browse and edit custom tables
QtopiaAlarmClock	A simple travel alarm clock for the Zaurus
QtopiaAlarmClock-Advanced	Adds several features absent from the original Alarm Clock
WAP11GUI	SNMP management for the Linksys WAP11 wireless AP
Win4	Now available in english and german.
ZaurusCheck	Constantly checks the connection of a Zaurus to a Linux PC

Office

OBM	Intranet management and contact application
ProjectAssistant	A CRM, Project Management, and Paperless Office system
SystemsManagementInterface	Web-based Enterprise Management System
Taskware	Task/project scheduling

Server

ApacheFileManager	An Apache mod_perl Web server filemanager
BotproofEmail	Hides the contact email on a Website from spambots
ELOGElectronicWebLogbook	Electronic logbook with a Web interface
EmailRBLcheck	RBL checker for gmail and sendmail
EnemiesOfCarlotta	Mailing list manager
Face2	Java-based one-to-one chat system for ecommerce sites
JustForFunNetworkMgmt	Open Source network management system in PHP
mod_spam_die	Apache 2 module for trapping spam crawlers
Noffle	News server optimized for low speed dial-up connections
NotFTP	Clean and simple PHP-based HTTP-FTP gateway
Postfix	The Postfix MTA
PostgreSQL	A robust Object-Relational DBMS
PostgreSQLSessionHandler	A PostgreSQL session save handler for PHP4
SpamAssassin	Mail filter that identifies spam using text analysis

Sound

AlsaPlayer	PCM (audio) player for Linux/ALSA
AudioBurn	Commandline audio CD burning application
EQXMMSGraphicalEqualizer	XMMS plugin that equalizes whatever is being played
LAoE	Graphical audio sample editor

System

Apt4RPM	Create an apt repository from an RPM repository
FileUploadWebminModule	A small module to allow file uploads via Webmin
KPortage	Graphical frontend for the Gentoo Linux portage system
LFS-install	Script for automatically compiling a Linux system
OpenMosixClusterForLinux	Kernel extensions for transparent, scalable clustering
RPMrebuild	Rebuilds an RPM file from an installed package
TimosRescueCDSet	A kit for creating bootable rescue CDs
Usermin	Web-based interface for UNIX users
Webmin	Web-based interface for Unix system administration
WebminNetworkUtilities	Common Network tools with Webmin look and feel



Fun and addictive, that's the general consensus of Frozen Bubble, see for yourself.

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 Jonathan Cole

41 MIDLANDS

URL <http://midlandslug.port5.com/>
Contact Pete Thompson

42 CUMBRIA

URL www.cumbria.lug.org.uk
Contact Jamie Dainton

43 DORSET

URL www.dorset.lug.org.uk
Contact John 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**Email N.J.Murison@exeter.ac.uk**74 SUNDERLAND**Email thomas.croucher@sunderland.ac.uk**75 EAST YORKSHIRE**Email sharkonline@whatemail.com**76 CLEVELAND OPEN SOURCE GROUP**Email openlug@digitalmedia.co.uk**77 BEVERLEY**Email vladimir_lukyanov@hotmail.com**REVISED
DETAILS****78 DUNDEE & TAYSIDE**URL www.dundee.lug.org.uk**79 SUSSEX**URL http://sussex.lug.org.uk/**80 WIGAN & ST HELENS**Email paulfjohnson@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@ntword.com**92 MENAI**URL www.menai.lug.org.uk**93 ABERDEEN**Email mark@arricc.net**94 SHETLAND**URL www.shetland.lug.org.ukEmail c_s_s_butler@yahoo.com**OTHER GROUPS:****B-LUX**URL www.b-lux.org/Email blux-admin@lineone.net**FUENGIROLA GNU/LUG**URL www.fuengirola.lug.org.uk/Email fuengirola@lug.org.uk**SCHOOLS**URL www.schools.lug.org.uk/**TR LINUX – TURKISH LINUX**URL www.trlinux.lug.org.uk/Email c_s_s_butler@yahoo.com**YOUNG LINUX**URL www.young.lug.org.uk/

LUG OF THE MONTH

NYLXS

New York GNU Linux Scene is dedicated to providing resources to the New York Linux Community and developing Open Source leadership. Supporting www.lxny.org and www.nylug.org, but bound to no specific local NYC or NY Metro Area organisation, its members support the NY Linux and Free Software scene with manpower, technology, money and time. Its goal is to empower NYC Area Linux users, Free Software users, and the NY

population in general through their lugs, schools, businesses, and government agencies. NYLXS helps with installfests, lectures, tutorials, mailing lists, scheduling of events, political lobbying, training, publicity, journey, Internet radio shows, and educational support to public schools and libraries.

Membership is determined by both paying dues and hours of volunteer work given to NYLXS-sponsored events. The Membership

Committee is chaired by Richard Weinberg and is currently formulating specific rules for membership with the help of our Treasurer Joe Maffia and President Ruben Maffia. NYLXS believes that the advent of Free Software and Linux is not down to accident, or technical achievement, but is primarily a political activity. Visit us on the Net at www.nylxs.com and <http://fairuse.nylxs.com>



Worldwide Linux User Groups

Free Software users across the globe

Africa

EGYPT

URL www.linux-egypt.org

Contact Hesham Bahram

GAUTENG, SOUTH AFRICA

URL www.glug.org.zaEmail glugmin@revolution.org.za

Australia

ADELAIDE

URL www.linuxsa.org.auEmail mtippet@anu.edu.au

ALICE SPRINGS

URL www.aslug.org.au

MELBOURNE, VICTORIA

URL www.luvasn.auContact luv-committee@luvasn.au

PERTH

URL <http://plug.linux.org.au/>

SYDNEY

URL www.slug.org.au

Europe

AUVERGNE

URL www.linux-arverne.orgEmail Cyril.Hansen@wanadoo.fr

COSTA DEL SOL (English speaking)

URL www.fuengirola.lug.org.uk

DENMARK

Aalborg www.alslug.dkEsbjerg www.eslug.dkFyns www.flug.dkMidt-og Vestjylland www.mvjlug.dkNordjylland www.njlug.dkSkåne Sjælland www.sslug.dkTrekantsområdet www.tlug.dkVest-fyn www.haarby-net.dk/vflugÅrhus www.aalug.dk

EIRE

URL www.linux.ieEmail root@linux.ieURL www.dilu.orgEmail glossary@dilu.org

GOTHENBURG

<http://nain.oso.chalmers.se/LUG/>

Middle East

ISRAEL

URL www.iglu.org.il/IGLU/Email webmaster@iglu.org.il

PALESTINE

URL www.lugps.orgEmail 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.orgEmail aftyde@balug.org

PAKISTAN

URL www.linuxpakistan.netEmail tux@clug.org

HYDERABAD, SINDH, INDUS VALLEY

URL www.geocities.com/slug_pk/

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.jionline.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.freeservers.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

TCO issues often result in Linux being chosen over pricier proprietary solutions. Jono Bacon counts the pennies...

Recently in this little series of articles we have discussed that Total Cost of Ownership (TCO) has a factor in the adoption of a particular software system, in businesses in particular. This month we will focus on this issue some more and look at the separate entities involved.

TCO is a tricky issue to work out due to the different methods and requirements of each practitioner, but the first general step is to identify the needs of the business. This needs to include what they want to do, how big a need it is and which needs are more important. From these needs we can begin to craft a proposal of hardware and free software and how the separate components fit together. This is mainly the job of a system architect, so some of this work may be already done.

With the proposed solution set, it is now time to figure out the general costs of the solution. At this stage it is often wise to consult a single vendors price list and keep to all prices from

that vendor, but these differences can often be minimal. It is important at this stage to include all costs including hardware, licences, staff costs to set up the solution, running costs, upgrading costs etc. You need to provide a cost estimate NOT for the initial set up but for the entire proposed life-span of the set-up; the life-span of a solution is obviously an area where Linux scores due to it being free in nature and having a free upgrade path.

Next month we will delve further into these issues and look at how the initial cost estimate can be further improved and compared against competing vendors. It is this comparison that can have an overriding effect on the suitability of Linux and free software as a solution, so we will concentrate on it in some detail. There's an exhaustive piece on how TCO relates to schools on the Kolkata chapter of the Indian GNU/Linux Users' Group www.ilug-cal.org/node.php?id=68

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 matthewnailon@futurenet.co.uk**Editorial Contributors** Jono Bacon, Neil Bothwick, Chris Brown, David Cartwright, Andy Channelle, David Coulson, Dave Cross, Chris Denton, Richard Drummond, Hoyt Duff, Alan Ferguson, Marco Fioretti, Chris Howells, Maurice R Kelly, Jon Kent, Biagio Lucini, Patrick K O'Brien, Mike Saunders, David Taylor, Bill Von Hagen, Tom Wilkinson**COVER CD PRODUCTION****CD Editor** Neil Bothwick**Senior CD Editor** Kate Hadley**ART CONTRIBUTORS****Art assistance** Martin Mulchinock**Photography** Louise Parker, Getty Images, Corbis, Photodisc**Illustration** Paul Bateman, Chris Winn, Shane Collinge, Andrew McLaughlin**ADVERTISING SALES****Senior Portfolio advertising manager** Kelly Corten
01225 442244 kellycorten@futurenet.co.uk**Key account manager** George Gill george.gill@futurenet.co.uk**Business development manager** Charlotte Donovan**Senior Sales Executive** Stuart Quinn**Ad sales executives** Ruth Huchrak, George Lucas, Laura Cullingworth**Classified Sales Executives** Helen Tullett, Zahid Patel**MARKETING AND PROMOTIONS****Marketing Manager** Fiona Tully**Subscriptions Executive** Marie Spicer**PRODUCTION****Production Co-ordinator** Diane Ross**Group Production Manager** Clare Tovey**Advertising Production Co-ordinator** Jo Crosby**MANAGEMENT****Group Publisher** Dave Taylor**Publishing Director** John Weir**Managing Director** Colin Morrison**DISTRIBUTION AND CIRCULATION****Circulation Manager** Jamie Malleyjamie.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 Wearsimon.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 96**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 TPL

© Future Publishing Ltd 2003

LINUX is a trademark of Linus Torvalds, GNU/Linux is abbreviated to Linux throughout for brevity. All other trademarks are the property of their respective owners

Future Publishing Ltd. is part of The Future Network plc.

The Future Network produces carefully targeted specialist magazines and websites for groups of people who share a passion. We aim to satisfy their passion by creating titles that offer superb value for money, trustworthy information, multiple ways to save time and money, and are a pleasure to read or visit. Today we publish more than 80 magazines and over 30 magazine websites and networks from offices in five countries. The company also licenses 32 of its titles resulting in over 60 local editions in a further 23 countries.

The Future Network plc is a public company quoted on the London Stock Exchange (symbol: FNET).

Non-executive Chairman Roger Parry

Chief Executive Greg Ingham

Chief Operating Officer & Managing Director, UK Colin Morrison

Group Finance Director John Bowman

Tel +44 (0)1225 442244

BATH • LONDON • MILAN •
NEW YORK • PARIS •
SAN FRANCISCO • WROCLAW

www.thefuturenetwork.plc.uk



Media with passion

NEXT MONTH

BREAKING NEWS:
SCO vs Linux – full report on the claim filed by SCO against IBM for alleged misuse of UNIX code. Details, interviews and comment next issue!

LINUX FORMAT AWARDS 2002

AWARD WINNERS!

The votes are in, the polling boxes closed, and in next month's issue we'll reveal the winners of the most prestigious awards for Linux. Did your favourite project, app or company win? Did YOU win? Join us next time when we recognise and reward those who gave the most!

LINUX FOR MAC USERS

Why would you use Linux, what would you use it for, which version would you use and what problems might you have with it? Find out with our brief guide to installing Linux on Mac hardware.



ON TEST

The latest and greatest hardware and software for Linux lines up to be counted. In among our tests find out how well the latest Gigabyte server performs, take a look at some customised IP programming libraries and design your latest hardware project with VariCAD 9.

PLUS:

Don't forget that the only place you can get the new A4 *Linux Pro* and its 24 pages of real-world Linux for IT professionals is with *Linux Format*.



Issue 40 on sale Thursday 24 April

DON'T MISS YOUR COPY OF LINUX FORMAT!
Get it delivered to your door every month – subscribe on page 96

The exact contents of future issues are subject to change

Welcome

Twenty-four pages of real-world Linux for IT professionals

Welcome to another issue of our 'super-sized' *Linux Pro*, and thanks very much for the feedback we received last month. We normally get quite a few messages, but I was impressed with both the quality of the feedback and its quantity! Many good suggestions were made for topics we should cover – please forgive us if we didn't reply to you all individually, but keep it up! Your reward is a magazine that reflects what you want to read about.

This month we have dipped our toes back into the waters of high Performance Computing. SGI's latest supercluster, the Altix 3000 is a real-life Linux record-breaker, and we're pleased to be able to give you some more information on what it's capable of, and how it has been developed with close co-operation with the Linux community. Even if superclusters aren't really your bag, I think you'll find it interesting to find out how companies like SGI are contributing massively to the development of Linux itself on the worldwide stage.

On the security front we have a short piece illustrating what many sysadmins secretly know but aren't really sure what to do about – the biggest security threats to your network are usually on the inside. Good policy-making and detection can mitigate the risks though, as this piece by David Taylor attempts to demonstrate.

Also on the security side of things, we have the concluding part of our *iptables* series, which looks at options for extending *netfilter* support as well as some solutions for common routing problems. The feedback for part one has been tremendous, so I expect this will be pretty popular.

We're also finishing off the *Apache* rewriting and *Tomcat* tutorials this time, but we'll have some more web-centric material popping up soon, as well as an extended storage section.

And again, if you do have ideas for things you'd like to see, or comments on the material we have featured, drop us a line.

Nick Veitch Editor
nick.veitch@futurenet.co.uk



“I hope you will find it interesting to discover how companies such as SGI are contributing massively to the development of Linux as a whole.”

Contacts

EDITORIAL

Editor

Nick Veitch nick.veitch@futurenet.co.uk

Art editor

Julian Jefferson julian.jefferson@futurenet.co.uk

Production editor

Matt Nailon matthewnailon@futurenet.co.uk

Contributors

David Taylor, David Coulson, Chris Brown

Photography

Louise Parker, Eyewire

Group Publisher **Dave Taylor**

ADVERTISING

SeniorPortfolio Advertising Manager Kelly Corten
01225 442244 kelly.corten@futurenet.co.uk

Key Account Manager George Gill
george.gill@futurenet.co.uk

Business Development Manager Charlotte Donovan

Senior Sales Executive Stuart Quinn

Ad Sales Executives Ruth Huchrak, George Lucas,
Laura Cullingworth

Classified Sales Executives Helen Tullett, Zahid Patel

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

CONTENTS

SGI ALTIX 3000

SGI's record-breaking Altix **p4**

SECURITY

The greatest threat is internal **p8**

STORAGE

External Maxtor 5000XT **p11**

APACHE REWRITING

Rewriting with other modules **p12**

TOMCAT

Architecture, access & logging **p16**

IP TABLES

DNAT, SNAT and extensions **p20**



SGI's Linux

Record breaker!

SGI's latest server launch turned heads at the recent LinuxWorld Expo in New York. It wasn't the groovy rack case or the brightly coloured trim that attracted visitors, but the promise of the best performing Linux system that money can buy that piqued interest.

The Altix 3000 family of servers and clusters is the culmination of several years hard work by SGI and the Linux community in general.

Designed for high performance tasks, the Itanium 2 based servers have a number of unique features which have enabled a leap in performance and opened up new avenues for high performance computing on Linux.



NICK VEITCH breaks his own records to bring you a report on the SGI Altix.

High performance computing (or HPC) is a big area for Linux. With commercial software and operating systems often being licensed on a processor basis, building large clusters has been Linux territory for some time. Sideways scaling (running large numbers of independent but interconnected nodes) is one thing, but Linux was never built for vertical scaling, and performs badly when run as a single instance on a large multi-processor system. At least, that's what Linux's detractors have always tried to have us believe. But the Altix 3000 series may just change all that. So, how does it work?

Memory

One of the key considerations in multiprocessor/cluster environments is memory handling. Since any memory location on the system may be required by any individual processor, this is usually one of the performance bottlenecks – the memory may be local to processor number 1, but if processor number 12 needs it, the time taken to transfer the data is usually far greater than the time to perform the actual processor operation. The usual system architecture solution is NUMA – Non-Uniform Memory Access, which allows for the fact that some data is stored in 'local' memory, while some is stored elsewhere on the system. SGI have addressed this problem with their NUMAflex technology. Data from remote memory nodes still has to be transferred, but SGI has redefined the interconnect fabric which connects the processor units together. The new NUMAlink interconnect fabric carries memory data and network information between units and nodes with latencies as low as 50 nanoseconds – a figure so low that to all intents and purposes, all memory on the system can be considered part of a global pool.

Numaflex isn't a new technology – it was first introduced to the SGI Origin series on MIPS 64-bit processors in 1996, but this is the first time it has been used on an Itanium system. The same principles apply to the construction of a system. Components are packaged together into 'bricks' which are then put together to build a system.

XFS

The Altix range of systems also implements the XFS filesystem developed by SGI. Who needs another filesystem? We already have ext2, ext3, Reiser, JFFS and many more. XFS was developed by SGI to provide a high performance 64-bit compatible filesystem that was journaled, but also gave a similar performance to a traditional filesystem like ext2.

On massively multi-user systems, the journaling overhead becomes more of a hindrance, and ext2 easily outperforms ext3, Reiser and others by a wide margin. By contrast, XFS can actually outperform ext2 in some circumstances, and gives better overall performance than ext3, while still providing robust journaling.

There are also further useful additions in functionality. The XLV volume manager allows flexible attribution of partitions on the system, with a number of additional options such as concatenation (pooling disk

space), striping (optimising disk access for speed) and plexing or mirroring data (for extra security). There is support for dynamically altering the XLV space – very useful for high availability systems that are running out of storage!

To help programmers, XFS can also operate in a guaranteed I/O rate mode. Rather like USB2 and IEEE1394, this means that a certain data bandwidth can be guaranteed by the system for a time, enabling the confident ability to handle streaming media or data acquisition software whenever it's needed.

For backup purposes, XFS also supports active disk dumping – you can dump an image of the disk without unmounting it, and get a consistent, useable image, even when the device is actively being used. This certainly makes backups an easier task, and SGI provide utilities to make best use of this feature.



“Numaflex isn’t a new technology – it was first introduced to the SGI Origin series in 1996, but this is the first time that it has been used on an Itanium2 system”

Bricks

Each brick is a 1 or 2U box containing one of five different system building components. Usually the most populous of these in any given system is the ‘C-brick’, which contains processors and memory. A single C-brick can contain up to 4 Itanium 2 processors and 32GB of shared memory. Importantly, it also contains two specialist ‘SHUB’ ASIC devices, used to communicate between the processors and memory within the brick, and other C-bricks within the node or in remote nodes. As each SHUB has a maximum of two processor located on its Front Side Bus (FSB) it is capable of providing a bandwidth of up to 6.4GB/s. Limiting the number of processors on the SHUB FSB was a design decision to maximise the available bandwidth on each C-brick and throughout the system to minimise the effect of bottlenecks.

The other brick types provide storage and interconnect functionality. The IX brick is the base I/O module, the R-brick provides router interconnect facilities, the larger D-brick provides hard disk storage and the PX-brick adds space for up to 12 PCI-X hotswappable cards.

The flexibility of the brick system means that servers can be customised for their tasks. Some applications may largely require processing power, in which case a single rack can house 32 processors and the require interconnect. For a system where large volumes of storage are required, a 24-processor system could be built with space for a D-brick in a single rack.

While the NUMA interconnect system blurs the line between nodes and components of the same node in terms of connection speeds, there is a limit of 64-processors for running a single system image - ie for the purposes of processes and the OS, a 64-processor system can be seen as a single entity.

Linux doesn’t scale?

A 64-processor Linux system seems a little adventurous. Aren’t we supposed to believe that Linux doesn’t scale well? While it is still true of any useful operating system that you can’t get linear scalability in real performance, Linux isn’t that far off the mark these days.

Linear scaling would be if a task on a single processor machine could be performed 64 times faster on a 64-processor machine. Due to the overheads of inter-processor communication, this is an ideal rather than an achievable result. But improvements to the kernel can make a difference to scalability, and a lot of work on this has already been done, thanks to the profile of scalability as an issue. SGI has been able to contribute to kernel development here, as well as build on the work of groups such as the Linux Scalability Effort. 

COVER FEATURE SGI SERVER

« The LSE is really a collection of subprojects that cover matters such as scheduling and support for NUMA. Kernel support for NUMA systems is important, because scheduler awareness of the physical layout of the memory is crucial to making sure that processes are dispatched to CPUs adjacent to the memory they will access. This obviously results in a more efficient system, where processes can utilise memory directly through a local SHUB. Without this kind of system being present, the process could be running on a processor far removed from the memory allocated to it. In a system like the Altix the memory latency is very low, even in this sort of scenario, but it is still an inefficient use of the resources and will undoubtedly lead to bottlenecks as more and more processes are allocated.

SGI has further expanded on this functionality by adding kernel additions such as CPUMemSets, which allows processes to be directed to run and use the resources of specific groups or single processors. This allows for a finer grained management of the system, with priority given to the processes that require it.



The 'entry level' Altix 3300 server is a half rack device which can accommodate up to 12 Itanium 2 processors and 96GB of pooled memory.

As well as functionality improvements, SGI has also implemented a number of monitoring techniques that are of special interest to high performance computing. One of these is the Linux Kernel Crash Dump module, which is aimed at providing more useful crash information. The LKCD software will save the kernel memory image in the event of a crash and analyse it on restart to determine the exact nature of the failure.

Interestingly, the construction of 64-processor systems gave SGI to the opportunity to test a lot of other system software and it's suitability for this sort of high stress work. For example, the Altix systems implement the devfs device filesystem for managing device names. One of the key advantages here is that the device names are persistent across reboots, and don't vary with whatever interface was registered first – a great boon when dealing with tens of disks spread across different fibre channels and controllers.

One further improvement to Linux on the driver side is the implementation of the SGI SCSI subsystem, originally developed for IRIX. While much work is going on in reinventing the SCSI subsystem for Linux to get over some limitations (especially when it comes to multiprocessor systems); the IRIX SCSI system running on the Altix delivers much better performance overall.

DEVELOPING SOFTWARE

Having the fastest Linux box in the world isn't much use if you haven't got any software to run on it. Fortunately, there is a rich environment of development tools available. Intel produced their own compiler for Linux, which will optimise software for the Itanium series. The good old workhorse GCC can also compile optimised code to an Itanium2 target. Both of these will also work with Fortran source code, which is still in wide use in the field of computational mathematics.

As well as compilers there is a choice of debuggers. The Intel debugger partners

the range of compilers and provides support for MPI and OpenMP systems debugging. The Gnu debugger is also available and will work equally happily on 64-bit code.

A range of support libraries is also provided by SGI. These are essential for making use of some of the features of the Altix environment, such as the MPT, CPU sets, array services for building clusters, and the SGI SCSL (Scientific Computing Software Library) which provides a huge range of optimised, useful scientific and mathematical functions.

ITANIUM 2

Work on the Altix servers began on Intel's original Itanium processor. The original Itanium was bulkier, more costly and generated a lot more heat, so building huge superclusters out of them wasn't a particularly easy task. Although the Itanium never sold in great volumes, it was seen very much as a developer release to give hardware builders and software developers the 64-bit Itanium environment to work with. Although some Itanium-based servers did appear, it wasn't until the launch of the Itanium 2 in 2002 that most server manufacturers began to introduce devices based around this processor.



Available in 900MHz or 1GHz clock speeds, the Itanium 2 features three levels of cache – a 32K level 1 cache, 256K level 2 and up to 3MB integrated level 3 cache. Combined with a 6.4GB/s system bus bandwidth, this gives the Itanium 2 impressive performance, twice the computing power of the original Itanium. New technologies in the Itanium 2

also include a Machine Check Architecture, which handles various system errors intelligently to enable highly available systems.

As a result, the Itanium 2 features widely in high-performance solutions from suppliers such as NEC, HP, Toshiba and Unisys, as well as SGI themselves.

Linux and the Itanium have a long history. Intel encouraged the development of 64-bit tools for the Itanium and kernel support for the processor long before launch. That's now paying off, as the Itanium 2 is able to leverage a host of tools and software already ported to 64-bit Itanium Linux.

Consequently, reported throughputs are up to 1000MB/s, around five times faster than the current kernel drivers.

Breaking records

So, what records have been broken? Well, the result of all this effort is the best-performing single instance of Linux currently available. In SPECfp_rate2000 tests, a 64-processor Altix 3000 recorded a score of 862. By comparison, HP's Superdome, a PA 8700 based supercomputer, scored 267. That's not breaking a record – it's smashing it spectacularly. Similarly, a 32-processor based Altix outperformed IBM's eServer to record a SPEC CFP2000 score of 443, more than 1.5 times the performance of the IBM machine.

The Altix set another record for memory bandwidth, recording an amazing 125GB/second on a single instance of Linux running across 64 processors. That's over 4 times the performance of the Superdome.

Of course, benchmarks are benchmarks, and the real world is a completely different kettle of bottlenecks; but the Altix range is already able to address High Performance Computing demands with real-world applications. SGI have tested the Altix with a range of software commonly used in HPC environments and have been pretty pleased with the results. The STAR-CD computational fluid dynamics software for example (used in the automotive industry) managed a 45 times speedup when running on 64 processors compared to a single processor running the task.

All SGI has to do now is watch the order book fill up – there's no better advert than good results. In the meantime



The Altix 3300 is designed for smaller tasks, as a supercluster node or development platform for the 3700.

“As well as functionality improvements, SGI has implemented monitoring techniques of interest to high performance computing.”

they have not only proved the Altix concept, but also thoroughly discredited the notion that Linux doesn't scale for high performance tasks. ■

I thought my network was secure!

Hopefully, if you've ever found yourself saying this, it was before an intrusion or security audit informed you of a failure. If you have, you're not alone. Many systems administrators have been faced with this stark reality when discovering that network security is much more than simply protecting from the 'outside-in'. While securing the network with demilitarized zones (DMZs), gateways and intrusion monitoring software are all excellent intrusion counter-measures, they do not address the security of the network from the 'inside-out'. For example, if an administrator issues a command causing a clear-case password to be transmitted, the network is at risk. Or if an applications administrator passes around a generic applications password, no matter how much external intrusion detection hardware and software are installed, the security of the network could fail, causing assets to be at risk.

DAVID TAYLOR
Managing Director
EMEA of Symark
Europe Ltd,
stresses that most
threats to your
network security
actually come
from inside your
company.

Much focus on network security often 'faces outwards' and is well-managed. But our IT practices sometimes circumvent this security profile where we are most vulnerable – from *within* the network, by 'trusted' users and administrators. As Sysadmins, we need to think about protecting ourselves from ourselves. Even in today's distributed, heterogeneous, n-tier environments, comprehensive security is readily available in usable software packages. These products focus on securing the network from the inside, by providing granular control over who can access systems, from where and when systems can be accessed, and by restricting or prohibiting action that could compromise data assets.

Password resets

Near the beginning of the article, I posed two potentially disastrous scenarios. In the first, a systems administrator

FOCUSING ON AN 'INSIDE-OUT' THREAT

Decide your security policy

	RDBMS		Web		Financials		ERP			PowerPassword		
	System1	System2	System3	Web 1	Web 2	Fin01	Fin02	Fin03	MFG001	logistics01	Allowed login method	IP access from
REGULAR USERS												
realuser2				✓	✓				✓	✓	ssh,ftp	any
realuser3				✓	✓				✓	✓	ssh,ftp	any
realuser4				✓	✓	✓	✓	✓			ssh,ftp	any
SYSTEMS ADMINISTRATORS												
realuser5	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	telnet,ftp,X-term	any
realuser6	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	telnet,ftp,X-term	any
realuser7	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	telnet,ftp,X-term	any
NETWORK ADMINISTRATORS												
realuser8	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	ssh,ftp	any
realuser9	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	ssh,ftp	any
RDBMS ADMINISTRATORS												
realuser10	✓	✓	✓								ssh,ftp	any
realuser11	✓	✓	✓								ssh,ftp	any
FINANCIALS ADMINISTRATORS												
realuser12						✓	✓	✓			ssh,ftp	any
realuser13						✓	✓	✓			ssh,ftp	any
realuser14						✓	✓	✓			ssh,ftp	any
APPLICATIONS ADMINISTRATORS												
realusername15									✓	✓	ssh,ftp	197.206.187.25
realusername16									✓	✓	ssh,ftp	197.206.187.26
realusername17									✓	✓	ssh,ftp	197.206.187.27

issues a command that causes a clear-case password to be translated across the network. A common example of this occurs when the sysadmin (or help desk) is called to reset a password that a user has forgotten. The admin simply resets the password without considering the implications of such an act. What's more, in today's distributed computing environment, a password-sharing facility is often used, so that the user can log in to many systems, sharing the same username and password. In this example, when the admin resets the password, it is transmitted in clear-case to the server master. This potentially allows an intruder snooping your network to capture important passwords, compromising valuable data.

A more secure alternative uses encrypted password synchronisation, available in some of the best-of-breed security products available today. Here's how this works: When security software is installed, multiple configurations are possible. These include everything from flyweight, thin-client programs, to systems requiring kernel modifications and multiple module configurations. It is important to choose a product that allows configurations to encrypt data for network traffic (and potentially files) between the systems. Once in this configuration, arrange each (non-'master') system as a 'slave'. This typically means the system holds a cached copy of the access system – be it a password and history database, login and password policy file(s), or a combination of both. This method enables password changes from local systems to be written via

ABOUT SYMARK

Symark is exhibiting at Infosecurity Europe, Europe's largest and most important information security event. Now in its 8th year, the show features Europe's most comprehensive FREE education programme, and over 200 exhibitors at the Grand Hall at Olympia from 29th April – 1st May 2003. www.infosec.co.uk

encrypted network traffic to the master server, which then propagates the copy to all slaves (also using encryption). No clear case passwords are ever transmitted. However, care must be taken to keep the network key (and potentially the files key) in a safe place, as they must be equivalent on each system for communication to take place.

Password history

The R.O.I. for this type of system is immediate and considerable. A rich set of password history and quality rules requiring minimum length, punctuation, and alpha/numeric standards (including a comprehensive set created for importing all new accounts) is set in place on the network.

Adding further security, these rules include a history component that prevents users from reverting back to previous passwords any time they are prompted to make a

“Our IT practices sometimes circumvent this security profile where we are most vulnerable – from *within* the network.”

password change, and can also include a 'three strikes' rule, disabling a login after three (or some other easily configurable number) login attempts.

Combine this with a login policy that specifies the systems users can access, the devices and IP addresses from which users can gain access, the specific days and times of day when access is permitted, and the ability to inform your security staff in real time of repeated access denials, and you have added significant security from the 'inside out'. If a powerful account's password has been compromised (the *Oracle* account, for example), protection still exists by restricting *Oracle's* login to a specific system and IP address. Security goals are thus achieved by allowing staff access to perform their duties, while thwarting and reporting errant or malicious access attempts before vital assets are compromised.

In the second scenario, a powerful generic account and password is passed around by applications administrators. This is common among many of the multi-module RDBMS-based applications for ERP, clinical trials management, financials management, manufacturing, human resources, and others. There is an old adage that states, "A secret is something that is hushed around from person to person." Experienced administrators understand this truism only too well.

While these practices are seemingly convenient or even necessary for business to function normally, they can cause serious damage to the system. What if the RDBMS administrator mistakenly runs the wrong stored procedure, or connects to the wrong (but similar) database instance before dropping tables? In a perhaps less-likely example, an applications administrator could gain access to a system which contains medical patient information, causing a potential breach of HIPAA policy or other audit guidelines. While infrastructure security teams may not always view these accesses as in scope, or even

Power Broker				
Days to login	Hours to login	UID	UserName	Run this command with root (or privileged account) permission
m-f	07:00-19:00	2000	user10	none
m-f	07:00-19:00	2001	user11	none
m-f	07:00-19:00	2002	user12	none
any	any	3000	user16	all commands
any	any	3001	user17	all commands
any	any	3002	user18	all commands
any	any	100	user31	all commands
any	any	101	user32	all commands
any	any	200	dbadmin1	all database commands
any	any	201	dbadmin2	all database commands
m-f	any	300	fin01	financial applications commands
m-f	any	301	fin02	financial applications commands
m-f	any	302	fin03	financial applications commands
any	any	400	appadmin01	applications administration commands
any	any	401	appadmin02	applications administration commands
any	any	402	appadmin03	applications administration commands



view the access as a security breach at all, assets have been placed at risk.

Security enhancements

Defending data assets from the 'inside-out' makes the network more secure, without causing disruption to business best practices. More and more businesses are aligning their business model with their IT infrastructure by translating business use-case scenarios into comprehensive policy and access management. By determining who should be able to log into systems, execute certain commands, switch user id's to a more powerful account, and access data with certain file systems, internal security is greatly enhanced without putting expensive development or implementation teams through the security wringer.

This kind of security is taking its place in technology's acronym stew, assuming monikers like Rules Based Access Control (RBAC), Policy Based Access Control (PBAC) and Roles and Rules (R&R). Think of it as protecting ourselves from ourselves, or protecting the network from the 'inside-out'. Here's how to make it work:

Take time to map out your infrastructure business processes, as you would when developing any use case (or business case) scenario. The model begins when someone attempts to access a system, or execute a command. The pre-thought and pre-planning spent here can not be over-valued – it is pennies on the pound compared to the value of the assets you are protecting. Use any visualisation tools necessary to make this easier, and translate them into a framework on which your policies and restrictions (or roles and rules, etc...) will be based. The table on the previous pages is an example (using a spreadsheet), including systems, accounts, users, applications, commands, and login methods. Note that commands on the spreadsheet are only generically illustrated. For your own spreadsheet you would potentially list complete path names to real commands, as well as real system names.

Admittedly, this looks a little thick, but it contains information you need to create comprehensive access and restrictions from the 'inside-out'. In this spreadsheet, the user community is on the left axis grouped by business role or duty. Systems types, access times and systems, and allowed commands are listed across the vertical axis. From this spreadsheet, we can determine:

- Who can log in
- From where and when they can log in
- To which system(s) they can log in
- What method of access is allowed to the system(s)
- Which commands are allowed by users

When creating your own spreadsheet use your best methods to visualise your use cases in a complete matrix and evaluate potential policies and restrictions. Of course, this example is limited – it does not show allowed login environments, full path names to allowed commands, directory limitations that are placed on users (if any), and much more. In creating your own spreadsheet, you will think of your own scenarios based on your organisation's business practices. The key is to create a comprehensive 'inside-out' set of policies that do not limit your company's ability to do business.

“Many of the products on the market offer GUIs to help users create policy without having to be a syntax expert”



Next, move your spreadsheet (or equivalent) into practice using the security software (and languages) of your choice. Migration of your spreadsheet to systems practice is usually accomplished from 'simple-to-elaborate', using a basic access case, then building additional restrictions as you become comfortable with the product's language. Many products offer graphical user interfaces (GUIs) to help you create policy without having to be a syntax expert.

The resultant policy should meet each of the use cases you defined in your spreadsheet. For example, the login policy restricts access to systems by evaluating the username, login method, originating IP address, day, time, and UID. If any of the values in the spreadsheet do not match, the login is rejected, the event is logged, and perhaps a SMTP e-mail event notification is sent to the infrastructure group in real time for investigation. In this example, we have focused on the threat from the 'inside-out', and we have been inside the company LAN or WAN the entire time.

Indelible auditing

Additionally, the ability to provide indelible auditing of sessions is not only critical for security, it is also useful for monitoring outside consultants as well as for training. The ability to create a comprehensive log file, containing every keystroke, as well as standard out and standard error messages, allows security teams and administrators to determine who did what to the system in case of a disaster, or to provide training to new staff for tasks like configuring databases, firewalls, or any activity you choose to monitor in your role and rule set. Some products can track sessions across multi-system access, so if a user becomes another user, or gains access to another system, auditing continues, reporting activity from there as well.

While making the environment more secure without inhibiting normal business is the central goal, other steps can further secure the environment. You can restrict the commands allowed by the users that have already passed our login policy by using additional security configuration.

For example, users who are allowed system access are further restricted to the limited sets of commands that constitute their normal business practices. This can be done by giving users a menu-style interface to the commands, silently creating iolog files for each allowed command and/or even the entire user session for auditing or training purposes. Using this type of security helps protect ourselves from ourselves. If a regular user, database administrator (DBA), or even a network administrator issues a command we have determined to be outside normal business practice, the configuration policy simply rejects the command, without allowing damage to the system. In the event an account has been hijacked by an intruder, the attempted execution of forbidden commands can also trigger real-time notification of misuse, before any damage is done.

Pre-planning these policies and configurations before attempting to implement them will not only provide a more comprehensive, less intrusive security, but it will also minimise the frustration level for everyone concerned.

External access protection is an important part of any comprehensive IT security plan, but don't ignore the (statistically) larger threat...from the inside. ■

EXTERNAL HARD DRIVE

Maxtor 5000XT

An external storage device gets the once-over from **NICK VEITCH**.

While much of the time we may think of disk storage devices in terms of internal drives, the use of external hard disks, while not as

**INFORMATION**

PRICE £251
SUPPLIER Maxtor
WEB www.maxtor.com

SPECIFICATIONS

TYPE IEEE1394/USB2 external
RPM 5400
SEEK TIME 11ms
CAPACITY 250GB
BUFFER 2MB

Two ports should be supplied as standard on all IEEE1394 devices to allow chaining!

**DANTZ BACKUP**

One of the features of this drive is the automated backup feature. Pressing the button on the front of the drive causes a Dantz backup server on the host to kick into life and begin an unattended backup of any data you may have specified. Obviously a 250GB drive is quite suited to this task, and since the backup requires no user

prevalent as it once was, hasn't yet faded into obscurity. There are plenty of benefits to using external drives – for a start, they can easily be switched between systems, to provide storage where they are needed.

The Maxtor *Personal Storage* range has been designed very much with this in mind. Portability has obviously been a key design notion, and the result is a self contained drive that is barely larger than a 5.5" bay device

interaction other than the initial setup, this makes the drive useful as a mobile backup solution.

Unfortunately, although Dantz do support Linux on the client side, there is no Linux server software as yet, so this feature is not implemented on a Linux system, though if you also use Windows or Macintosh machines, it may be of interest.

(41 x 152 x 219 mm are the exact dimensions if you are interested).

With a weight of 1.22Kg (and a little more if you factor in the weight of the custom power supply) it is certainly portable enough to slip into a laptop case with big pockets. The case isn't particularly ruggedised – plastic corner protectors and curved rubber feet may provide some protection, but don't expect it to survive too many journeys through baggage reclaim. The drive itself seems well protected inside the case though, which is perhaps just as well, as the snap-together case design deters opening it up for any guarantee-endangering tinkering.

Capacity

Maxtors *Personal Storage* range includes a number of devices, including the 5000LE, 5000DV and the 5000XT (reviewed here). The main difference between them is merely the drive unit installed inside, and hence capacity, with the LE starting at 80GB, through to the 250GB offered in the 5000XT, reviewed here. These are unformatted capacities – for our tests we set up an ext3 partition which managed a formatted capacity of 232GB – still plenty big enough for most uses. Bear in mind that you may get different results with other formatting options.

Connectivity

The *Personal Storage* drives come with both Firewire (IEEE1394) and USB2 connection options. The USB connection is backwards-compatible with 1.1, useful if you want to share the device between a number of computers which might not all have the most up-to-date connections.

Although USB 2 support is included in the latest development kernel, in previous tests we have found it hard to get it to work reliably, and therefore we didn't include this as part of our tests. As well as the USB connection,

the rear of the unit sports two six-pin IEEE1394 ports. This is a bonus, as many external FireWire devices only bother with one port. Having two means you can chain devices, such as another 5000XT drive if required, or perhaps more usefully if you only have one port available, other devices such as DV cameras.

Although IEEE1394 support is still classed as experimental in the current stable kernel series, it has proved reliable enough, especially for generic SBP-2 devices like hard drives. As with USB, the IEEE1394 device will be patched through the kernel SCSI modules, so the unit will appear to the Linux system as a SCSI hard drive.

Performance

The actual drive in the 5000XT unit is a Maxtor Diamond Max Ultra ATA133 unit, with a 5400rpm spin speed. We tested the real-world speed by copying files from a RAMdisk to the drive and reading them back, and using the xhdbench utility.

For sequential I/O, sustained speeds of up to 28MB/s were achieved – a little lower than the quoted 37.5MB/s. The actual bottleneck here may be the kernel implementation of SBP-2 rather than the drive, but nevertheless, it gives results similar to internal ATA drives, and it's certainly more consistent.

If you are worried about drive noise, don't be. It isn't going to be terribly noisy at 5400rpm anyway, and much of the noise seems to be soaked up in the case – you'll find yourself checking the light to make sure it's working. Any noise that does escape is sure to be blotted out by the noise of the computer it's attached to anyway. ■

VERDICT**MAXTOR 5000XT****FOR**

- + Good value for 250GB storage
- + Quiet and portable

AGAINST

- Not the fastest way to boost storage capacity

Value 10/10

Features 10/10

Performance 8/10

Overall 9/10

public_html/\$2

This configuration would map a subdomain of example.net to a user's public_html directory, by virtue of a RewriteRule. If we accessed <http://www.david.example.net/index.html>, we would end up rewriting to /home/david/public_html/index.html. We've first used a RewriteCond to check that the hostname is actually valid before proceeding to stick the HTTP_HOST variable on to the front of our document path. Now, rather than performing the rewrite against purely the document path sent by the client, we can rewrite the entire URL, creating the opportunity for a dynamic hosting system.

We're only going to perform a rewrite on the actual HTTP_HOST variable at the moment, so the actual content path will not be modified. We can quickly perform a rewrite based on %{HTTP_HOST} using a mapping:

```
RewriteMap host-map txt:/usr/local/apache/conf/hosts.txt
RewriteEngine On
RewriteRule ^(.+)    %{HTTP_HOST}$1    [C]
RewriteRule ^(www.)*?([^\+]+)(.*) /home/${host-map:$2}
/personal_web/$2/htdocs/$3
```

Our hosts.txt would contain a list of domains and the user who runs that site:

```
davidcoulson.net    david
example.net        bob
```

Note that we've got a option **www.** string in our RewriteRule, so if www.davidcoulson.net was access, it'd automatically be handled as **davidcoulson.net**. It's rare for the domain and the www RR to point to different sites, so this solves many problems. We also need to setup another mapping, using an internal function of **mod_rewrite**, to generate lowercase HTTP_HOST variables, otherwise **Davidcoulson.net** and **davidcoulson.net** would be handled differently.

We can extent upon these rules by including capabilities for CGIs, as well as for our /icons folder, which usually has to point at /usr/local/apache/icons, for when a client performs a directory listing.

```
RewriteMap lowercase int:tolower
RewriteMap host-map txt:/usr/local/apache/conf/hosts.txt
RewriteEngine On
RewriteRule ^/icons/(.+) /usr/local/apache/icons/ $1 [L]
RewriteRule ^(.+)    ${lowercase:%{HTTP_HOST}}$1 [C]
RewriteRule ^(www.)*?([^\+]+)cgi-bin/(.*) /home/${host-map:$2}/personal_web/$2/cgi-bin/$3 [T=application/x-httpd-cgi,L]
RewriteRule ^(www.)*?([^\+]+)(.*) /home/${host-map:$2}
/personal_web/$2/htdocs/$3
```

We now have the issue that the web server will log everything all in one lump, so it's pretty much impossible to figure out which site the log entry refers to. The easiest way to do this is to define a couple of variables as part of the rewrite, allowing us to include these variables in the log. We can then write a parser to look through the logs and chop it up so each user gets their own.

```
RewriteMap lowercase int:tolower
RewriteMap host-map txt:/usr/local/apache/conf/hosts.txt
RewriteEngine On
RewriteRule ^/icons/(.+) /usr/local/apache/icons/$1 [L]
RewriteRule ^(.+)    ${lowercase:%{HTTP_HOST}}$1 [C]
RewriteRule ^(www.)*?([^\+]+)cgi-bin/(.*) /home/${host-
```

A simple virtual hosting system can be developed using mod_rewrite, making for a very flexible installation.



```
map:$2}/personal_web/$2/cgi-bin/$3 [T=application/x-httpd-cgi,L,E=VHOST:$2]
```

```
RewriteRule ^(www.)*?([^\+]+)(.*) /home/${host-map:$2}
/personal_web/$2/htdocs/$3 [E=VHOST:$2]
```

We can now modify our log file script to use the VHOST variable as part of the log line:

```
CustomLog "/usr/local/apache/logs/access_log"
"%{VHOST}e %h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\""
```

The actual functionality of CustomLog does not change, and since we would want to run it through a parser, it's a good idea to create separate log files for each day:

```
CustomLog "/usr/sbin/cronolog /usr/local/apache/logs/%Y/%m/%d/access_log" "%{VHOST}e %h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\""
```

This is all with a static text file as a rewrite mapping, but we can provide access to a backend database by writing a simple script to interface between Apache and the database. We can have a 'program' mapping type, which spawns a single process from the webserver which can access a database, read from files and really do anything we like. It just needs to do everything on stdin and stdout, so it can be written in something as simple as **bash**. Apache will feed us a mapping key on stdin and we need to spit one back to it on stdout. By using a Perl or PHP script, it is trivial to provide access to a MySQL or PostgreSQL service:

```
httpd.conf:
RewriteMap host-map prg:/usr/local/apache/conf/hosts.php
hosts.php:
#!/usr/local/bin/php -q
<?php
mysql_connect("10.1.1.2","username","password");
$fdin=fopen("php://stdin","r");
$fdout=fopen("php://stdout","w");
set_file_buffer($fdout,0);
while($i=fgets($fdin,256)) {
    fputs($fdout,key_lookup($i)."n");
}
function key_lookup($key) {
    $res=mysql_query("SELECT mapval FROM web.hosts "
        ."WHERE mapkey='".$key."'");
    if(@mysql_num_rows($res)) {
        return @mysql_result($res,0,0);
    } else {
        return $key;
    }
}
```

While `mod_rewrite` gives us many possibilities with virtual hosting, `mod_perl` is far more powerful and this is what we'll be looking at using instead.

URL rewriting with `mod_perl`

`mod_perl` is an *Apache* module which allows Perl code to be executed directly by the `httpd` process, rather than by an external Perl interpreter. While this has numerous benefits for those using CGIs, as we no longer have to spawn a perl process to actually run the Perl code, it also allows us to configure *Apache* using Perl code embedded within our `httpd.conf` file.

Building `mod_perl` is a little different to most other *Apache* modules, as the installer is written in Perl, rather than *bash*. Anyone who is familiar with Perl modules, such as those distributed through CPAN, will be able to jump straight in and install `mod_perl`. However, anyone who hasn't played with Perl modules in the past will likely find the installation sequence a little odd.


We're going to build `mod_perl` as a loadable module through the use of *apxs*, so *Apache* needs to have been built with the module **so** available, otherwise we would be unable to load a module into the server dynamically. It is possible to compile `mod_perl` into the server, although this does require that we recompile *Apache*. Statically compiling `mod_perl` into *Apache* is often a good choice on highly loaded servers, but on the most part, it's not worth recompiling *Apache* for `mod_perl`. Everything will still work as expected when `mod_perl` is compiled as a DSO, plus it allows us to rebuild `mod_perl` or *Apache* independently of each other. As one would expect, this makes life easier when we come to upgrade one or the other, as we don't end up recompiling everything following a small change.

Downloading `mod_perl` is straightforward, and it can either be downloaded and built via CPAN, or from <http://perl.apache.org/download/index.html>. While there are both 1.0 and 2.0 releases of `mod_perl`, these relate to the *Apache* distribution for which the `mod_perl` code is capable of working with. As we're using *Apache* 1.3, we need to download the tarball from http://perl.apache.org/dist/mod_perl-1.0-current.tar.gz. Once this has been downloaded, we can untar it and set it off compiling:

```
# tar xzf mod_perl-1.xx.tar.gz
# cd mod_perl-1.xx
# perl Makefile.PL \
  USE_APXS=1 \
  WITH_APXS=/usr/local/apache/bin/apxs \
  EVERYTHING=1
# make && make test && make install
```

We've told `mod_perl` that we want to build it using *APXS*, then supplied it with the path to our *apxs* binary. The **EVERYTHING=1** option enables all of the Perl call back hooks for use with *Apache*. It's worth compiling them in just in case a particular script desires a call back to be present within the `mod_perl` build.

Once `mod_perl` has compiled it's compilation, it will perform a number of tests to ensure that everything is working correctly. It will then install the `libperl.so` module into our `libexec` directory, and modify `httpd.conf` so that the `mod_perl` module is loaded when *Apache* starts up.



mod_perl is a hugely popular Perl engine for *Apache*, which allows Perl code to be executed by the `httpd` process.

At this point, before we configure `mod_perl`, we restart *Apache* and ensure it comes up properly with `mod_perl` available to the `httpd` process. Anything which returns the *Apache* version, such as the `error_log` file from the startup of *Apache*, a 404 page, or even using **netcraft.com** should state that `mod_perl` is built into the web server.

Currently, we've got `mod_perl` built into our server, but nothing will actually use it. We need to define a directory which we will put our Perl scripts into and then specify that we want them handled by the `mod_perl` engine. Unlike PHP, it's not simply a matter of using an **AddType** directive. Of course, we can use a directory, location or file definition for `mod_perl`, so there are quite a few ways to define which documents are handled by `mod_perl`. Generally, one would provide a `/perl/` location, which is an alias to a directory outside of the `htdocs` directory. Anything people put into this directory is handled as a Perl script.

To achieve this, we would add a **Location** section for `/perl` to our `httpd.conf` file, and set the handler to `Apache::Registry`.

```
<Location /perl>
  SetHandler perl-script
  PerlHandler Apache::Registry
  Options ExecCGI
  Allow from all
  PerlSendHeader On
</Location>
```

The **PerlHandler** is actually a Perl module provided by `mod_perl`, which is part of a large number of other modules forming part of the `mod_perl` API. However, before we can use `Apache::Registry` as a `PerlHandler`, we need to load it. Normally with Perl, one would use the **use** directive, but with *Apache* we need to use the **PerlModule** directive:

```
PerlModule Apache::Registry
```

This **PerlModule** definition should be put before anything tries to use it in the `httpd.conf`, but after the `LoadModule` to load `mod_perl`. To get this to work requires yet another restart of *Apache* is needed, using **apachectl**. With luck, everything will come back up happily and we can actually start to write Perl code and have it handled by our `mod_perl` installation.

With `mod_perl`, we can create our own *PerlHandlers* and deal with scripts as we want. These are basically Perl modules with a **handler** function, which is what deals with the actual request. We can expand on this by using a *PerlHandler* with any request, even if the requested file has nothing what so ever to do with Perl.

We can start writing our own Perl module and then have *Apache* use it as a handler. For URL rewriting, we want to use a **PerlTransHandler**, which is use for URL translation. There are many other handlers, for various locations of the request process, but we're only going to use a

PerlTransHandler.

A simple example of a translation handler in `mod_perl`,

“While both `mod_perl` and `mod_rewrite` are very powerful, one should look at the exact need before getting over-complex.”

which we would normally just use `mod_rewrite` for, would be to redirect a specific location to a query string in a CGI or a PHP script. The Perl code to do something like this would be:

```
<Perl>
package My::Trans;
use Apache::Constants qw(:common);
sub handler {
    my $r = shift;
    my $uri = $r->uri;
    my ($id) = ($uri =~ m!^/news/(.?)!);
    $r->uri("/news/index.php");
    $r->args("id=$id");
    return DECLINED;
}

package Apache::ReadConfig;

</Perl>
```

PerlModule Apache::Registry PerlTransHandler My::Trans

We can either create a separate Perl module which is loaded into *Apache* via a **PerlModule** directive, or we can write the code directly in our `httpd.conf`. Which is chosen depends upon the amount of code and how easily it can be organised. We've set our server up to use the **My::Trans** Perl module, and *Apache* will use the **handler** function within the module to perform the actual translation. We can, of course, define additional functions depending upon the functionality of our module. The handler function gets an *Apache* request object as it's first argument, which we can read from, as well as write to. After we've performed the translation, we return with a `DECLINED` value, which tells *Apache* to continue with the rest of the request process before sending it back to the client. We can at this stage send back a `FORBIDDEN` or another request error.

To tie this in with a database, we simply need to have **DBI::mysql** available, or indeed any module to connect us to our database of choice.

```
<Perl>
package My::MySQLTrans;
use Apache::Constants qw(:common);
$dbh=DBI->connect("DBI:mysql:web:10.1.1.2","user",
"pass");
sub get_hosting_info {
    my $h = shift;
    $h =~ tr/[A-Z]/[a-z];
    if($h =~ /^www\.(.+)$/) {
        $h=$1;
    }
    $sth=$dbh->prepare("SELECT mapval FROM hosts
WHERE type='alias' AND mapkey='".$h."'");
    $sth->execute;
    if($row=$sth->fetchrow_hashref) {
        $h=$row->{mapval};
    }
    $sth->finish;
    $sth=$dbh->prepare("SELECT mapval FROM hosts
WHERE type='vhost' AND mapkey='".$h."'");
    $sth->execute;
    if($row=$sth->fetchrow_hashref) {
        $user=$row->{mapval};
    } else {
        $user="";
    }
    return ($user,$h);
}

sub handler {
    my $r = shift;
    $u = $r->uri;
    if($u =~ /^\/icons\/.+/) {
        $r->filename("/usr/local/apache".$u);
        return OK;
    }
    ($user,$h)=get_hosting_info($r->hostname);
    if($u !~ /^\/(perl/cgi-bin)\/.+/) {
        $u="/htdocs".$u;
    } elsif($1 eq "cgi-bin") {
        $r->content_type("application/x-httpd-cgi");
    }
    $r->document_root("/home/".$user.
"/personal_web/".$h."/htdocs");
    $r->filename("/home/".$user."/personal_web/".$h.
.$u);
    $r->subprocess_env(VHOST=>$h);
    return OK;
}

</Perl>
```

There are numerous other ways to expand this script, depending upon the exact requirements of the service. This code performs almost exactly the same function as the `mod_rewrite` rule we built earlier using a PHP script, but this also supports site aliases. Lots of example code for `mod_perl` can be found at <http://perl.apache.org/docs/1.0/guide/snippets.html>, if you don't want to start from scratch. As we have seen, in true Linux style, there are many ways to do the same thing. While both `mod_perl` and `mod_rewrite` are very powerful systems, one should look at the exact need before jumping in with something overly complex. ■

Configuring the Servlet Container Tomcat



Last month's tutorial looked at the role of *Tomcat* as a container for Java servlets, and went through the process of installing it along with the underlying JDK. We also saw how to write a simple web application by extending the Java class `HttpServlet`, and how to deploy the application, both as an unpacked hierarchy of files, and as a .WAR file. We were able to do all this using the default configuration of *Tomcat*.

This month, we'll be looking at the component architecture of *Tomcat*, and explore how it can be configured. We'll see how to serve multiple hosts (sites) from one instance of *Tomcat*, how to define "contexts" in which web applications can run, and how to define users and roles and use them to control access. We'll also explore the use of components called 'valves' which allow us to do access control and logging.

Tomcat's Configuration Files

The files which control the overall configuration of Tomcat live in the directory `$CATALINA_HOME/conf`. If you didn't read last month's piece (or even if you did, for that matter ...) `$CATALINA_HOME` refers to the top level directory in which Tomcat is installed. In our installation, this was `/opt/tomcat`. The configuration files are shown in a table below as **figure 1**.

Most of our attention in this article is directed at the main configuration file, `server.xml`. This file is roughly analogous to the file `httpd.conf` used to configure the *Apache* web server, and on the assumption that some of our readers may already be familiar with that file, we'll draw a few comparisons along the way.

As you'll guess from the name, this is an XML file. Being the wrong side of fifty, over the years I've seen quite a number of markup languages come and go (remember troff macros, anyone?) and thus I've been slow to get excited about XML. But given the general love affair of the

PART 2 DR. CHRIS BROWN examines component architecture, access and logging.

Java community with XML, it's not surprising to see it used extensively within *Tomcat*. And it has to be said that the hierarchy made explicit by the structure of *Tomcat*'s XML configuration file nicely reflects the hierarchical nature of the component structure within *Tomcat* itself. You don't need to understand the formal details of XML to work with this file, you just need to get your head round the idea of a nested hierarchy of tags, with associated attributes.

Those of you familiar with the *Apache* configuration file will have come across the idea of "container directives" (such as `<Directory>` and `<VirtualHost>`) that define a kind of "context" within which inner, nested directives apply. You'll also be familiar with the idea that settings established at the top level of `httpd.conf` (**LogFormat**, for example) are inherited by containers unless explicitly over-ridden inside the container.

The *Tomcat* configuration model takes these ideas further by defining multiple levels within the hierarchy within the more formal structure of an XML document.

Figure 2 tries to show the overall hierarchical structure of *Tomcat*. It won't all make sense at once, but should help you get the big picture of the configuration file's structure as we work through it. **Figure 3** shows how this structure looks as XML; please note that this isn't a complete example, many of the details have been elided out.

Peeling back the layers

It's probably easiest to start with the outer layers of the structure and work inwards. At the top level is the `<Server>` element. There must be exactly one of these in the file and it must be the outermost element. You can define a port number on which the server will listen for a "shutdown" command, and the exact command it will listen for. (The default is the string "SHUTDOWN").

Inside a server you can place one or more `<Service>` elements. The name "service" here is somewhat misleading; a service element is really just a container into which you can put `<Connector>` and `<Engine>` components. The service can be given a name, which is used in messages logged by that service, and also shows up in the graphical configuration tool, if you choose to use it. Complete examples of `<Server>` and `<Service>` elements, along with their attributes, are shown in **figure 3**.

A service must contain one or more connectors and exactly one engine, so in pidgin-XML it looks like this:

```
<Service name="My Standalone Service">
  <Connector ... />
  <Connector ... />
  <Engine ... />
```

Figure 1: Tomcat's configuration files and what they do.

File	Description
<code>server.xml</code>	The main configuration file which controls the hierarchical structure of <i>Tomcat</i> 's components, described in detail in the text.
<code>catalina.policy</code>	Establishes a default set of security policies to be used by the underlying java execution environment. This file is only relevant when <i>Tomcat</i> is run under the java security manager.
<code>web.xml</code>	Defines system-wide default parameters for web applications. Individual applications can override these by supplying a <code>web.xml</code> file of their own.
<code>tomcat-users.xml</code>	A simple file in which you can define users and give them passwords and roles that can be used to control access to web applications served by <i>Tomcat</i> .

</Service>

Next, a **<Connector>** element defines a component which listens for incoming connection requests. There are several connector types available which understand different protocols. The most important is the HTTP connector; which understands HTTP requests (as might be made by a browser) and allows *Tomcat* to function as a stand-alone web server. The current implementation of the HTTP connector is called the **Coyote** connector. Alternatively, if your service is using some other web server (such as *Apache*) as its front end, you'll need to use a connector which speaks one of the protocols used for server-to-server communication – either AJP13, or WARP. It's possible for a service to have several connectors; perhaps one for HTTP and one for AJP13, or maybe several HTTP connectors listening on different ports.

The HTTP connector has quite a lot of attributes, some of which relate to the TCP/IP configuration of the underlying socket which the connector creates. For example, the port attribute specifies the TCP port number, the **acceptCount** attribute specifies the length of the queue of (unserved) connection request, and the address attribute specifies the IP address on which the connector will listen. (This is relevant only on machines with connections to more than one network.)

There's one other component you must include inside a service, and that's the **<Engine>** element. An engine represents the request processing machinery of the service. The attributes of an engine include a name (which, as for the **<Service>** element, simply serves to identify it in log messages and in the graphical admin tool) and a defaultHost. The notion of a default host is bound up with the concept of virtual hosting and is explained in the next section. A simple **<Engine>** element might look like this:

```
<Engine name="standalone"
defaultHost="info.myco.com">
```

Virtual Hosting

Like *Apache*, *Tomcat* is able to serve multiple virtual hosts. For example, an organisation might host what appears to the user as three separate sites, **www.myco.com**, **info.myco.com** and **downloads.myco.com**, on a single server. The virtual hosting mechanism looks at the host name in the HTTP request header, and forwards the request to the correct virtual host. For this to work, the DNS must resolve all of these host names to the IP address of our server. Each host is defined by a **<Host>** element within an Engine, and there can be as many of these as you like.

A **<Host>** element essentially defines a mapping between a host name (like **www.myco.com**) and the name of the top-level directory where the web applications and content for that host are to be found. This is analogous to putting a DocumentRoot directive inside a **<VirtualHost>** container in *Apache*. The host name is specified as the **name** attribute and the directory is specified as the **appBase** attribute, like this:

```
<Host name="www.myco.com" appBase="mycodir">
```

Here, **mycodir** is relative to `$CATALINA_HOME`; I could also have specified an absolute path name here.




Figure 2: *Tomcat* component hierarchy.

It's time to complete our unfinished business with the **<Engine>** element. The defaultHost attribute specifies which host should be used if an HTTP request arrives for a

“Like *Apache*, *Tomcat* can serve multiple virtual hosts. What users see as three separate sites can be from a single server”

CONNECTIONS

Any port in a storm?

The default port for the HTTP connector is 8080. If you want *Tomcat* to act as your main web server you might be tempted to change this to port 80, the standard port for HTTP. If you do, keep in mind that in the linux world there's a rule that you can only listen on ports below 1024 if you're running as root. There are obvious security worries if your servlet engine (and the servlets deployed in it) runs with super-user privileges, and you probably don't want to go there. *Apache* overcomes this by starting up as root, then dropping to a less privileged identity after establishing the socket on port 80. *Tomcat*, as far as I can tell, lacks this facility.

host for which there is no **<Host>** element explicitly defined. For example, for the three-site scenario in the previous paragraph, how should the server deal with a request destined for the host **uploads.myco.com**? The host nominated as the default host must actually exist within the engine.

Finally, within our nested hierarchy of components, we come to the **<Context>** element. This element defines a web application running on the associated virtual host. Specifically, it maps a context path onto a document base. When I first read the documentation for this, it was far from clear what either of these terms meant, so I'll try to explain. The context path means, in essence, the part of the incoming URL that follows the machine name; for example, the context path specified by the URL **http://info.myco.com:8080/training/catalog.html** would be **training**. The document base means the top level directory where the application is deployed, or the name of the .WAR file in which the application is deployed. (In last month's example, this directory was named `formexample`. Because we put it under the `$CATALINA_HOME/webapps` directory, it was automatically deployed onto the context path `'formexample'` when *Tomcat* started up, so we didn't have to create a **<Context>** element for the application.)

Here's an example of a **<Context>** element:

```
<Context path="/catalog" docBase="cat2003">
```

The **docBase** is relative to the **appBase** of this virtual host.

From URL to Application

Putting all this together, **figure 4** attempts to show how the various components of a URL work together to select the appropriate web application. The last

◀ step of this mapping, which selects the actual class file to be loaded as the servlet, is controlled by entries in the application's web.xml file. This is (as they used to say in the textbooks whenever they got to a bit they didn't want to cover in very much depth) left as an exercise to the reader!

Bells and Whistles

We've now met all the key elements in the config file that are needed to establish a working server. We've missed out many details, such as the debug attribute, which can be used in most elements to control the level of detail that each component will log. There are three other component types – loggers, realms, and valves – which can be placed at various levels in the component hierarchy to extend

“A ‘realm’ in *Tomcat* is a database of user names, passwords and roles, rather than a resource set controlled by a single log in.”

Tomcat's capabilities. The basic idea is that these components are inherited by any components nested inside the element in which they are defined. For example, if a **<Logger>** element is defined within an **<Engine>**, that logger will also be used by any virtual hosts defined within that engine, unless a virtual host over-rides this by defining its own. This is very similar to the way properties defined in the **main** server in *Apache* propagate down into the container directives, unless over-ridden.

A logger defines a destination for logging, informational, and error messages. By itself, a logger doesn't generate any output, it just delivers logging information generated by other components. We'll meet some components whose sole purpose is to log information, when we discuss valves. There are three types of logger, which log to standard output, standard error, and a named file, respectively. The first two are rather trivial to define. The third is marginally more complicated because you have to provide attributes called prefix and suffix which determine the name of the log files. *Tomcat* starts a new log at midnight each day, and uses the date as the middle component of the file name. For example, a logger which defines a prefix of “mycolog” and a suffix of “.txt” would write to the file mycolog.2003-01-30.txt on 30 January. There's also a boolean attribute called timestamp which controls whether a timestamp will be included with each logged item.

I know that high volume sites like to rotate their log files more often than once a day. There doesn't seem to be any way to do that with *Tomcat*, though I don't doubt that with a bit of ingenuity (and a suitably named named pipe, perhaps) some kind of mechanism could be devised to chop the logs up into smaller pieces.

Realms ain't what they used to be

I had a lot of trouble getting my head around realms in *Tomcat*, until I realised that I needed to forget what I already thought they were. In ordinary English usage, a realm means a kingdom, or a domain or subdivision of

some sort (as in, for instance, “the realm of science fiction”). Within the authentication mechanisms of the HTTP protocol, a realm means a “protection space” – a set of related resources to which access is controlled via a single log in. The word is used in a similar sense in the the authentication service *Kerberos*.

It would appear that the word was adopted by the designers of *Tomcat*'s config file without regard for (or possibly in ignorance of) prior usage of the word. In any event, a realm in *Tomcat* means a database of user names, passwords, and roles (which are a bit like Linux groups) which can be used as the basis of an access control mechanism in *Tomcat*. To explain how this works, let's give a specific example.

Maintain database

Let's say a web application has been written to maintain a database of pre-booked fares for a taxi firm. A servlet accessed via the context path **getbooking** is used by the cab drivers to retrieve details of their next fare. A second servlet, accessed via the path **addbooking** is used by the dispatchers who take telephone bookings from customers. So, when the application is deployed, entries are placed in the deployment descriptor (the web.xml file) to specify security constraints which effectively say: “getbooking is only accessible to users who have the drivers role; addbooking is only accessible to users who have the dispatchers role.” Within the realm, the users Bob and Tinytim are assigned the drivers role, and the users Tracey and Jenny are assigned the dispatchers role. Tracey drives her own cab at weekends so she has the drivers role, too.

When access is attempted to a resource protected by a security constraint, *Tomcat* will initiate a standard HTTP authentication exchange to collect a name and password from the user. These are checked against the data supplied by the realm, and if that succeeds, a check is made to see if the user has the required role. If so, access is allowed. If not, *Tomcat* returns a message to say that access is denied. This mechanism is called “container managed security” because ..., well, er, because its managed by the container, rather than by code written into the application itself.

Tomcat offers several types of realm (i.e. several sources of user account information). Some of them are intended

Figure 3: Skeleton server.xml file.

```
<Server port="8005" shutdown="SHUTDOWN" debug="0">
  <Service name="Tomcat-Standalone">
    <Connector className="...HttpConnector" port="8080" />
    <Connector className="...Ajp13Connector" port="8009" />
    <Engine name="Standalone" defaultHost="localhost">
      <Logger className="...FileLogger" ... />
      <Realm className="...MemoryRealm" />
      <Host name="localhost" appBase="webapps">
        <Logger className="...FileLogger" ... />
        <Context path="/examples" docBase="examples" ... >
          ...
        </Context>
      </Host>
    </Engine>
  </Service>
  <Service name="Tomcat-Apache">
    ...
  </Service>
</Server>
```

to allow tomcat to integrate into systems that maintain user accounts in an external database. The JDBC database realm, for example, takes account data from a relational database accessed through a JDBC driver, whereas the JNDI Directory Realm uses an LDAP directory accessed through a JNDI driver. These realms are relatively complex to configure, with lots of attributes that determine how tomcat will connect to and log in to the database, and the names of the tables and columns within which the user, password, and role information is to be found.

The simplest type of realm is the Memory Based Realm, which reads user account data from a simple XML file when *Tomcat* starts. By default, the file is `$CATALINA_HOME/conf/tomcat-users.xml`, but if you want, you can define several memory-based realms (effectively, several user communities) and use the `pathname` attribute of the `<Realm>` element to identify the file. Here's an example of `tomcat-users.xml` based on our taxi scenario; the format is pretty self-explanatory.

```
<tomcat-users>
<user name="bob" password="qwertyuiop"
roles="driver" />
<user name="tinytim" password="xyzyzy" roles="driver"
/>
<user name="tracey" password="mmds"
roles="driver,dispatcher" />
<user name="jenny" password="sepulveda"
roles="dispatcher" />
</tomcat-users>
```

Valves

Finally, we come to valves. No, not those fragile little glass bottles we used to use for heating radios and audio amplifiers. These valves are software components that can be inserted into the request/response stream flowing into an element within the *Tomcat* component hierarchy. *Tomcat* provides four types of valves: access logging, single sign-on, request filter, and request dumper.

Access logging valves intercept requests into an engine, host, or context (depending on where the valve is placed) and write an entry into a named log file for each request/response. By default, the format of these entries are compatible with the "common log format" (as used by *Apache*) which means the logs can be processed with standard log file analysis tools such as *analog* or *sawmill* to summarise page hit counts, and so on. You can define custom log formats by specifying a pattern attribute for the valve, in much the same way that the `LogFormat` directive is used in *Apache*. It even uses the same format codes that *Apache* does, which makes life easier. By choosing the correct level in the hierarchy to place the valve, you control the specifics of what gets logged. For example, a `<Valve>` element within a `<Context>` container logs all traffic to that specific application.

Request filter valves implement a crude form of access control based either on the IP address or the name of the client machine. These are similar to the `allow` and `deny` directives in *apache*. For example, a valve declared within a `<Host>` container like this:

```
<Host name="info.myco.com" appBase="infodir">
```




Figure 4: Mapping a url to a resource.

```
<Valve classname="org.apache.catalina.valves.
RemoveAddrValve"
allow="192.168.0" />
...
</Host>
```

would allow access to the **info.myco.com** host only from clients on the 192.168.0 network (the local network, perhaps?) You can filter on client host name as well (for example, to reject all requests from the **sales.myco.com** domain), but you'd need a second valve.

Single sign-on valves are used to cache user credentials at the server end, so that, for example, once a user has authenticated to a specific virtual host, his credentials are available to all subsequent requests he makes to that host.

Finally, request dumper valves can be used to aid debugging of the interactions with your server. They record (via the associated logger) a rather detailed record of the HTTP requests and responses (including all the headers). As for the other valves, the level they're placed at controls the specificity of what gets logged. The output is pretty voluminous, so this isn't something you'd want to leave in place on a production server.

“Define custom log formats by specifying a pattern attribute for a valve, in much the same way as `LogFormat` in *Apache*.”

Overall, *Tomcat's* configuration is clean and well-documented. The XML file looks a bit intimidating at first but isn't too bad once you get used to it. If you prefer a fill-in-the-form approach to an edit-the-XML approach, there is a graphical configuration tool (in similar spirit to the *SWAT* tool for configuring *Samba*), but as with all graphical tools, you still need an understanding of what the elements actually mean in order to make productive use of the program. In light of the fact that *Tomcat's* main purpose is as a reference implementation, it's perhaps surprising (and commendable) that the *Tomcat* team should take the trouble to write a graphical configuration tool at all. Happy servleting! ■

Network security iptables

PART 2 In the final part of our series on *iptables*, **DAVID COULSON** looks at advanced filtering and *netfilter* extensions.

Last month we looked at the basics of *iptables* with the Linux 2.4 kernel, as well as example configurations for simple firewalls and gateways. However, for a production firewall implementation, or situations where we need to finely tune the packet filtering, we have to look at some of the more advanced capabilities of the *netfilter* system and how we can make use of those with *iptables*.

Destination NAT

On a large network with many systems running public-facing services, such as web and mail servers, giving each front-end system its own IP is often rather expensive. Each system may be running one or two services, but requires a single IP to perform that function.

Instead, another option is to use the Destination NAT, or DNAT, capabilities of *iptables* in order to redirect packets from the front-end IPs onto the servers using internal IP addressing. This permits us to run many services on a single IP address, even if these services are actually provided by more than one physical box. Each system will have an IP in the 10.0.0.0/8 range, allowing structured networks to be created, offering a DMZ and secure back end networks for databases and file serving. On the outside, all we need to do is have our gateway system ARP for the external IP addresses, then we can proceed to modify the packets using DNAT.

If we assume that our public range is 192.168.1.0/27, we can use a number of DNAT rules to rewrite the packets so that they hit our internal systems:

```
iptables -t nat -A PREROUTING -i eth0 -d 192.168.1.5 -p tcp --dport 80 -j DNAT --to 10.11.2
```

This does, of course, assume that our web server is listening on 10.11.2, using port 80, and that the eth0 device on our gateway is the interface connected to the public Internet. By expanding upon this rule, we can easily forward packets from external IP addresses:

```
iptables -t nat -A PREROUTING -i eth0 -d 192.168.1.5 -p tcp --dport 80 -j DNAT --to 10.11.2
```

```
iptables -t nat -A PREROUTING -i eth0 -d 192.168.1.5 -p tcp --dport 25 -j DNAT --to 10.11.2
```

```
iptables -t nat -A PREROUTING -i eth0 -d 192.168.1.5 -p tcp --dport 53 -j DNAT --to 10.11.8
```

“Over a large network with public-facing services, giving each front-end system its own IP is often rather expensive”

```
iptables -t nat -A PREROUTING -i eth0 -d 192.168.1.5 -p
udp --dport 53 -j DNAT -- to 10.1.1.8
```

```
iptables -t nat -A PREROUTING -i eth0 -d 192.168.1.5 -p
tcp --dport 22 -j DNAT -- to 10.2.1.3
```

If we don't specify a port following the `-to` argument to DNAT, *iptables* will rewrite the destination IP and keep the port the same. Now, rather than wasting four IPs to run services on different systems, we have one IP which handles four different services, for four different boxes. We can also quite easily modify the external IP address block of the network without having to reconfigure any of the internal systems.

This type of setup does present a number of rather interesting issues. While it will work fine with an external connection, if we want a system on the internal network to access services via their external IP address, we need to add a few extra *iptables* options. The problem occurs when a system, such as 10.2.1.8 makes a connection to 192.168.1.5:80. Our firewall will modify the destination of the packet to 10.1.1.5 and will then send a packet back to the source. However, the source address is still 10.2.1.8, so rather than routing the packet back via the gateway machine for correct reversal of the DNAT operation, 10.2.1.8 will see a packet coming back to it from 10.1.1.5, rather than 192.168.1.5, and will assume it's a bogus packet and drop it.

This can be solved by performing a SNAT on the packet when it enters our gateway, so that the source address is set to the internal IP address of our gateway. The returning packets from the server will be routed via the gateway and will end up back to our client correctly. The only issue with this is that all clients will appear to come from 10.1.1.1, rather than the unique IP address of each system on the internal network. This can present some problems, although by using a proxy server such as *Squid*, one can get the original IP address via the `HTTP_VIA` or `HTTP_FORWARDED_FOR` headers from within a CGI or PHP script. To allow systems from our 10.0.0.0/8 range to access external IPs correctly, we would use a *iptables* configuration as follows:

```
iptables -t nat -A POSTROUTING -s 10.0.0.0/8 -d
192.168.1.0/27 -j SNAT --to 10.1.1.1
```

Of course, if we have both a front-end and back-end internal network, we can DNAT the packets to an IP from either range, as long as the IP belongs to the gateway:

```
iptables -t nat -A POSTROUTING -s 10.1.0.0/16 -d
192.168.1.0/27 -j SNAT --to 10.1.1.1
```

```
iptables -t nat -A POSTROUTING -s 10.2.0.0/16 -d
192.168.1.0/27 -j SNAT --to 10.2.1.1
```

Note that we have not specified a interface, since if we have many internal interfaces,

With an internal network with this type of configuration, we also need to ensure that packets appear to come from the correct IP. This is especially important with services such as DNS, which do IP-based authorisation for notification and transfers. With a master/slave DNS configuration, we would configure our slave to use 192.168.1.5 as the master, since that is the IP which our DNS server is accessible via. However, when our system sends a notify following a zone update, if it comes from the wrong IP address and the slave will ignore it. Naturally, if DNS servers out of sync for the refresh interval, then it



Bifrost handles *iptables* configuration over HTTP, but it's not freely distributable so is unlikely to become hugely popular.

might cause problems. We can perform SNAT on specific packets, based on what TCP or UDP service it applies to:

```
iptables -t nat -A POSTROUTING -o eth0 -s 10.1.1.8 --
dport 53 -p tcp -j SNAT -- to 192.168.1.5
```

```
iptables -t nat -A POSTROUTING -o eth0 -s 10.1.1.8 --
dport 53 -p udp -j SNAT -- to 192.168.1.5
```

```
iptables -t nat -A POSTROUTING -o eth0 -s 10.0.0.0/8 -j
SNAT -- to 192.168.1.2
```

In this configuration, we have set our default SNAT address to be 192.168.1.2, however any packets heading to port 53 will have it's source address set to 192.168.1.5, in line with our incoming DNS configuration.

Mangling

As well as performing filtering and NAT, we can also 'mangle' packets. Mangling simply allows us to modify packet options, such as the TOS field or to mark the packet with a value. The 'mangle' table in *iptables* has five different chains, so we can mangle packets at almost every stage of their processing and routing. However, to actually make use of mangled packets in something other than *iptables*, we need to compile in certain options, such as the ability to route based on **fw mark** or to perform QoS matching based on the packet mark.

If we want to mark all packets which belong to ssh sessions so that we can route them out over another link, we need a mangle entry:

```
iptables -t mangle -A PREROUTING -p tcp --dport 22 -j
MARK --set-mark 1
```

This won't mark packets which head out of our firewall box, so we need to duplicate this in the OUTPUT chain:

```
iptables -t mangle -A OUTPUT -p tcp --dport 22 -j MARK
--set-mark 1
```

We can then perform routing, or QoS, based on this mark. As neither of these support packet matching based on protocol or port, we have to use marking to handle routing. This type of routing is known as 'policy routing', as it is based on something other than packet destination, and we need to use the `iproute2`

IP TABLES

packet in order to configure the correct rules and create a separate routing table for these packets:

```
ip ru add table 10 fw 1
```


```
ip ro add table 10 default via 192.168.2.1
```

Of course, this will require an appropriate SNAT rule to modify the source address appropriate to the outgoing interface. We can send packets out into the Internet with the incorrect source IP, but this may cause issues of the ISP blocks packets not belonging to its own network to avoid spoofed packets heading onto the Internet.

Graphical IPTables Configuration

The command line is always nice and welcoming for die-hard Unix users and those who really want to know what's happening behind the scenes. However, if you've just got one or two boxes and it's more important to get it all working quickly and easily, there are a number of different GUI front-ends to *Netfilter* available. Of course, one still needs an understanding of firewalling, as it's still just as easy to create a completely broken configuration where the firewall isn't of any use.

Possibly one of the most popular GUIs for *netfilter* is *knetfilter*, from the KDE project. This is a fairly simple, yet comprehensive, tool which allows all aspects of the iptables system to be configured. The name *knetfilter* is somewhat inappropriate, as it won't work with anything other than the



***Knetfilter* is a simple KDE-based iptables configuration tool for those who don't want to get their hands dirty with the command line.**

iptables modules, so if we're trying to use *ipchains* or *ipvs* with *netfilter*, it's not going to work with those. More information on *knetfilter*, along with documentation and source code, can be found for your perusal at <http://expansa.sns.it/knetfilter/>

For headless servers, having a non-X GUI interface is very useful, and more often than not, it's a web interface

NETFILTER PATCH-O-MATIC

Getting the latest firewall capabilities for the Linux kernel can prove complex, but *patch-o-matic* makes life easier.

As *netfilter* has its capabilities provided by loadable kernel modules, there are a great number of third-party kernel modules available which allow our firewall to do somewhat obscure things. Of course, among the nonsense, there are a number of very valuable contributions to the *netfilter* project. All of these patches can be downloaded from the iptables.org site, and we can happily apply each of them in turn. However, this is very tiring and boring, not to mention if a patch fails, we have to decide if we should continue with other patches, or just give up. Fortunately, the *netfilter* project has a very useful suite of utilities which makes all of this much easier.

Netfilter provides a system known as *patch-o-matic* which has a nifty system to manage patches and ensure that you don't end up with a half patched kernel that won't compile properly. *Patch-o-matic* can be downloaded either as a tarball or from *netfilter* cvs, depending upon how stable you

want your firewall to be and if you need some of the new exciting netfilter modules. Both tarball releases and cvs checkouts contain the two main patch groups: **base** – The current 'stable' netfilter module release **extra** – Development modules

Patch filter also includes the netfilter patches for the last few stable kernels, so if you're still using 2.4.18, it will automatically upgrade you to the current netfilter release in the stable kernel tree before applying lots of other things. To apply the *patch-o-matic* patches to our kernel tree, we use the aptly named **runme** command passing it our kernel source path in the **KERNEL_DIR** environmental variable:

```
KERNEL_DIR=/usr/src/linux ./runme extra
```

This will step through each patch in turn and attempt to install it. If for some reason the patch won't apply, *patch-o-matic* gives you the option of testing to see if it works if it is applied anyway.

```

david@macha:~/cvs/netfilter...ilter/patch-o-matic (pts/12)
Welcome to Rusty's Patch-o-matic!

Each patch is a new feature: many have minimal impact, some do not.
Almost every one has bugs, so I don't recommend applying them all!

Already applied: submitted/01_2.4.19
                 submitted/02_2.4.20
                 submitted/03_ip_conntrack_proto_tcp-lockfix
                 submitted/04_neumat-ude-helper
                 submitted/05_REJECT-fuspotting-phrack60-fix
                 submitted/06_ftp_conntrack-msg-fix
                 submitted/07_ECN-tcpchecksum-littleendian-fix
                 submitted/08_ftp_conntrack-debugftp
                 submitted/ipt_ULOG-mac-len-fix
                 submitted/ipt_multisport-invfix
                 pending/06_early_drop-backwards
                 pending/08_mangle_input_norouteneharden
                 base/HL-IP6G
                 base/IPV4OPTSSTRIP
                 base/METLINK
                 base/METMAP

Testing...

```

The *patch-o-matic* system is a must for those who need some of the more advanced *iptables* capabilities.

Frequently *patch-o-matic* will complain about missing files, but the patch will actually apply perfectly fine. For a 'base' patch, this isn't much of a problem, but with an 'extra' patch, it gets very frustrating pressing 'T' and 'Y' every few patches. So that *patch-o-matic* is useful to people who don't have time to watch what it's doing, there is a **--batch** option which will test patches which it has problems with. It will stick as

about patches which cause rejects, but these are rare, and generally depend upon the state of your kernel tree, rather than the *netfilter* patches themselves.

Patch-o-matic contains a wide variety of patches, so here's a quick rundown of the best:

fuzzy

This option adds **CONFIG_IP_NF_MATCH_FUZZY**, which allows you to match

which proves popular. *Bifrost* fits this need, by having a very simple web interface which runs via *Apache's* `mod_cgi` system. Unfortunately, it is not available under a free license, so if we want full control of your firewall, cash has to be duely coughed up. It's open to the individual as to whether *Bifrost* is worth the cash, but it does what it does well. *Bifrost* can be downloaded from <http://bifrost.heimdalls.com/>

“Administrators can save a great deal of time by using *Firewall Builder*, as opposed to handling each system individually”

However, the LXF vote goes to *Firewall Builder*, which is a cross-platform configuration builder for firewalls, which supports iptables to Cisco PIX, with a whole selection of things in between. It is a vice nicely implemented system, which allows multiple firewalls to be configured from the same front-end. While it may be a little more complex to install than other systems, such as *knetfilter*, for a large network where nearly all machines are firewalled, administrators can save a great deal of time by using *Firewall Builder*, as opposed to handling each system individually. Plus, as an added advantage for large networks,

Firewall Builder integrates with `ucd-snmp`, so that monitors and logging software can be developed to watch and tune the firewall in real-time. *Firewall Builder* can be downloaded from www.fwbuilder.org/ As it can handle a large number of other platforms as well as *iptables*, it's certainly worth trying to see if it fits your needs.

Going further with netfilter

In this series we've only looked at the iptables modules which can be used with netfilter. As *netfilter* is a completely modular system, there are a wide variety of netfilter modules to provide varying network capabilities to a Linux system. While most people will know of the `ipfwadm` and `ipchains` modules for backwards compatibility with 2.0 and 2.2 kernel firewall systems, the popular Linux Virtual Server has been ported to 2.4 using the netfilter subsystem. This makes it somewhat more attractive, as it does not require any patches to be applied to the kernel. As always, the more patches which a kernel is using, the longer it takes to get the kernel up and running after a new kernel release is made.

As always, all the latest information on *netfilter* and *iptables* can be found at <http://netfilter.org/>. There are also some excellent mailing lists there covering *netfilter* and *iptables* configuration, so anyone wanting to seriously use *iptables* should subscribe. ■

packets according to a dynamic profile implemented by means of a simple Fuzzy Logic Controller (FLC) .

Supported options are:

-- **upper-limit** => Desired upper bound for traffic rate
-- **lower-limit** => Lower bound over which the FLC starts to limit traffic

iplimit

This adds `CONFIG_IP_NF_MATCH_IPLIMIT` match allows you to restrict the number of parallel TCP connections to a server per client IP address (or address block).

Examples:

```
# allow 2 telnet connections per client host
iptables -p tcp -- syn -- dport 23 -m iplimit --
iplimit-above 2 -j REJECT
# you can also match the other way around:
iptables -p tcp -- syn -- dport 23 -m iplimit ! --
iplimit-above 2 -j ACCEPT
# limit the nr of parallel http requests to 16 per
class C sized
# network (24 bit netmask)
iptables -p tcp -- syn -- dport 80 -m iplimit --
iplimit-above 16 \
-- iplimit-mask 24 -j REJECT
```

NETMAP

This adds `CONFIG_IP_NF_TARGET_NETMAP` option, which provides a target for the nat

table. It creates a static 1:1 mapping of the network address, while keeping host addresses intact. It can be applied to the `PREROUTING` chain to alter the destination of incoming connections, to the `POSTROUTING` chain to alter the source of outgoing connections, or both (with separate rules).

Examples:

```
1.2.3.0/24 -j NETMAP -- to 5.6.7.0/24 -j
NETMAP -- to 5.6.7.0/24
iptables -t nat -A POSTROUTING -s 5.6.7.0/24
-j NETMAP -- to 1.2.3.0/24
```

mport

This module is an enhanced multiport match. It has support for byte ranges as well as for single ports.

Examples:

```
# iptables -A FORWARD -p tcp -m mport -- ports
23:42,65
```

Up to 15 ports are allowed. Note that a portrange uses up 2 port values.

quota

This option adds `CONFIG_IP_NF_MATCH_QUOTA`, which implements network quotas by decrementing a byte counter with each packet. Supported options are:

- **quota**
The quota in bytes.

ROUTE (Development/Works for me)

This option adds a `ROUTE` target, which enables you to setup unusual routes not supported by the standard kernel routing table.

For example, the `ROUTE` lets you directly route a received packet through an interface or towards a host, even if the regular destination of the packet is the router itself. The `ROUTE` target is also able to change the incoming interface of a packet. This target does never modify the packet and is a final target. It has to be used inside the mangle table.

ROUTE target options:

```
--oif ifname
Send the packet out using the ifname
network interface.
--iif ifname
Change the packet's incoming interface to
ifname.
--gw ip
Route the packet via this gateway.
```

SAME

This adds `CONFIG_IP_NF_TARGET_SAME` option, which is similar to `SNAT`: it takes a range of addresses (`-- to 1.2.3.4-1.2.3.7`) and gives a client the same address for each connection. It has a `--nodst` option to make it not use the destination-ip in the calculations when selecting the new source-ip.