LXF34 Linux Format December 2002 ● RED HAT 8.0 ● SUSE 8.1 ● MINI DISTRO ROUNDUP ● PHP ACCELERATORS ● MAN PAGES EXPLAINED ● DOXYGEN ● LINUX EXPO LONDON ● NETRAVERSE ●

LINUX FORMAT

www.linuxformat.co.uk

The ultimate scanner set up guide and expert tips! p62

SCANNING

DISTRO WARS!

Brand new releases, SuSE 8.1 and Red Hat 8.0 reviewed inside

PLUS: Eight specialist mini distros compared (and on your DVD!)

AND: SmoothWall and Lycoris on test!

Man Pages Demystified!

Find out how to use, edit and create these most useful resources p68

Going faster with PHP

Accelerator engines compared – could your site be going faster? p54

LINUX EXPO 2002!

All the releases, debates, rumours and gossip from the UK's biggest Linux event p12

//Any Comments?

Use the latest documentation tools and your work will be cleaner, smarter and much easier to maintain p82

CD issue also available Printed in the UK
LXF34 DECEMBER 2002 £6.49
the future network Media with passion
9 771473 968012 12
(£6.49 outside UK and Republic of Ireland)

THE UK'S BEST-SELLING LINUX MAGAZINE

Interesting times...

While the annual Linux Expo in London is a very good event in its own right, with plenty of enthusiasts, new products, exciting demos, bizarre mascots and the occasional free beer, it is also a rather important barometer of industry and user feeling towards Linux as a platform.

A few years back, there were ostentatious stands and large crowds. But the crowds were mostly of people who had only a peripheral interest in Linux, and many of the companies behind the ostentatious stands are no longer with us. This year, the crowds (though larger than last year perhaps) wouldn't exactly constitute a throng, but they were people who for the most part were serious about Linux. Serious developers, community members, vendors – they may have had different opinions on which direction Linux was or should be taking, there may have been concerns over training, education, support and other things – but overall these people were friendly and enthusiastic. And serious about Linux. You can find out more about the events and non-events of the show in our full report starting on page 12.

Of course, some of the exhibitors were there touting the latest versions of the distributions. The almost simultaneous release of several major updates prompted us to give a distro theme to this issue. While SuSE and Red Hat battle for the desktop and Lycoris caters for the migrators, there are also many specialist distros for specific tasks, so we have a special roundup this issue featuring some of these useful tools (and indeed, you can find them on your coverdiscs this issue).

If you're happy with your current distribution, then perhaps you'll be more interested in all the other stuff we have for you. There's a complete guide to PHP accelerators, including some thought provoking benchmarks, a look at the man page system and of course, tutorials galore. We hope you enjoy it!

Nick

Nick Veitch EDITOR

The Red Hat 8 controversy, SuSE gets even better, Lycoris updates and a hatful of mini distros p22,36,42

Everyone loves a show! Get the lowdown on what went on p12

PHP is good, but it's even better when it's accelerated p54

12

54

22, 36, 42

LINUX FORMAT

Aims of the magazine

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

» To provide the most accurate, unbiased and up to date information on all things Linux.

» To promote the use of Linux in business and the home, for servers and on the desktop.

» To support the Open Source community by providing a resource of information, and a forum for debate.

» To help all readers get more from their Linux experience by providing insightful and useful tutorials.

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

Richard Smedley
Rich's favourite thing at the show was an HP 'lightsabre' keyring. Hmmm bright lights, sparkly things...

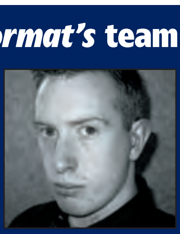

David Coulson
Our Answers guy is a networking and security guru with plenty of sysadmin experience.

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

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

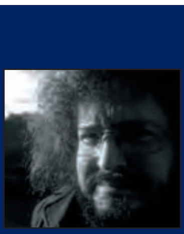

Charlie Stross
Master of Perl, Charlie has been writing about Linux for more years than probably even he can remember.

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

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

Richard Drummond
With more time on his hands, Rich is no doubt developing a Java app to write his reviews...

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

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

Contact us

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

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

Contents

LINUX FORMAT
LXF34 December 2002

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

REVIEWS

TUTORIALS

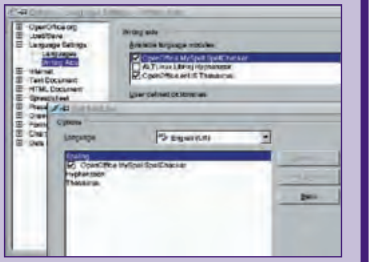

REGULARS

FEATURES

Coverdiscs

DemoLinux Linux without installing; **Unreal Tournament 2K3** demo of the new UT (see p101 for system requirements); **RUTE** teach yourself sysadmin; **OOoDictionaries** and installer; **Note Editor** notation editor; *On the DVD:* **KDE3.1** beta2; **SmoothWall2** latest GPL edition; **GCompris** games and education for young Linux users; **Scribus** rapidly-improving DTP; **GNUPrivacyGuard** protect your emails; **Movix** bootable movie discs

Please read the coverdisc instructions starting on page 100 before installing from the coverdiscs!

On the CD and DVD

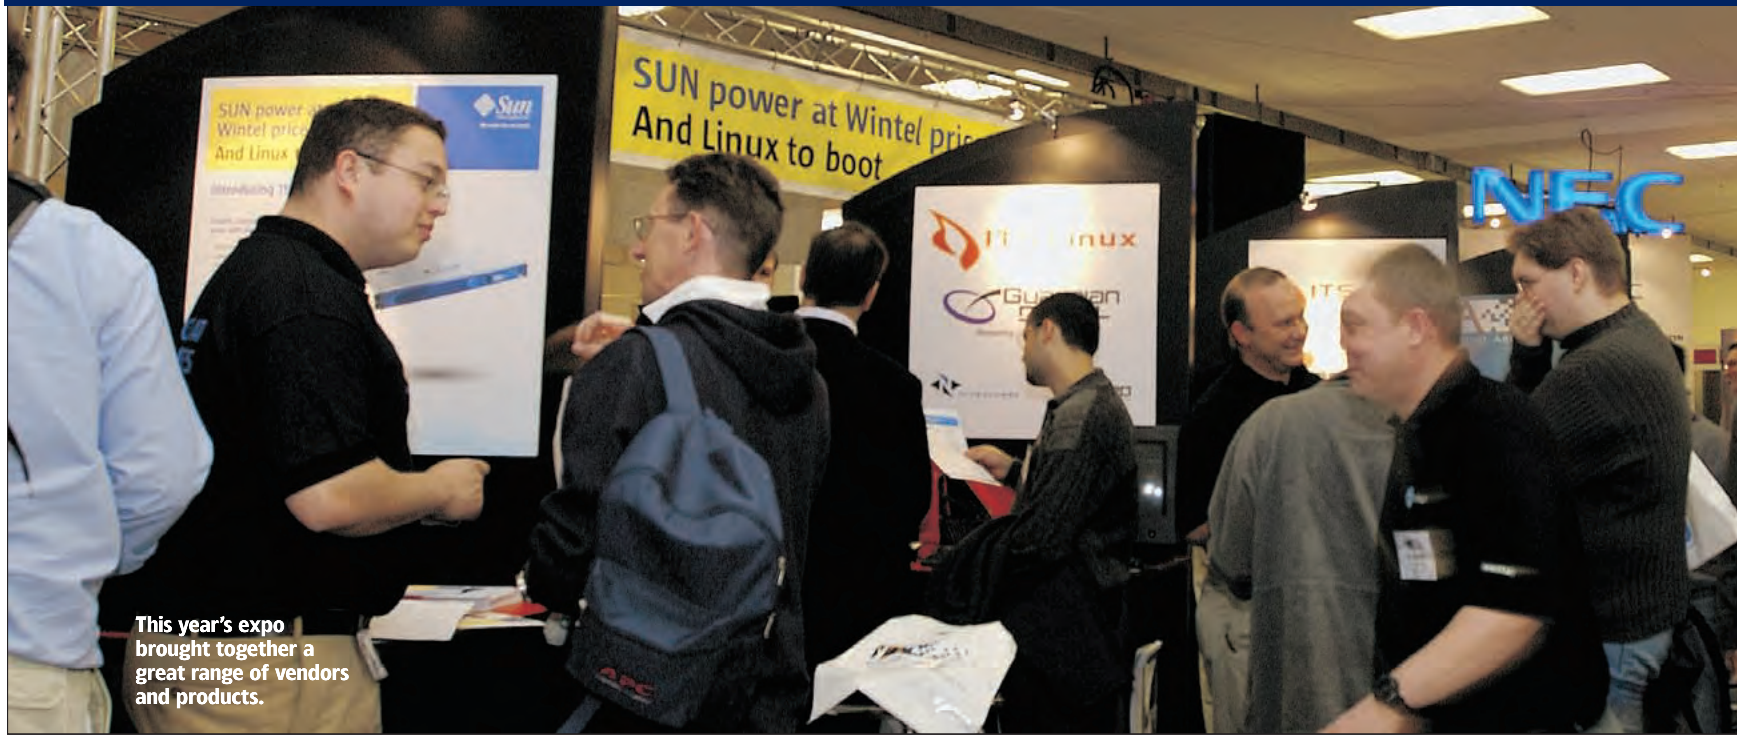

This year's expo brought together a great range of vendors and products.

LINUX EXPO UK 2002

SPECIAL REPORT Braving the dangers of the big city, the LXF crew attended the expo in force to bring you all the news, gossip and rumours.

NEC had everything from laptops to large servers for the enterprise.

The Linux Expo at Olympia is an important event for *Linux Format*. Not only is it the place that many of the exhibitors choose to reveal new products or pricing structures, and the coming together of some of the most knowledgeable Linux experts from the UK and further afield, but it's one of the few events we get to meet our readers in great quantities. Final attendance figures for the show haven't been calculated yet, but it certainly seemed a lot busier than last year's event judging by the crowds of people we had to fight our way through to get to the bar (*surely you mean the* LXF *stand – Ed*).

It was great to meet you all. Thanks for stopping by the stand and giving us your perspective on the magazine and the world of Linux in general. We picked up a lot of good ideas, and were quite gratified that nobody seemed particularly disgruntled with the magazine. Sorry if some of you stopped by the stand and didn't get to speak to one of the team, but please feel free to email us any time!

There were other people there apart from *Linux Format* though. In fact, maybe it was just our imagination, but there seemed to be a greater number and more diverse range of exhibitors than ever. If you wanted to buy cheap Linux distributions or books, attend seminars on all sorts of topics, or learn about the latest clustering/ backup/code development technologies, you were well catered for.

It's always interesting to see which distro vendors turn up to the show, and this year we were a bit light on big names. Red Hat declined to appear again, though their latest 8.0 distro appeared on many stands. Mandrake didn't have a presence. SuSE were there though, with the latest 8.1 release, and seemed to attract much interest at the IBM stand. The newly renamed SCO where there, with a flashy looking stand that had

Downstairs, Web Solutions 2000 turned up some interesting Linux servers and admin tools.

Corporate branding made fun as the Positive Internet Company trumpet their Linux success.

During the "battle of the mascots" the ISPs' representatives remained on good terms. Shame.

"I'll have one of those, please." Everyone is impressed to see the new Cyclades terminal server sprout a PCMCIA card for wireless potential.

Macs run *nix, too! Apple show off their entry into the world of grown-up OSs.

apparently been shipped all the way from the US for the event. They were keen to promote United Linux, and certainly interested most visitors enough to take away the Beta3 CDs they were handing out.

Perhaps not best known for their support of Linux, Apple were there under the guise of Toucan, trying to lure the faithful with some cool looking hardware running Mac OS 10.2, "Jaguar". The easy to use interface certainly seemed to be impressing some, though of course questions were asked about the legitimacy of charging so much for what is largely Free Software (*e.g.* the standard compiler is *GCC*). Still, the 'Unix with a nice interface' story was a compelling reason to have a play at least.

Gordano, creators of *GLMail*, which is now known under the title *GMS Mail*, where attracting a lot of interested visitors to their booth. Armed with a copy of *LXF33* where their software scored 9 out of 10 in our massive roundup, they were talking through the new additions to the suite, which now includes an SMS messaging service as well as the usual Antivirus, antispam, webmail interface and other goodies we've come to expect. Some more

»

IBM dispensed Linux knowledge and enthusiasm to many a visitor.

Assorted Linux personalities debate the issues of the day, while editor Veitch grabs a quick nap.

« modules are due out soon including a mailing list option and some other stuff that's still secret! We will be featuring a full review soon, but for the moment, you can take our word for it that this is the best proprietary mailserver about.

Cyclades, makers of the remote terminal servers we have reviewed recently, had a large stand and a big message to get across. When we featured the terminal server, some readers seemed to be at a loss as to what it could be used for. The answer is out of band management. Connecting up the serial ports of servers, routers and other devices, then connecting a modem or other device to the box gives you a secure way of monitoring and accessing servers without using always-open Internet ports. Remote management tools are useful, but if your 'Net connection goes down, you're in for a long trip to the data centre. No doubt that's why Cyclades has continued to grow rapidly in the past three years, and accounts for the interested group of admins clustered around their stand!

While others may go in for grand visions and lots of talk, HP's approach to Linux is more nuts and bolts. One of the more interesting recent developments is their series of Itanium 2 servers. The original Intel 64-bit server processor received a mixed reaction, but its successor is much faster, smaller and, in some ways, simpler to deploy. Itanium 2 servers are already shipping in reasonable numbers from HP and, of course, the choice of 64-bit OS to run on them is almost always Linux.

Backing up your system isn't sexy, but it is important. And it has got a bit sexier, thanks to Arkeia. At their stand visitors were treated to the visual feast that is Arkeia v5. Earlier versions proved to be solid and reliable, but that wasn't always the impression people got when they saw the GUI interface. This has been completely overhauled, and we are told the application code has been around 80% rewritten for this version. But looks aren't everything. One of the reasons for the rewrite is that Arkeia now supports a plug-in architecture to help deliver an easier to use backup experience. The first notable plug-in is for *MySQL*. With a few simple clicks, you can configure the Arkeia system to make backups of all your database tables and files, without having to remember where they all live. Look forward to a review really soon.

Downstairs, the Web Solutions 2000 event was taking place simultaneously, and it was good to see there was a lot of relevant crossover here. Visitors who managed to negotiate around the Microsoft .Net stand could find many big name

Different perspectives

Andrew Channelle ambles round the stands

The best looking device award (if they'd had one) at this year's Linux Expo would have gone to the Eye-box Max from Rightvision; an Internet appliance which looks like a very well-fed Aibo. The little fella is, in fact, a Red Hat-based network server and integrated hub capable of handling the Internet, email, file/printer sharing and security needs of up to 25 users. As well as a small LCD face, the device has an innovative – but perhaps a little too twee – web-based management interface based on the concept of a physical desktop; click on the desk calender to access the calender etc. There was a bigger 150 user unit, but it wasn't as cute.

Ever since it was announced, the Linux for Playstation2 package has been an object of desire for me, but having had my first real encounter with the device at the show, I have to say I'm glad the temptation to shell out for one wasn't acted upon. The whole kit looks lovely, and the gadget-points would have been considerable, but this is really a developer unit, not for serious use. A lack of expandability and readily available software (it all needs to be build-specific for the PS2) make it cool, but not for everyday use.Still the tiny Sony stand was busy throughout the day and the demonstration of the machines graphical prowess was quite impressive.

While marvelling at the PS2's puppet demo, the sounds of a competently-played guitar drifted across the hall from the .org section. The source was the *Rosegarden* stand where Richard Bown and his fellow developers were showing off version 4 of their music software. The previous version (2.1) was interesting but lacked too many features to challenge the likes of *Cubase* or *Cakewalk*. Version 4, however, looks like it's in an entirely different league. Let's hope the development team can secure some funding to allow them to achieve *Rosegarden's* potential.

Hoping to confuse hapless punters, the members of team Debian were selling free software and giving away free beer. Little wonder this was one of the more popular stands. As the first day wore on, the volunteers running the stand were becoming more adept at 'reading' the punters. "It's pretty liberating," one told me, "to know within seconds whether someone is ready for Debian, or if we should send them off to buy Mandrake."

Coding on the Playstation2 is even more fun than playing games.

In between Debian and GNOME, the KDE stand was empty for much of the day and became a sort of rest station for Debianites trying to get a break from the crush. Perhaps they thought that the fact that the vast majority of hardware at the show was running KDE was advertising enough. It would have been nice to see them though...

The Association For Free Software, despite having the plainest stand on Earth, played host to a steady stream of

Debian's Phil Hands spoke up for Free Software and free beer.

hosting companies. One of the newer faces there was Jool, a company with some shiny boxes and some cunning ideas.

Jool launched their 'Netserv' range of servers at the show, which are aimed at everyone from SoHo users to Enterprise customers. Apart from some nice looking hardware, the really cool thing about Jool was their software. The servers run a customised version of Linux which has been stripped down and optimised for speed and security. Jools's MD and founder Anjula Perera told us that this delivered a significant performance increase over a standard Red Hat installation.

Then there's *Radmin*. This fantastic looking remote management tool is designed to make administering Jool servers as easy as possible. As well as the standard features you might expect, there are loads of automated tools for common jobs. There are even special admin tools for setting up game servers!

Sun also had their new LX50 servers at the show, which looked even better in the flesh and we saw many a drooling passer-by. Sun have taken some criticism for their lateness to the Linux party, and the less than polished version of Linux they are shipping with the server (solid and reliable, but a bit unfinished looking). However, they certainly talked a good talk and many came away with changed attitudes.

Many wondered whether the various web hosting companies might stage a 'Battle of the mascots' sideshow. It seemed more than likely that The Positive Internet Company giraffe would get into some sort of turf war with Rackspace's impressively proportioned representative. There was much debate as to which had the most padding, but our efforts to run a book on who might be the victor of such a titanic struggle came to nought, as they both (wisely) determined to remain on friendly terms. The show was, they agreed, big enough for the three of them.

Of course the show wasn't just about exhibitors. There were a few theatres running, and a lot of interesting presentations ranging from "What is Free Software" to more complex topics like testing code and automating cluster management. Once again, the 'Great Linux Debate' proved a popular draw, with representatives from SuSE, SCO, IBM, HP, Sun, Debian and, ahem, *Linux Format* fielding a variety of questions on the current state of Linux. Initially many of the questions were on familiar themes – "why are there few training courses" and "is Linux ready for Enterprise," which certainly gave some the impression that not a lot of progress has been made. There was a fair amount of posturing by the big companies represented, and some (we hope) lighthearted jokes at the expense of Sun. Sun's Simon Tindall, gamely accepted a special award for their contribution to the Linux market so far. The final debate ended on a positive note though, with the panel and audience in agreement that Linux had made significant progress since last year, and they believed it would continue to do so. If the main protagonists failed to agree on specific strong areas for Linux, at least the general consensus was that it was moving forward. In a way, we were left thinking that it didn't really matter whether Sun and HP could agree on the scalability question, or whether SCO thought there was a future in a desktop market. The main strength of Linux is that it's a community, not controlled by a single corporate entity, so ultimately it will go where the people who use it want it to. That's you by the way.

By the end of the show, we at least were convinced that there were more attendees this year, and that the general knowledge about and use of Linux had become much more mainstream than before. An interesting trend – but we'll all have to wait and see what happens in the year ahead. LXF

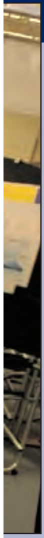

punters hoping to discuss the difference between free and really free, and find out what could be done about the seemingly inevitable imposition of software patents in Europe. Tip for next year though: get some chairs. It's more conducive to decent conversation, especially if you've spent the best part of the day stood outside the Enterprise Linux Theatre straining ears to hear words of wisdom from Sun and IBM.

XIOS, provider of support, training and system integration, used the old 'Win a Zaurus!' ploy to draw in passers-by. And very successful it was too, though the three men running the stand must have got tired of being asked "what exactly do you do?" Still, it was probably water off a duck's back for a company which aims to 'bridge the gap between the technical world of Linux and management understanding.'

Situated on the route from the main show floor to the 'Deli', GenaWare had a pretty good shot at catching the eye of hungry attendees. They had a very attractive display too, though it's understandable that a lot of the passing trade thought it was something to do with the genetic/pharmaceutical industry, and not a Geographic Information Systems (GIS) provider. The GenaWare product is surprisingly complete, and their presence at the show was a pretty good demonstration of Linux making inroads into almost every facet of computer use. And like many others, GenaWare cited user demand as one of the main reasons for their decision to go beyond simply supporting Windows.

I ended my day investigating the unglamorous world of firewall appliances on the Rockstone stand. Like its name, Rockstone's presence at the show radiated Roy Keane style hardness – from the enormous chains tethering their imposing hardware to the firmament to the names of their products (*Fortress*, *Imperial*, etc.). But despite the talk of rackmount cases, multiple Ethernet interfaces and "legendary Linux stability" my eye kept being drawn back to the dog-like lines of the Eye-box thing just across the walkway.

Newsdesk

New distro releases from Red Hat, SuSE and Mandrake; SuSE become SAP partner; Bundestag funding for KDE PIM; Inland revenue considering StarOffice; China's CPU; Apache worm goes wild; Zend speed up.

GRAND DISTRO ROLLOUT

The next round of distro

A glut of new products has been cluttering the release calendar of late, bringing the latest versions of KDE, GNOME and not a little controversy.

Red Hat and Blue

America's biggest Linux seller finally launched Red Hat 8.0, which heralds the company's long-predicted push for the desktop market. The release inspired plaudits, complaints and even the resignation of one KDE hacker who refused to be associated with a company that was actively turning the KDE desktop into 'crippleware'. The cause of KDE developer Bernhard 'Bero' Rosenkraenzer's ire was Red Hat's attempt to unify – or nullify – the two competing desktop environments to create a more 'convenient, organised and user-friendly desktop' under its own Bluecurve brand.

While still giving users the opportunity to opt for KDE or GNOME as their desktop of choice, Red Hat engineers have tried to make both systems look the same, and operate in a similar fashion. So far so sensible. But in the process of creating Bluecurve, Red Hat have removed many of the distinct characteristics of both KDE and GNOME, as well as most of the visual branding, and as the company has favoured GNOME in the past, the perception is that KDE has been altered to fit the GNOME look and feel. Rosenkraenzer said his decision to leave Red Hat was sort of mutual. "I don't want to work on crippling KDE, and they don't want an employee who admits RH 8.0's KDE is crippleware," he wrote on the KDE developers' mailing list.

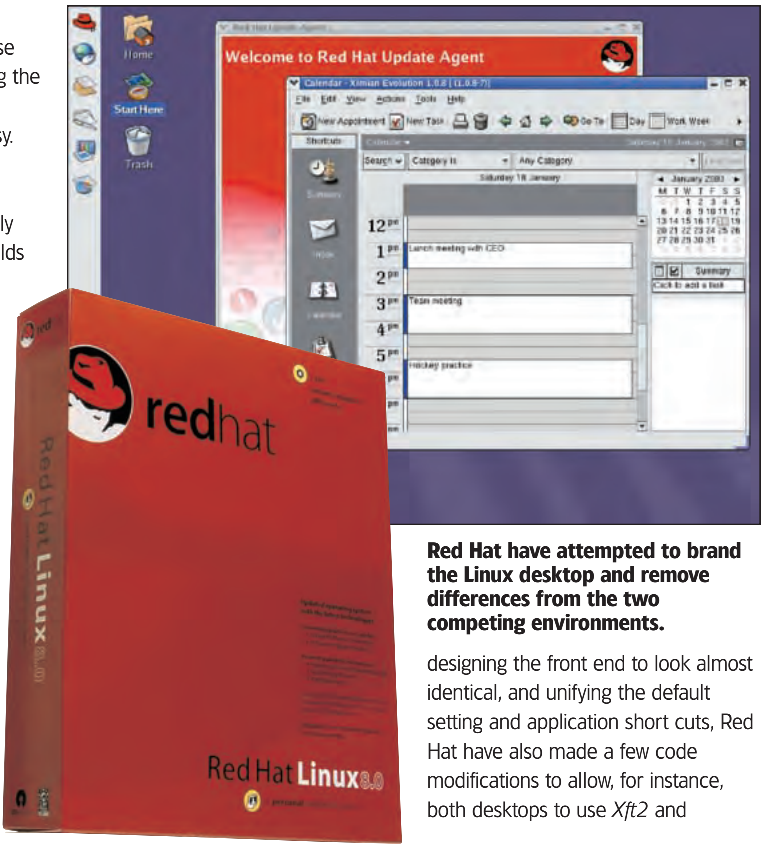

Red Hat have attempted to brand the Linux desktop and remove differences from the two competing environments.

User benefits

Red Hat said the aim behind Bluecurve was to improve the user experience and also avoid unnecessarily duplicating the development effort. As well as designing the front end to look almost identical, and unifying the default setting and application short cuts, Red Hat have also made a few code modifications to allow, for instance, both desktops to use *Xft2* and *fontconfig* for font rendering. These decisions, Red Hat said, made sound business sense. "Creating two sets of configuration tools, two websites, and two boxes isn't feasible or desirable. So we have to make the desktop fit in with the rest of the product instead of making the rest of the product fit in with the desktop." But the key goal is making Linux (and of course Red Hat) more enticing to the casual user.

Erik Troan, Red Hat's Director of Product Marketing, said the new distribution offered functionality to hardcore and professional Linux users, but also offered a more friendly experience for mainstream users. "With this release we have given more attention than ever before to usability," he said.

Mandrake's great leap

Mandrake – just – managed to steal the march on Red Hat with the announcement of the fully LSB compliant Mandrake Linux 9. Hoping to build on its reputation for cutting edge but stable releases, Mandrake 9

UnitedLinux beta

"Global enterprise OS"

The UnitedLinux project got a new general manager – Paula Hunter – and also put out an initial Beta release. In one of her first outings as the public face of UnitedLinux, Hunter stressed the importance of broadening membership of the group on a global basis. "Our open beta release demonstrates our collective commitment to UnitedLinux as the global enterprise operating system of the future," she said.

"Because it's being developed to unify rather than to fragment Linux offerings, UnitedLinux will help vendors support a single high-value Linux offering rather than many different versions."

UNITEDLINUX

ENTERPRISE PARTNERSHIP

SuSE bring SAP certification to Linux

SuSE are the first Linux vendor to become a SAP Global Technology Partner. SuSE CTO Boris Nalbach said the partnership was a significant milestone in the advancement of Linux in the enterprise sector. “SuSE is the first Linux enterprise OS to enable SAP clients worldwide to run SAP’s leading e-business applications on Linux,” he said. The partnership enables companies to use mySAP.com on SuSE Linux Enterprise Server.

releases

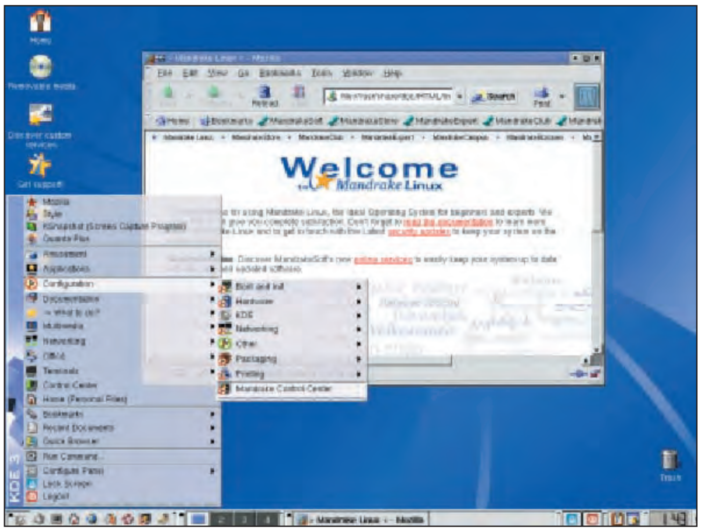

Mandrake continue to offer a range of cutting edge packages.

features Linux Kernel 2.4.19, KDE 3.0.3, GNOME 2.0.1, *OpenOffice 1.0.1*, *Mozilla 1.1* and *GCC 3.2*. One of the most attention-grabbing innovations in the distribution is its dynamism in both software and hardware recognition. For instance, when software is installed or removed the application menus are updated dynamically updated to reflect the changes (so far, so Debian) – and this method is carried over to hardware, so if you plug in a USB camera or storage device, its icon should magically appear on the desktop. Mandrake 9 also supports USB 2.0 out of the box and, it’s claimed, is also the first mainstream distribution to natively support NTFS partitions and the WebDAV protocol. New security features include the ability to send and receive encrypted email seamlessly from *Mozilla*, *KMail* or *Evolution*.

As ever, Mandrake 9 is available in a range of packages including a download version.

SuSE-friendly?

Either by choice or circumstance, SuSE’s latest distro release was decidedly low-key, which is odd as SuSE are billing it as a tenth anniversary commemorative launch.

As ‘desktop season’ is in full swing, SuSE are also trumping ease of use advances including hotpluggable USB 2 and firewire; a new ‘comfortable’ CD Burner; a new installation guide aimed at home users and a preinstalled edition of *Wine* and *fakeWindows*. *YaST 2*, SuSE’s proprietary configuration tool, also has a new addition intended to smooth the installation of joysticks and other ‘home specific’ hardware.

NEWSBYTES

- After a brief legal skirmish users can purchase an **Apple iPod** secure in the knowledge that their new toy will play happily with their Linux box. See **www.blinkenlights.ch/cgi-bin/fm.pl?get=ipode** for more information on the *gnuPod*.

- Get in before the **MPAA**. *DVDSynth* is a small app that can be used for developing ‘virtual’ DVD/CD drives and also virtual content for them. More interestingly, *DVDSynth* can apply filters to pre-existing discs. These could, for instance, add subtitles, and strip out region coding and annoyances such as blocked fast-forwarding. It’s available for Windows only at the moment but the source code is at **www.roundelay.net/dvdsynth/prerelease.html**.

- After reports last month of users running SuSE on an Xbox, **Microsoft** have apparently modified the hardware so that the current generation of mod chips won’t work. An MS spokesperson said small changes had been made to the consoles internals to increase security and reduce costs. Meanwhile the original project have managed to get Windows 2000 running within an emulator on their system.

- **IBM** are to supply two enormous retailers with Linux-based cash registers. The deals, with Regal Entertainment Group and Brazil’s Casas Bahia covers 5,000 devices.

- A buffer overflow exploit has been discovered in PDF viewers *gv*, *kghostview* and *ggv* which could enable an attacker to run arbitrary code on your machine. Disclosing the vulnerability, security company **iDefense** said that, though any malicious code would only be able to take advantage of the users file permissions, it wasn’t unusual for less-savvy users to check email when logged in as root. The other *caveat* is that the bug is not exploitable if the ‘infected’ doc is opened in the traditional manner. It has to be opened from the cli.

- Eminent self publicists **Lindows** managed to grab headlines again recently by announcing a new licensing deal with Netscape/AOL which would effectively turn a Lindows-based PC into an ‘AOL Computer.’ The deal turned out to be just a standard license to bundle the *Netscape* browser with version 2 of the OS. So one of the things you can’t do with this new ‘ AOL Computer’ is sign up with AOL as your ISP as they still don’t support Linux.

David Cartwright

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

COMMENT

Freeing Windows

Although I spend as much time as I can persuading people that Linux is most definitely an alternative to Windows, I do live in the real world and so I have to do the MS thing quite a lot, not least because my main client right now is very heavily into Bill’s products. I was asked recently to look at the .NET development environment and report on how it’s different from the old way of developing Windows applications. It was an interesting little voyage of discovery, not least because of finding the fact that you no longer need to go out and buy yourself a compiler to write Windows software.

Sounds like someone is taking a leaf out of the *nix world on this one. Linux has, of course, always had *GCC* in there somewhere to let you write software without spending loads of cash on compilers and IDEs. Even the likes of Sun have made free development possible – although their compilers have traditionally been chargeable items, you’ve generally been able to get *GCC* for nowt (and all the “include” files and key libraries are on the installer CD, you just need to tick the ‘developer’ box when you install).

So maybe the Windows world has gone full-circle (remember the old days, when your PC used to ship with a BASIC interpreter?), and ended up right back where Linux has been all the time. The downside, of course, is that now that you no longer have to shell out several hundred quid for your copy of *Visual Studio*, people will develop more on Windows. But it’s nice to see that the world’s biggest software company has decided to start doing things the Linux way.

LIBRE SOFTWARE UBER ALLES

German government commissions KDE

The KDE development team have been commissioned by the German government to create a single groupware product from the various PIM applications that ship with the desktop. The task, from Germany's Federal Agency for IT Security, is to create a "heterogenous environment and provide email, contacts, appointments and tasks lists." On the server-side, *Kroupware* will based around *OpenLDAP*, *Cyrus IMAP* and *Postfix*, while the client-side will be tackled by extending *KMail* and *KOrganizer*.

Matthias Kalle Dalheimer said the project would stick to the standard open development method, but that the deadline "forbids us from following the normal KDE development schedule." Instead of bringing potential instabilities into the main *KMail* and *KOrganizer* at a crucial stage in the KDE3.1 development cycle, Kalle Dalheimer said the project would create a **kroupware_branch** in the modules *kdelibs*, *kdebase*, *kdenetwork* and *kdepim*. "We aim to retrofit the changes on this branch as timely as possible back into the **HEAD** branch; this will happen in close collaboration with the maintainers of the affected projects," he said.

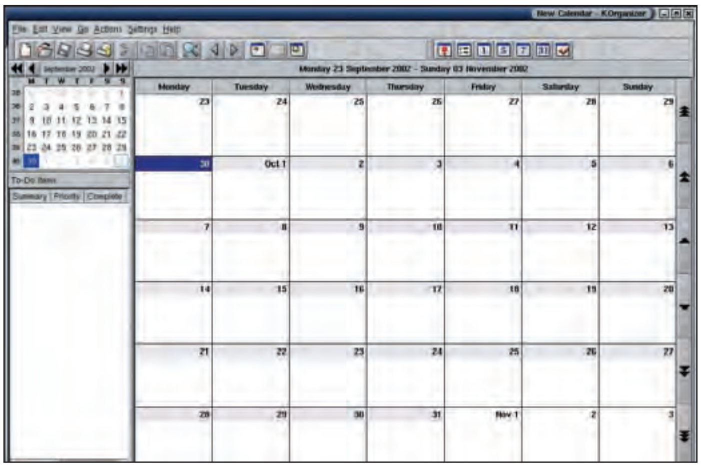
New Calendar - KOrganizer
Monday 23 September 2002 - Sunday 03 November 2002
Monday
Tuesday
Wednesday
Thursday
Friday
Saturday
Sunday
To-Do Items

The German government want KDE to organise your life.

Linux Web Watch/

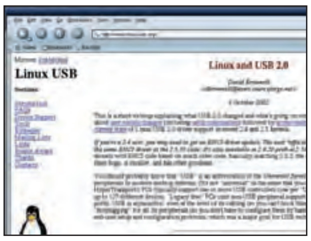

USB 2.0 – now in the kernel.

1394 – watch your chipsets!

Bluez – for all your PAN info.

O'Reilly – homebrew recipes.

Just connect...

RS232 and Ethernet are not the only fruit.

A PC that isn't connected to something else is essentially a bulky, useless metal thing. Fortunately there are loads of different things you can plug into your computer to make it more useful. And as devices proliferate, so too do the different ways you can actually hook them up.

USB has gained a lot of ground when it comes to connecting mice, joysticks, scanners and cameras and it makes the thought of using a serial or parallel port almost painful. But it's not good enough for really data intensive jobs such as mass storage or full-screen video capture. Luckily the USB 2.0 standard is faster than its little sibling by a factor of 40, and since 2.5.39 (or 2.4.20-pre8-ac2) it has been part of the main kernel as part of the *EHCI* driver. Go to www.linux-usb.org/usb2.0 for more information.

IEEE1394 is better known as Firewire (if you're a Mac fan) or iLink (if you like Sony-speak), and has become a standard for connected digital video (DV) cameras. Support under Linux is patchy, but if you're looking for 1394 supported hardware you'll need to look for chipsets from Texas Instruments (PCILynx/PCILynx2) or those that are *OHCI* compliant. Sony's own chipset won't work. Apple, by the way, have committed to open sourcing their Rendezvous network configuration system (http://developer.apple.com/macosx/). www.linux1394.org

Bluetooth has had a slow start but is now making its presence felt. It has been part of the Linux kernel since 2.4 and the official Bluetooth stack, originally developed by Qualcomm, is housed at http://bluez.sourceforge.net/ Bluetooth technology may bring about the advent of PANs (personal area networks) whereby your phone can talk to your watch and your watch can talk to your fillings (or something).

The other big wireless standard – and one with a bigger immediate impact – is WiFi, or more correctly IEEE802.11b. This is the wireless networking standard that basically does ethernet without the wires. It's perfect for home networks, and getting cheaper all the time. There are thousands of websites and HOWTOs devoted to this phenomenon, but one of the more useful ones is O'Reilly's recipe for creating a Linux-based WLAN. See www.oreillynet.com/pub/a/wireless/2001/03/06/recipe.html

NEWSBYTES

■ The **Isreali government** is set to stump up a significant amount of funding to help with the Hebrew localisation of *StarOffice*. IBM will also be involved in the project.

■ **The Instant Messaging Standards Board (IMSB)** have met with representatives of all the major IM technology providers to try and hammer out some sort of interoperability standard. IMSB spokesman Navin Rajapakse said the intention is not to create a new set of standards, but to nudge existing standards into an interoperable state. “We want instant messaging to be like email,” he said.

■ **The Liberty Alliance**, who are devising a standard for federated online identification, have announced that the next iteration of the standard may be compatible with MS's Passport. Speaking at the Catalyst Conference in Munich Paul Madison, from LA member Entrust Inc., said interoperability could be a part of the 1.1 spec. “We see opportunities for interoperability between Passport and Liberty Alliance.”

■ The idea of skipping from number 2.4 to 3.0 for the next major **kernel release** has been mooted on the main development mailing list. The idea received few positive responses until Linus said that he didn't have a problem with it. Ingo Molnar said: “I consider the VM and IO improvements some of the most important things that have happened in the past 5 years – and it's definitely something that users will notice. I think due to these improvements if we don't call the next kernel 3.0 then probably no Linux kernel in the future will deserve a major number.”

■ **NetOp Remote Control** is an established app which allows users to connect and control applications remotely (hence the name) and the latest version supports a range of devices and OSs including Linux, Symbian, Windows and Solaris. We should have a review in an upcoming issue.See **www.blinkenlights.de/interactive.en.html** for look at one spectacular use for a small Linux PC.

■ **YAFFS** is the first GLP'd flash filesystem designed specifically for NAND flash. Most flash RAM to date has been NOR (see the article in this month's *Linux Pro*), but NAND flash is increasingly being used. Smartmedia cards are the most obvious example. Tools for generating YAFFS images, & using YAFFS in a test environment without real flash are included. www.aleph1.co.uk/armlinux/projects

TAXMAN SEEKS REAL VALUE

Inland Revenue prep Sun migration

The UK Inland Revenue service is allegedly in talks with Sun about migrating 70,000 desktops away from Windows and MS *Office* to Linux and *StarOffice*. The move is said to be in response to Microsoft's decision to drop support for Windows NT4 next year.

The organisation says there are no plans to upgrade to Windows XP or 2000 and as they intended on installing 30,000 new desktops next year, it would provide the ideal opportunity to begin migrating to alternative desktop and productivity software.

Sun backtrack on x86

A concerted campaign to convince Sun to release Solaris 9 for the x86 architecture has apparently borne fruit with the company now saying it will ship an 'unbundled' edition in the next four weeks. The package will cost $99 for a single CPU license.

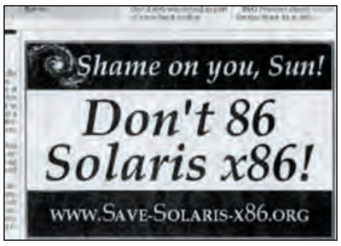

Shame on you, Sun!

Don't 86 Solaris x86!

www.Save-Solaris-x86.org

The Save-Solaris on x86 campaign involved adverts in a number of newspapers including *The Guardian*.

Solaris users launched a campaign (www.save-solaris-x86.org) back in January when Sun, citing development costs, said it would no longer support non-Sun hardware. Talking to eWeek, John Loiacono, vice president of the operating platforms, said Sun had listened and responded to the Solaris community. “”We always said we would continue to talk to the community and find a viable business model that would support it unbundled. We've now done that,” he said.

CHINESE CHIP FAB

Enter the Dragon

After establishing its own OS project, the Chinese government has announced and launched a dedicated CPU program in order to cater for the boom in the countries IT sector. It's estimated that China's current growth rate will require 17 billion chips before 2005, making it the second biggest market in the world.

The chip is said to run an embedded version of Linux, called Shuguang, and a dedicated web server called Soaring Dragon. The initial offering from the Computer Institution of the Chinese Academy of Science is said to be on par with 486 performance, but gains are expected to be fast. The government hope to achieve Pentium III standard by 2003.

The government are thought to be keen on developing it's own technology to avoid relying on imported product.

“Information security is our first line of national defence,” a spokesman told China's People's Daily. “However, the line was built on the foreign technology, which is an obvious worry.”

Jono Bacon

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

COMMENT

Organisation? Bah...

“ As a humble reader of *LXF*, you may have noticed the wide variety of people involved in the Linux and Open Source effort. Within these varieties of people, there are some geniuses, oddballs, dictators and followers – safe to say there are all sorts indeed.

Although we have lots of different people in this community, one wonderful thing is that very few people make demands of your time: in many cases a project's group of developers or contributors often work in an *ad-hoc* fashion that suits everyone. There may indeed be some critics of this *ad-hoc* way of working, but it is important to remember that there is more than one way of running a community.

Take LUGs. Regular readers of my incessant ramblings will note that I am a bit of an advocate of getting involved with them. The man behind the LUG effort is a chap named Mark Lewis who has devoted his time and expertise into forming an infrastructure for setting up and running a LUG. Now, want to know the great thing about how Mark runs this? Well, it is not the fact that he provides all of the services that you need such as mailing lists and web servers...but merely the fact that he makes no formal demands on your time. Some may see it as *ad-hoc*, some may see it as simplicity that works well.

The moral of this story is that the world is not an average place; there is no typical user, no typical developer and no typical method of doing something. In a community such as ours, it is important to remember the visceral beliefs and practices on display and to learn and share with them. Oh, and go and join a LUG. :) ”

APACHE WORM

Infected by a slapper

A worm targeting *Apache* web servers has been quietly attempting to build a peer-to-peer (P2P) network capable of launching distributed denial of service attacks. Symantic noted as many as 3,500 infections in the few days following identification and released an advisory notice saying it was "in the wild and actively attacking other servers."

Named *Slapper*, the worm uses a well-documented and patched hole in *Apache's OpenSSL* implementation, so will only be able to infect older unpatched systems. It's first DDoS target appeared to be a major Internet service provider.

For more on its detection, footprint and removal see the Internet Storm Centre at **http://isc.incidents.org/analysis.html?id=167**

PHP ACCELERATION

Fastest PHP Performance

Zend have just announced the release of *Zend Performance Suite 3.0*, their next-gen PHP acceleration product for organisations that have high-volume web sites powered by PHP.

Through the use of three core optimisation components – code acceleration, content caching, and data compression – Zend hope to increase the PHP acceleration stakes by providing an estimated speed-up of up to 25 times that of unaccelerated PHP, enabling IT managers to serve more content using less hardware.

New and enhanced features in this version include automatic compression of output to supporting clients including caching of compressed content, auto identification of scripts that would benefit the most from caching, and configurable cache-hit determination.

The Zend Performance Suite Enterprise Edition will be launching for about £650 ($1000) for one CPU. An Accelerator edition, which is just the code acceleration component by itself, is priced starting at £370 ($575). The product is available at www.zend.com.

Embedded Linux News

- Sandia National Labs have created a four node Linux cluster using off-the-shelf small form PC/104 hardware. Does this bring supercomputing to the mobile crowd? Mitch Williams, engineer at Sandia's Embedded Reasoning Institute (ERI) said that the design managed to integrate all the standard features of normal-sized rack-mounted clusters, including individual keyboard, video and mouse access to every node into a device barely over a foot tall and five inches wide.

- Sharp have given developers a sneaky peak at the next generation Zaurus machine which, like the SL-5500 will be built on Linux. The device is said to be a clam-shell design with one half housing a keyboard and the other a new CG Silicon display which is capable of packing in twice the pixels of a traditional TFT screen. The screen will apparently run at 640x480 resolution.

DIAL-A-DONATION

Support as you surf

UKFSN is a new ISP with a difference – *all* of the profits from the operation will be donated to fund Free Software projects in the UK. Each month UKFSN will release a public statement clearly showing how much income was raised during the preceding month and how much is cost to run the service. The surplus will be paid to the Association for Free Software (**www.affs.org.uk**) who have agreed to handle the distribution of the money raised.

The UK Free Software Network offers all the standard services one expects from an ISP including 0845 dialup access, email, web hosting and usenet. These standard offerings are beefed up with the inclusion of Perl, PHP and python, and *MySQL* support for the web hosting and extensive SPAM control measures in the email service.

The service is not just for geeks however. As well as subscription free Internet access and services UKFSN.ORG offers professional hosting options for businesses and domain registration services.

UKFSN is being run by Jason Clifford, the person behind Definite Linux – the now defunct British Linux distribution. **www.ukfsn.org info@ukfsn.org**

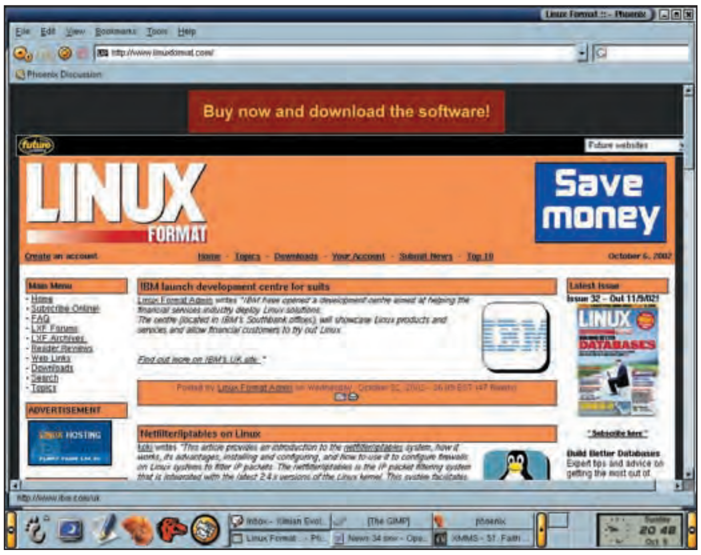

Buy now and download the software!

Create an account · Home · Topics · Downloads · Your Account · Submit News · Top 10 · October 6, 2002

IBM launch development centre for suits

Linux Format Admin writes "IBM have opened a development centre aimed at helping the financial services industry deploy Linux solutions. The centre (located in IBM's Southbank offices) will showcase Linux products and services and allow financial customers to try out Linux.

Find out more on IBM's UK site."

Netfilter/iptables on Linux

kiki writes "This article provides an introduction to the netfilter/iptables system, how it works, its advantages, installing and configuring, and how to use it to configure firewalls on Linux systems to filter IP packets. The netfilter/iptables is the IP packet filtering system that is integrated with the latest 2.4.x versions of the Linux kernel. This system facilitates

Latest Issue: Issue 32 – Out 11/9/02!

Build Better Databases: Expert tips and advice on getting the most out of

***Phoenix* is a *Galeon*-like browser built in XUL and *Komodo*. (below) ActiveState's *Komodo IDE* reaches version 2.**

FASTER BROWSER

Small, fast Mozilla!

Mozilla's core development team have begun an offshoot project to create small fast browser in the *Chimera* (OS-X) mould for Windows and Linux. *Phoenix* drops the mail, composer and chat elements of the main project, features a Scaleable *Gorilla*-like UI and compatibility with many of the current *Mozilla* add-ons. The first release lost the sidebar, but that has now returned along with an extension of the address-bar autocomplete feature, which does the same job on web forms.

The latest official *Mozilla* release, meanwhile, has a new Typeaheadfind feature which allows you to activate a link on a page by typing its name. If a number of links have the same initial first letter, subsequent key-presses will cycle through them all. Ideal for the musophobics out there.

ActiveState's *Komodo* – an Integrated Development Environment and one of the most successful XUL projects to come from the *Mozilla* application framework – has just hit version 2.

The IDE is available for both Linux and Windows in personal and professional editions. Designed for working with various 'open' languages including Perl and Python, *Komodo* has a new XSLT debugging tool, GUI builder and support for the latest web services technology.

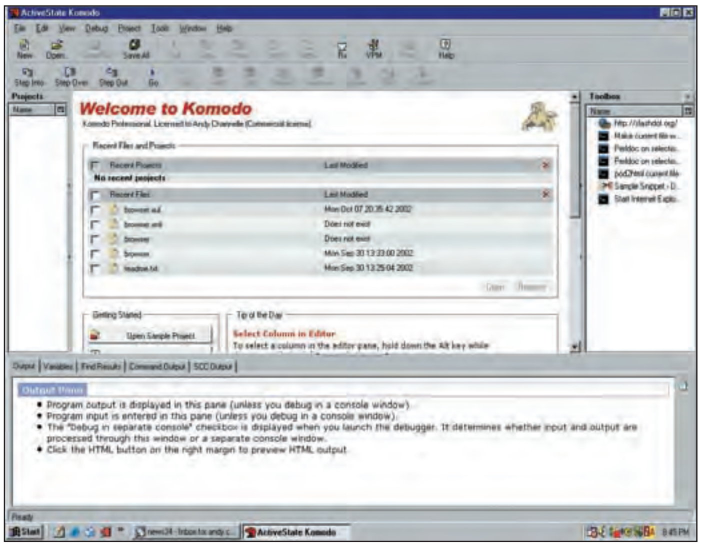

Welcome to Komodo

Komodo Professional. Licensed to Andy Channelle (Commercial license)

No recent projects

Recent Files	Last Modified
browser.xul	Mon Oct 07 20:35:42 2002
browser.xml	Does not exist
browser	Does not exist
browser	Mon Sep 30 13:33:00 2002
readme.txt	Mon Sep 30 13:25:04 2002

Select Column in Editor

To select a column in the editor pane, hold down the Alt key while

Output Pane

- Program output is displayed in this pane (unless you debug in a console window).
- Program input is entered in this pane (unless you debug in a console window).
- The "Debug in separate console" checkbox is displayed when you launch the debugger. It determines whether input and output are processed through this window or a separate console window.
- Click the HTML button on the right margin to preview HTML output

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

A change as good as rest

"Admittedly, I have been slavishly devoted to a Mandrake/KDE desktop. I felt comfortable with it. I learned the tips and tweaks, the ins and outs, the glitches and the gotchas.

But I got an assignment to write about Red Hat and participate in their beta development program, following it all the way through to the final release. I hesitated to accept because I knew I would have to spend time in the GNOME desktop since that is the Red Hat default.

Dual booting until I could learn my way around the "new" OS, I found it had problems, as do all betas, and I found myself booting back into Mandrake and KDE to get my work done. I was not learning anything about my assignment.

Taking the plunge, I backed up critical data and installed Red Hat Limbo (the 2nd beta) as my only OS, also making a commitment to use only the GNOME desktop until the next release.

What I discovered was that my initial discomfort disappeared after a week or so. The two distros may be similar, but are not identical and GNOME is very different in use from KDE, even in the new Red Hat version of the desktops.

Once comfortably immersed, however, I wound up filing a few *Bugzilla* reports and participating in mail list discussions. I have enjoyed my experience and learned a great deal from it. I'm now less inclined to be zealous about OS/desktop advocacy as I got my work done on either system, in either desktop. That ideology has become unimportant to me.

I suggest you make a similar change for a month. Try a new distro or desktop. Don't dual boot – go all the way."

Mailserver

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

★Letter of the month

This month's winner receives a copy of **SuSE Linux Personal Edition**

Which Filesystem?

I wonder if it would be possible for you guys to review some of the journaling file systems in a future issue please. As an *ext2* user, I like the idea of being able to migrate painlessly to *ext3* and then back to *ext2* if I ever need to. On the other hand, if one of the other systems such as *ReiserFS*, *jfs* or *XFS* offers significant advantages over *ext3* then I could be tempted to use one of them instead. Benchmarks that I am interested in include:

1. Reliability – code stability and code maturity are important but so is the reliability of the underlying technology. Do any of the journaling file systems use forward error correction techniques to improve reliability in a system which is beginning to suffer from hard read/write errors?

2.Tools – for disk tuning, disaster recovery etc.

3.Performance – does latency and peak bandwidth performance vary between filesystems or is this primarily down to the performance of the underlying hardware?

4.Bootability – less important as it's always possible to use an *ext2* /boot.

Steve Roper, *via email*

That's a really good idea. We have done a feature in the past about the different systems and also a tutorial on migrating to *ext3* before, but a more direct comparison may indeed be useful. In terms of tools and compatibility, *ext3* will probably come out best as it is really an extension to *ext2*, and most *ext2* tools will work happily with *ext3* (though there are exceptions).

In recognition of your good work, We'll be sending you a copy of SuSE 8.1 Personal, this month's star prize.

Save our CD!

I just noticed the other week that one of your competitor magazines has dropped their coverdisc because of feedback from their readers. Ultimately this has meant that the magazine price has dropped as well to compensate, which I is fair enough.

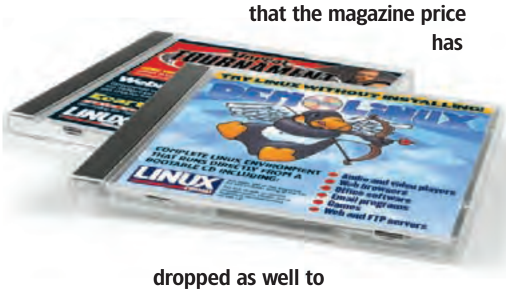

If *Linux Format* is receiving similar feedback from it's loyal readers then all I can say is ignore them completely! Don't listen to a word they say! They must have gone temporarily insane! You can drop the cover price if you want though.

Since changing employer I no longer have access to a stupidly large amount of Internet bandwidth and have to rely on a 56k modem, which isn't really ideal for the likes of the Linux kernel source or apps like *OpenOffice.org*, let alone downloading ISO images of beta releases of Red Hat etc.! I now rely solely on my subscription to *LXF*, and the attendant DVD, for all those essential bits and pieces that I can no longer download for myself.

On a totally different point, can I make another request: can we have *more* adverts please. Not adverts for hosting companies though, I'm sure these are useful to some people but not to me, or maybe not for a lot of home users if it comes to that. No, what we need are adverts for computers, bits and pieces and other useful stuff to buy.

There are, I believe, quite a few companies that sell Linux friendly systems or components and it would be really nice to have a magazine with a nice range of adverts for these, as well as the excellent articles, features, etc that we already enjoy.

Keep up the excellent download service, er, I mean magazine.

Andy Bianchi, *via email*

You will be pleased to know that we have no intention of dropping the CD or DVD at the moment. Some people do suggest that they might prefer a cheaper magazine without the CDs, but to be honest, in many ways it wouldn't cost that much less. We would certainly want to keep a disc version anyway, and then there would be three versions of the magazine!

You may or may not have broadband, but isn't it more convenient to have the stuff on a disc – it still saves time. I have broadband myself, but I'm not likely to try and download the latest Mandrake distro at home!

Also, this rather ignores the fact that the disc and the magazine can work together, not just for the programs we mention, but for source files and scripts from the magazine, and extra information on particular topics. What do the rest of you think?

Knoppix etc.

The Knoppix distribution you provided on the cover DVD of *LXF 32* was brilliant. As a fairly experienced Windows user who is new to Linux, it gave me a great opportunity to try Linux for the first time without actually installing anything. It's a pity that some of the more mainstream distros don't have 'demo' versions on discs like that, as it would be a great way to try them out. The article on Palladium was very interesting, although I can't exactly see Europe introducing laws that give US companies control over what you can do on your own computer. I am also curious as to why Intel and AMD appear to be so keen on it. If anything TCPA, Palladium, etc. will restrict what users can use their computer for, and hence decrease demand for PCs and in turn processors. Also, adding "anti-piracy chips" to processors will surely increase their cost of production and make them less competitive. Finally, would it be possible to include more introductory Linux topics each month, as the majority of the current guides are at a little too high a level for new users, perhaps you could include an intro guide on the discs each month.

Keep up the good work on the magazine

James Porter, *via email*

More TCPA

I thoroughly appreciated your article on Palladium/TCPA. However,

Knoppix – a great introduction for the hesitant Linux newcomer.

one paragraph is missing and that's the one explaining what you can do to boycott or otherwise complain about Palladium. The European Parliament seems to be one key factor that could stop Palladium – exerting pressure on them (as EU citizens) seems to be the thing to do.

After all, this is a purely US 'solution' being shoved down everyone's throat. The EU is not exactly in the US jurisdiction, or is it?

Gunnar Mortier, *via email*

Well done on giving prominence to the issue of TCPA/Palladium – I had already come across it, but the more publicity and de-FUDing the better!

I agree that TCPA/Palladium is a serious threat, but I see some major flaws in the concept, other than the ones raised in your recent article.

1. I can see how the TCPA could enforce their "standards" on the home and business desktop, where Windows is already established, but what about servers? Presumably it won't be necessary for all servers involved in, say, distribution of DRM-enabled music and videos to be TCPA-compliant, unless the TCPA propose winding back the clock to the days of private networks like Prodigy – the Internet is just too *big* now to impose such a far-reaching technological change on a realistic timescale. If servers don't have to be TCPA-compliant, why change them? And which OS is used on the majority of servers...?

2. If a TCPA-compliant box won't run uncertified software, does that extend to the shell scripts I write to help me run my webserver more efficiently? [I don't use Windows on servers, so I don't know whether batch files might also fall foul of this one.] Not a problem if non-compliant servers are feasible, but otherwise there's going to be a lot of unhappy sysadmins around!

3. It will also significantly slow software development and increase costs, if every beta has to be certified before it can be tested on a typical user desktop – not to mention requiring extra machines for developers, who will need a TCPA box to run their nice M$ development tools and a non-TCPA one to try out their pre-certified binaries?

Don't get me wrong - I think that one can't publicise the bad aspects of TCPA/Palladium enough, and I dread the prospect of having to go back to Windoze full-time (I'm quite happy to visit occasionally as long as I don't have to live there!) – but I'm not yet convinced that it will destroy Open Source Software totally.

Anne Parker, *via email*

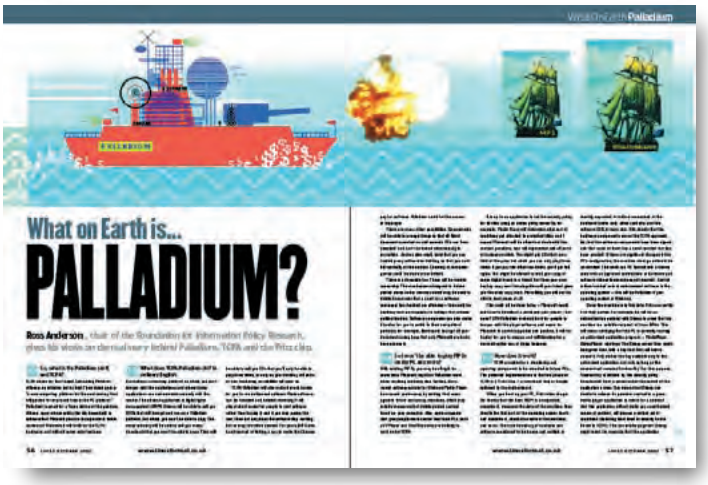

Palladium – removing real ownership from the computer user.

Thanks for these letters and the many other emails and calls we've had about TCPA. Recently exponents of the scheme said that it wouldn't block the running of OSes such as Linux, but didn't really clarify that one would still be able to enjoy your OS of choice for activities like online banking, watching DVDs, listening to music and all the other things a desktop OS is currently being used for.

This really is a classic case of a solution without a problem (at least from the user point of view). As we have seen it will do nothing to diminish SPAM or halt viruses, but will restrict the users freedom to use their computer (for legal as well as illegal purposes).

I think everyone can quite understand the desire on behalf of the owners to stamp out copyright theft, but in many ways this proposed solution goes too far, and effectively removes real 'ownership' of the computer from the user.

"A classic case of a solution without a problem (from the user point of view) as we have seen that it will do nothing to diminish SPAM or halt viruses"

More tips

In *LXF31*, your tip-of-the-month was on the subject of kernel boot parameters. You missed one great tip, however:

After your kernel has started, /proc/cmdline contains the kernel boot parameters that the kernel was started with. Since the kernel will ignore any parameters that it doesn't recognise, you can take advantage of this to pass any boot-time options you care to dream up to your startup scripts.

Mandrake's /etc/rc.d/rc.sysinit make extensive use of this facility in code fragments such as this:

```
if grep -iq nopnp /proc/cmdline
 >/dev/null 2>&1 ; then
   PNP=
else
   PNP=yes
fi
```

Month by month, your mag is getting better and better. Keep it up!

Neal Crook, *via email*

Thanks for this great tip. For those of you a little confused, the example shows how a script can read the supplied kernel parameters, even ones the kernel does not do anything with, to set further runtime system parameters.

GUI standards

As a new user it has been interesting as well as laughable the bickering between GNOME and KDE coders. The single most important factor keeping Linux at bay from garnering a greater share of computer installations is that it does not have a standard gui interface.

It requires a C++ OOP library base standard on which a combination of full fledged applications with stripped options are available for complex, simple work and thin client use.

A complete CD recording program, and multimedia integration of programs like *Mplayer* plus a basic video editing package would leave windows with the game playing market only.

It would allow a GUI install/uninstall program, without which Linux will remain a niche player.

Dr Matt Despot, *Australia*

»

« An interesting point of view. Interestingly there are already several GUI install programs that don't rely on either desktop system, though they aren't in common use. Does anyone else think that KDE vs. GNOME is good or bad for Linux? Does having two projects hasten the pace of development or hold us back? Do we like what Red Hat have done in 8.0? Let us know your thoughts by emailing or writing to us!

Configuration

Justin Eastham (*Answers*, page 93, *LXF 31*) has hit the nail right on the head with his comment "Perhaps I'll go back to Slackware – at least I knew where I was with plain ASCII config files ;-)" As a professional Unix/Linux sysadmin, I find myself constantly fighting recent versions of Red Hat (and, to a lesser extent, SuSE) – a small change to a simple ASCII config file is often all that's needed to reconfigure something yet if I do this with these distros I find it overwritten by *linuxconf*, *et al.* at the next boot!

I know the creators of these distros have worked hard to make them easy to use and configure for novice users (often those from a MS Windows background) but it's really frustrating for experienced Unix types to work with these systems. Can we ask that a simple or 'lean' config mode is made available on these distros for those who are familiar with standard Unix/Linux config files, bypassing *linuxconf*, etc?

This is one of the reasons 'Unix-like' Linux distributions such as Slackware, Gentoo and, of course, FreeBSD, NetBSD, etc. continue to be so popular and I was pleased to see Slackware 8.1 on your coverdisc!

Andy Thomas, *via email*

Some very good points here. Is it really the case that ease of use should always equate to a restriction of flexibility? Or does it just have to be applied in a smart way. In many respects 'ease of use', at least for common config tasks and desktop use, is a fairly new concept to Linux – perhaps it just needs time to mature. In the meantime, try not to frustrate those who are comfortable with more arcane methods!

Do user-friendly config tools cause you more headaches than they fix?

Nuggets

I have been reading *Linux Format* since issue one. I enjoy the tutorials and reviews but, I get the most pleasure when I discover a really good open source package. Your *HottestPick* (October) *ntop* is just this sort of nugget. I run a small network (about 150 IP addresses) and I was looking for a network monitoring tool. *ntop* is fantastic, and being open source I can add any requirements that I may have to it, if I can think of any. I can't praise *ntop* enough, even though it did take me nearly two hours to install it (the installation notes are not located in one place) the effort was worth it. I was shocked by the visual appearance and ease of use of the system. In your article you said that the web interface does not update, well my installation did, it refreshes every two minutes.

Finally, I installed *ntop* on an old PC that had been retired because it was too slow to run Windows 2000 and the system runs perfectly. Has anybody ever installed a disk-less Linux system *i.e.* a PC without a hard disk, *e.g.* using compact flash instead of a hard disk. I could then plug this device into the network and forget about it as it will monitor the network forever!

Please keep on finding nuggets like these.

Regards,

Dave Walsh, *Eire.*

Thanks for the kind words. In many ways, the unearthing of such nuggets is the whole point of hotpicks. So often it is easy to loose really useful tools or utilities in the sheer volume of software that is released or updated each month. As well as perhaps sparking ideas for new uses of Linux, and keeping you abreast of what's happening, *HotPicks* can bring the best of these to your attention.

The tops

Thank you for the very nice write-up about *ntop* – turning the pages, it was a nice surprise to see it. As one of the (minor) developers of *ntop* (By no stretch of the imagination am I the author of ntop), it's nice to see glowing words about one's work.

First the good news – since *v2.1* which was reviewed – we have released *v2.1.2* and now most recently (2002-09-15) *v2.1.3*, with additional bug fixes and some minor improvements. The current version (and Red Hat RPMs) are available on SourceForge. I know *LXF* is a Linux magazine, but in addition to almost any flavour of Linux you can imagine, *ntop* also

Helpdex

shane_collinge@yahoo.com

BY SHANE COLLINGE

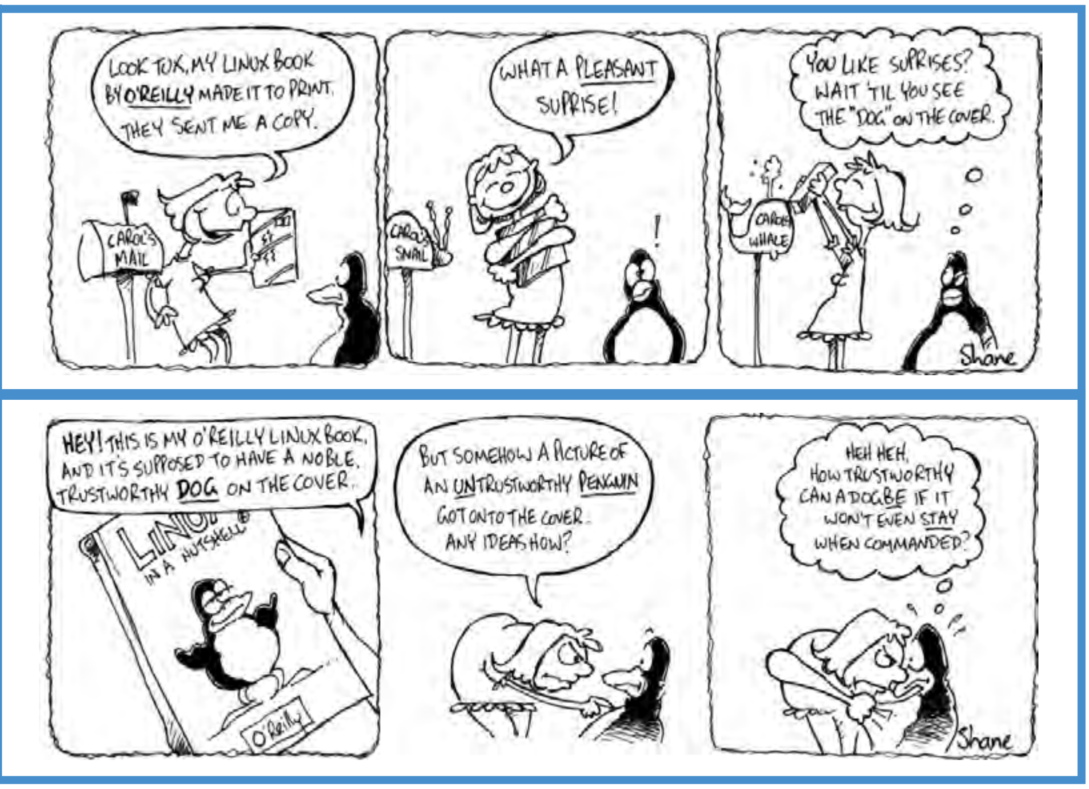
LOOK TUX, MY LINUX BOOK BY O'REILLY MADE IT TO PRINT. THEY SENT ME A COPY.
CAROL'S MAIL
WHAT A PLEASANT SUPRISE!
CAROL'S SNAIL
!
YOU LIKE SUPRISES? WAIT 'TIL YOU SEE THE "DOG" ON THE COVER.
CAROL'S WHALE
Shane
HEY! THIS IS MY O'REILLY LINUX BOOK, AND IT'S SUPPOSED TO HAVE A NOBLE, TRUSTWORTHY DOG ON THE COVER.
LINUX IN A NUTSHELL
O'Reilly
BUT SOMEHOW A PICTURE OF AN UNTRUSTWORTHY PENGUIN GOT ONTO THE COVER. ANY IDEAS HOW?
HEH HEH. HOW TRUSTWORTHY CAN A DOG BE IF IT WON'T EVEN STAY WHEN COMMANDED?
Shane

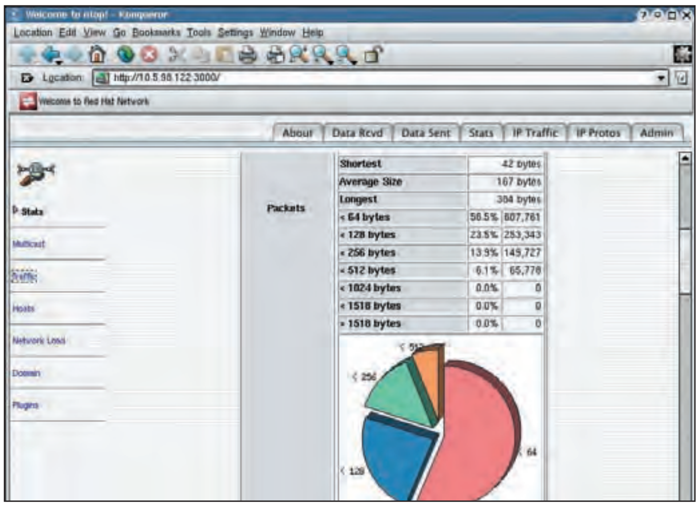
Welcome to ntop! - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help
Location: http://10.5.98.122:3000/
Welcome to Red Hat Network
About
Data Rcvd
Data Sent
Stats
IP Traffic
IP Protos
Admin
Stats
Multicast
Traffic
Hosts
Network Load
Domain
Plugins
< 512
< 256
< 128
64

Packets	Shortest	42 bytes	
	Average Size	167 bytes	
	Longest	384 bytes	
	< 64 bytes	56.5%	607,761
	< 128 bytes	23.5%	253,343
	< 256 bytes	13.9%	149,727
	< 512 bytes	6.1%	65,778
	< 1024 bytes	0.0%	0
	< 1518 bytes	0.0%	0
	> 1518 bytes	0.0%	0

ntop **– a great network monitoring tool.**

runs under Mac OS (Darwin), Windows32, Solaris and FreeBSD. We're working with a user trying to get it to run under HP-UX and I personally have it running (albeit single threaded) under NetBSD.

Unfortunately, I have two issues with the review.

Contrary to the reviewer's assertion, the web interface does in fact automatically refresh its pages. There are a small number of (administrative and control) pages that do not automatically refresh, but all of the data pages refresh automatically. The refresh rate (in seconds) is set via the -r parameter. The default is 120 seconds (2 minutes), so perhaps the reviewer was simply impatient.

A large portion of the review is really about *intop*, which is obsolete. All of our efforts since *v1.3* have focused upon expanding and improving the internals and the html interface. The individual who developed and maintained *intop* has left the project to pursue other interests. This means that *intop* has not been updated in functionality since the *ntop1.3* release.

We maintain it to the bare minimum to ensure it compiles and starts up, but that's all. As your reviewer points out, the web interface is "well laid out," "perhaps easier to use" and "provides you with all the info that the command line interface does and then some." The only fault I can find with that is "perhaps".

However, those are minor points. Thanks for the review!

If anyone would like to help with *ntop*, y'all come on over – *especially* if you're willing to do a little (creative) writing or willing to maintain *intop*.

Don't worry if English isn't your native language, we're a pretty fluent bunch – active people on the mailing list this week include a couple of Italians (Luca Deri, the Author of *ntop* is at the University of Pisa), some Brazilians, a Norwegian, a Finn, a German, a gent from the Netherlands, an Australian and a couple of Californians. (And those are just the people who's email address makes it easy to figure out). And there is always Babelfish!

Burton M. Strauss III
Texas, USA (Burton@ntopsupport.com)

Thanks for the feedback. I'm glad you thought (overall) that the piece was good! Reviewers do tend towards impatience, so that may well have been the problem. Thanks very much for getting in touch and for producing, if you'll excuse the pun, some top software! LXF

Submission advice

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

WHAT WE DON'T WANT:
- Technical question – direct those to our Q&A pages!
- Random abuse
- Nonsense rants
- 200 pages of meandering diatribe

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

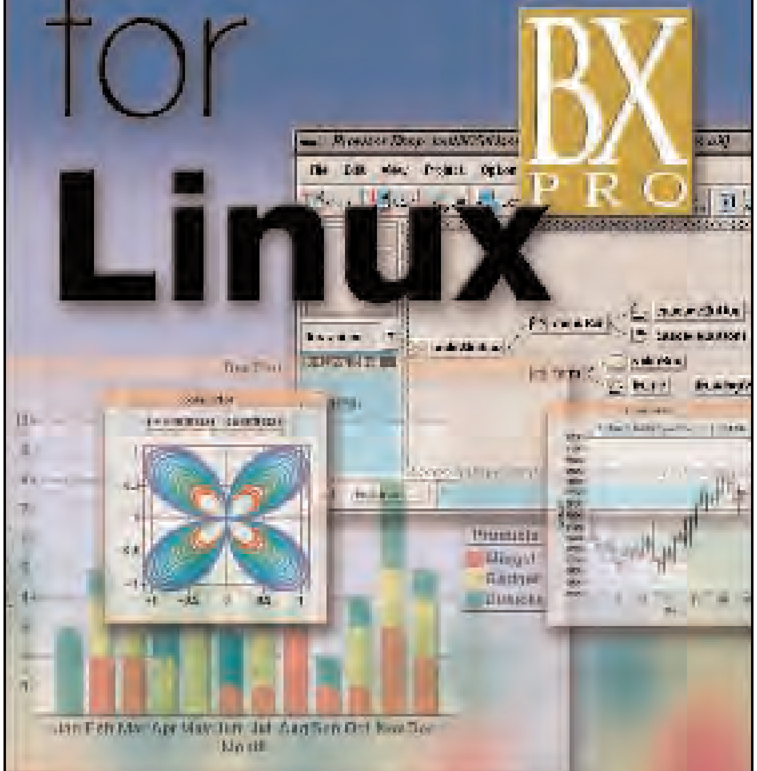

Reviews »

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

LXF verdict explained

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

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

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

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

Value for money: Does it have a competitive price?

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

10 The close to perfect product.

8-9 Good, but has a few niggles.

6-7 Does the job, but needs work.

5-4 Average.

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

The Top Stuff Award

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

THIS MONTH...

SuSE Linux 8.1

More packages, easier installation, no hardware problems – happy 10th birthday, SuSE **p36**

Red Hat 8.0

Have the Bluecurve modifications made the desktop more or less usable? **p22**

Netraverse Terminal Server

Serve your Windows desktops from Linux, to Linux, and reduce administration headaches **p26**

Snapgear PRO+ firewall »

Embedded Linux,web-based configuration and a powerful feature list **p29**

Smoothwall

Is the commercial version of everyone's favourite firewall worth shelling out for? **p30**

Lycoris Desktop

Continuing after the Windows desktop market, Redmond Linux adds a powerful update tool – but when will they deal with the hardware problems? **p34**

Books

Go to bed early and read: *XML in a Nutshell* (2nd edition); *MySQL Reference Manual*; *Linux Kernel Programming* (3rd Edition) **p40**

COMING UP SOON...

Mandrake 9.0

The much anticipated LSB 1.2-compliant release of the *LXF* readers' favourite distro features *GCC 3.2*; *Glibc 2.2.5* and the latest desktops, Internet tools and multimedia apps. We look at the installation, configuration and security improvements.

NetOp

The commercial, cross-platform, fast, scalable and secure remote control application.

Webmin 1.0

The incredibly versatile, web-based remote admin tool reaches its milestone release.

UT 2003

The latest FPS reaches Linux – and this time it's available off-the-shelf in a simultaneous release.

Xandros Linux

The project to resurrect Corel Linux nears completion.

DESKTOP DISTRIBUTION

Red Hat Linux 8.0

Hoyt Duff, co-author of Red Hat Linux 8.0 Unleashed, examines the latest offering from Red Hat, a distribution targeted at the corporate desktop.

Interesting features and a surprise or two. Will Red Hat be a trendsetter?

- **DEVELOPER** Red Hat
- **WEB** www.redhat.com
- **PRICE** Personal–US$40; Professional–US$150 (includes DVD); free GPL version reviewed here.

cover feature

By now, you've read the Internet reviews and probably participated in the heated debate caused by the release of the eighth version of what is arguably the world's best-known Linux distro, Red Hat Linux. If not, we'll recap and illuminate the discussion a bit:

THE BUZZ: The *Anaconda* installer is now prettier because it uses a native video driver or alternatively defaults to a VESA driver.

OUR TAKE: It looks beautiful, but some video cards are not supported; those will need to use the text install.

THE BUZZ: The Bluecurve Theme is gorgeous and looks the same whether you choose KDE or GNOME. Detractors claim it looks like Windows XP.

OUR TAKE: It's gorgeous. The name Bluecurve is trademarked and the Bluecurve icon package can't be redistributed without a name change, giving rise to the false rumour that Red Hat has gone proprietary.

THE BUZZ: GNOME users are pleased to see so many GNOME apps as the default choices.

OUR TAKE: Disappointingly, the desktop runs on the *metacity* window manager (*Sawfish* is provided, slightly broken) and GNOME 2 lacks a menu editor. It doesn't include several popular GNOME apps and uses the *VTE* terminal emulator (whose cut and paste behaviour may be odd to you) instead of *Zvt*. Double-click is the default mouse behaviour. *GDM* is the default login manager, but can be changed.

THE BUZZ: KDE users are angry that *Konqueror* is not the default browser in KDE and they suspect that Red Hat has "crippled" KDE.

OUR TAKE: Many KDE apps are missing and none of the KDE apps show an "About KDE" selection. *Qt* has been modified to work with AA font support as implemented by Red Hat. *Konqueror* is all but hidden and "virtual folder" support from GNOME is hacked into KDE.

THE BUZZ: Red Hat has removed MP3 and MPEG capabilities from the distro.

OUR TAKE: Red Hat justify this on legal grounds; they could be wrong, but are playing it safe. We don't blame them.

THE BUZZ: Red Hat dropped my favourite application.

OUR TAKE: Yes, because of copyright or patent issues, or the app didn't play well with UTF8 character encoding. There are also apps installed that didn't get added to the menus, so you might have just overlooked your favourite.

RHCE: Rated #1 in overall quality of programme by Certification Magazine (January 2002). Register now for a Red Hat Training course at www.europe.redhat.com/training
redhat
redhat

THE BUZZ: Red Hat has ruined what was a fabulous Linux distribution. Red Hat is evil. Red Hat is Microsoft. Red Hat hates KDE.

OUR TAKE: That's the controversy, isn't it? Fortunately, Open Source Software can be modified to suit your tastes; Red Hat doesn't stand in your way.

What the Internet rabble fails to realise is that they are examining a watershed event on the Linux timeline. Red Hat have produced a foundation for their enterprise-class offerings; it's their answer to United Linux. The design criteria is strict and they are exercising their rights in an Open Source market. They have winnowed the field of offered applications to a recommended few to facilitate the effective support, and integration, of their distribution. With SuSE and Mandrake looking like the Swiss Army knife of distros (and expensive to fully support at the corporate level), Red Hat is attempting to be fiscally prudent in its minimalism.

Installation

If you have performed an installation of a 7.x version of Red Hat, you'll notice that little has changed other that the appearance. A new installation profile, "Personal Desktop", has been added and the "Workstation" profile has been modified to provide a software development environment.

Goodies

Thank Heaven for Open source! Travel to www.freshrpm.net where you'll find a number of replacement and enhancement apps. The first step should be to install *apt* and *Synaptic*, however, as a wonderful front-end for package management. After install additional software, you'll be enjoying MP3s and DVDs. We suspect that full KDE and GNOME package will be available soon.

Missing In Action

***Netscape*, alien, *fvwm2*, *gnorpm*, *gphoto*, *ical*, *kaffee*, *pump*, *rpmfind*, *taper*, *xbill*, and others have been removed. *LPRng*, *Lilo* and *sndconfig* are scheduled for removal.**

Substitutes

***eog* replaces *ee*; *privoxy* replaces *junkbuster*; *lftp* replaces *ncftp*, among a few others.**

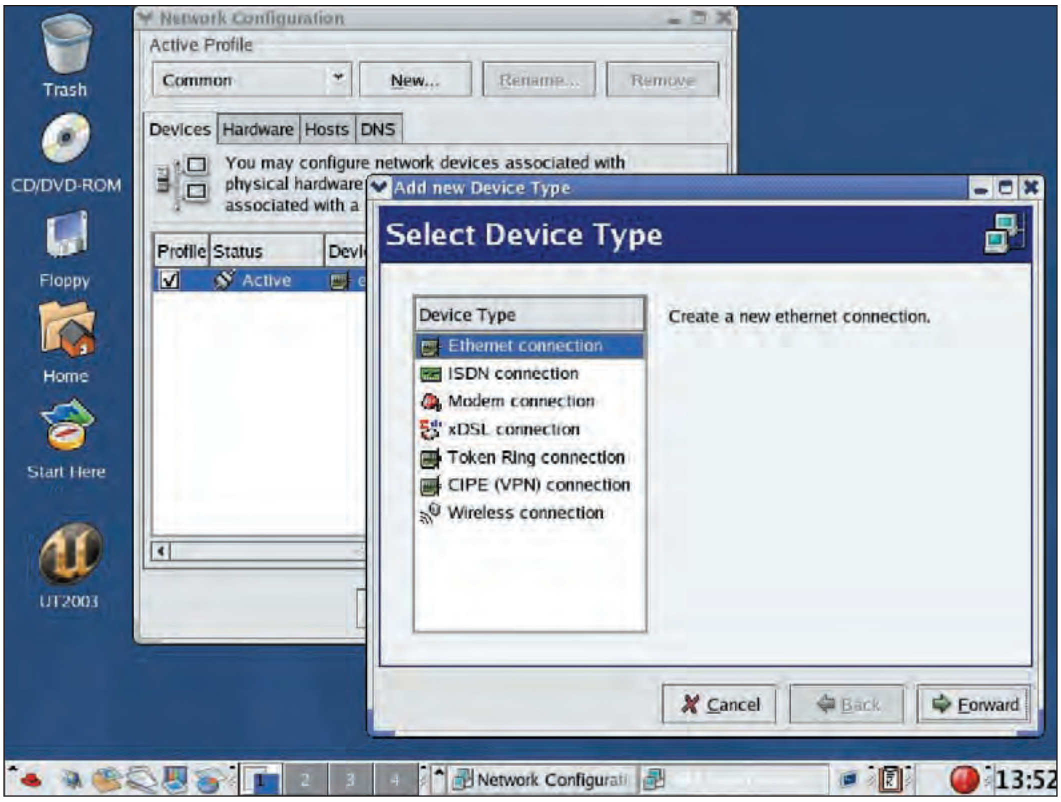

You may configure network devices associated with physical hardware associated with a

Create a new ethernet connection.

Two of the many configuration tools provided by Red Hat: the Network Connection Wizard and the Network Configuration GUI; they could use some context-sensitive Help dialogs.

A Hat of Your Own

Grappling with an anaconda

The average user is unaware that the Red Hat installer, *Anaconda*, is mostly a collection of Python scripts that are Open Source. This means that an enterprising individual could customise the installation of Red Hat to add new or updated RPMs to the CD, change the default options, or add even new functionality. The only thing standing in the way is a deliberate lack of documentation for the process. One can understand Red Hat's reluctance to flaunt the inner workings of the code that is a significant part of the value that they add to a Linux distribution.

Fortunately, Tony Nugent has done the difficult work for us and has assembled some of the disparate documentation at www.linuxworks.com.au/redhat-installer-howto.html. His information focuses on substituting updated RPMs and generating new CDs, also providing a list of additional resources. As well, Miguel Freitas documents his experiences of doing the same thing at http://cambuca.ldhs.cetuc.puc-rio.br/RedHat7-CDs-HowTo.html. It's a task not that difficult to accomplish, only requiring a good reading of the HOWTO and the ability to search the *Anaconda* mail list archives for solutions to the inevitable errors you'll encounter.

Little has changed in the workings of *Anaconda* for Red Hat 8.0 other than the video support described elsewhere and the CD boot method. Red Hat now employs H. Peter Anvin's ISOLINUX boot method as described at http://syslinux.zytor.com/iso.php. The application allows CDs to boot in non-emulation mode, obviating the need for a separate floppy-sized boot image. The Red Hat Release Notes contain some information as to making bootable CDs with ISOLINUX.

The option to install everything loads nearly 5GB of software on your hard drive. We wish that there were a "Minimal Install" option to provide a base install for servers. It would also be nice if Red Hat enabled the **Num Lock** key by default during the installation phase.

Several early adopters have encountered installation issues related to unsupported video cards, but have used the text installation method successfully.

Modifications

Part of the controversy surrounding this release has been the dearth of multimedia applications and missing MP3 and MPEG functionality. Several websites address these perceived inadequacies by providing drop-in replacements and additional applications.

http://psyche.freshrpms.net/ (*apt* and multimedia)

www.math.unl.edu/~rdieter (KDE 3.0.3)

http://neomundi.com.br/usuarios/mab/redhat8.html (Flash, JRE, *Acrobat*, MP3 for *XMMS*)

http://cachalot.ods.org/RPMS/ (*kio_fish*, *unace*, *jikes*)

Feature rundown

Media Five CDs comprise the GPL set; the last two contain the source code. A DVD is provided with the Professional boxed set. The documentation (available from their web site) has been rewritten to address the many changes.

Installation Methods The *Anaconda* installer is little changed from its previous incarnations except that it will probe for and utilise a native video driver. The Red Hat engineers put that to good use by fully implementing *GTK* and previewing the now-famous Bluecurve icons. What results is a classy, sexy and beautiful display to watch as you proceed through essentially the same install steps you did with Red Hat 7.3. We like the built-in ability to verify the install media since so many errors arise from faulty media and so few people take the time to do an md5sum check. There are a few undocumented and unsupported install options. While at the boot: prompt, typing **linux reiserfs** or **linux jfs** will enable the choice of those filesystems during the install. Supported network install options are accessed by typing **linux askmethod**.

Logical Volume Management *LVM* is fully supported in Red Hat 8.0; it is the ability to dynamically resize hard drive space. Essentially, *LVM* is a layer between what you know as the disk partitions and what you see in a directory listing. If you run out of disk space in /tmp, for example, just install an additional hard drive and add it to the logical volume group. *Caveats*: you can only edit *LVM* groups when in the GUI install mode, not in text mode and the /boot partition can't be on a logical volume.

The X Server The older *XFree86 3.x* of is no longer provided. The default config file is now named /etc/X11/XF86Config (note, no -4). A USB mouse is configured during installation just in case you add one later, very handy for our laptop. The old *Xconfigurator* tool is gone, replaced by the awkwardly named *redhat-config-xfree86*; *xf86config* is still available as is **XFree86 --configure**. The latter useful if you use a dual-head display since the Red Hat tool doesn't support that (that functionality is a low priority for them). If your video card is not supported in *XFree86 4.x*, Red Hat defaults to the generic VGA driver. While much has been said about Red Hat being unfriendly toward third-party drivers (like nVIDIA), they have made changes in the *Mesa* packages to accommodate third party driver vendors who supply their own GL

»

libraries. It's not that Red Hat doesn't like third-party drivers, they don't want to be their support staff.

Fonts Red Hat has made a great effort to make Anti-Aliased fonts available throughout the X applications, no small feat in GNOME and *GTK*. The fonts in *Mozilla* and *OpenOffice.org* look beautiful. In an effort to support internationalisation, Red Hat uses UTF8 character set encoding, causing many problems with applications that don't play well with it (and those apps get deprecated by Red Hat).

Under the Bonnet The big news is the switch to *GCC 3.2* and the inclusion of *glibc 2.3* and an updated *libstdc++*. Kernel modules compiled with older kernels will not work, but might if forced to load. Experimental prelink support is included which allows binaries to start faster and occupy less memory (try **prelink -avm**); it seems to be working well for us. Perl 5.8, *Apache 2.0*, *Mozilla 1.01* are included. DMA for CDROM drives is turned off by default and the way to turn it on is different now: add the line **options ide-cd dma=1** to /etc/modules.conf. The much ballyhooed Access Control Lists (a finer-grained support for permissions) seen in the beta kernels has been removed because it was found to be unstable.

The RPM functions for version 4.01 have been split into two packages: *rpm* and *rpmbuild*. The former does everything the old *rpm* did except build and create RPM packages, the latter now handles that function. RPM can now verify digital signatures when installing packages and now can suggest packages to resolve dependency problems. The GUI for RPM is called the *Package Management Tool*; it's nice, but not yet as full-featured as *Synaptic* or *URPMI*.

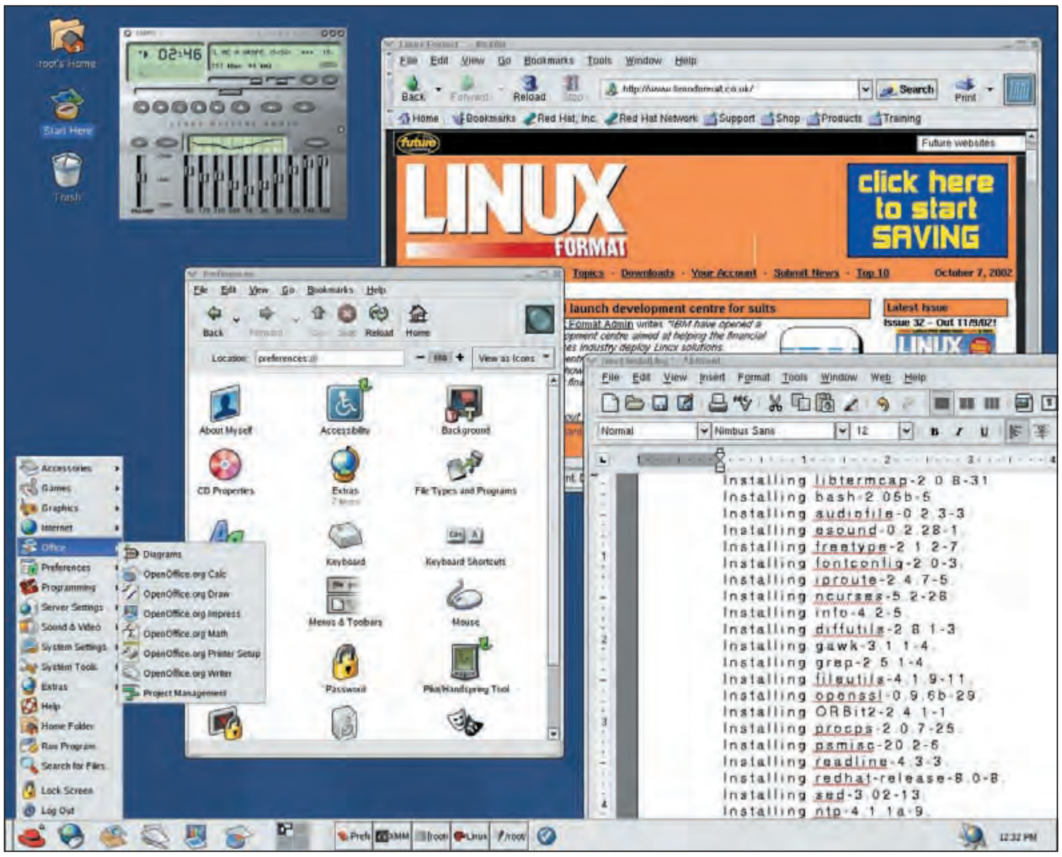

```
Installing libtermcap-2.0.8-31
Installing bash-2.05b-5
Installing audiofile-0.2.3-3
Installing esound-0.2.28-1
Installing freetype-2.1.2-7
Installing fontconfig-2.0-3
Installing iproute-2.4.7-5
Installing ncurses-5.2-28
Installing info-4.2-5
Installing diffutils-2.8.1-3
Installing gawk-3.1.1-4
Installing grep-2.5.1-4
Installing fileutils-4.1.9-11
Installing openssl-0.9.6b-29
Installing ORBit2-2.4.1-1
Installing procps-2.0.7-25
Installing psmisc-20.2-6
Installing readline-4.3-3
Installing redhat-release-8.0-8
Installing sed-3.02-13
Installing ntp-4.1.1a-9
```

The default desktop is GNOME2. Click on "Start Here" to easily access the Preferences (we agree that makes no sense). Just look at those beautiful icons!

Since this distro is now targeted at businesses, multimedia support is slight. No MP3 or MPEG support is provided and the provided apps are crippled in this regard. You have: *XMMS*, *gnome-cd*, *Grip*, *Kaboodle*, *Noatun*, *Timidity*, *KMid*, *Kmidi*, *Gtoaster*, *KonCD,* and *xcdroast*.

OpenOffice.org is the centrepiece of the Red Hat office suite, which includes a few apps from GNOME *Office*. *AbiWord* is provided, as is most of the KDE *Office Suite*.

Evolution is the default mail client (*Balsa*, *KMail*, *pine*, *mutt* and *mail* are provided as well). The default browser is *Mozilla* (it provides its IRC and Mail functions). *Konqueror* and *Galeon* makes an appearance.

Third Party Problems: *Win4Lin* doesn't like some of the kernel patches used. Codeweaver's *Crossover* doesn't get along with the menus. *VMware* also has problems. These companies all promise fixes for their products. nVIDIA "fixed" their driver incompatibility by suppressing the error messages. Many third-party apps will likely have problems with UTF8 and fonts. There may be some compatibility problems with the shared libraries being compiled with *GCC 3.2*; statically compiled apps should have no difficulty. The Sun JDK 1.3/1.4 appear to work. The Macromedia Flash plug-in is reported to work. *Acrobat* needs a little fix mentioned in the Release Notes (well worth reading, by the way). *StarOffice 6* needs an update patch.

Conclusion

We found the installation uneventful and the setup and configuration flawless. Few problems have been reported from early users, the result of a longer beta test program. Once third-party apps catch up, 8.0 will undoubtedly have a significant impact on the Linux world.

In this outing, Red Hat displays great technical excellence with far fewer problems associated with a x.0 release than we've ever seen. While their corporate orientation is executed well, the power user will be disappointed with the default installation. Fortunately, a few well-chosen package upgrades will make the power user experience enjoyable. Red Hat has made a statement with this release; it will be interesting to see who listens. LXF

A Hat of a Different Colour

Open Source improvements

Kevin Fenzi employs the "build a better mousetrap" approach to customising Red Hat. Found at www.tummy.com/krud, Kevin's Red Hat Uber Distribution (known as KRUD) provides a paid subscription service where they supply a customised version of Red Hat every month that seamlessly includes all the errata and updates as well as a few (195 at last count) applications Kevin feels Red Hat left out of the recipe.

As well as their own software updating tool, *krud2date*. They plan to began shipping their iteration of Red Hat 8.0 in November and offer free support, paid support and mailing lists. Notable is a single-CD server-only version, KRUDserver, that strips out the unneeded RPMs that include X and any workstation tools. We have looked at the KRUD CD set before and have found the enhancements worthwhile and the service excellent.

LINUX Format VERDICT

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

As a power user desktop, too many tweaks are required and too many replacement applications need to be installed. Users are forced to use feature-impaired tools and software choices are limited. As a business desktop however, it excels.

LINUX Format **RATING**

7/10

WINDOWS ON LINUX

Win4Lin Terminal Server 2.0

Want to get rid of your Windows desktops? **Chris Denton** looks at some software that can help.

Server version of *Win4Lin 4.0* – enables multiple users to connect to their own personal Windows installation over a network.

- **DEVELOPER** NeTraverse
- **WEB** www.netraverse.com
- **PRICE** See box

Previously called *NeTraverse Server Standard Edition*, *Win4Lin Terminal Server* has been renamed to position it more closely with the celebrated Windows emulation software from which it is derived. *Win4Lin Terminal Server* (*WTS*) is primarily aimed at small to medium businesses, although it may be suitable for enterprise environments in some circumstances. The basic concept is that multiple copies of MS Windows are installed onto a Linux server using *Win4Lin* emulation. These can then be accessed concurrently by client boxes running X, or suitable redisplay software.

This potentially means that Linux can become the desktop OS – with any necessary but Windows-only software run from *WTS* sessions, rather than physical machines. However, before anybody gets too excited, it should be pointed out that there are some fairly big limitations to setting things up in this way. First of all, because *WTS* is built on *Win4Lin 4.0*, it shares the same restrictions as that program. Windows 95,98 and ME are supported, but a lot of potential customers will want to use NT or 2000 and these are not compatible – nor is XP. *Win4Lin* doesn't work with *DirectX*, and that means that a lot of apps, including practically all games, are similarly no go. Some USB devices (keyboards, mice) will be okay but most are not going to get recognised. If any of this is going to present a problem then *WTS* is not for you.

On the other hand, there are many positive features of *Win4Lin* that make *WTS* surprisingly alluring. A decent list of applications that have been tested and are supported includes *Internet Explorer* (or *Netscape* if preferred) MS *Office*, Adobe *Photoshop*, *Dreamweaver* and Norton *AntiVirus*. Furthermore, the most recent versions of these apps should all be okay, in addition to any legacy stuff you might have hanging around. Since programs are added and removed at a Windows level these tasks work just as they do on a normal MS box, and there is no desperate struggle to install any software such as can be experienced with *Wine*. In fact, there is a good chance that any app you might require for business use will install and function perfectly. Basic needs like printing, networking, CDROM use and file-sharing are all seamlessly catered for.

Web Install

Before you attempt the *WTS* install it is necessary to give some thought as to how your system is going to be set up. It requires some kernel modifications – potentially involving quite a bit of effort. Fortunately, the NeTraverse graphical installer will automatically install a pre-compiled kernel, but only if you are using one of a few supported distros. These are Caldera/SCO OpenLinux, Connectiva, Mandrake, Red Hat and SuSE. In addition, only the 2.4.x kernel is supported by this version of *WTS*. Systems with 2.2.x will only work with previous releases. If you don't want to run any of the specified distros but have a 2.4.x kernel you will be able to try manually building a working kernel as instructions and code are available on the website, although not yet for *WTS 2.0*. If you want to go down that route then good luck to you, but for the purposes of this review I installed a fresh Red Hat 7.3 system and used that.

The *WTS* software is available on CD out of a box, but it is at least as easy to perform the whole installation from the web. The first thing to do is download the manual, which is a 1.3MB PDF file, and in the main extremely clear and well-written. After reading this doc, or at least the introduction and installation chapters, you should log in to your account on the NeTraverse website. You should get full instructions on this bit when you hand over your money, but suffice to say if all goes smoothly you will be presented with the option of downloading the 2.4MB tarred and gzipped *Win4Lin Terminal Server 2.0* installer. Once unzipped and extracted you will end up with an executable called *wts-install*, and kicking this off starts the graphical installer.

Licensing Costs

5 users – £414
10 users – £637
25 users – £1436
additional user – £51
All prices exclude V.A.T.

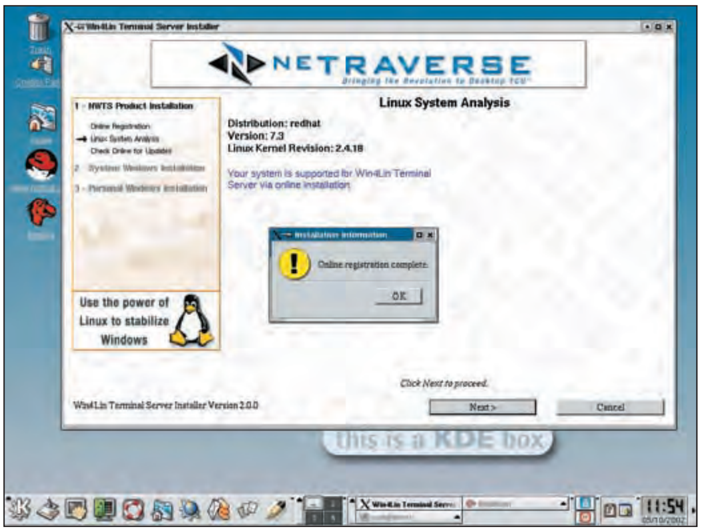
Win4Lin Terminal Server Installer
NETRAVERSE
Bringing the Revolution to Desktop TCO™
1 - WTS Product Installation
Online Registration
Linux System Analysis
Check Online for Updates
2 - System Windows Installation
3 - Personal Windows Installation
Use the power of Linux to stabilize Windows
Installation Information
OK
Next >
Cancel
this is a KDE box
Win4Lin Terminal Serv
11:54
05/10/2002

Linux System Analysis

Distribution: redhat
Version: 7.3
Linux Kernel Revision: 2.4.18

Your system is supported for Win4Lin Terminal Server via online installation

Online registration complete.

Click Next to proceed.

Win4Lin Terminal Server Installer Version 2.0.0

Finding a supported distro makes the *WTS* installer very happy.

Hardware Requirements

10 Concurrent Users:
866MHz Pentium III or equivalent processor
768MB Memory
30GB Disk space on SCSI/RAID
25 Concurrent Users:
866MHz Pentium III or equivalent processor
2GB Memory
60GB Disk space on SCSI/RAID

Once the GUI fires up, it will ask for some registration details before taking you through the installation process. At this point you will come to grief if you are connecting to the Internet in any way other than directly (*e.g.* using an HTTP proxy) because the installer doesn't seem to recognise them, or if it does it certainly can't make use of one. Still, not to worry, as long as you can meet this condition, and NeTraverse have actually heard of you, which they presumably will have done if you've got this far, the registration will take a minute or two and then the real business can begin.

There are three distinct stages to an initial install of *WTS*. Firstly, the Linux software is installed. Then what is known as a "System-wide" Windows installation is performed as root. This in effect copies the contents of your Windows OS CD onto a local filesystem. Finally, a "Personal" Windows installation takes place. This is actually the standard Windows install but run from within *WTS* so as to keep the guest OS happy. This step should be run from an unprivileged user account, and must then be repeated for each instance of Windows you need. If that last bit sounds rather arduous then you'll be glad to know that standard profiles, called recipes, can be implemented to save going through the whole Windows installation procedure every time a new user is required.

Each part of the install process seems to work extremely quickly and

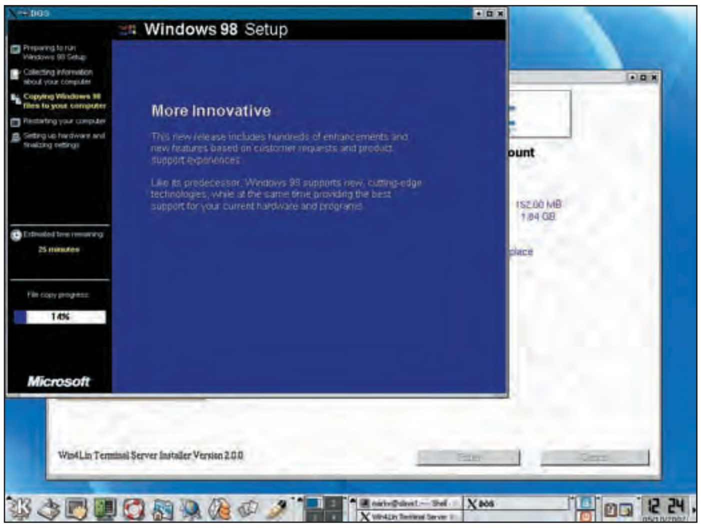

Windows 98 Setup

More Innovative

This new release includes hundreds of enhancements and new features based on customer requests and product support experiences.

Like its predecessor, Windows 98 supports new, cutting-edge technologies, while at the same time providing the best support for your current hardware and programs.

Estimated time remaining: 25 minutes

File copy progress: 14%

Win4Lin Terminal Server Installer Version 2.0.0

Preparing to run Windows 98 Setup
Collecting information about your computer
Copying Windows 98 files to your computer
Restarting your computer
Setting up hardware and finalizing settings
Microsoft
152.00 MB
1.84 GB

Fortunately, the time remaining estimate is wrong by about 22 minutes.

efficiently. The first section is potentially a nightmare as the kernel and bootloader are both modified, but in actual fact it was a breeze. The second stage takes practically no time at all and the third is basically an MS wizard, so won't tax a seasoned sysadmin unduly. However, there is an extra step where you must decide whether to use either WinSock or VNET as your networking option. NeTraverse suggest the former as a default but actually I would caution anyone against it. With WinSock you don't get access to *Network Neighbourhood* or *Exchange*, but with VNET you do and all you need in the way of extra configuration is a spare IP address and some network info.

Open Windows

The proof that everything has been setup comes when Windows boots for the first time. This happens with creditable swiftness, and seeing it for the first time can be pretty startling. The performance on my test machine was excellent, much faster than if it had been running natively. I was able to logon to my NT domain without any problems and got unreasonably excited upon discovering that the Windows key was supported. A special driver is provided to utilise the server's local Linux printer, but I had no problems in adding a Windows one from the network, which seemed more logical.

Bringing up *Explorer* reveals that the partitioning is a bit different. The C: drive is known as "~/win", there is a D: drive for local storage labelled "~/mydata", and also a useful J: drive that's for filesharing between the *WTS* clients. The A: drive for floppy and N: drive for CDROM give you full control over the corresponding resources on the server. Choosing VNET pays off here because the whole *Network Neighbourhood* gets displayed. Access to Windows and *Samba* shares works just as it should.

Connecting to a Windows session is handled by X on the local machine. Obviously, this performs well on the server itself, but normally users would be connecting from client PCs on the LAN. The documentation gives very good instructions on doing this using *SSH*, or if that's unavailable, *Telnet*. This is all fine as far as it goes, and would probably be suitable in most situations. It does however, rather assume that all the users are running Linux as their primary OS since X must be running on the client machine.

The *WTS* docs do not include instructions on connecting to the server with redisplay technology, although it does mention that it is possible. I contacted NeTraverse about this, and asked them how exactly connection without X could be achieved? They came back with two choices. If you want to do this at no extra expense, VNC can be used. VNC is not officially supported, but nor is it discouraged and indeed it seems to function well enough, if a bit inelegantly for this purpose.

With a bit more money to spend, Tarantella Enterprise software can be purchased and configured to deliver *WTS* sessions over a Java-enabled browser. Tarantella is a versatile product and it's even supported by NeTraverse if you buy it through them. However, going down this road could effectively double your outlay, so it's only for those with deep pockets. It seems to me that NeTraverse are taking a bit of a risk by pushing Tarantella at all, considering in some situations it could well be used instead of their product rather than in combination with it.

Even without any add-ons, *WTS* is far from cheap. In addition to the cost of the NeTraverse licensing, you still need all the relevant MS licences, a fairly hefty server to run it all on and a decent network infrastructure to deliver the sessions. Faced with the choice between the bill for this or keeping a number of Windows desktops and hoary old hubs going, many managers are not going to elect for the latter. This is a shame because *Win4Lin Terminal Server* is a fundamentally sound piece of software. If they ever get round some of its more glaring restrictions, it has the potential to make a very large impact indeed. LXF

LINUX Format VERDICT

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

Not particularly economical and there are some big limitations that will put off a lot of people. Nevertheless, if you think you can use it, you probably should.

LINUX Format **RATING** 8/10

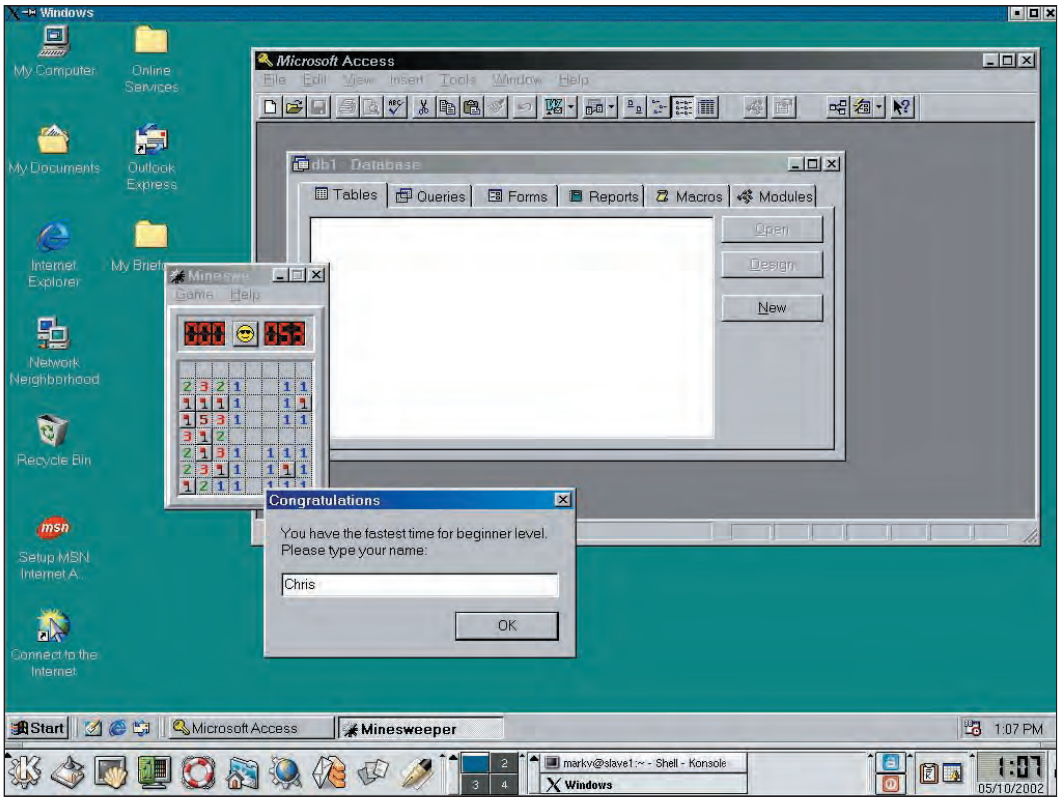

Congratulations

You have the fastest time for beginner level. Please type your name:

Chris

OK

Here we see arguably Microsoft's most popular ever application. Behind it is MS *Access*.

FIREWALL/VPN APPLIANCE

Snap Gear PRO+

The best things come in small packages. Richard Drummond asks whether that includes firewalls, too.

Firewall appliance using embedded Linux. Rivals include Lightning Multicom Speedsurf.

- **MANUFACTURER** Snap Gear
- **WEB** www.snapgear.com
- **PRICE** $749

When it comes to choosing a firewall, there are two principal types of solution: software, which will install on any PC – such as SmoothWall, reviewed on page 30 – or a dedicated appliance, such as Snap Gear's range of devices. The former has the advantage of flexibility (it's running on a general purpose computer, after all), the dedicated firewall device can be smaller, quieter, consume less power and be easier to install and maintain.

The Snap Gear PRO+ is the big brother of the LITE+ device we reviewed in *LXF 24*, targeted at larger networks. Snap Gear firmware has seen many significant revisions since our review.

The Snap Gear PRO+ is driven by an embedded AMD x86, aided by a dedicated cryptographic coprocessor. This extra processing power offers better throughput, especially on VPN links. The PRO+ offers two 10/100 Mbit Ethernet ports, for internal and external networks, a single standard serial port, and a built-in analogue modem. It can thus make use of a wide variety of devices for connection to the Internet, including cable modems, ADSL, external modems or you can use the internal modem. The modem can also be configured for dial-in. A useful feature is the ability to fail-over from a broadband connection to a modem link to maintain availability. That is, if your broadband network goes down, the the PRO+ can automatically establish a dial-up link as a back-up.

Embedded Linux

This entirely solid state device, based on a 2.4.17 kernel, is flash-upgradable in the field with firmware updates from the Snap Gear web site. The software offers an *iptables* firewall, NAT, DHCP server, DNS proxy, web proxy, VPN functionality (supporting both Ipsec and PPTP protocols) and intrusion detection. The web proxy optionally features URL-blocking and content-filtering, but this requires a licence purchase from Snap Gear (prices start at $49 for 5 users).

Like the LITE+, the PRO+ is easy to install. Simply plug it in to your network and run the set-up wizard from a Windows host on the network. Further config can then be performed with the PRO+'s web-based admin console. If you don't have Windows, the PRO+ will get it's IP number from a DHCP server on your network, and then you can log in via the web-based interface and perform the rest of the config.

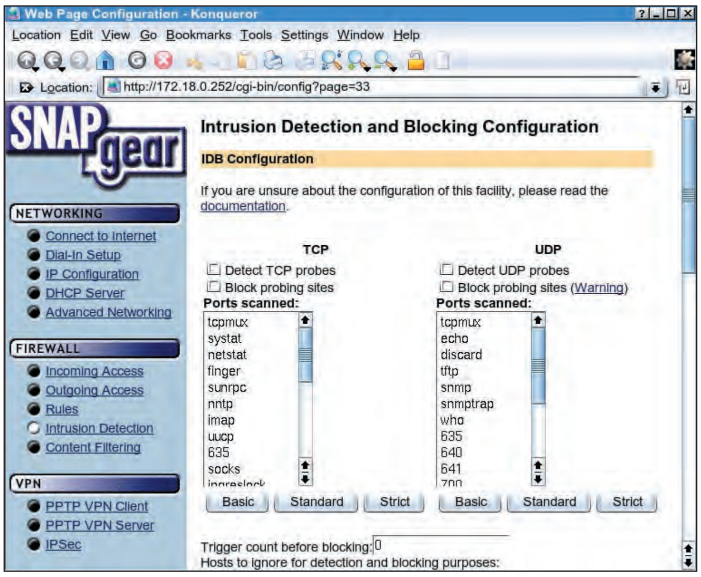

Location: http://172.18.0.252/cgi-bin/config?page=33

Intrusion Detection and Blocking Configuration

IDB Configuration

If you are unsure about the configuration of this facility, please read the documentation.

TCP
□ Detect TCP probes
□ Block probing sites
Ports scanned:
tcpmux
systat
netstat
finger
sunrpc
nntp
imap
uucp
635
socks
ingreslock

UDP
□ Detect UDP probes
□ Block probing sites (Warning)
Ports scanned:
tcpmux
echo
discard
tftp
snmp
snmptrap
who
635
640
641
700

Trigger count before blocking: 0
Hosts to ignore for detection and blocking purposes:

The web-based interface is simply to use, yet offers a surprising amount of control over the PRO+'s more advanced features.

The web-based console makes admin and monitoring of the Snap Gear device simple. From here you can config the network, the firewall, any services that are running and also monitor the device. It's mostly straightforward to use, although the memory constraints mean there is little online help. The printed docs supplied with the device are minimal, but fuller and more useful docs are supplied as a PDF on the companion CDROM. The web-based interface strikes a good balance between making common tasks quick and simple to perform, while still giving you the power and flexibility to do more complex tasks. For example, you may add or remove firewall rules by specifying in the web interface which services are to be allowed to and from which IP ranges. However, the interface also gives you the ability to directly add commands to the device's *iptables* script or replace it entirely. Similarly, you can edit various key config files directly from a text-dialog within the interface.

The PRO+ has limited features for monitoring. You can view system logs and various diagnostic data, but the device provides no traffic accounting. There's no alarm or support for mailing of alerts, either. With limited storage space (16MB of RAM), storage of logs could prove to be a problem, external logging is supported via *syslog*.

The Snap Gear PRO+ is an attractive proposition for medium enterprises, who don't want the hassle of maintaining a PC-based firewall. It offers a surprising amount of punch in a small footprint and compares favourably on price with its competitors and with many software-based solutions, especially when you consider, apart from the content-filter, there are no extra per-user licenses to be bought. LXF

LINUX Format VERDICT

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

A powerful firewall/VPN solution that's easy to install and maintain, yet flexible enough for most needs.

LINUX Format **RATING** 8/10

FIREWALL DISTRO

SmoothWall Corporate Server 2.0

With the SOHO market sown up, this popular firewall looks for bigger fish to fry. Richard Drummond investigates.

Competitors include Astaro Security Linux and SecurePoint.

- **DEVELOPER** SmoothWall Ltd
- **WEB** www.smoothwall.co.uk
- **PRICE** £120+VAT.

When first released in July 2000, SmoothWall revolutionised the entry-level firewall market by giving users the ability to turn a redundant, low-spec PC into an easy-to-set-up, secure, Linux-based firewall. The original SmoothWall was free software, but the creators soon recognised the commercial value of their product and set up a company to market a supported version.

The new SmoothWall Corporate Server 2.0, the product on test here, is a whole lot more than just a retail version of the current SmoothWall GPL release, however. It has been expanded and re-targeted at the medium enterprise market. It features more flexible networking tools, support for SCSI disks for the first time, a revamped admin interface, and a new modular and scalable architecture. Corporate Server can be extended with add-on services – such as more powerful VPN functionality and content-filtering – and hence you only pay for services that you need. Despite all of this, it remains as easy to use as ever.

Smooth installation

One of SmoothWall's big selling points is its ease of installation. Simply boot from the install disc, answer a few questions, reboot your machine and you have a functional and secure firewall. By default it sets up the firewall rules so that all incoming traffic is blocked and all outgoing traffic is allowed. This policy will work for most, so most users will find that little configuration is required – unlike with some firewall products.

SmoothWall is simplest to install from the bootable CD. If your machine can't boot from CD – this is designed to work on redundant Pcs, after all – images are supplied on the CD to create boot floppies. You can even install on a machine *sans* CDROM drive, via the network.

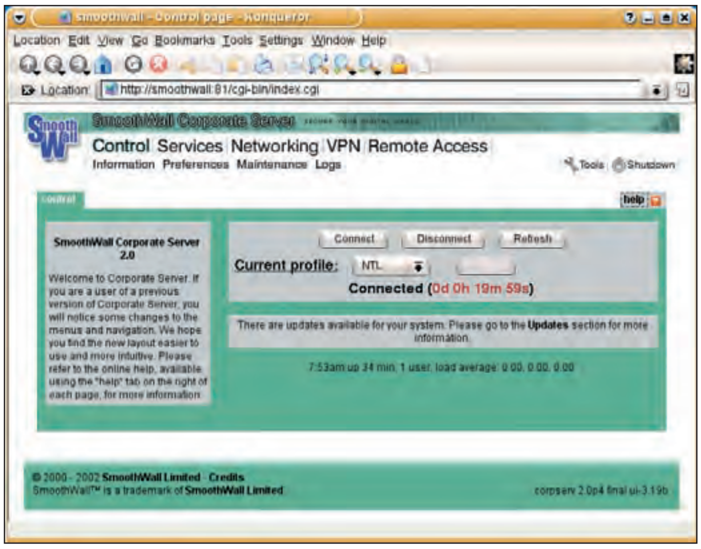
smoothwall - Control page - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help
Location: http://smoothwall:81/cgi-bin/index.cgi
SmoothWall Corporate Server
Control Services Networking VPN Remote Access
Information Preferences Maintenance Logs
Tools Shutdown
control
help
Connect
Disconnect
Refresh

SmoothWall Corporate Server 2.0

Welcome to Corporate Server. If you are a user of a previous version of Corporate Server, you will notice some changes to the menus and navigation. We hope you find the new layout easier to use and more intuitive. Please refer to the online help, available using the "help" tab on the right of each page, for more information.

Current profile: NTL

Connected (0d 0h 19m 59s)

There are updates available for your system. Please go to the **Updates** section for more information.

7:53am up 34 min, 1 user, load average: 0.00, 0.00, 0.00

© 2000 - 2002 SmoothWall Limited Credits
SmoothWall™ is a trademark of SmoothWall Limited

corpserv 2.0p4 final ui-3.19b

Smoothwall Corporate Server has a new, professional-looking administration interface.

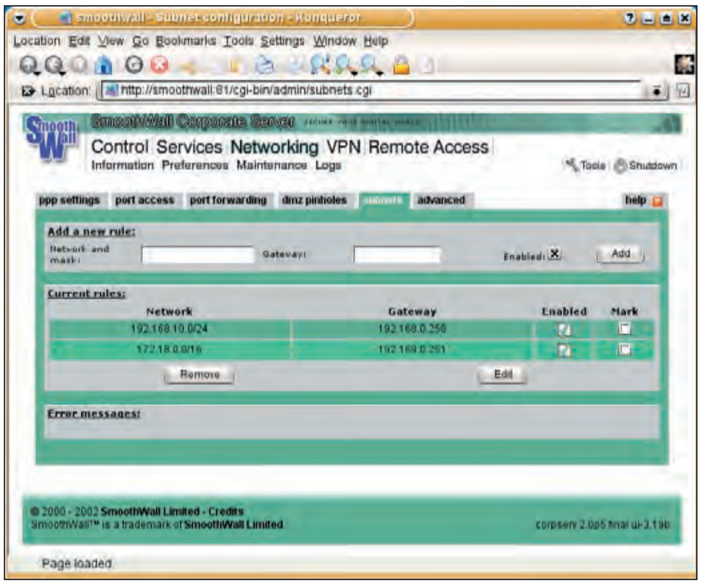
smoothwall - Subnet configuration - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help
Location: http://smoothwall:81/cgi-bin/admin/subnets.cgi
SmoothWall Corporate Server
Control Services Networking VPN Remote Access
Information Preferences Maintenance Logs
Tools Shutdown
ppp settings
port access
port forwarding
dmz pinholes
subnets
advanced
help
Add
Remove
Edit

Add a new rule:

Network and mask: Gateway: Enabled: ☒

Current rules:

Network	Gateway	Enabled	Mark
192.168.10.0/24	192.168.0.250	☑	□
172.18.0.0/16	192.168.0.251	☑	□

Error messages:

© 2000 - 2002 SmoothWall Limited - Credits
SmoothWall™ is a trademark of SmoothWall Limited

corpserv 2.0p5 final ui-3.19b

Page loaded

It supports more complex networking, such as the ability to define multiple subnets within your local network.

The installer itself is based on the familiar Red Hat text-console dialog-based system, although this has been vastly simplified and asks you for a minimum of info. After picking the install type – from CD or network – it prompts you that the hard drive is to be repartitioned. SmoothWall takes over an entire disk, deleting previous contents. It creates root, boot and swap partitions and and a /var/log partition for storing log files. SmoothWall now uses the *ext3* filesystem throughout, and this – and the separation of log files – should ensure that if your SmoothWall box does go down while writing to the disc, damage will be a minimum and quick and easy to recover.

Next up is network config. Like the previous incarnations of SmoothWall, Corporate Server supports three network zones: Green, or the local network; Red, the Internet; and Orange, the DeMilitarized Zone. The first step is to set-up your Green interface – which must be an Ethernet adaptor. SmoothWall can probe for a NIC automatically, or you can manually select a kernel driver to load. You just then need to specify the IP address and netmask for the interface.

Support for a fairly comprehensive range of NICs is included but, perhaps, this is the only area of SmoothWall that the user will notice and care that SmoothWall is actually based on a 2.2 kernel. The reason: driver support. For example, my test machine had a recent Acton Technology Tulip-based card which failed to be detected by SmoothWall and is unsupported by SmoothWall's tulip driver (this card works in 2.4.x). To be fair, the vast majority of NICs will work – and, certainly, all the others I tried did. However, the info about precisely what hardware Corporate Server supports isn't available on the SmoothWall website, although it is included on the documentation on the install disc (which is excellent, by the way).

Configuration of the Red interface is similar if you choose this to be a NIC rather than a dial-up device such as an analogue modem or ISDN adapter – although here DHCP is supported for determining IP address, etc. PPP over Ethernet is included, so you can use a wide variety of routers, cable modems, etc,. on your Red interface. If

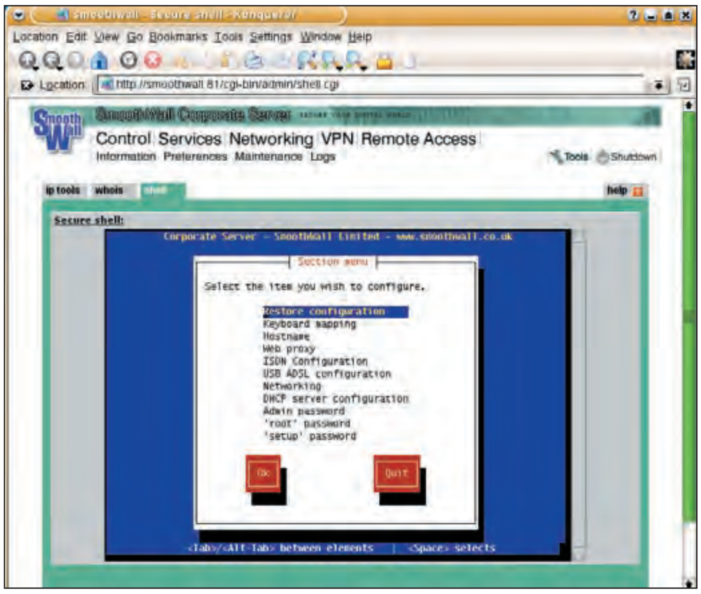
SmoothWall - Secure shell - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help
Location http://smoothwall:81/cgi-bin/admin/shell.cgi
SmoothWall Corporate Server
Control Services Networking VPN Remote Access
Information Preferences Maintenance Logs
Tools Shutdown
ip tools whois shell
help
Secure shell:

```
Corporate Server - SmoothWall Limited - www.smoothwall.co.uk

Section menu

Select the item you wish to configure.

Restore configuration
Keyboard mapping
Hostname
Web proxy
ISDN Configuration
USB ADSL configuration
Networking
DHCP server configuration
Admin password
'root' password
'setup' password

<Ok>    <Quit>

<Tab>/<Alt-Tab> between elements   |   <Space> selects
```

The web-based interface now gives access to ping and traceroute tools, and a Java-based remote shell, shown here running the setup program.

you are using ISDN, SmoothWall can probe for your ISDN adapter – and drivers are included for an impressive range of internal and external devices – or, again, you load a driver manually. The other option for the Red interface is the Alcatel USB ADSL modem. The connection settings for any kind of dial-up device must be configured from the web interface post-install.

The rest of the installation is fairly routine. You choose your keyboard layout, hostname, and passwords. Usefully, you can also configure SmoothWall's DHCP server here. Another handy feature is that the installer gives you the option of restoring your firewall settings from a back-up floppy, and this will save you a lot of work should you need to re-install.

Smooth interface

When you boot up your newly installed SmoothWall machine, you will have a working firewall. A limited amount of config can be performed locally, via a dialog-based set-up tool, but most admin will be performed via SmoothWall's web interface. This can be accessed via a browser pointed at port 81 and a secure connection is available at port 445. The web interface has been redesigned in Corporate Server, with a more business-like look. The interface is split into eight pages according to task (and most of these are further sub-divided into 'tab panes'), so navigation is easy. Online-help is available on each page.

The networking page is where you can set-up your security policy and configure your dial-up connection. Most of this will be familiar from previous SmoothWall releases. The 'Port access' section lets you open up specific ports for incoming traffic to your local network or DMZ; 'Port forwarding' lets you forward incoming connections from a port on the Red interface to a port on a specified host within your DMZ, useful when you only wish to expose a particular service on a host in the DMZ; and 'DMZ pinholes' allows you to open up connections from the DMZ to the Green network, useful if your web server needs access to a database in the local network, say. New in this section is support for multiple sub-nets within the Green network. Here you can define the address ranges of any sub-nets and the IP address of the router via which SmoothWall can access them. This features allows SmoothWall to protect much more complex local networks.

The services section should also prove familiar. Here you can configure SmoothWall's DHCP server, web proxy and DNS cache. The DHCP section is straightforward and lets you add blocks of IP addresses for dynamic assignment and to configure static IPs based on MAC address. The web proxy page is for simple config of *Squid*. Interesting features here include forwarding to a remote proxy and 'transparent' operation. The latter enables caching on all web connection via the firewall, not just those via the proxy port, and eases the setting up of client machines. Static DNS is where you add static IP numbers and their host names to SmoothWall's DNS cache, and Dynamic DNS provides automatic support for various dynamic DNS services – such as no-ip.com and dyndns.org – which let you assign a world-visible host and domain name to your firewall when it's on a dynamic IP address. The final page is for the intrusion detection system, *Snort*, and this just turns it on or off.

After your firewall is configured, the pages that most users will spend their time viewing are in the Information and Logs sections. The former provides status info, and detailed traffic graphs and statistics. These display the traffic on each of the interfaces and VPN traffic separately. The logs section lets you monitor the behaviour and security of your firewall in more detail. As well as a general activity log, this provides access to the web proxy, firewall and intrusion detection logs. The firewall log displays IP packets dropped by the firewall (you cannot configure which packets are logged), and provides a look-up function to do a reverse look-up on the source address of any dropped packets. All the various logs can be saved out as files for archival or analysis purposes, but SmoothWall doesn't support the automated external storage of logs, emailing of logs, or any kind of alarm system. The upshot is, you do have to monitor these pages regularly to maintain the health of your firewall.

The other aspect of maintenance is ensuring your software is up-to-date. SmoothWall checks online for available updates and these can be viewed in the maintenance section of the web interface. Using this interface, you can download updates to your admin host, and then upload them to your firewall where they will be installed. Simple enough, if a rather roundabout way of doing it. Add-on modules can be installed in a similar manner.

Add on modules

Need you firewall to do more? Then just ask.

SmoothHost
Price: £100+VAT
Extends SmoothWall to support multiple IP addresses on the external interface, which is useful, for example, when you need to give each web server in the DMZ its own IP address.

SmoothGuardian
Price: £200+VAT
Provides URL blocking, regular expression-based URL blocking, MIME-type blocking and advanced content-filtering.

SmoothTunnel
Price: £200 + VAT
Provides a VPN gateway supporting multiple VPN tunnels and more extensive features, such as certificate management and Road Warrior Support.

SmoothNode
Price: £50 + VAT (5 connections)
Allows SmoothWall to connect to a VPN gateway such as *SmoothTunnel*.

SmoothRule
Price: £100 + VAT
Provides policy-based control of outgoing Internet traffic to regulate what service network users can access.

Conclusion

SmoothWall Corporate Server builds on the strengths of the free SmoothWall to create a more scalable product. The increased hardware support, the more flexible networking tools – such as the support for multiple subnets – and the module, extensible architecture will certainly widen its appeal to larger businesses. I have two main worries. One is that SmoothWall is based on a 2.2 kernel and so support for newer hardware isn't as good as if it had been based on a 2.4 kernel (I know it was meant for redundant hardware, but scaling up to larger networks requires more horsepower, after all.) My second concern is the lack of any automated tools for firewall monitoring. It should really be able to email intrusion detection alerts to an admin, rather than requiring somebody to constantly monitor log files. Having said all this, SmoothWall Corporate Server does what it sets out to do and, at this price, is an outrageously good bargain. LXF

LINUX Format VERDICT

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

As simple to use as the original SmoothWall, but those extra features and scalability means that it can tackle bigger networks. A bargain.

LINUX Format **RATING**

9/10

DESKTOP DISTRO

Lycoris Desktop/LX Personal Amethyst, Update 2

Can Linux really replace Windows on the desktops of the world? Richard Drummond tests a distro that could prove it can.

Other Windows-like distros include Xandros and Lindows.

- **DEVELOPER** Lycoris
- **WEB** http://www.lycoris.com/
- **PRICE** $29.95 (Deluxe edition $39.95)

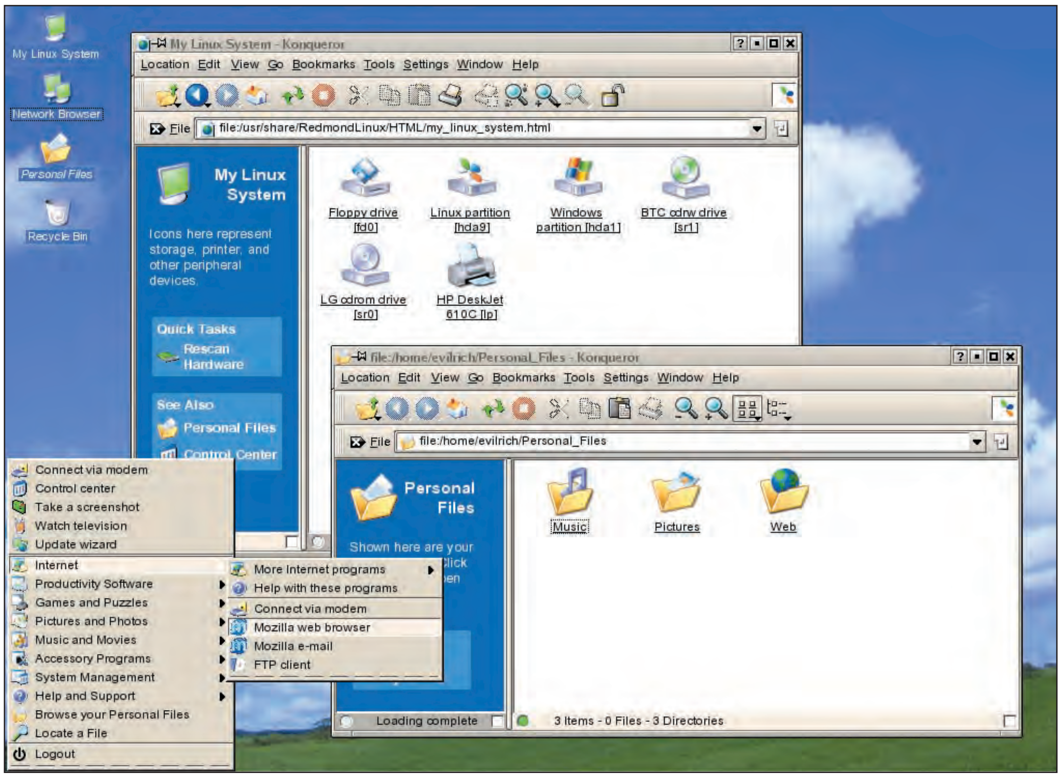

Love it or hate it, that XP like and feel is bound to be re-assuring for the Windows convert.

Pundits have been claiming for some time, each year, that this will be the year that we will see Linux attack Microsoft's dominance of the desktop market. Well, it hasn't happened yet. But, with the big boys like Sun and Red Hat at last taking an interest, and with distros like Xandros and Lindows that are friendly to the Windows user, Linux is better-placed than ever to offer a credible desktop solution in the home and the office. One of the more successful attempts at bringing Linux to Windows users to date is Lycoris's Desktop/LX (*née* Redmond Linux). When we last reviewed this product in *LXF25*, we were impressed by its easy installation, its automatic support for Windows networking and its well laid-out, easy-to-navigate desktop. Here we take a look at the latest version, known as Amethyst Update 2.

Solid base

Lycoris is based on Caldera's OpenLinux – and adds things like a simpler and streamlined installation, a custom build of the KDE 2.2 desktop that should prove familiar to the Windows user, and a carefully chosen selection of apps. It ships as a single CD or a 3-disc set (Deluxe edition). The single-disc version should suffice for most users since the second and third discs contain only source code and development tools. Lycoris may be downloaded for free, but you need to buy a licence for installation support (although free support is available from the Lycoris community website at **lycoris.org**) and to use their new web services. Existing users can get Update 2 using the built-in software updater, and this will download and install the required packages automatically. The update service is free for all and requires no registration.

Installing Update 2 from the CD is easy, since it uses the same graphical installer as last time. It still offers no advanced partitioning tools, such as partition-resizing, but it is one of the most pleasant installers to use. Hardware detection threw up no obvious problems during installation (but see the *X trauma* box for post install problems), and, perhaps, best of all for new users, there's no package selection stage. You don't have to wade through a list of strange sounding software; everything to make a full-featured desktop is installed for you automatically.

As with the previous release, the Lycoris build of KDE is tweaked to create an environment comfortable for Windows users. A lot of attention as been paid on aesthetics, and they really have done a good job on making the desktop theme look like Windows XP (whether you like that or not). Anti-aliasing of fonts is enabled by default, as is a lot of KDE eye-candy like menus that fade out. More importantly, they have overhauled the KDE menu. Programs are named by task rather than software name, making it much easier for those unfamiliar with Linux to find there way around.

Borrowing more from the Windows world, Lycoris puts three icons on your desktop in addition to the trashcan. The first of these is 'My Linux Box', which when opened brings up a browser showing various hardware attached to your system, and it acts as a launch pad for hardware-related tasks. (*E.g.*, clicking on a CD writer loads up the CD burning application *KonCD*, while clicking on a printer shows that device's printer queue.) Lycoris has built-in, automatic support for Windows networking – via *Samba* – and the 'Network Browser' icon lets you access any shared folders. Your own home directory is exported automatically. The third icon is called 'Personal Files' and is simply a link to a folder in your home directory for storing your data. This has been sub-divided into 'Music', 'Pictures' and 'Web' – but it's really only a cue for you to organise yourself, since no apps default to saving files into these folders.

The 'My Linux Box' browser has been improved in Update 2, and – besides hard drive partitions, CD drives,

X Trauma

Not so newbie-friendly

By and large, hardware detection works well in Lycoris, but it's X configuration still leaves a lot to be desired. On the three machines I tested none of them had a working 3D card upon installation. One had an nVIDIA card, which requires that you download and install the drivers separately. Lycoris do provide their own build, and these install smoothly, so this is a simple job. The other two were more problematic, however. One box I tried had an ATI Rage 128 card, which seemed to be correctly detected and configured by Lycoris, but hardware-accelerated 3D just didn't work. A quick look at the X log file revealed that the wrong version of the kernel DRM modules had been included (Lycoris uses *XFree86 4.2.0*, but the DRM modules were for *4.1.0*).

The only solution there is to rebuild the kernel, not something the target user for Lycoris is going to want to do (especially since the basic install doesn't even include *GCC*). Lycoris haven't provided an updated kernel yet.

My third machine, with a Matrox G200, caused even more grief. Despite me pointing this out last time around, Update 2 still chose to use *XFree 3.3.6* SVGA server for this card. No prospect of anti-aliased fonts or direct rendering there, and this server seemed to have lots of problems with the KDE window manager. In a moment of folly, I tried to tackle this problem as the average user might, and use Lycoris's X config tool in the Control Center to fix this. Pick the MGA driver to force it to use *4.2.0*, and restart X. Simple? No. X refused to start, because the config file was broken. A moment of grappling with *vi* (there's no *emacs*), and we were in business.

Although anti-aliasing now worked, Truetype fonts were being rendered on screen at an inch high or more. The X log revealed that X thought my display was 2 metres wide and had a resolution of 11 DPI! Quite why it believed this to be the case, I don't know, but experiment revealed the problem was due to Matrox's HAL driver. Lycoris includes Matrox's X drivers, which includes the binary-only HAL library needed for setting up certain features on some cards. I renamed this, so that it wouldn't be loaded, and my screen size was reported more sanely and fonts became readable.

and printers – it can show a lot more types of hardware, such as Zip drives, USB storage devices, cameras and TV cards. On the systems I tested, however, it failed to mention my ATAPI Zip drive, Kodak DC-3600 camera or WinTV card. Interestingly, though, my Rage 128 All-in-Wonder was reported here, and clicking that brought up *motv* for me to watch TV. Neat! (Lycoris include the GATOS ATI drivers for Xvideo support on ATI cards. See **http://gatos.sf.net/**.)

Being in control

The *KDE Control Center* in the previous release of Lycoris had been given a good make-over. As well as integrating third-party tools – for configuring network interfaces and services, adding users, etc. – the structure of the *Control Center* had been re-organised and components re-named to make navigation easier. All of that is in Update 2, but it's hidden from view. The default *Control Center* is now HTML-based, again with a nod to XP. The first page presents groups of configuration tasks, such as 'Video and Monitor' and 'Internet and Network'; clicking one of these brings up a page with links that will launch the individual configuration tools. While all of this look pretty, I'm not sure that this system is really any easier to use. It will make the Windows people feel more at home, though.

The set of admin tools is the same as last time, with two additions. The first is a basic firewall tool, which simply turns the kernel's firewall on or off (you cannot configure the rules). The second is the software installer, grandly named *IRIS*, or *Internet Rapid Installer for Software*. This provides access to an online gallery of additional software, which can be browsed and installed with a web-based interface. Use of *IRIS* requires that you have a licence, but, at the moment, this service alone isn't worth the licence fee – the selection of software currently available is poor.

I really want to like Lycoris, because it does go a long way to make Linux accessible to all. This really is a distro that I would recommend to my Grandmother. However, while from a user-interface point of view Lycoris excels, the under-pinnings are a little shaky. Lycoris need to spend as much time making sure that that their distro runs on stable foundations as they do making it pretty and easy-to-use. Perhaps they should implement some better quality control mechanisms and at least a proper bug-tracking system on their website. Another issue is that the versions of key software components are lagging behind. Okay, so Update 2 includes kernel 2.4.18 and *XFree86 4.2.0*, but other vendors are now shipping systems built with *GCC 3.2* and KDE 3.0. Lycoris may need to do some catching up.

These gripes aside, Lycoris provides a usable and comfortable desktop system, and offers excellent value for money. Update 2 – although largely a cosmetic update – brings some important software revisions and extra hardware support. Existing users will want to upgrade, and, if you are new to Linux, it's well worth a try. LXF

LINUX Format VERDICT

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

Update 2 means that Lycoris is now friendlier than ever, but they've failed to many fix hardware configuration issues, and introduced some new ones.

LINUX Format **RATING** 7/10

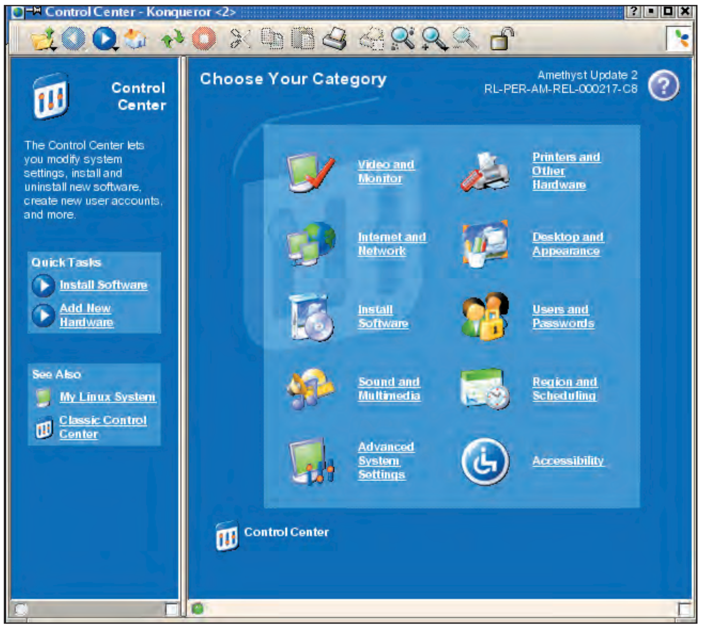

The Control Center lets you modify system settings, install and uninstall new software, create new user accounts, and more.

The new HTML-based front-end looks pretty, but KDE fans may prefer the traditional Control Center.

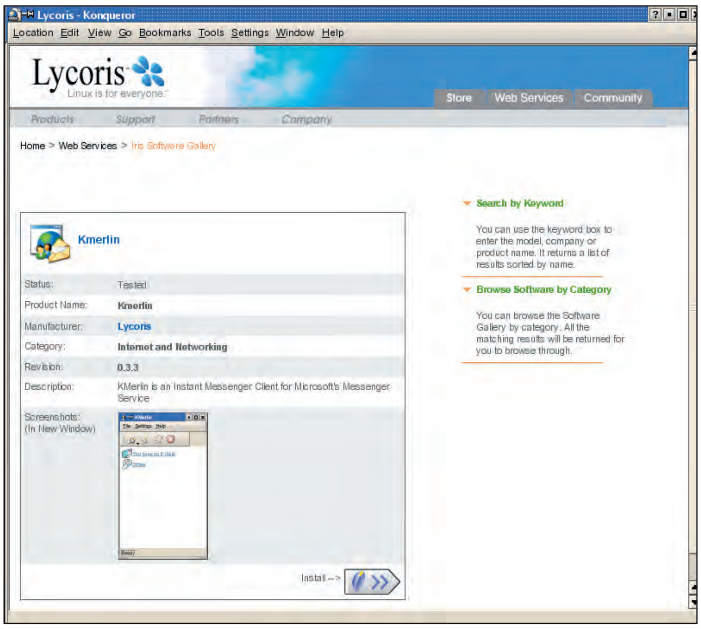

Status: Tested
Product Name: Kmerlin
Manufacturer: Lycoris
Category: Internet and Networking
Revision: 0.3.3
Description: KMerlin is an Instant Messenger Client for Microsoft's Messenger Service
Screenshots: (In New Window)

You can use the keyword box to enter the model, company or product name. It returns a list of results sorted by name.

You can browse the Software Gallery by category. All the matching results will be returned for you to browse through.

The *IRIS* software gallery is a good idea, but it needs to be stocked with more, up-to-date software.

LINUX DISTRIBUTION

SuSE Linux 8.1 Professional

The ten year history of SuSE's Linux distribution has seen some amazing developments. Has a decade of experience given SuSE something to boast about? Rob Fenwick finds out…

Also consider: Linux Mandrake is good for desktop users, while Red Hat is popular on servers and desktops alike.

- **PRICE** £59.00 (£39.00 for personal edition) inc VAT
- **SUPPLIER** SuSE Linux Ltd
- **PHONE** 020 8387 4088
- **WEB** www.suse.co.uk

cover feature

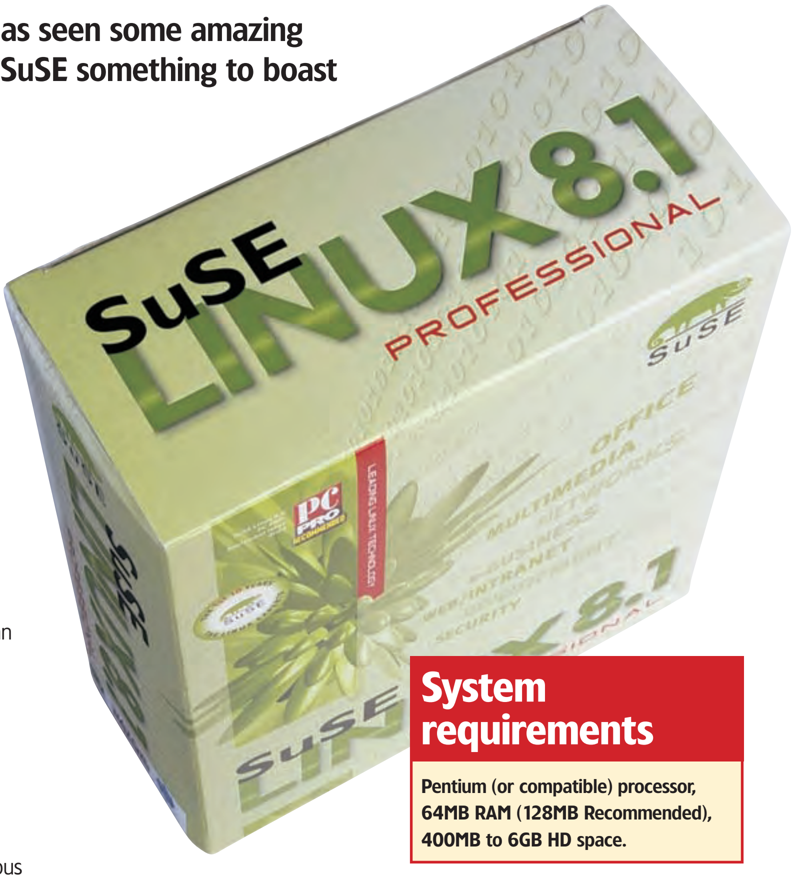

We installed SuSE 8.1 on a Compaq Presario 700 laptop. SuSE have reduced the installation to such an extent that you need to make very few clicks to install with the default settings. You can still, however, heavily customise the installation, or choose from one of numerous profiles to select your software. We were particularly glad to see a choice of screen resolutions to work in during the installation – screen resolutions less than 1024x768 do not scale well on Compaq (and countless other) laptops.

As ever, SuSE's drive partitioning is still something of an abstract process – the graphical representation of the drive layout you're building is basic, and isn't interactive. Red Hat and (particularly) Linux Mandrake, still outstrip SuSE when it comes to ease of partitioning. Our system was already adequately partitioned, however, and SuSE had no problems automatically assigning the most appropriate mount points to the various partitions – the *YaST* installer's analysis of existing partitions is extremely intelligent.

It is worth mentioning, at this stage, that the traditional installation interface that most of us are used to – a step-by-step wizard (or, if you prefer, druid) is not a feature of the SuSE installation. The installation has, to all intents and purposes, been made to look like a webpage instead. The entire installation process is laid out in front of you on one 'page', and you simply click on the sections you want to customise. While it sounds like a good idea in practice, the implementation of it may need some polishing.

We altered the setup profile to install all available binary packages, but either through software or human error, only the basic packages were actually installed. We then had to return to *YaST2* after the setup was completed to install the rest, which was a simple enough process.

Accessible to all

SuSE's installer still supports numerous braille interfaces, making SuSE one of the best choices of Linux distribution for the visually disabled. SuSE are to be commended for this – it perfectly demonstrates their commitment to make Linux accessible to all, regardless of commercial pressures.

It is in tailoring the distribution for desktop users that SuSE has really excelled. Of course, like all distros, if you want to stick to a strictly server install, SuSE has strengths there, which we will come on to, but it is the slickness of the experience for desktop users which sets SuSE apart from the myriad other distros on the market.

Many of the differences between SuSE 8.1 and 8.0 are in improvements to the admin tools. It took some time, but *YaST2's* graphical incarnation has come of age. That said, although *YaST2* has progressed dramatically, there are still parts which are only regurgitation of the plain text config files, with little interpretation. In these components, you may as well edit the file directly, as the basic interface thrown around it is perhaps not likely to be too much help – *see screenshot, left*.

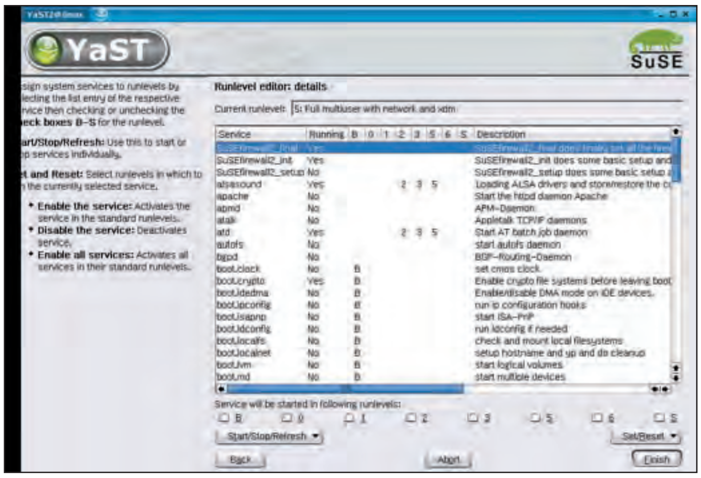

sign system services to runlevels by lecting the list entry of the respective rvice then checking or unchecking the eck boxes B–S for the runlevel.

art/Stop/Refresh: Use this to start or p services individually.

t and Reset: Select runlevels in which to n the currently selected service.

- **Enable the service:** Activates the service in the standard runlevels.
- **Disable the service:** Deactivates service.
- **Enable all services:** Activates all services in their standard runlevels.

Runlevel editor: details

Current runlevel: 5: Full multiuser with network and xdm

Service	Running	B	0	1	2	3	5	6	S	Description
SuSEfirewall2_final	Yes									SuSEfirewall2_final does finally set all the firew
SuSEfirewall2_init	Yes									SuSEfirewall2_init does some basic setup and
SuSEfirewall2_setup	No									SuSEfirewall2_setup does some basic setup a
alsasound	Yes				2	3	5			Loading ALSA drivers and store/restore the c
apache	No									Start the httpd daemon Apache
apmd	No									APM–Daemon
atalk	No									Appletalk TCP/IP daemons
atd	Yes				2	3	5			Start AT batch job daemon
autofs	No									start autofs daemon
bgpd	No									BGP–Routing–Daemon
boot.clock	No	B								set cmos clock
boot.crypto	Yes	B								Enable crypto file systems before leaving boot
boot.idedma	No	B								Enable/disable DMA mode on IDE devices.
boot.ipconfig	No	B								run ip configuration hooks
boot.isapnp	No	B								start ISA–PnP
boot.ldconfig	No	B								run ldconfig if needed
boot.localfs	No	B								check and mount local filesystems
boot.localnet	No	B								setup hostname and yp and do cleanup
boot.lvm	No	B								start logical volumes
boot.md	No	B								start multiple devices

Service will be started in following runlevels:

□ B □ 0 □ 1 □ 2 □ 3 □ 5 □ 6 □ S

The runlevel editor in YaST2 is a pretty basic regurgitation of the raw configuration file.

System requirements

Pentium (or compatible) processor, 64MB RAM (128MB Recommended), 400MB to 6GB HD space.

Key Software

What's in the box?

SuSE 8.1 was relatively up to date at the time it hit the shelves. Here are some of the key version numbers, as they are at installation time.

Kernel: 2.4.19
GNOME Desktop: 2.0.6-60
KDE: 3.0.3
Window Maker: 0.80.1
glibc: 2.2.5
XFree86: 4.2.0
Apache: 1.3.26
BIND: 9.1

SuSE have had the decorators in to KDE3, and produced a default theme which is slick, friendly and easy on the eye. Grey and blue combinations are most certainly the new beige. There

Tip Your postman

Manual delivery

If you mail order SuSE Linux 8.1, spare a thought for your postman – thanks to the manuals, it weighs a fair amount!

In addition to the hefty manuals provided by SuSE, which help to get you started and cover many of the most common problems encountered by new users, there is also an extensive on-line support database, with articles dating as far back as the time of SuSE Linux 04/95, before the version numbering system we have today.

Up until as recently as eighteen months ago, English users of SuSE's documentation could feel that they were getting a relatively poor deal, compared to their German counterparts. Poor translation and vague advice, while not being common place, were nevertheless frequent. That's changed, now.

The quality of the English language support docs and on-line knowledge base is better than it has ever been, and if this interviewer could speak German, French, Korean, or Hungarian, I would comment on the quality of those translations. For now, though, it is enough simply to say that those translations are there if you need them.

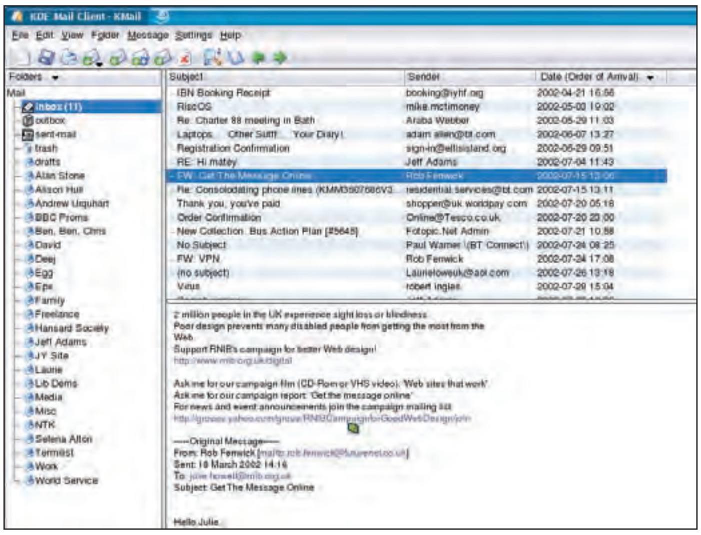

Subject	Sender	Date (Order of Arrival)
IBN Booking Receipt	booking@iyhf.org	2002-04-21 16:56
RiscOS	mike.mctimoney	2002-05-03 19:02
Re: Charter 88 meeting in Bath	Araba Webber	2002-05-29 11:03
Laptops… Other Stuff… Your Diary!	adam.allen@bt.com	2002-06-07 13:27
Registration Confirmation	sign-in@ellisisland.org	2002-06-29 09:51
RE: Hi matey	Jeff Adams	2002-07-04 11:43
FW: Get The Message Online	Rob Fenwick	2002-07-15 13:06
Re: Consolodating phone lines (KMM3507686V3	residential.services@bt.com	2002-07-15 13:11
Thank you, you've paid	shopper@uk.worldpay.com	2002-07-20 05:16
Order Confirmation	Online@Tesco.co.uk	2002-07-20 23:00
New Collection: Bus Action Plan [#5645]	Fotopic.Net Admin	2002-07-21 10:58
No Subject	Paul Warner \(BT Connect\)	2002-07-24 08:25
FW: VPN	Rob Fenwick	2002-07-24 17:08
(no subject)	Laurieloweuk@aol.com	2002-07-26 13:18
Virus	robert inglas	2002-07-29 15:04

2 million people in the UK experience sight loss or blindness
Poor design prevents many disabled people from getting the most from the Web
Support RNIB's campaign for better Web design!
http://www.rnib.org.uk/digital

Ask me for our campaign film (CD-Rom or VHS video): 'Web sites that work'
Ask me for our campaign report: 'Get the message online'
For news and event announcements join the campaign mailing list
http://groups.yahoo.com/group/RNIBCampaignforGoodWebDesign/join

-----Original Message-----
From: Rob Fenwick [mailto:rob.fenwick@futurenet.co.uk]
Sent: 18 March 2002 14:16
To: julie.howell@rnib.org.uk
Subject: Get The Message Online

Hello Julie

***KMail* is a capable and easy to use email app, and the excellent import utilities allow you to transfer for your mail from MS apps if you need to.**

aren't all that many themes provided, and so if you have unconventional tastes then you will have to resort to the numerous KDE themes sites online.

The installer automatically detects and creates mountpoints for partitions belonging to other OSs where it is able to. In addition, be you root or a regular user, you will find icons to these other volumes on your KDE desktop. As ever, for safety, NTFS partitions are set up for read only access, and it is recommended that you do not alter this.

Strength in diversity

The customisation treatment doesn't end with KDE. In previous releases, there appeared to be something of a reluctance on SuSE's part to recognise the presence of the GNOME desktop environment in their distribution. Things have moved on a little, now. The default GNOME appearance has also been customised by SuSE, and given a very corporate, professional 'feel'. Again, desktop icons to other volumes have been preconfigured, and there are easy 'shortcuts' set up to *Mozilla*, and *YaST*, but unlike in KDE, no link to the office apps. While things have undoubtedly moved on, there is still a heavy bias towards KDE.

For those who find KDE and GNOME a little too much to bear, SuSE have also included *Window Maker 0.80.1*, and have added the SuSE menu to this environment, to make all of the apps available to GNOME and KDE users also accessible to *WM* users. The Window Manager selection list on the login screen contained options to launch *mwm* and *twm*, but neither of these options successfully launched the respective Window Manager for us.

SuSE has chosen a broad selection of packages for the CD, and the result is a diverse starting point which demonstrates much of the power of Linux. We were suprised, however, to see that SuSE have already dropped the 'locate' databases (and corresponding apps: *locate*, *updatedb*, *bigram* and *code*) from their default installation. Despite the slightly cumbersome nature of this system, it is still popular with many.

Keep it safe

Historically, the Red Hat distribution has been ahead of SuSE in providing easily accessible updates to the OS – the Red Hat Network is an excellent piece of engineering, particularly for businesses. Under older versions of *YaST*, online updates were cumbersome, but *Yast Online Updates* (*YOU*) is much easier to use now than it has ever been. Keeping your system up to date with the lastest updates is essential, and SuSE have now joined the ranks in making on-line updating easy. *up2date*, Red Hat's update software is still a few steps ahead of SuSE's *YOU*, but SuSE are catching up fast – we'd be particularly keen to see *YOU* install with a KDE panel applet to proactively warn you of updates, in the same way that *up2date* integrates with GNOME.

»

Easy as 1-2-3

Adding new hardware is simple, as we discovered when we connected our USB colour deskjet printer

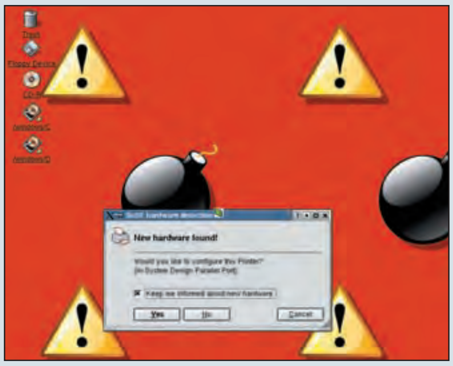

STEP 1

We plugged the printer in to the USB of the laptop, and logged in. We were immediately greeted by a pop-up box informing us that new hardware had been found. After clicking 'yes' to configure the printer, *YaST2* was launched.

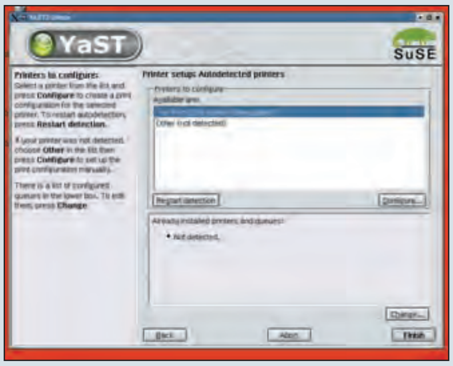

Printers to configure: Select a printer from the list and press **Configure** to create a print configuration for the selected printer. To restart autodetection, press **Restart detection.**

If your printer was not detected, choose **Other** in the list then press **Configure** to set up the print configuration manually.

There is a list of configured queues in the lower box. To edit them, press **Change**.

Printer setup: Autodetected printers

Printers to configure
Available are:
Other (not detected)

Already installed printers and queues:
- Not detected.

STEP 2

***YaST2* highlights the new hardware available to be configured, but the 'Configure' button is in an unusual place – we would normally expect to see the main action button at the very bottom-right of the screen.**

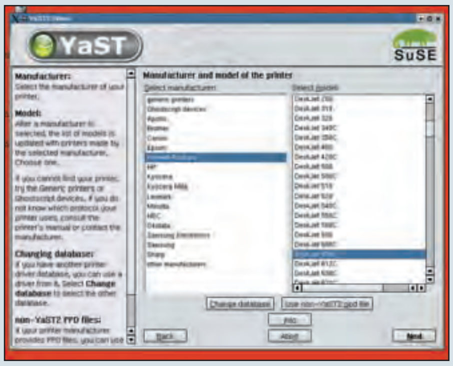

Manufacturer: Select the manufacturer of your printer.

Model: After a manufacturer is selected, the list of models is updated with printers made by the selected manufacturer. Choose one.

If you cannot find your printer, try the Generic printers or Ghostscript devices. If you do not know which protocol your printer uses, consult the printer's manual or contact the manufacturer.

Changing database: If you have another printer driver database, you can use a driver from it. Select **Change database** to select the other database.

non-YaST2 PPD files: If your printer manufacturer provides PPD files, you can use

Manufacturer and model of the printer

Select manufacturer:
generic printers
Ghostscript devices
Apollo
Brother
Canon
Epson
Hewlett-Packard
HP
Kyocera
Kyocera Mita
Lexmark
Minolta
NEC
Okidata
Samsung Electronics
Seiko
Sharp
other manufacturers

Select model:
DeskJet [illegible]

STEP 3

Select the printer in the list, and the next screen invites you to select options (colour vs greyscale, dpi, etc.) and print a test page. Our test print worked flawlessy, first time – turning what was once an incredibly laborious task in to a two minute job.

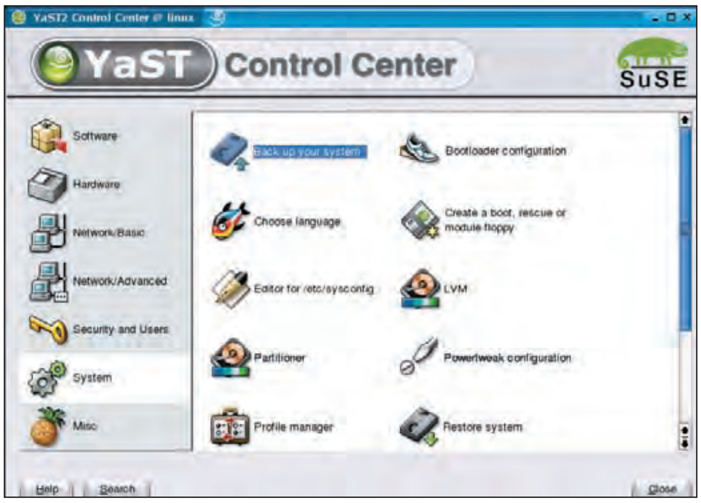

***YaST2* has come of age. It is friendlier and more professional than ever before.**

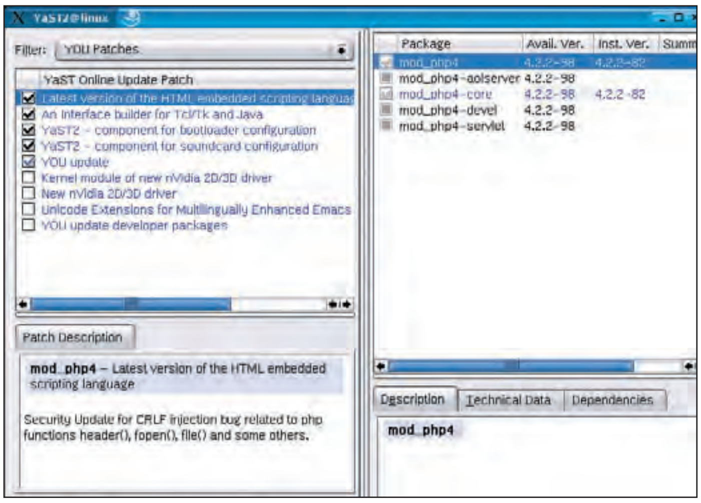

YaST Online Update Patch

- ☑ Latest version of the HTML embedded scripting languag
- ☑ An interface builder for Tcl/Tk and Java
- ☑ YaST2 - component for bootloader configuration
- ☑ YaST2 - component for soundcard configuration
- ☑ YOU update
- ☐ Kernel module of new nVidia 2D/3D driver
- ☐ New nVidia 2D/3D driver
- ☐ Unicode Extensions for Multilingually Enhanced Emacs
- ☐ YOU update developer packages

mod_php4 – Latest version of the HTML embedded scripting language

Security Update for CRLF injection bug related to php functions header(), fopen(), file() and some others.

Package	Avail. Ver.	Inst. Ver.	Summ
mod_php4	4.2.2-98	4.2.2-82	
mod_php4-aolserver	4.2.2-98		
mod_php4-core	4.2.2-98	4.2.2-82	
mod_php4-devel	4.2.2-98		
mod_php4-servlet	4.2.2-98		

mod_php4

Online updates are now easy for SuSE users, though the package manager's interface can take some getting used to.

« SuSE have a history of using the /etc/sysconfig/ directory as a repository for system configuration, much more heavily than other distributors – whether or not you like this is a matter of personal choice. Many new users find the 'English language' variables in sysconfig to be easier to understand than some of the more complex and detailed files directly under /etc. Again, the sysconfig editor under *YaST2* is a module which only provides the most simple interpretation of the raw text files, but it is at least convenient for navigating around the numerous individual files, as configuration options are grouped together by type – mail, networking and so on.

As with most distributions now, SuSE have moved from using the *Lilo* boot loader by default, to *GRUB*. Kernel 2.4.19 is packaged in the box, and *YaST2* provides easy access to the boot logs, and so discovering what your system was up to during the boot process is simple. The final significant under the bonnet change is that *CUPS* is the default print system, replacing *lpr*.

SusE need to pay a little more attention to detail when it comes to advanced power management – our test system suffered from a partial ACPI installation, rendering battery monitoring useless – Red Hat's *apm* installation has worked sucessfully with our laptop, out of the box, for the previous two versions.

In the machine room

Up until this point we have focussed on the experience for desktop users – but SuSE is of course suitable for the server room, too. A typical Internet server installation will deliver *Apache 1.3.26*, with PHP 4.2.2 for web, *MySQL 3.23.52* for databases, *Bind 9.1* for DNS, and *Squid* for proxying. All of these services can also be used on a LAN as well as on the Internet, and in addition *Samba 2.25* can be used for file sharing.

We installed the latest version of *webmin* (www.webmin.com) on our system, as SuSE do not provide a strong array of remote admin tools. We then purposely enabled *telnet*, and configured the *SuSE Firewall* to block the port, and test the firewall. The *Firewall* is simplicity itself to configure and enable, though we would have preferred to see SuSE 'push' this feature a little more aggressively to the user. As it stands, the user has to go and purposely dig out the firewall configuration module in *YaST2* to protect their system.

Wired World

The Internet runs deep in to the development and use of Linux – the history of the kernel, surrounding OS and key software is thoroughly entwined with the Internet, so how do SuSE cater for net users in this release?

Netscape has been dropped from the default installation, with *Mozilla 1.0.1* taking its place. *Konqueror* and *Galeon* are also installed. Thoughtfully, SuSE have installed and configured the Macromedia flash player plugin – how widely this will be welcomed is something of a contentious point. In line with most distributors, *lynx* isn't included by default.

E-mail is handled by *KMail*, Ximian *Evolution* or *pine*. Each of these tools caters for a slightly different audience, but they each go to show that we have truly been dragged out of the dark ages of flakey, featureless e-mail clients.

KNode has come on in leaps and bounds, for NNTP users, and again for the text based users, *pine* is available. FTP is handled by *Konqueror* and *KBear*, with the latter being a particularly powerful FTP client. *Realplayer* is installed, and *xmms* is great for MP3 streaming.

SuSE have produced a great Internet ready package, out of the box, with just about every taste catered for – however, there is a heavy bias towards the graphical tools, with console friendly users often left to do their own thing.

10 years on

Over the last ten years, SuSE's distribution has matured into a polished product for the professional user. It is the fine detail which sets this distribution apart from the rest – the little touches that make the system easier, and more enjoyable to use. 10 years ago, few dreamt that SuSE, and Linux, would be where they are today – so, with that in mind, SuSE can look forward to the next ten years with a solid product base on which to build. LXF

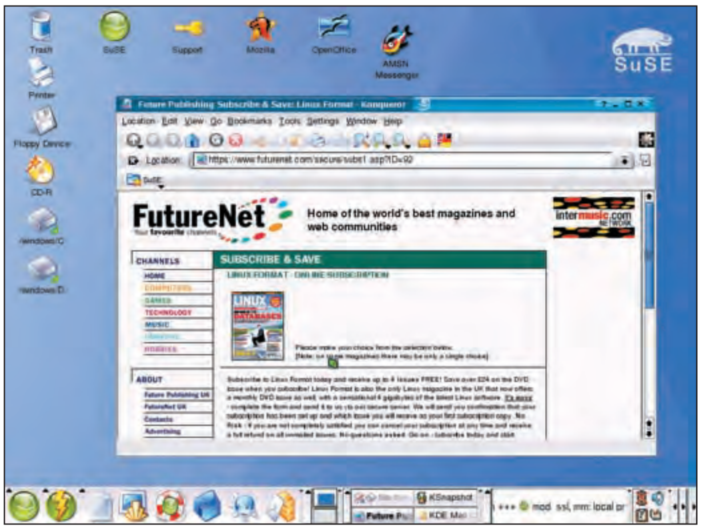
Trash
SuSE
Support
Mozilla
OpenOffice
AMSN Messenger
Printer
Floppy Device
CD-R
Windows C
Windows D
SuSE
Future Publishing Subscribe & Save: Linux Format - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help
Location https://www.futurenet.com/secure/sobe1.asp?ID=92
FutureNet
Home of the world's best magazines and web communities
intermusic.com
CHANNELS
SUBSCRIBE & SAVE
LINUX FORMAT
ABOUT
KSnapshot
Future P...
KDE Mai
mod_ssl, mm: local pr

The *Konqueror* browser has come so far, so quickly, and is an excellent and capable web, FTP, and file browser.

LINUX Format VERDICT

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

Overall comment: An excellent distribution, particularly for the home user.

LINUX Format **RATING**

9/10

Linux Kernel Programming

(3rd Edition)

Richard Drummond tackles a manual for the wannabe kernel hacker.

- **PUBLISHER** Addison Wesley
- **AUTHOR** M Beck, H Boehme, *et al*.
- **ISBN** 0201719754
- **PRICE** £34.95

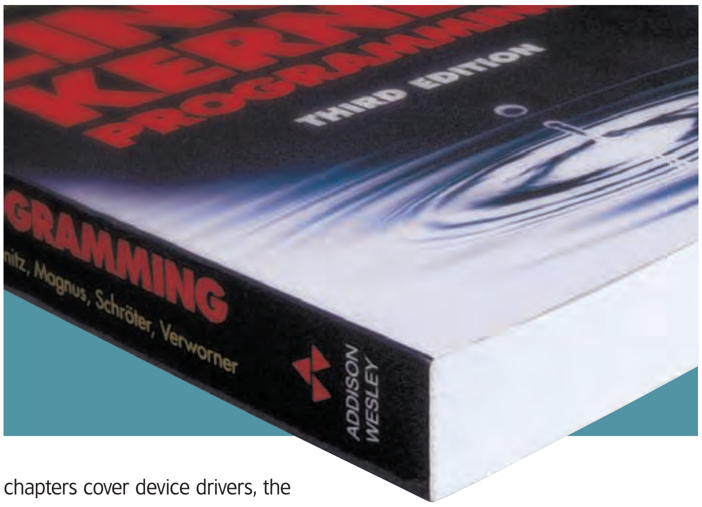

You would be forgiven for thinking that the Linux kernel – with current versions weighing in at 160MB of source code and documentation – is incomprehensible to mere mortals. That is not true at all. Strip out all the hardware and filesystem drivers and you are left with a core that is elegant and modular. But, while the kernel's modular design aids understanding, it is still useful to have a guide to show you the way around. *Linux Kernel Programming* is just such a guide. This latest edition has revised and updated to cover kernel 2.4.

This books assumes a knowledge of C and a grounding in operating system theory, but, other than that, it is approachable by anybody with a grasp of Linux. Chapter 3 gives a useful general overview of the kernel's architecture and main components, and will help get you up to speed. The following three chapters discuss the kernel's principal subsystems in detail, namely memory management, interprocess communication and the virtual filesystem switch. Other chapters cover device drivers, the network stack, kernel modules and multiprocessing. The appendices bring up the rear with handy reference material listing system calls, kernel-related commands and the structure of the /proc filesystem.

Linux Kernel Programming will be a boon to the fledgeling kernel hacker or to anybody else who desires a deeper knowledge of the inner workings of Linux. It is meticulous in detail and well illustrated with diagrams and excerpts of kernel code.

My only serious complaint is that this book has suffered in translation from its original German, and the text can sometimes be belaboured or ambiguous.

Linux Format VERDICT

An authoritative guide to the architecture of the Linux kernel.

LinuxFormat **RATING** 8/10

XML in a Nutshell

(2nd edition)

Dave Coulson sings the praises of a thorough guide to all things XML.

- **AUTHOR** Harold & Means
- **PUBLISHER** O'Reilly
- **PRICE** £28.50
- **ISBN** 0-596-00292-0

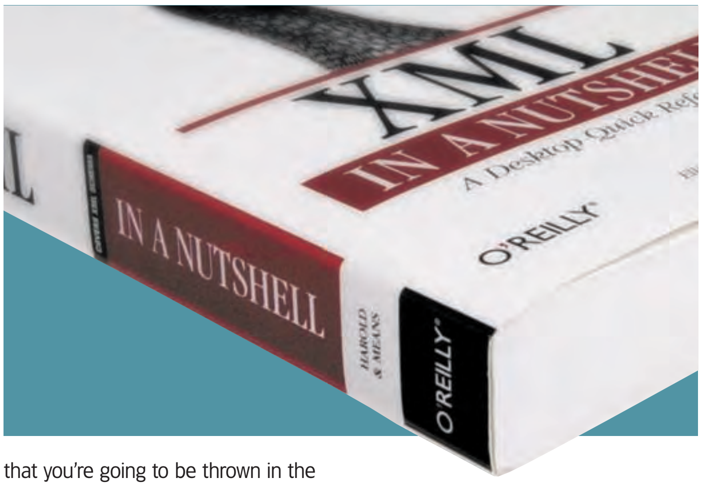

XML is certainly one of the buzz-words of the twenty-first century, and while it may frequently be hyped up by e-business people, it is actually a very popular and useful standard for a great deal of applications. The XML specification can be found at w3.org, but don't expect an easy time if you are going to try to learn XML from that. This, the second edition of the popular O'Reilly desktop reference, covers everything from the basic XML standards up to the Document Object Model (DOM) standards.

Of course, writing XML is one thing and actually doing anything useful with it is altogether different. There are many examples of writing XML documents, strictly sticking to the standards along the way, but there are very few examples of processing XML, and those which do exist are in Java. Not quite the end of the world, but a clear insight into the XML parser which you will be using, be it PHP, Perl or *Expat* in C, is going to be needed to make use of XML. Adapting the examples are not too difficult, but don't expect your hand held when you write your own routines to handle your XML. O'Reilly, and other publishers, have books covering XML with specific languages which may be a preferred choice.

Everything involving XML is covered in this book, from basic XML standards, all the way through to XSLT, XPath and CSS. Fortuanately, it does start with the most simple of examples, looking at individual tags and attributes, so you need not worry that you're going to be thrown in the pool head first at the deep end. However, once the XML specification is understood, then going straight to the section of the book which is of interest to you does not present too many problems. Should you already know what you're doing, then a quick reference is provided in the back of the book to jog your memory.

If you're looking at using XML for something specific, such as XML-RPC or SOAP, then a more detailed book covering that standard is probably preferable, as a great deal of this guide will be redundant. Those who are interested in using XML directly, and want a handy reference to lean on when they run into trouble, then this is an ideal publication and deserves pride of place on a bookshelf of anyone using XML.

Linux Format VERDICT

An excellent guide to XML, although you're going to need to know how to write your own parser.

LinuxFormat **RATING** 9/10

MySQL reference Manual

Big software needs a big manual. **Nick Veitch** checks out 780 pages of MySQL knowledge.

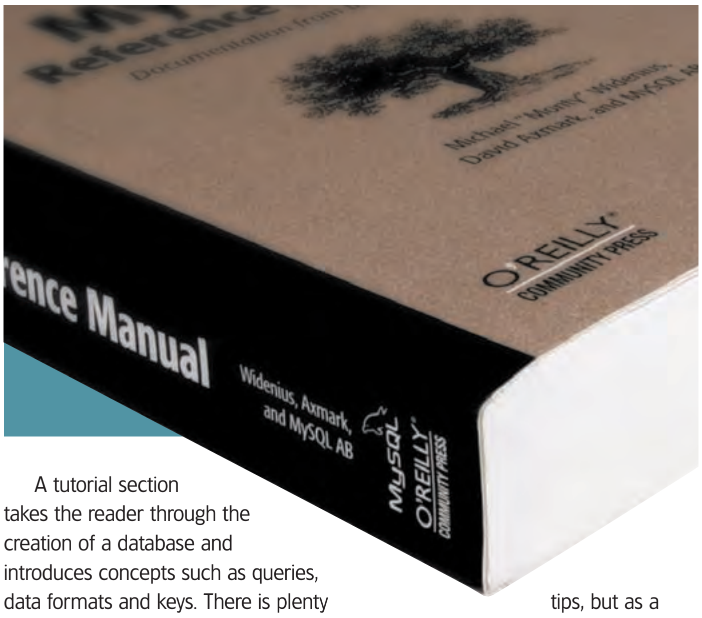

- **AUTHOR** Michael "Monty" Widenius, David Axmark
- **PUBLISHER** O'Reilly
- **ISBN** 0-596-00265-3
- **PRICE** £23

MySQL is probably the most commonly used database on Linux servers. In combination with the PHP language, it certainly powers a huge percentage of websites, including the Linux Format site. However, until recent years it hasn't been very well documented, and most people aren't using it as cleverly and efficiently as they might be.

This reference manual goes some of the way towards helping site developers and other database users get the most out of this terrific software. Written by *MySQL* experts including Monty Widenius who is the co-founder and CTO of MySQL AB, the book is split into nine logical sections.

Although the technical information is here, and the book is very thorough, it manages to be useful without assuming any knowledge of the *MySQL* software.

While the hardcore database programmers are going to get more out of the chapters on Optimisation and APIs, everyone who uses *MySQL* can get something from this book. The early chapters explain what *MySQL* is, what standards it supports, how the company supports development, an honest appraisal of how it compares to other database systems and plenty of links to web resources.

A tutorial section takes the reader through the creation of a database and introduces concepts such as queries, data formats and keys. There is plenty of discussion about security.

Perhaps most impressive is the way the book doesn't treat *MySQL* in isolation, but looks at how *MySQL* will interface with other applications and scripts, which – of course – is how you will make use of *MySQL* in the real world.

It may have been possible to do some more thought provoking examples, or to tackle good design tips, but as a reference work this is essential reading for *MySQL* developers and users. LXF

Linux Format VERDICT

A thorough and well written reference book – essential reading for database admins and coders.

LinuxFormat **RATING** **8/10**

Roundup »

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

Mini distros

Jon Kent puts the space saving mini Linux distros under the microscope.

cover feature

Our selection at a glance

- Demo Linux
- BasicLinux
- LNX-BBC
- tomsrtbt
- floppyfw
- LOAF
- Trinux
- Freesco

Linux distributions seem to be becoming bigger and bigger with each release. Debian's lastest release spans seven CDs – and that is just for the binaries, not the source code. So, what if you have a fairly old PC lying around that you want to put to good use. Well, you could either upgrade some of the components in your old PC, or look at some of the mini Linux distributions that are around. Some are designed to

What is a RAMDISK?

And why do most mini distros use one

A RAMDISK is a portion of memory that you allocate to use as a partition. Or, in other words, you are taking memory, pretending to treat it as a hard drive, and you are saving your files to it. Most of the distribution here are designed to work without a hard drive, and therefore use a RAMDISK to act as a pseudo partition. In addition, as some use a compressed file to store that actual filesystem, you can unpack this file into the RAMDISK and have more binaries to work from. Of course, once you switch off your system, any changes you have made to the files in the RAMDISK are lost, so you need to be careful to back these onto a floppy, or hard drive partition, if you wish to save anything.

Aside from the usefulness to mini distros, RAMDISK file systems are quicker to read and write to, so in some situation, such as a web server, it may make sense to use a RAMDISK to speed up your systems response times.

be used on floppy disks whilst others have been designed to run from a CDROM. Some mini Linux distributions have been created to serve a specific function such as a firewall. Nearly all only require a very small memory footprint to run.

The other, extremely useful, function of some of these distributions is the ability to use them as rescue disks should you be unable to access your normal Linux system for some reason, such as a corrupt master boot record (MBR) on your hard drive. I always make sure that I have one at hand should the worst happen, and it has saved me from having to rebuild a system many a time when I have accidently made my Linux system unbootable. It is always a good idea to have access to a mini Linux distribution that you can use, should the worst occur.

The are three main media types of mini Linux distribution: floppy-based, CDROM-based and hard drive-based. The hard drive distributions are perhaps the strangest ones as they are designed to be located within a DOS or Windows partition, and booted into using *LOADIN*. How useful this is to have is debatable, especially as the most recent version of Windows has no DOS access available, which is required for these types of distributions. For this review we are only going to cover floppy- and CDROM-based distributions, as these provide the most flexibility, and are accessible to all users.

Obviously these are cut down versions of Linux but, as you will see, there is still a surprising amount that can be fitted even onto a 1.44MB floppy disk.

How BusyBox works

The little binary that mimics 54 others

Lineo developed this useful utility for embedded Linux systems, but it has also found its way into the various rescue and mini distributions.

It manages to pack 54 of the essential commands into a single binary of just over 110KB. *Busybox* is used via symbolic links. The name of the symbolic link dictates which binary *Busybox* emulates. The --help option provides a brief description of the command, eliminating the need for man pages as well.

DemoLinux

Impress your colleagues with a demo of Linux.

■ **VERSION** 3.0 ■ **WEB** www.demolinux.org

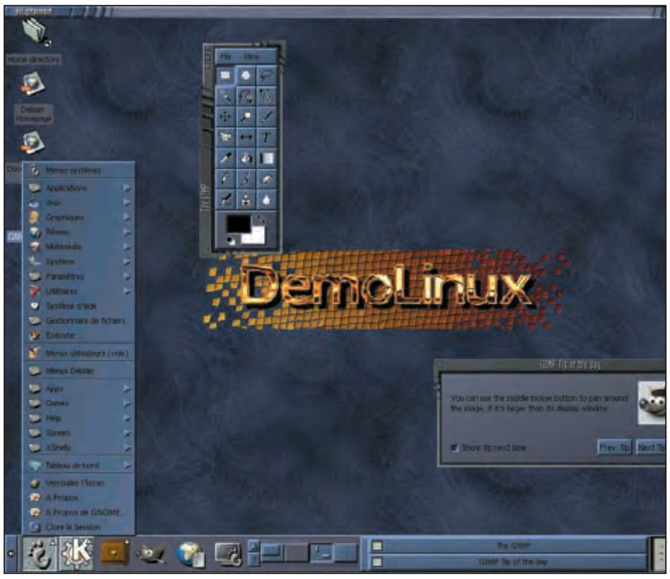

The snazzy French distro may run from a CD, but it looks every bit as slick as its permanently-installed counterparts.

Quick Overview table

Minimum System	Minimum Memory	Kernel	Network card Support	X	ppp
386	32MB	2.2.18	Large support base	Yes	Yes

DemoLinux is a distribution designed to use a bootable CD, nothing that unusual there. However it also comes complete with XFree86, and complete with networking apps (including *Netscape*), games, tools and loads of additional programs.

DemoLinux is currently based on Debian 2.2, running kernel 2.2.18. Obviously this is a big distribution, and takes a long while to download the very large ISO image, which is almost 700MB in size. There is also a compressed image, which comes in at around 500MB, but this is still a large download.

Requirements

The minimum requirements are higher than with the other distributions here. You need to have a Pentium based computer with a CDROM drive, the faster the CDROM drive the better, and at least 64MB of RAM. DemoLinux also comes with *StarOffice*, but if you want to run that you will need at least 128MB of memory.

After you've download one of the ISOs and burned it to CD, all of which takes quite some time even if you have an ADSL link, you are ready to boot into DemoLinux. When its boots it looks for various pieces of hardware, such as SCSI cards, disks and disk partitions, and so on.

Usefully, if it finds a hard drive it will assign it to hda1 and so on if you have multiple partitions. DemoLinux also looks for a network card, of which it supports many. If a network card is found it asks you if you have a DHCP server it can use, or if you wish to manually configure it. This is simple enough, except the DHCP option does not seem to work at present, so you need to manually start the DHCP client to automatically configure your network interface.

Once DemoLinux has booted up, you are presented with the familiar *KDM* login screen. As you see on the *KDM* login there is only one account, called demo, although you can log in as root if you prefer. Neither the demo or root accounts have password set, so once you press return KDE will start up. From here you can access many of the KDE applications and utilities that you can normally, with the added bonus of *StarOffice*. Surprisingly KDE is not the only window manager you can choose from, there is also GNOME and *twm* to use instead.

Performance

The one thing you will notice is that DemoLinux is slower than you may be used to. It is still surprisingly fast, you just notice that applications take slightly longer to start. This is because it is using the CDROM drive, which is far slower than your hard drive is. If you have a hard drive on the system, you can speed DemoLinux up by using the *anchor* program, which is available as an icon on the desktop. This creates a small area on you hard disk which DemoLinux uses for swap and configuration files. Just adding the swap file can help speed up DemoLinux quite considerably.

DemoLinux is a very interesting distribution that proves that you can fit a surprisingly complete distribution onto one CDROM. Yes, it is slower than you would be used to, but the main purpose behind it is for it to be used to demonstrate Linux (no, really) and it does this very well indeed. You can also use this as a temporary Linux desktop when you have no access to Linux. It is well presented, and is bound to get people asking what it is, which is the whole purpose of the distribution.

LINUX Format VERDICT

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

A full distribution on one CDROM including KDE and GNOME. Amazing but true!

LINUX Format **RATING** **8/10**

BasicLinux

Complete (old) distro on two floppies.

■ **VERSION** 1.7 ■ **WEB** www.volny.cz/basiclinux

Quick Overview table

Minimum System	Minimum Memory	Kernel	Network card Support	X	ppp
386	8MB	2.0.34	Limited	No	Yes

Do not let the name mislead you BasicLinux is anything but basic. It is based on Slackware version 3.5 and supplies the usual command line tools via the *Busybox* utility. Network connectivity is provided through both Ethernet and dialup, and you can also configure a firewall using *ipfwadm*. Those who have been around Linux for a while may notice something here. *ipfwadm*? Yup, BasicLinux uses the 2.0.x kernel. That is not a bad thing as this is a solid and well proven kernel, but then so is the 2.2.x series.

There are two main ways of running BasicLinux, via DOS or from two floppies. DOS mode is only available to you if you are running Windows 95/98, otherwise you will have to use the floppy route as later Windows versions do not have a genuine DOS mode. Out of the box, BasicLinux does not provide many modules for network cards, but luckily you can download these modules and either add them to the compressed image *baslin17.zip* or put then onto another drive. It's at this point when you start to realise that BasicLinux needs a bit more effort to make it completely useful, depending upon your requirements.

As you would expect there is no X available so you need to be fairly used to using the command line if you want to use BasicLinux. To be fair, this is the case for nearly all the distributions reviewed here and is not a negative point. You do have utilities like *fetchmail*, *links* and our old friend the *Bash* shell. *ppp* is available ready to be used, once you have configured it. For some strange reason BasicLinux also comes with two text editors, *e3* and *pico*, but no *vi* – which seems a little odd. Lastly, by default you are also provided with three console logins, although you can increase this if you wish.

There is nothing particularly wrong with BasicLinux, it does a lot in very little space, but perhaps it tries to do too much. As a rescue distribution it falters as you need two disks – whereas others will quite happily use one, which makes them slightly more portable than BasicLinux.

Unfortunately, it is not really strong enough in the security space to use as a dial out server to the Internet, and again, others do it better. In a nutshell, it is not bad, but needs to have more focus.

LINUX Format VERDICT

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

Tries to do too much, does not excel in any area, and needs two floppies Look elsewhere.

LINUX Format **RATING** 6/10

LNX-BBC

Credit card-sized graphical rescue distro.

■ **VERSION** 1.618 ■ **WEB** www.lnx-bbc.org

Quick Overview table

Minimum System	Minimum Memory	Kernel	Network card Support	X	ppp
386	16MB	2.2.19	Good support base	Yes	Yes

The Linuxcare bootable business card has a history that goes back to 1999, when four Linuxcare employees designed a bootable CD to be used for rescuing damaged systems. The distro was designed to be compact enough to fit on the small business card-sized CDs, hence the name BBC (Bootable Business Card). The early versions were fairly basic, providing a minimal system, but today's version of LNX-BBC now includes XFree86 and *Blackbox* to provide a more usable system. There are some useful apps included, such as *BrowseX* (a web browser), *lynx* and *Ethereal*.

The current version of LNX-BBC is a fork from the original code. The original developers of the Linuxcare BBC have left Linuxcare but have continued to work on the project with help from other developers, creating the LNX-BBC project. Installation is simply a case of downloading the fairly small ISO image and burning it to a CD.

LNX-BBC is aimed to be used by current Linux users and is not intended to be used as an easy introduction to Linux like DemoLinux. Although X is available this is not started automatically, and if you have any network cards, and other such devices, these need to be manually loaded into the kernel either using the modules supplied (of which there are many), or using your own modules. Obviously, if you have to use your own modules you need to locate these either on a floppy, another CD, or on the local hard drive of the system you are using. This is not newbie territory.

LNX-BBC is still aimed as a tool to aid in fixing a malfunctioning system, although you can use it as a temporary system, if required. As this is aimed at experienced users, there are no gui tools to use, so everything has to be done at the command line, even in X. Interestingly, for a distro that aims to help you fix your main system, *GRUB* is not available – but *Lilo* is. This is not a major problem as you could, hopefully, use *GRUB* from your main system to repair any MBR problems, but it does seem odd that it is missing.

When you boot LNX-Linux it first asks you for a screen resolution that you would like to use. You can also select a text option, but if you do so you will be unable to run X. The remaining options invoke framebuffer support, which allows you to use X. Once the operating system starts, you can log in as root using the instructions that appear on the screen. If you want to start X Windows, you can simply use the usual **startx** command and X starts with the *Blackbox* window manager. As mentioned previously, if you have a network card that you need to use, you need to install a module into the kernel. You can either do this by hand, which is easy enough, or use the *trivial-net-setup* script provided which will perform this for you. If you require dhcp, this script will also start the daemon for you.

A nice little feature of LNX-Linux is that upon bootup it will try and mount your local hard disk. This is not done in any intelligent manner, it just simply goes though a list of likely partitions and tries to mount them. Nevertheless this is one less step to have to go through. The only point to note is that all the partitions it finds and mounts are mounted read-only, so if you do need to make any changes you need to remount the partition which is perhaps a little annoying but does safeguard your partitions from accidental damage.

LNX-BBC is a useful distribution, with many tools to help you fix a problematic system. Being CD-based allows it to have many more utilites than a floppy-based system, and to have full version of these utilites. It is also still supposed to be able to fit onto a business card-sized CD, although I have not tried this, so you could put it in your wallet for future use. A very useful distribution to have.

LINUX Format VERDICT

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

Excellent CD-based rescue distribution – and it comes with X to boot!

LINUX Format **RATING** 10/10

tomsrtbt

Ubiquitous single-disk rescue distro.

■ **VERSION** 2.0.103 ■ **WEB** www.toms.net/rb

Quick Overview table

Minimum System	Minimum Memory	Kernel	Network card Support	X	ppp
386	8MB	2.2.20	Good support base	No	No

And now we come to tomsrtbt, or Toms as it often referred to as. Toms is probably the best known and most widely used recovery disk version of Linux. There are around 200 modules and utilities that make up Toms and it all fits onto one boot floppy. How on earth does that all fit on to one floppy disk you may ask. Well it sort of does not, because Toms is actually 1.722 megabytes in size. This is achieved because most 1.44 floppy drives can actually support floppies written in a 1.722MB format. Obviously the older the drive the more probable it is that you will not be able to use Toms, but so far I have not had a problem.

Installing Toms is straightforward. You simply download the tar file, or zip file is you are installing from Windows, unpack it and run the installation script. The only thing you need to make sure if that you have an error free disk as Toms need all of the capacity, so any errors will mean that you cannot install it.

Now you can put the floppy into the required PC. During the bootup you will be asked to identify a video mode to use and the keyboard to use (40 is the UK keyboard layout). Like most other floppy-based distros, Toms uses a RAMDISK so once bootup is complete, you can eject the floppy. You log in as root using the default password **xxxx**, which you can change if you wish. As this is a single floppy based distribution there is no X, so you need to know you way around a system from the command line.

Functionality provided is the minimum required for a Linux system, so you can copy files, run *Lilo* (useful if you MBR is corrupt), mount drives and so on. There is also TCP/IP support, but this is limited to a few network cards. You also have several network utilities such as *wget*, *telnet* client and daemon, *rsh* client and daemon, a DHCP client, with many more available. However, because of the limited support for network cards, if you need networking support you would probably be better off with a different distribution – although you can add support for your card if you wish. There's a wide range of additions that can be made if you need to add additional hardware support, but whether it is better to use them or to use another distribution is debatable.

Toms main use is as a rescue distribution and at this task it is very good, as it provides most of the functionality you would require to restore a system, all on one floppy. If this is your main reason for using a mini-distribution, Toms is definitely worth looking into. If you want more from you mini-distro though, you are probably better off looking elsewhere.

LINUX Format VERDICT

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

Been around for ages and for a good reason as it is a superb system recovery floppy distribution.

LINUX Format **RATING** 9/10

floppyfw

IPtables firewall-on-a-floppy.

■ **VERSION** 1.9.9 ■ **WEB** www.zelow.no/floppyfw

Quick Overview table

Minimum System	Minimum Memory	Kernel	Network card Support	X	ppp
386	12MB	2.4.19	Good support base	No	No

floppyfw is a single floppy distro that supplies static routing and firewall capabilities – albeit just a packet filtering firewall. In addition, floppyfw can be configured for IP-masquerading (or Network Address Translation - NAT), port forwarding, dial out and a DHCP server. The current stable release (1.0.x) provides *ipchains*, and is based on the 2.2.x series of the Linux kernel, while the beta tree (1.1.x), reviewed here, provide *iptables* and is based on the 2.4.x series (2.4.19 for version 1.9.9). The kernel support provided by floppyfw makes it one of the most up to date distributions reviewed here.

floppyfw's main aim is to provide firewall capabilities on a redundant PC. The PC will need two network interface cards in order to work correctly, which is a consideration as sourcing supportable network cards for older PCs can be problematic. If you would like to set up a firewall using floppyfw (or any other Firewall solution), it is probably best to use a Pentium based PC as these usually support the PCI architecture, making sourcing cheap network cards easy. Installation of floppyfw is simply a case of downloading the lastest floppy image and *dd* this to a floppy.

Configuring floppyfw is a straightforward task. The easiest method is to mount the floppy drive on your Linux system, or you can access the floppy via Windows if that is all you have access to. You can use Windows to access floppyfw as it uses the FAT format. So if you are mounting the floppy under Linux you will need to do a mount command like this:

```
mount -t msdos /dev/fd0 /mnt/floppy
```

floppyfw does not use a compressed filesystem for its config files, so changing the config is simply a case of going to the correct directory and directly editing the file using your favourite editor. This certainly makes life simpler. There are five floppyfw config files:

```
config (main configuration)
firewall.ini (filter rules)
modules.lst
 (additional ip_masq modules)
syslinux.cfg (kernel boot parameters)
syslog.cfg (syslog config, such as
 /etc/syslog.conf)
```

The only configuration files that are you are likely to need to modify are the *config* and *firewall.ini* files. Within the config file you modify (if required) settings such as which interface is connected to the Internet; the various network settings for the internal facing interface; DNS caching; turning on or off *syslogd*; and so on. The *firewall.ini* file that floppyfw comes with only sets up basic masquerading and rejects a couple of ports, and does not take advantage of the full firewalling capabilities that Linux provides. If you want to add additional firewall rules this is the file you would make the changes to. By default *syslog* is set up to log everything to the console, so you may wish to change this. Remember that you cannot write anything to the floppy as there is not enough space so you will need to set up a loghost on another system and configuration *syslog* to send all logs to that system.

It is very easy to setup, as you have seen, and seems to work very well. floppyfw is not bleeding edge firewall technology, but it is still strong enough to secure your internal networks from most unwanted intruders. If you are used to using *ipchains* or *iptables* you should feel right at home with floppyfw. floppyfw is a very good starting point if you need to setup a firewall between your network and the Internet, or any other untrusted network. It sets out to do one thing well, and certainly achieves this.

LINUX Format VERDICT

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

Sets out to be a firewall distribution on a floppy, providing easy setup. Good for small networks.

LINUX Format **RATING** 8/10

LOAF

Sparse floppy distro for network use.

■ **VERSION** 1.2 ■ **WEB** www.ecks.org/projects/loaf

Quick Overview table

Minimum System	Minimum Memory	Kernel	Network card Support	X	ppp
386	4MB	2.0.36	Good support base	No	No

LOAF is a single floppy distribution which consists of the kernel and a small collection of utilities. The intention of LOAF is for it to be used as a client for various network protocols and it provides various utilities such as *lynx*, *ftp*, *telnet*, and *ssh*. From a hardware perspective, LOAF's requirements are very low – needing only a 386SX CPU, 4-8MB of RAM, and a floppy drive. As LOAF is designed to be used on a network you will also need a NIC to get the best from it. As with most other floppy-based distros, installation is straightforward and only requires downloading the chosen image and *dd'ing* this to a floppy.

As LOAF is designed for network use you will not find any utilities for recovering a system from a failure, so this distribution will be of no use to you in that situation. Surprisingly though, for a network-centric distribution, there is no support for *ppp*, so if you need to use a modem you will have to look elsewhere. However, within the console-based world of the local network and the Internet, LOAF does work very well, if you can live without a gui front-end. Why you would want to do this is another matter, but it you like a very sparse interface to the Internet, LOAF will happily provide that for you.

Another missing feature is support for DHCP, so even if you do have a DHCP server available LOAF will not use it. So you need to manually configure your network interface and add any routeing entries that you may require. Also, unlike some other distributions, you have to select the correct floppy image that supports your NIC in order to ensure LOAF can use the interface. Most of the others in this roundup supply a good selection of kernel modules for network interface cards on just one disk.

In use though, LOAF is a fairly fast distribution, and the addition of *ssh* is a nice touch. However LOAF is a very limited distribution, beyond simple text based web surfing and use of ftp, telnet and *ssh* there is very little else that you can do with this. The lack of ppp support is LOAF's real let down, as this limits it usefulness even more. The latest version may have addressed some of these issues, but there are no working releases yet available. All in all, it promises much but delivers too little to be of much practical use and does need improvement.

LINUX Format VERDICT

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

Not as full a distribution as the name implies. With limited uses, this is a disappointing distro.

LINUX Format **RATING** 5/10

Trinux

Network analysis and a package system!

■ **VERSION** 0.80rc2 ■ **WEB** http://trinux.sourceforge.net

Quick Overview table

Minimum System	Minimum Memory	Kernel	Network card Support	X	ppp
486	12MB	2.4.5	Large support base	No	No

Trinux is another RAMDISK-based mini(ish) Linux distro, that you can install on a bootable floppy or CD. Trinux' main focus is as a network analysis OS, to check security on your local network. It contains latest versions of popular network security tools for port scanning, packet sniffing, sniffer detection, session-hijacking, backup/recovery, vulnerability scanning, packet construction, network monitoring, active/passive OS fingerprinting, computer forensics, intrusion detection and more – even *OpenSSH* support for secure remote management of your Trinux box.

There are three floppy images. The Network image has a stripped-down kernel with support for common NICs, but no support for IDE devices or filesystems (except FAT and minix). The PCMCIA image contains the smallest kernel, with PCI and ISA support removed, and PCMCIA kernel module packages available. The IDE image adds support for IDE hard drives, CDROMs and NTFS. There is also an ISO image which amalgamates the floppy images.

Trinux will run on a 486 or better with 12-16MB of RAM, depending on how many packages are loaded. Trinux is developed using Slackware 7.1, supports Linux kernel 2.4.x, and utilises *BusyBox* to replace many common Unix utilities.

When you have created your boot floppy or CD from the selected image and booted up you may notice the complete lack of network tools available. This is where we come to an interesting difference between Trinux and the other distros here. The boot floppy, and even the CD, only provide enough to boot the kernel, build the RAMDISKs, extract a minimum number of necessary utilities, and configure the network, usual via DHCP. To gain access to the network tools you need to download the relevant package from the centralised webserver using *snarf* (much like *wget*). Why does Trinux do this? Well the main reason given is that "This saves the trouble of building and updating floppies", but this also ensures that you have the latest version of the tools ready to use, without have to build another boot floppy or CD. Coming from a Debian background I find this feature quite natural, but to some this may seem strange. As Trinux is designed to troubleshoot networks, you would not normally need to run this type of distro at home, where your bandwidth maybe limited. I like the fact that Trinux uses HTTP for file transfer as this should not cause any problem from within a corporate network.

With Trinux' packaging system you can choose which packages you wish to use, so you do not have to download the complete set, which is quite comprehensive. Of course, once you switch off the system, all the packages you have downloaded will be lost, and you will have to download them again next time you use Trinux. But as this ensures that they are up to date and as you would not expect to use Trinux every day this is not a disadvantage.

As with all RAMDISK distros, Trinux config files and data are non-persistent, as they only reside in RAM. Luckily it is possible to save config info on the boot floppy using the **savecfg** command which saves the contents of /etc/tux to the boot floppy. You can also load or save your home directory using the **savehome** and **gethome** command from a remote Unix system using ftp or *SSH*.

Trinux is very impressive and comes highly recommended for network analysis or intrusion detection. It provides every tool you could possible want, from *tcpdump* and *nmap*, through to network load generators to allow you to perform stress testing. The ability to download packages from the Trinux servers is a definite plus, allowing Trinux to fit on one floppy. Loaded up on a portable PC there are no ends of use for this distro. Just plug it into a suspect switch and then access it remotely to gather the network logs.

LINUX Format VERDICT

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

Very impressive network analysis distribution with a packaging system to boot.

LINUX Format **RATING** 10/10

Freesco

Handy router/firewall for small networks.

■ **VERSION** 0.2.7 ■ **WEB** www.freesco.org

Quick Overview table

Minimum System	Minimum Memory	Kernel	Network card Support	X	ppp
386	6MB	2.0.38	Good support base	No	Yes

Freesco, which stands for FREE ciSCO is designed to be a replacement for commercial low end routers, and can also be used as a firewall, dialup server and dialout server. Currently Freesco is based around Linux kernel 2.0.38, can run from one floppy or from a hard drive, and supports up to three ethernet/arcnet/token_ring/arlan network cards and up to two modems. The design goal of Freesco is for it to be easy to set up and use, and to provide equal reliability to its commercial couterparts – all within a small disk footprint.

The hardware requirements for Freesco are, as ever, very minimal. All that you need is a 386 or better PC with at least 6MB of RAM. Previous comments about the problems with availability of network interfaces should be taken into account though, so a realistic spec with be a Pentium PC to make network cards easy to purchase if required, or for replacements in the future. Installation of Freesco is straight forward – just requiring you to download a floppy image and transfer this to a floppy using *dd*.

Once you have booted up your PC into Freesco you can now start to configure it. You can configure Freesco from the boot prompt if you wish, but the easiest way is from the command line. This is performed by the setup utility, which provides an easy and helpful console-based configuration tool. The only slight problem with this tool is that you cannot go back if you make an error, but that is only a slight annoyance. From here you need to choose the type of setup that you want to configure, with options such as dialup router, ethernet bridge, remote access server or ethernet router to choose from.

You choice will dictate which options appear next, but a common requirement is the need to know the I/O address and IRQ number of you network card. Other options that may need to be completed are the 'phone number of your ISP, login name and so on, and any local network information, such as address ranges. For setting up a home network you may also wish to set up the Freesco box as a DHCP server, which is also configured from within this setup utility.

Freesco also incorporates firewalling and NAT, which have been enabled within the Linux kernel. All of these features can be used in conjunction with each other or individually. Freesco also provides a web control interface so that you can administer the system from any web browser.

Although Freesco tries to emulate low-end Cisco equipment it probably is not really capable of replacing such equipment within a medium to large scale network. For use at home, or within a small network environment, Freesco could prove to be very useful. But if all you need is a firewall, floppyfw is probably a better distribution to use. Nevertheless, Freesco is very well thought out and very quick and easy to set up and use. The web interface is an interesting feature, which allows you to remove the keyboard and screen from the system.

However, there is no support for encryption over the web interface and *ssh* would probably have been a better solution. If you are looking to set up a small network, Freesco is worth evaluating to see if it can reduce the costs associated with Cisco and other such network equipment.

LINUX Format VERDICT

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

More useful within a home network, but it is well implemented and easy to administer.

LINUX Format **RATING** 8/10

MINI DISTROS THE VERDICT

This review of mini distributions has been an eye opener for me. I always used tomrtbt as a rescue disk and never really looked into any of the other distributions that are available. They have all impressed me, bar BasicLinux which really does not provide anything useful – and for its size it is surprising that it needs to use two floppies. Most of the distributions reviewed here are focused around one use, the ones that try to be general tend not to add any real value, because of the lack of focus.

Which distribution suits you depends upon what you want it to achieve and the media you want to (or can) use. If you want to build a firewall, floppyfw is the best one here for you – it is extremely easy to set up and use. If you want a rescue distribution both LNX-BBC and tomrtbt are good choice,s providing CD and floppy solutions respectively.

But which one is the best here? That was an easy choice, Trinux. It really impressed me as a network troubleshooting tool, but it can be also be used as a rescue distribution assuming that you use a network to connect to the Internet. It is extremely simple to use, binaries are kept up to date by downloading them from the Trinux servers and, of course, it is very portable. You can even choose between using Trinux on floppy or CDROM to suit your requirements. It does not have XFree86, but you do not need to on these type of distributions, it's just an added bonus if it's there.

So if you looking for a distribution to perform a niche function you may find that someone out there has already done it for you. And if they have not you could always put one together and release it into the wild. That's the whole beauty of Open Source, if you want something there is sure to be someone who does as well. LXF

Table of features

Name	Main Focus	Medium	Main features
BasicLinux	General purpose	2 floppies	Limited firewall abilities
DemoLinux	Full distribution	1 CDROM	Provides a full distribution with Xfree86
LNX-BBC	Rescue	1 CDROM	Multiple tools to repairing Linux systems
tomrtbt	Rescue	1 floppy	Lots of utilities on a floppy
floppyfw	Firewall	1 floppy	Good firewall capabilities
LOAF	General purpose	1 floppy	Console networking utilities
Trinux	Network troubleshooting	1 CDROM/floppy	Downloads analysis tools from the web
Freesco	Router	1 floppy	Router, bridging, and dial out server

HotPicks

The best new open source software on the planet!

Jon Kent
A Unix veteran and Real time information consultant.

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

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

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

HotPicks at a glance

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!

LOG MONITOR

Swatch

■ **VERSION** 3.0.4 ■ **WEB** www.oit.ucsb.edu/~eta/swatch

So you lock down your system and ensure that it is up to date with the latest patches. Your system is now, at the very least, a lot harder to break into. What about your systems logs files then? There is no point in performing all these steps without verifying that no one is trying to break in via one means or another. Equally, you would want to ensure that your system is running error free. You could manually check the log files every now and then, but that's tedious and inefficient. What you really need is something that will verify your logs file in real time and alert you, by one mean or another, should something unexpected appear. Well this is exactly what *Swatch* (*Simple WATCHer*) has been written to do for you.

Swatch is a Perl script that *tails* your log file, by default /var/log/syslog or /var /log/messages, and looks for predefined messages. You can use *Swatch* to monitor different log files by using the **--tail-file**, or **-t**, parameter. A single *Swatch* process can only monitor one log file at a time, so if you want, multiple log files monitored you will need to start a *Swatch* process for each.

Once you have downloaded and unpacked *Swatch* you run the usual **Perl Makefile.pl** command which will then ensure that you have the correct CPAN modules installed, and if not download and install these for you automatically. *Swatch* uses fairly standard CPAN modules which you may wish to install using your distro's package as opposed to downloading then directly, so it is worth verifying the CPAN modules required and seeing if you can use a package instead. Once this is completed a Makefile is generated which you use to build *Swatch* itself – with the usual **make**, **make test** (optional), **make install** – and **make realclean** to tidy things up.

Once installed from either source or from a package, you are ready to configure *Swatch* to monitor your system. By default *Swatch* uses the configuration file ~/.swatchrc, although it is probably sensible to place the config file in /etc and tell *Swatch*, via the **--config-file** (or **-c**) parameter, where you have relocated the file. If you want to monitor multiple log files it is probably better to create a *Swatch* config file for each log file you want to monitor, to avoid any erroneous error messages as you cannot create a *Swatch* configuration with the ability to monitor different events, depending upon the log file being monitored.

Swatch config

The config file syntax is relatively straightforward. It basically contains text patterns you want to match, using regular expressions. Each condition is followed by any actions you want to be taken if the pattern is matched. The following is an example of the syntax:

```
watchfor /user unknown/
   mail addresses=jkent\@thingy.com,
   subject=Unknown\ user\ error
```

This example is obviously simplistic, but shows how easy configing *Swatch* can be. Each event you want to monitor starts with the option **watchfor**, then the regular expression containing the event and then the actions you wish to take when the expression is matched. When creating a rule you need to bear in mind that this is a Perl script and thus some characters have special meanings in Perl. In the example above there are backslashes (****) before the **@** in the email address and before each space in the subject header to escape out these characters to ensure Perl does not try to interpret them. Of course you can use far more complex regular expressions, either using standard regex or Perl's extended regex options.

With the previous example *Swatch* will continuously try to match the patterns you have defined, however this is always be required for some log file. To facilitate this, *Swatch* also has a when option which accepts the syntax **range_of_day:range_of_hours**. You could then create the following rulebase:

```
watchfor /error/i
   mail addresses=jkent\@thingy.com
   subject=Error,when=2-6:8-17
```

This rulebase checks the log file for any line with error as an entry in any case (the **i** option in the regex). If this is matched an email will be sent only if the match is occurs between Monday to Friday and between the hours of 8 a.m. and 5 p.m. *Swatch* has many other options that can be used in addition to the examples above, such as the **exec** option that tells it to execute a binary or script when a condition is matched.

Running Swatch

As mentioned *Swatch* expects its configuration file to be in ~/.swatchrc. When *Swatch* is started is also creates a temporary file, and a dynamically created script, which it also places in this directory unless informed otherwise. As per the config file it is probably a good idea not to have *Swatch* do this but to locate these files elsewhere on the system, preferably in a directory that is not world readable. The following example shows how you can put this all together to run *Swatch*:

```
swatch -c /etc/swatch/swatchrc.errors
 --script-dir=/var/tmp/swatch
 -f /var/logs/errors &
```

This command uses the config file called /etc/swatch/swatchrc.errors, uses the directory /var/tmp/swatch to create the temporary file and script and monitors the log file /var/logs/errors.

Swatch is fairly easy to install and config, and provides good functionality. If you want to be able to monitor your log files in real time *Swatch* can provide that functionality. It may take a little effort to config all of the events you wish to monitor for, but once complete you can feel more secure in the knowledge that you'll be more aware of any issues that arise. So take control of your logs and install *Swatch*.

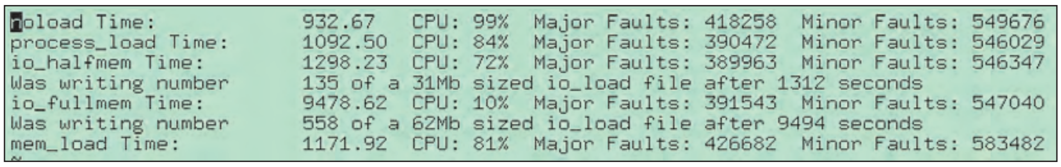

```
noload Time:            932.67   CPU: 99%  Major Faults: 418258  Minor Faults: 549676
process_load Time:      1092.50  CPU: 84%  Major Faults: 390472  Minor Faults: 546029
io_halfmem Time:        1298.23  CPU: 72%  Major Faults: 389963  Minor Faults: 546347
Was writing number      135 of a 31Mb sized io_load file after 1312 seconds
io_fullmem Time:        9478.62  CPU: 10%  Major Faults: 391543  Minor Faults: 547040
Was writing number      558 of a 62Mb sized io_load file after 9494 seconds
mem_load Time:          1171.92  CPU: 81%  Major Faults: 426682  Minor Faults: 583482
```

***Contest* provides the means to benchmark your kernel.**

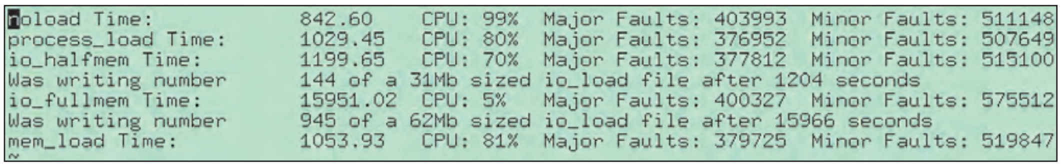

```
noload Time:            842.60    CPU: 99%  Major Faults: 403993  Minor Faults: 511148
process_load Time:      1029.45   CPU: 80%  Major Faults: 376952  Minor Faults: 507649
io_halfmem Time:        1199.65   CPU: 70%  Major Faults: 377812  Minor Faults: 515100
Was writing number      144 of a 31Mb sized io_load file after 1204 seconds
io_fullmem Time:        15951.02  CPU: 5%   Major Faults: 400327  Minor Faults: 575512
Was writing number      945 of a 62Mb sized io_load file after 15966 seconds
mem_load Time:          1053.93   CPU: 81%  Major Faults: 379725  Minor Faults: 519847
```

***Contest* statistics for a stock 2.4.19 kernel.**

KERNEL PERFORMANCE TEST

Contest

- **VERSION** 0.34
- **WEB** http://members.optusnet.com.au/ckolivas/contest/

One of the first things I tend to do once I install a distribution is to throw away the kernel that comes with it, download a kernel source from kernel.org and rebuild my own. I may apply some additional patches, but that depends upon what the system is to do. I do this as the kernels that come with most distributions are bloated as they tend to have support for a lot of devices to ensure that they will work on most systems.

Need for speed

So I create a kernel specifically for the system I have just built. It becomes interesting when you look at the various patches that are available and the modified kernel builds that you can use (like the *-ac* or *-ck* series of kernels and patches). How do you know if these patches or kernel modifications actually provide any benefit to you system? The same applies when a new kernel gets released. Do you just blindly upgrade because there is a new version and hope that there is some performance benefit or check the changes that have been made and only apply it if there are changes that are relevant to you?

Verifying that a new kernel or patch does actually help your system is now usually very easy to achieve, it can come down to "does the system actually feel faster or more stable?" – which is very subjective. To help you get some hard figures you really need a tool that can help you measure the performance on your system between different kernels and this is the design goal behind *Contest*.

Contest is written by Con Kolivas, the name behind the *-ck* kernel patch, which provides many patches, such as the Low Latency patch and Andrea Arcangeli's VM patch. *Contest* is designed to test system responsiveness by running kernel compilation under a number of different load conditions. It is designed to compare different kernels running on the same machine, without any hardware changes, to give you a good idea as to which kernel suits your system the best. The following four load conditions are tested by *Contest* (see table, *below*).

CONDITION	SUMMARY
Null load	Provides a benchmark
Process load	Fork and exec N processes, connected in a unidirectional ring by pipes. Insert M N chunks of data into the ring and pass them around.
Memory load	Repeatedly reference 110% of RAM in a pattern designed to cause cache misses.
IO Load	Copies /dev/zero continually to a file the size of half the physical memory then all of the physical memory.

Installation of Contest is very simple. Just download the source code, run **make** and then **make install**. The *Contest* program itself is just a shell script, which acts a a wrapper for the various test programs. I had problems with the script with it finding the *time* binary (it uses **which time** to find the location of *time*), but this is easily fixed by hard coding the path and binary name into the script.

Using Contest

If you do not have the source code for a kernel, download a kernel source tree, and extract it and configure the kernel as usual. You will need to you the same *.config* configuration file between the different kernels, so try not to switch on anything that is not available in a different kernel. In the documentation provided with *Contest* it recommends using the .config file that comes with the source, but I could not get this to work, you have more mileage with this .config file.

Once you have compiled and installed the new kernel, reboot the system into the kernel using single user mode. With *Lilo* you would type the following to achieve this:

```
[kernel name] single
```

If the partition the kernel source is located in does not have enough space to perform the IO Load test (see the table, *below left*) you can define an environment variable to inform *Contest* to write this file to another locate as follows:

```
export tmpfile=/path/to/file
```

When the system is up and running *cd* to the kernel source directory and then type:

```
contest [kernelname]
```

You do not have to enter the kernel name as *Contest* will, using the name from **uname -r** to determine this. *Contest* will now run a kernel compile under the different conditions and create six log files. One log file will list each test for that kernel and be called *[kernel].log*. The five other log files will be cumulative from any previous tests and are called:

```
noload.log
process_load.log
io_loadhalf.log
io_loadfull.log
mem_load.log
```

There are two numbers generated for each log, the time taken to compile the kernel and the average CPU percentage to do the compile. The [kernelname].log file contains more detailed information.

Reading the results

In a nutshell you are looking for lower times, the lower the better, and the higher the CPU percentage the better. What these numbers are is not critical here, but the differences between each kernel you test are. When the difference between noload and the loads is small it shows that on that kernel, the system is able to respond to requests for normal tasks (*i.e.* responsiveness). When the CPU percentage is high it shows that although the other loads are high, a cpu intensive application (the kernel compile) can still use as much as it needs.

Remember not to change the kernel tree you are using for the testing or the hardware settings. You would use *Contest* to verify if a patch, or group of patches, has a positive impact on you system. Any changes that you make between tests mean that the results will not give you any useful information. Of course after you run the first batch of tests you can then install or remove hardware to see if the addition, or removal, of hardware has a performance advantage by re-running the tests.

Contest is a useful utility which provides you with hard figures to base your kernel configuration upon, rather than relying upon configuration decisions that a distribution makes for you, or recommendations. You can be far more confident that your kernel choice and any patches you have applied, do indeed provide the benefit for which you are looking. It is easy to install and use and the resulting information is not provide in a difficult format. *Contest* is worth using to help you ensure that you are getting the best out of your system.

FTP ON STEROIDS

lftp

■ **VERSION** 2.6.2 ■ **WEB** http://ftp.yar.ru

No matter what you do, you will at some point have to transfer files from a website or ftp server. You of course use you web browser for this, point and click and download, download, download, freeze, stop and have to start again. This is the major problem when downloading via browsers, if the download stops for whatever reason, you have to start again as you cannot continue where the download left off with a browser. With command line tools you have the excellent *wget* tool, which does support this, and many other functions. But if you planning to download a file, but are not sure of the location, *wget* is not the best tool to use. If you are downloading from an ftp site you could just using your ftp client to perform this download, but not all client support continuation of a stopped download. For a big download you really need to have this just in case. However, *lftp* does, and has a lot of other features that make it a truly excellent file transfer tool that once you start using you will never want to go back to using other tools.

lftp is command line file transfer program (similar to *wget*) and an interactive shell (similar to most ftp clients). But don't let the name deceive you as *lftp* also supports ftp and http protocols. *lftp* is a ftp client on steroids it has many extremely useful features. The most important is that *lftp* is reliable so that any non-fatal error can be recovered from and the transfer is retried. So if you are downloading a 25MB kernel image and it stops, *lftp* can be restarted from that point automatically if you want, or you can continue the download later. To support the continuation of download the ftp server needs to support the **REST** command. However if the server does not support this command lftp will try to retrieve the file from the very beginning until the file is transferred completely, which is the worst case situation. Interestingly, and a point that needs to be noted, if you exit *lftp* with a file (or files) transfer not complete, *lftp* will go to **NOHUP** mode and continue in the background. This is both useful and a pain if you really *did* want to stop all file transfers, as you have to remember to stop this with the *lftp* tool, or kill the *lftp* process when you exit.

One of the most useful features of *lftp* is its shell-like command syntax. This allows you to launch several commands in parallel in background (**&**), group commands within **()** and execute them in background. All background jobs are executed in the same process which allow you to bring a foreground job to background with **Ctrl-Z** and back with command the **wait** (or **fg**).

To list running jobs, you would use the command **jobs** just as you would within a command line shell. Also some commands, such as *cat*, allow redirection to file or via pipe to external command. As per the *Bash* shell there is also support for autocompletion, which saves you from having to type in long file or directory names, and command history, accessible via the arrow keys. Lastly, commands can be executed conditionally based on termination status of previous command using the standard **&&** and || symbols.

Starting *lftp* is the same as for any other file transfer client:

```
jkent@smokey> lftp ftp.kernel.org
lftp ftp.kernel.org:/> (cd
 /pub/linux/kernel/v2.4/ && get
 kernel-2.4.19.tar.bz2 &)
lftp
 ftp.kernel.org:/pub/linux/kernel/v2.4>
 open http:/www.uk.debian.org
lftp ftp.kernel.org:/> (cd /debian/
dists/Debian3.0r0/main/
disks-i386/base-images-current &&
 get basedebs.tar &)
lftp ftp.kernel.org:/debian/dists/
Debian3.0r0/main/disks-i386/
base-images-current> quit
jkent@smokey>
```

The example above first connects to ftp.kernel.org changes to the latest kernel directory and starts a download of kernel2.4.19.tar.gz and puts this into the background. Then a new connection to www.uk.debian.org is made and the file basedebs.tar is downloaded and put, again, into the background. Then the *lftp* client is shutdown. As *lftp* is still running, even though the *lftp* interface has been stopped, these two files will continue to be downloaded, and if there are any connection problems the downloaded will continue once the connection has been re-established. As you can see from this example *lftp* additional features provide extremely useful and versatile capabilities.

lftpget

Of course you do not have to use the interactive interface to download a file, you can use *lftp's* **lftpget** command to download a file, much as you would with wget.

```
lftpget http://www.kernel.org/
pub/linux/kernel/v2.4/
kernel-2.4.19.tar.bz2
```

lftp is one of the tools that once you use it you never go back to using another. The additional functions are so useful that you will wonder why this approach is not used more often in a file transfer client. This little utility is a must have installation.

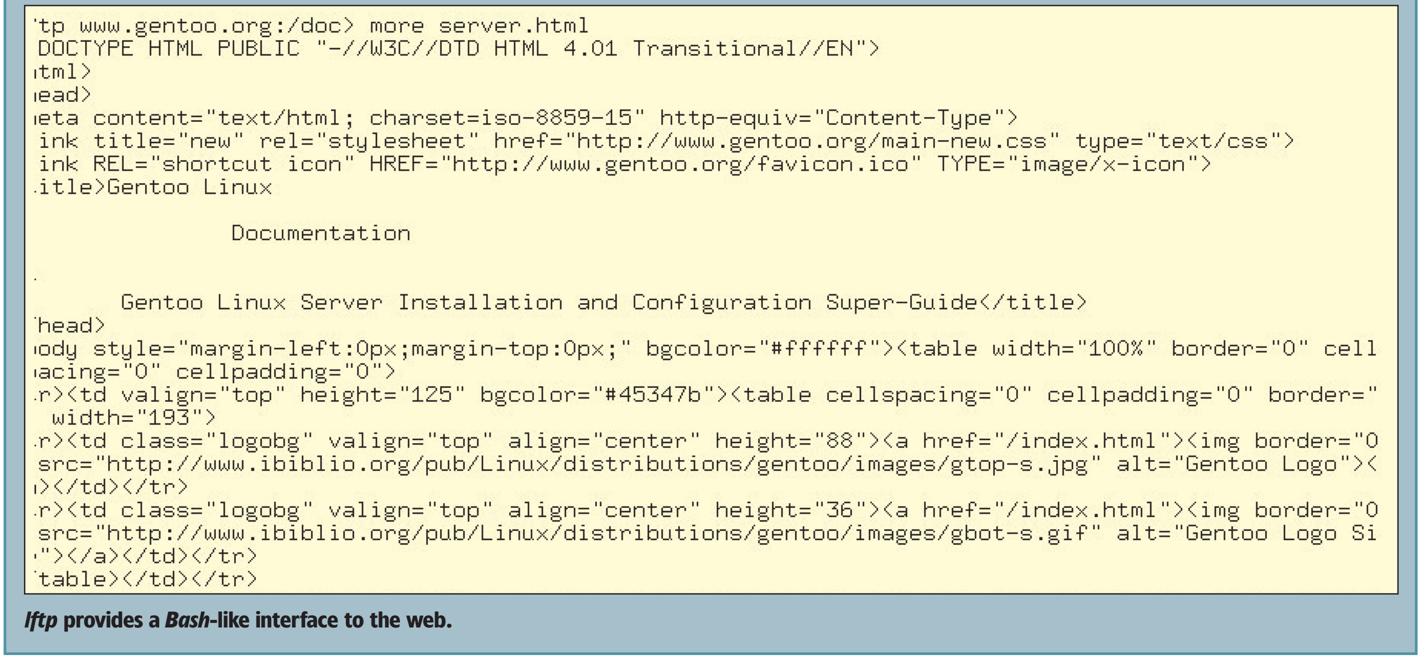

```
'tp www.gentoo.org:/doc> more server.html
 DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01 Transitional//EN">
ıtml>
ıead>
ıeta content="text/html; charset=iso-8859-15" http-equiv="Content-Type">
 ink title="new" rel="stylesheet" href="http://www.gentoo.org/main-new.css" type="text/css">
 ink REL="shortcut icon" HREF="http://www.gentoo.org/favicon.ico" TYPE="image/x-icon">
.itle>Gentoo Linux

                Documentation

.
       Gentoo Linux Server Installation and Configuration Super-Guide</title>
'head>
ıody style="margin-left:0px;margin-top:0px;" bgcolor="#ffffff"><table width="100%" border="0" cell
ıacing="0" cellpadding="0">
.r><td valign="top" height="125" bgcolor="#45347b"><table cellspacing="0" cellpadding="0" border="
  width="193">
.r><td class="logobg" valign="top" align="center" height="88"><a href="/index.html"><img border="0
 src="http://www.ibiblio.org/pub/Linux/distributions/gentoo/images/gtop-s.jpg" alt="Gentoo Logo"><
ı></td></tr>
.r><td class="logobg" valign="top" align="center" height="36"><a href="/index.html"><img border="0
 src="http://www.ibiblio.org/pub/Linux/distributions/gentoo/images/gbot-s.gif" alt="Gentoo Logo Si
ı"></a></td></tr>
'table></td></tr>
```

***lftp* provides a *Bash*-like interface to the web.**

GRAPHICAL MONITORING TOOL

GKrellM

■ **VERSION** 2.0 ■ **WEB** web.wt.net/~billw/gkrellm/gkrellm.html

As useful as utilities such as *top* or *vmstat* are they are not the easiest of tools to use to use to gain a quick snapshot of how your system is currently behaving. For example, *top* takes up too much space, as you need to run it in a terminal, and *vmstat* is time limited, although you could set a very long time period – and, again, run in a terminal. Sometimes it is also useful to be able to see how a remote system is running as well. *GKrellM* has been designed with this and a lot more in mind. *GKrellM* is a set of graphical utilities that is displayed in a column, with each block in the column displaying different statistics about your system in a graph or other formats. The bizarre name apparently stands for *GNU* (or *Gtk*) *Krell Monitors* (or *Meters*).

GKrellM can display the time and date, graphs for CPU activity, process load, disk I/O activity, network interface traffic, resource meters for memory, swap, and disk partitions, the number of new emails in your inbox, and the local system's uptime. Additionally there are many plugins available to extend the capabilities of *GKrellM* beyond these defaults, such as a remote system monitor, screensaver lock, an *XMMS* front end and a moonphase dial. Some of these plugins are can be very useful, whilst others are just plain strange, but they are there if you wish to use them.

When you first start *GKrellM* you are presented with a display with most of the options switched on. You need to configure *GKrellM* to work correctly with your system and to provide the details you are interested in seeing. Luckily configuring *GKrellM* is straightforward as you just right-click title bar and select the configuration option. Additionally, you can configure individual components by right-clicking on the component you will to modify and a configuration window will be displayed. The configuration windows are well laid out and make configuring the component straight forward. It takes no time at all to get *GKrellM* up and running monitoring your system in your preferred manner.

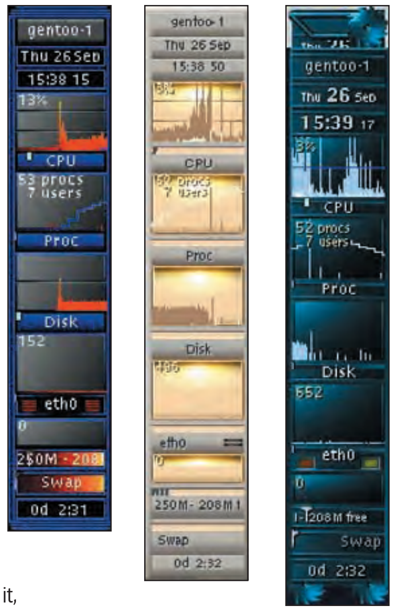

***GKrellm* graphically shows how you system is performing.**

Don't like the way *GKrellM* looks? Well that's easy to fix as it supports themes, of which there are many to download from the website, and from themes.org. You will easily find a theme that suits your desktop. On the *GKrellM* website you will also find many plug-ins to add, and there are other available on the web.

Once everything is set up you can also configure *GKrellM* to remember where you have located it, so the next time you start it you do not need to worry about relocating it to its proper location. This does not work will all window managers, but seems to work with all of the major ones (GNOME, KDE, *WindowMaker*, *Fluxbox*, *XFCE* and so on).

Although *GKrellM* could be considered eye-candy, it is in fact very useful to be able to quickly glance at your screen and get an overview of the system. I find this very useful as I tend to run with three Linux boxes through a KDM switch and another on a portable, so I find it useful just to glance quickly at the screen to get a good overview of the system's state. Even when I am running only one system, it is still useful to be able to quickly get this information. If I see something amiss I can then look into it in more detail. A very useful tool to have running on your system.

MULTIPLAYER TETRIS

Crack Attack

■ **VERSION** 1.1.8 ■ **WEB** http://qcd2.mps.ohio-state.edu/attack/

Obviously you cannot spend all your time working on your Linux systems, you also need time to relax. So in usual *HotPicks* fashion here is another game for you to try out, namely *Crack Attack*. *Crack Attack* is based on the Super Nintendo classic *Tetris Attack* which takes the classic game of *Tetris* and then adds several new layers of game play to it. In fact, *Crack Attack* probably owes more to the classic game *Columns* than *Tetris* as the rules are basically the same.

You can play *Crack Attack* either against your computer or against your friends over the network. In either mode the rules are the same. You must ensure that your stack of coloured blocks, which slowly grow from the bottom of the screen, do not reach the top of the screen, otherwise you will lose the game. To eliminate blocks from the stack you need to line up at least three blocks of one colour, horizontally or vertically. Once you do, those blocks disappear and you extend you life span.

Nice and simple so far, but occasionally chunks of large red blocks will fall from above which cannot eliminate directly. But if you eliminate any blocks which are touching the red block, the red block will shatter into normal, run of the mill blocks which you can, hopefully eliminate in the usual fashion.

Throughout the game your stack of blocks will grow from below – initially quite slowly, but gradually picking up speed. If the stack does reach the top it will pause and you have seven seconds to save yourself. It is possible, however, to get a momentary pause in the stack's ascent by eliminating a block, or by causing red blocks to shatter. Even if your stack has reached the top and you're about to lose, you will be given a little more time to save yourself. Luckily the time during which a block is disappearing or a chunk of red blocks is shattering does not count towards your seven seconds.

The cause of these red block is your opponent's way of trying to eliminate you and also your way of eliminating him. If you, or your opponent, eliminates four, five, or even more blocks at one time, you will not only be rewarded with a little light show, but red blocks will come raining down onto your opponent's stack. The bigger the elimination, the bigger the downpour of red blocks.

Beyond the multiple elimination method, there is yet another technique to generate these troublesome red blocks and that is via a combo. Whenever you eliminate blocks, it's likely to cause other blocks to fall as they were above the eliminated blocks. If, when these blocks fall, they cause a second elimination, this is referred to as a combo. If you're good enough to keep the combo rolling, your combo multiplier will increment for each elimination. When the combo is finally complete, a massive chunk of red blocks will descend upon your opponent, it's size proportional to your final multiplier.

Solo play works a little differently. Instead of trying to last longer than an opponent, now you're attempting to rack up as many points as possible before you die. The number of points you get for an elimination depends on how fancy it is. When you get a combo, all of the points you earned in the eliminations which made up that combo are multiplied by the combo's highest multiplier. However, in solo play, eliminations generate red blocks, except the red blocks falls on you.

Crack Attack is another simple, yet extremely additive game. It is great fun to play in solo mode, but if far more fun when you play against you friends. The graphics are well done and the game play has been well thought out. Cracking little game.

HARDENING SCRIPTS

Bastille Linux

■ **VERSION** 2.0BETA ■ **WEB** www.bastille-linux.org

Although *Bastille Linux* may sound like another distribution it is, in fact, a set of Perl scripts designed to help you secure your system. You can either run *Bastille* in interactive mode, select a template, or build you own *Bastille* configuration file and run *Bastille* against this configuration file. Of course, if you are already familiar with Linux security you can just use *Bastille* to save you time hardening your system.

This review covers the latest version, which is currently *2.0BETA*, although this should be production ready shortly. If you are using Red Hat 7.3 this version is the only one you can use as previous versions of Bastille do not support it. The distribution support list only covers Red Hat and Debian, although it should work, with some modifications, with other distributions – although the benefits of doing this are limited. Why is this? Well it's mainly down to location of configuration files and binaries, as these are still not standard across different distributions. Eventually issues such as this may be sorted out by the LSB, but don't hold your breath.

Interactive

The easiest way to get started with *Bastille* is via the interactive mode, which can be run from the command line (**bastille -c**) or from X using a *Tk* frontend (*bastille -x*), which is the default. If you have X, running the X mode is the easiest to use – however if you are hardening a server why have you got X installed? If you are using an version earlier than *2.0BETA, Bastille* is started by running **InteractiveBastille.pl**.

In interactive mode *Bastille* explains in reasonable detail what each question means and the actions that *Bastille* takes if you decide to implement that option. If you already know what you are doing you can run *Bastille* in **explain less** mode, which is a button on the bottom of the *Bastille* window.

Once you have interactive *Bastille* running you then have to go through a fairly extensive list of questions, which usually consists of just indicating if you wish to enable or disable a particular security function. These sections cover all areas you would expect, such as file permissions, firewalling, mail securing various services and so on. Once you have completed all the questions your configuration is saved to /etc/Bastille/config but is not actually implemented at this time. This gives you the opportunity to review your configuration and make any further changes that you may have missed. To implement your *Bastille* configuration you run **BastilleBackEnd** (or **BackEnd.pl** in pre *2.0BETA* versions).

If you prefer you can run *Bastille* using the **AutomatedBastille** script (**AutomatedBastille.pl** in pre-*2.0BETA*). This approach is more simplistic as you just choose a security baseline you want to implement and then this is applied to you system without having to run the **BastilleBackEnd** script. Obviously you have far less control over the changes that *Bastille* will make to your system, but you may find that this implements all the changes you require.

Bastille Distribution support information

LINUX DISTRIBUTION	VERSION OF BASTILLE SUPPORTED
Red Hat 6.2	1.2.0
Red Hat 7.0/1/2	1.3.0 and 2.0
Red Hat 7.3	2.0
Red Hat Advanced Server	2.0 (unofficial support)
Debian 2.2	1.2.0
Debian 3.0	1.3.0
Debian Testing	1.3.0
Debian Unstable	1.3.0 and 2.0 soon to be released
Mandrake 8.0/1/2	1.3.0

To verify that the **BastilleBackEnd** script has worked correctly it is a good idea to look at the logs that it creates. Two logs are created, action-log and error-log, both of which provide very verbose details of the changes that have been made and any errors that occurred. It is worth reviewing these logs to ensure that everything has been implemented correctly.

Continuing process

You may now be thinking that once you have installed *Bastille* you have a secure system. Unfortunately that is not completely true. *Bastille* does not disable all service that may be running on your system, so it worth reviewing which services are running and disabling those that you do not require. In addition you should also verify that /etc/passwd only contains user ids that your system requires as a lot of distributions load this file up with user ids that you do not need.

As system hardening is not necessarily a process that you perform regularly, *Bastille* provides a means to implemented a security policy on your system without having to dig out a book or surf the 'Net for hours on end. It is well implemented and easy to use, but just needs to expand its distro horizon beyond Red Hat and Debian. It cannot be expected to make your system 100% secure, you do have to do some work yourself and there are other issues, such as buggy apps, which need to be taken into account. However *Bastille* provides a solid base to build upon and is worth installing for the added piece of mind. LXF

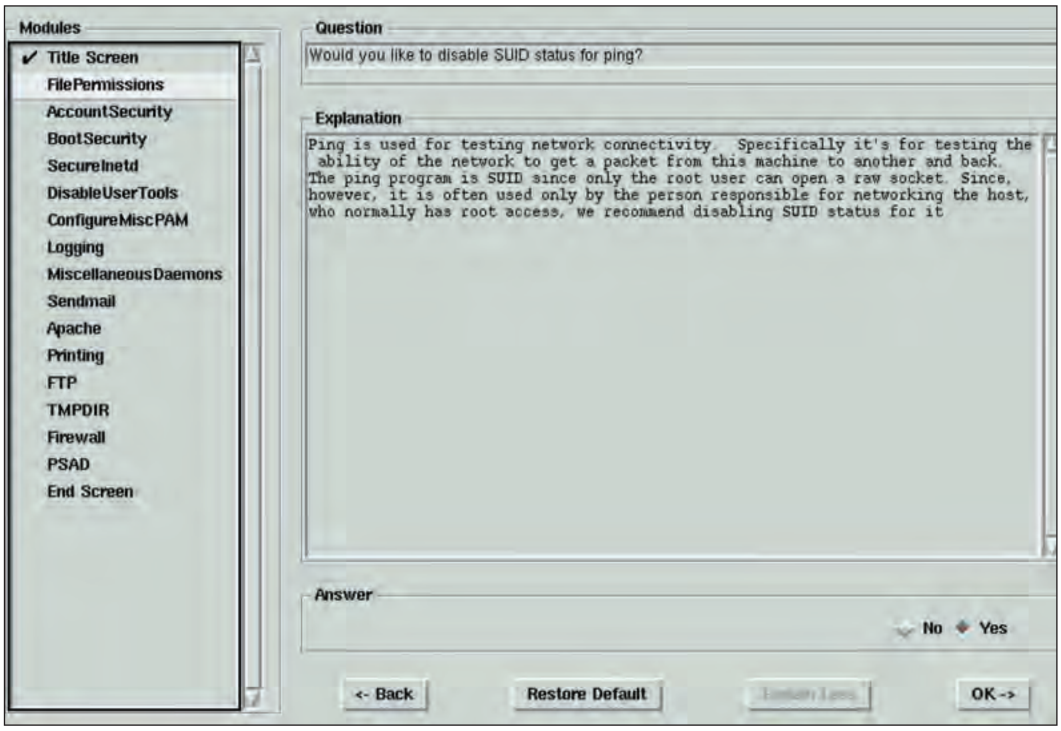

Would you like to disable SUID status for ping?

```
Ping is used for testing network connectivity.  Specifically it's for testing the
 ability of the network to get a packet from this machine to another and back.
The ping program is SUID since only the root user can open a raw socket. Since,
however, it is often used only by the person responsible for networking the host,
who normally has root access, we recommend disabling SUID status for it
```

Bastille **provides a friendly interface to securing your system.**

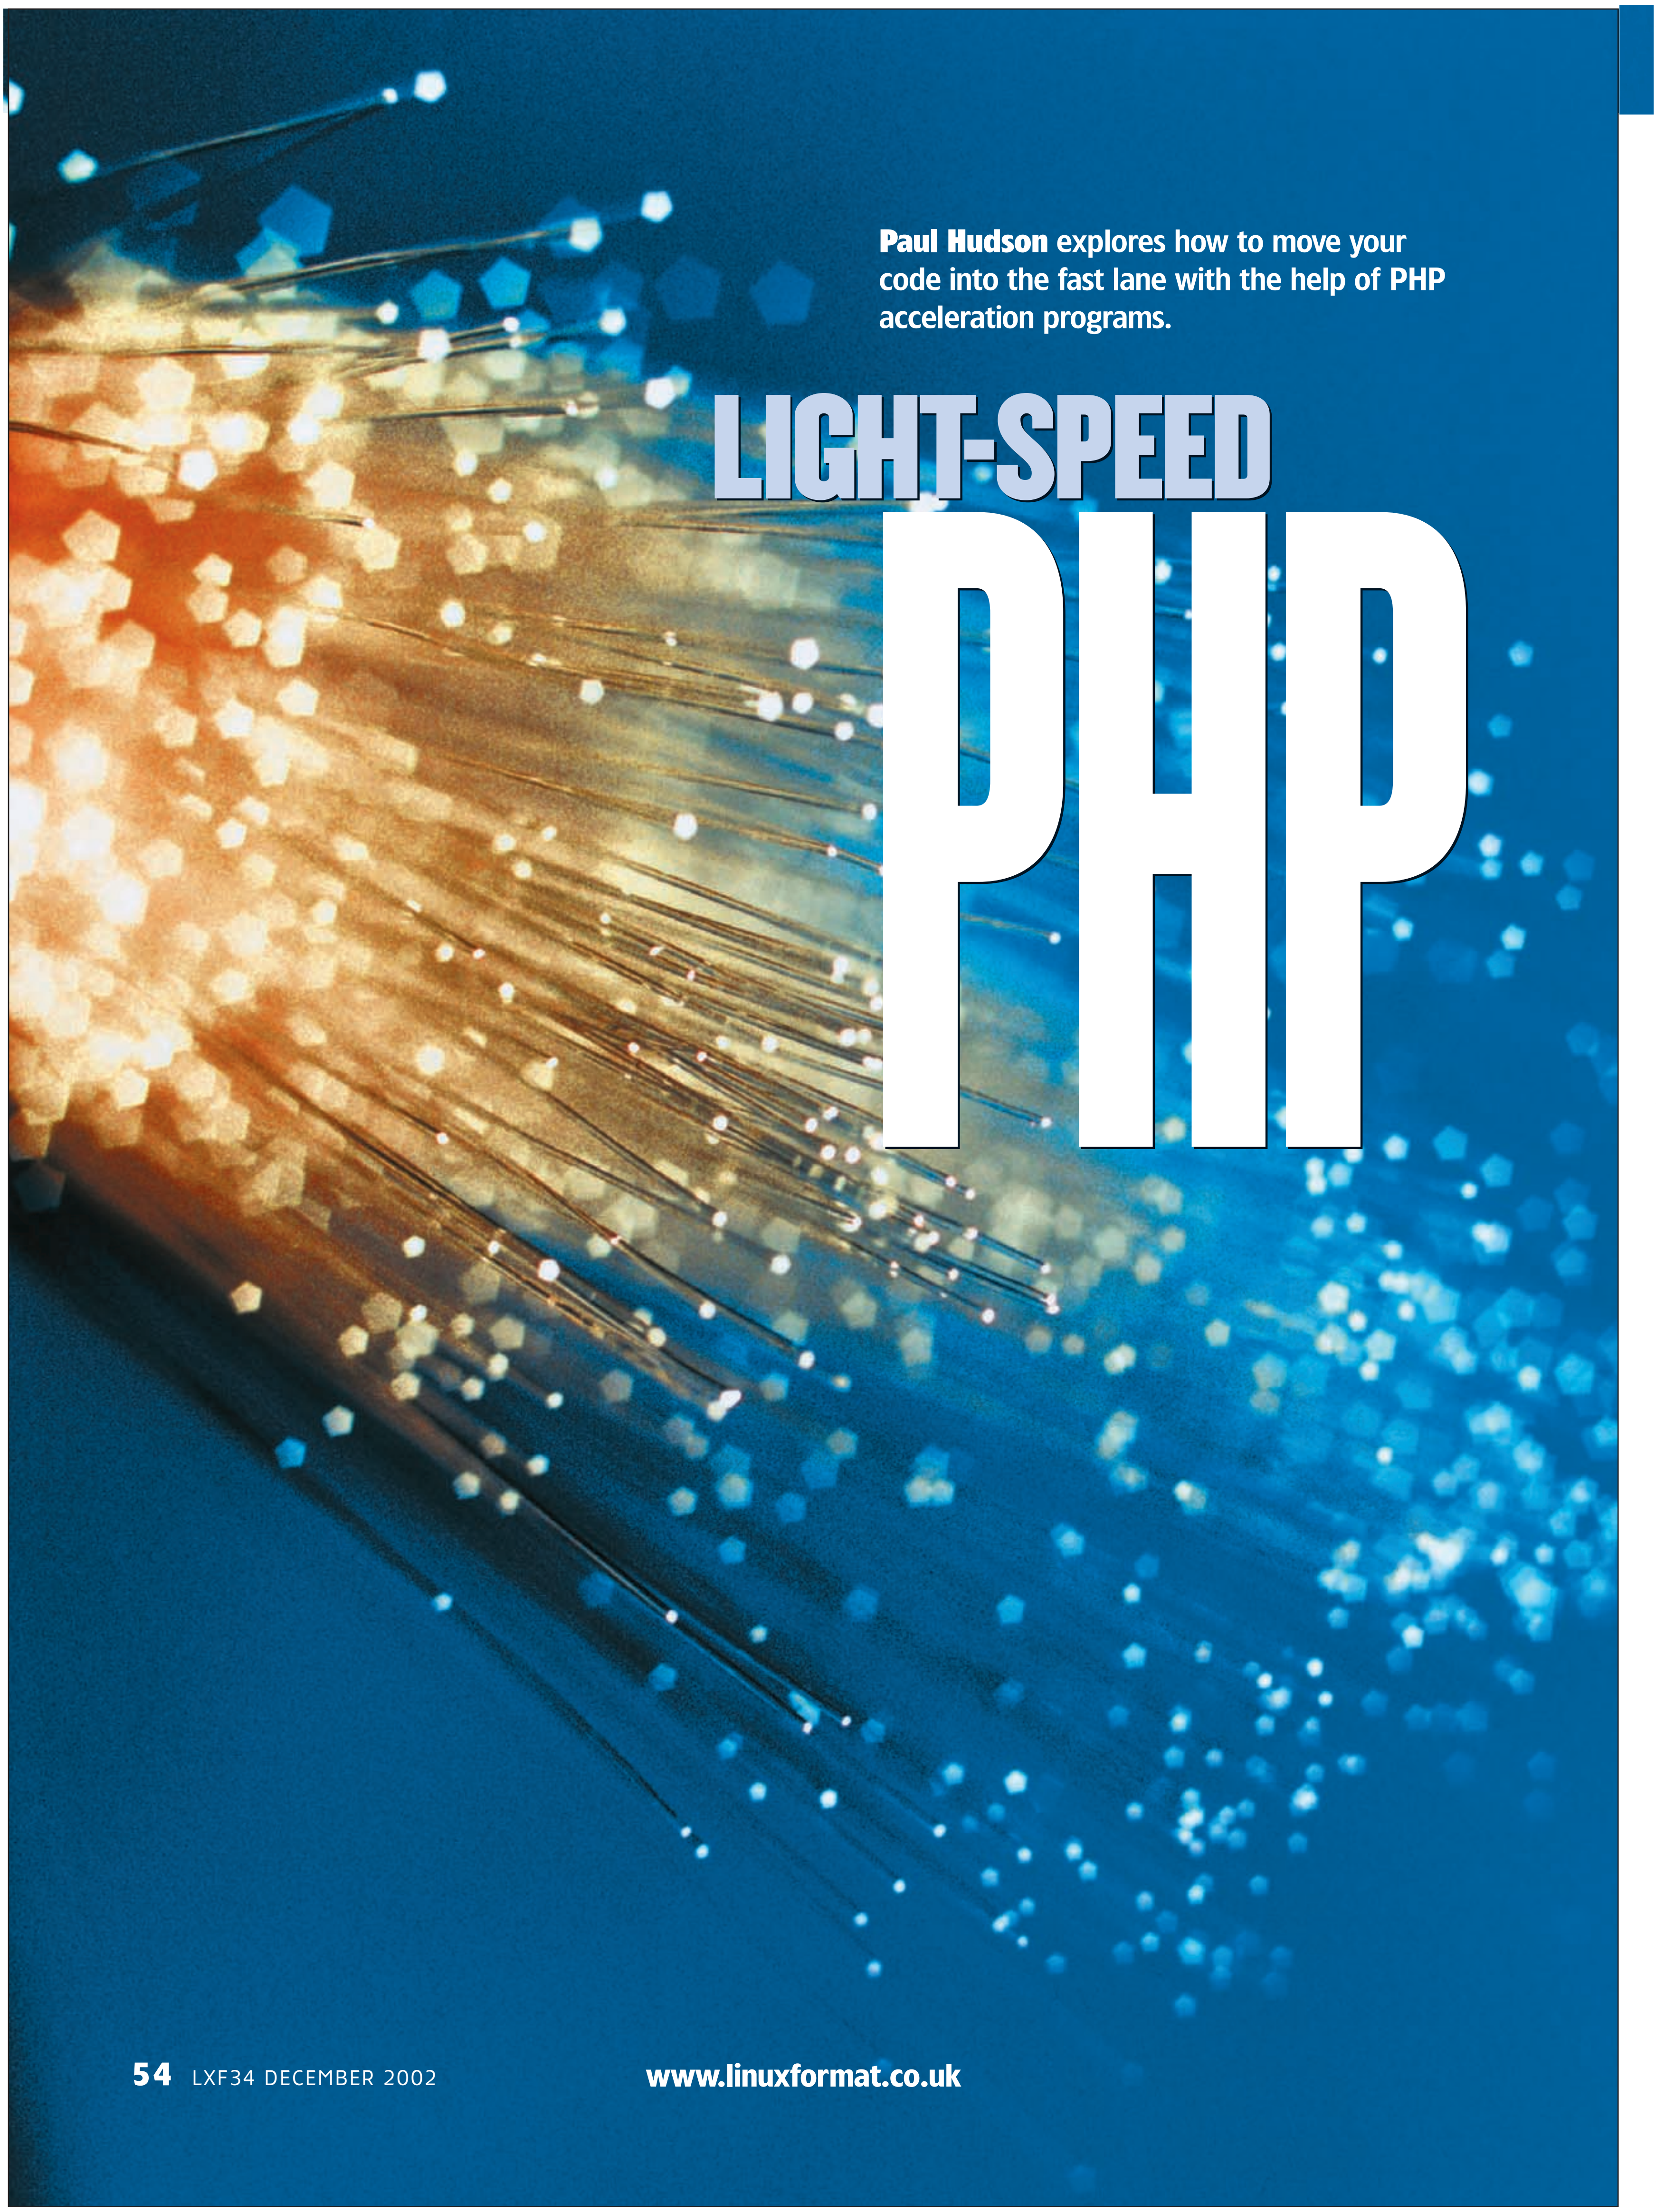

Paul Hudson explores how to move your code into the fast lane with the help of PHP acceleration programs.

LIGHT-SPEED PHP

As PHP continues to become increasingly popular with webmasters around the world, so the need increases for PHP-processed pages to be outputted faster in order to keep up with server load. Traditional methods of improving web server performance have invariably involved storing popular HTML pages in RAM so they can be served quicker than from disk.

Owing to PHP's dynamic nature, it's rarely plausible to cache its output – unless you want everyone to share the same shopping basket! To solve this problem, several programs have been created that cache your compiled code, not its output.

Before PHP 4 was released, the PHP parser executed each line of code as it was parsed. This led to slow execution times for scripts (especially within loops and functions), and the inability to cache any part of the process. PHP 4 split the process of reading source code into two parts, parsing and executing – take a look at the diagram, *above right*, to see what I mean. As a result of this split, code execution took a leap forward in speed, which was one of the most popular reasons for upgrading to PHP 4 from PHP 3.

The PHP engine converts PHP source code into compiled code, ready to be executed. At this point, database queries haven't been run, functions haven't been called, and so on – they have only been prepared. This compiled code is what is stored in RAM so that the next time a request comes in for the same page, the compiled page is executed from the cache – PHP doesn't have to re-read and re-parse your source code. This is important, because it means pages that are cached are *still* dynamic because they haven't been executed yet, whereas with normal caching (of the type popular for static-content servers) the PHP would be parsed *and* executed, with the output being cached.

Not content with eliminating one of the slowest parts of PHP processing, some programs also *pre-process* your compiled code, eliminating common performance hindrances and optimising your code.

While some people already cleverly eliminate most performance issues as they write their code (for example, pre-incrementing variables rather than post-incrementing them), this brings PHP a step closer to C++ compilers, which perform a massive variety of optimisations on your code.

In this article I'll be looking at the four PHP acceleration programs available on the market at the time of writing, comparing their features, performance, ease of installation, and price. If you were worried that accelerating programs may serve up old versions of your recently-edited pages, don't fret – all the programs covered here can be configured to automatically check for a newer copy of cached files.

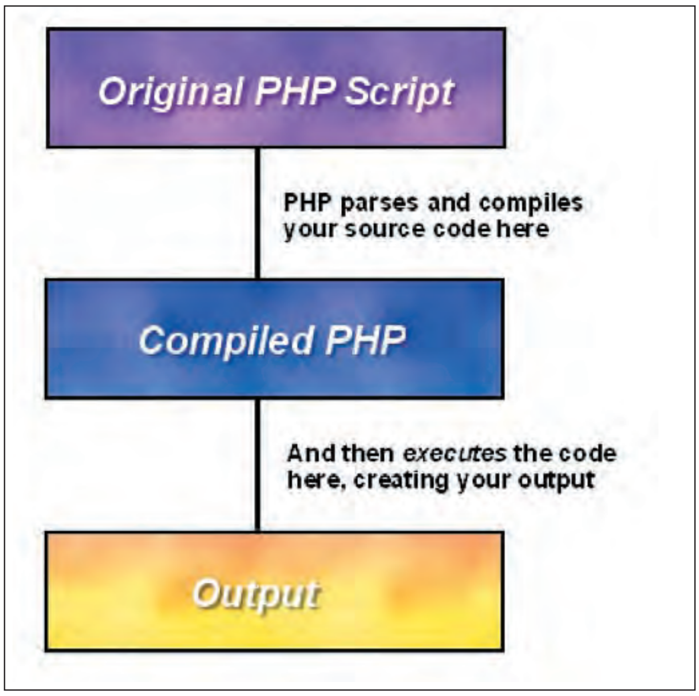
Original PHP Script
PHP parses and compiles your source code here
Compiled PHP
And then executes the code here, creating your output
Output

PHP 4 splits processing into two steps.

> **"Owing to PHP's dynamic nature, it's implausible to cache output unless you want everyone to share shopping baskets"**

ZEND ACCELERATOR

Zend, the makers of Zend Studio (reviewed in *Linux Format 31*), also produced the first caching solution for PHP. With Zeev Suraski, one of the lead developers of PHP, as the Chief Technology Officer at Zend, I had very high expectations of the *Zend Accelerator*, and certainly was not let down.

The installation was dialog-based, and, after a series of basic questions, Zend had detected my *Apache* installation and installed itself as a PHP extension. After restarting *Apache* (yes, it did *that* for me too), *Zend Accelerator* was installed and running.

It also installed its somewhat dubious online management tool that, whilst looking attractive to the eye, doesn't seem to have any figures that correspond to reality. For example, it told me that over half my scripts were sped-up by a factor of five. Given that I was using a fresh installation of Debian and *Apache* and the results certainly weren't *that* fast, I'm very skeptical about these numbers. However, the interface is pretty, for what it's worth ;)

With regards to features, *Zend Accelerator* caches your compiled script and also performs various code optimisations. You can also configure scripts you specifically want *Zend Accelerator* not to cache, how much memory it is allowed to use, and the maximum number of files it can cache.

On the performance front, *Zend Accelerator* consistently leads the pack – sometimes by quite a large amount. As can be seen in the graphs, *Zend Accelerator* really is currently the only choice for people serious about using *Smarty*, while also providing an excellent performance boost for all varieties of PHP operations.

In addition to being the fastest in the group, *Zend Accelerator* also manages to survive on only tiny amounts of memory, with memory usage being second only to the ionCube *PHP Accelerator*. One definite advantage to *Zend Accelerator* that should not be under-estimated is its close alliance with PHP development. As soon as a new

»

PHP release is out (pretty much to the day of release), *Zend Accelerator* is updated to work with it so that users can upgrade without problems. Zend are to be applauded for putting so much work into keeping up to date with PHP releases, as this helps take the compatibility sting out of upgrading. The downsides to *Zend Accelerator* are that it is proprietary, closed-source software, and that it comes with quite a hefty price tag – it's always an ominous sign when a company doesn't make prices obvious on their web site. The price starts at approximately $980 for a computer with 1 CPU, but other configurations come at different (higher) costs, which is more than most casual web developer have lying around and as such firmly places *Zend Accelerator* in the high-end/enterprise sector of the industry.

In some ways, being targeted at the enterprise sector would be a good thing for the *Zend Accelerator*. Firstly, there may be managers around who would find useful the pretty pictures and dubious bar charts in the management interface, and also increasing the throughput on each server to its maximum reduces the number of servers required, thus also reducing points of failure.

In the end, it pretty much comes down to this – are you willing to pay for the best performance available? Zend are banking that, for many people, the answer will be a resounding "Yes".

LINUX Format VERDICT

Ease of use	9/10
Features	7/10
Performance	10/10
Value for money	5/10

The fastest accelerator on the market, also with the smoothest installation. Ignore the management interface, though, and sit down before you see the cost for your setup.

LINUX Format **RATING** 8/10

The *Zend* web interface is very much designed for managers – while it looks pretty, the attraction is mostly skin deep.

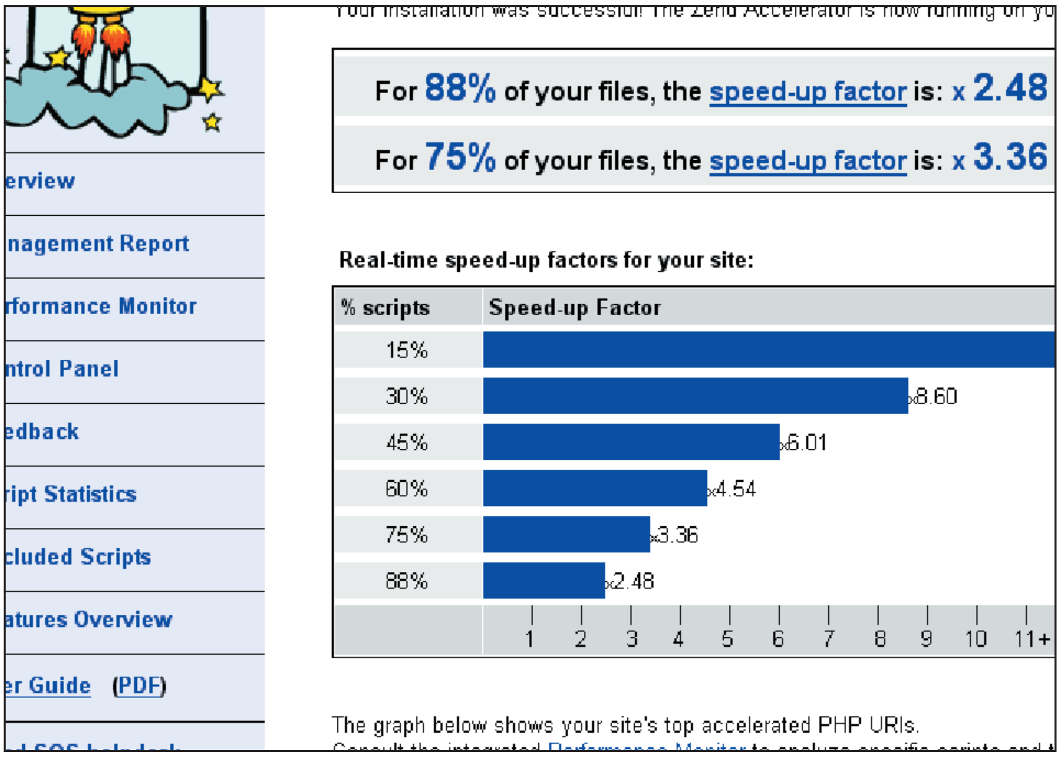

Your installation was successful! The Zend Accelerator is now running on yo

For 88% of your files, the speed-up factor is: x 2.48

For 75% of your files, the speed-up factor is: x 3.36

Real-time speed-up factors for your site:

% scripts	Speed-up Factor
15%	
30%	x8.60
45%	x6.01
60%	x4.54
75%	x3.36
88%	x2.48

The graph below shows your site's top accelerated PHP URIs.

APC

The Alternative PHP Cache (APC) was the second PHP acceleration solution to be released, but more importantly, it was the first *free* solution, and that gained it a large user-base in a comparatively short period. The advantage to having been around so long is that it arguably has a more mature codebase than the newer products, and has also had time to filter through into distro channels. For example, being a Debian user, I installed APC through Debian's *apt-get*. Beyond the smooth installation, however, things start to get a little rocky – not only was APC the worst performer of the four on offer, but the documentation (littered with such choice phrases as "we were too lazy to provide a workaround") was a long way from adequate. In some cases, *Zend Accelerator* was approaching twice the speed of *APC*, so, unless you really cannot use the other options for some reason, I wouldn't recommend you choose *APC* on the basic of performance.

One bright spot in the program is that it introduces several PHP functions to control caching. For example, calling **apc_reset_cache()** forces an update of all the cache content, whereas **apc_rm("/some/filename/here.php")** resets the cached copy of that one filename. This doesn't really make up for the lack of performance, but it definitely helps. Licenced under the QPL, the *APC* source code is available freely for download, and that's always an advantage. Don't get me wrong – *APC* is definitely heaps better than not using any accelerator at all, but it's performance is matched and beaten by all the alternatives. The only situation in which I would think of recommending *APC* is for administrators who want an easy installation of a free solution, because *APC* can often be installed with your package manager.

LINUX Format VERDICT

Ease of use	6/10
Features	5/10
Performance	3/10
Value for money	10/10

Been around for a long time, but still came bottom of the tests. Not recommended, other than for complete beginners.

LINUX Format **RATING** 6/10

BWARE AFTERBURNER

As the only free accelerator available for Windows, *Afterburner* has a larger user base than the other three covered here – but that's hardly an advantage for Linux users! What *is* an advantage, though, is its customisation options – it allows you to log its caching activity to check for problems, and also has a superior cache limiting system that allows you to be very precise as to which files you want cached and which not. On the performance front, *Afterburner* is a middle player, coming almost exactly half way between unoptimised PHP and the leading solution. Being the only GPL (well, LGPL) accelerator on offer, the satisfactory performance should be enough for most die-hard GPLers to be happy to use it over the other offerings. It isn't without problems, though. For example, don't expect immediate support for the latest versions of PHP. At the time of writing, the latest release supported for PHP 4.1.x, and was uploaded back in February. Secondly, the documentation is very poor indeed, which makes installation tricky at best. The options are listed in the accompanying README, and, being fairly self-explanatory, you can get *Afterburner* up and running with just a few edits. However, those of you less willing to edit configuration files may feel more comfortable with any of the other three programs, all of which come with installation guides. Whilst not the top performer, *Afterburner* does do very well for itself, and you could choose worse. If you run a mixed server farm (that is, Linux and Windows), it may be advantageous to you to run the same configuration on all the machines, in which case *Afterburner's* portability counts for a lot.

LINUX Format VERDICT

Ease of use	7/10
Features	6/10
Performance	5/10
Value for money	10/10

A middle player on the performance front, but is let down by less than adequate documentation. Good configuration options.

LINUX Format **RATING** 7/10

IONCUBE PHP ACCELERATOR

My initial impression of ionCube's *PHP Accelerator* (*PHPA*) was first class – the product's homepage is the best of the four, and includes performance information, an FAQ, forums, and more. Of course, if I were a less scrupulous writer, the fact that *Linux Format* is mentioned on the site would earn *PHPA* an extra point ;)

Installation was straightforward, and comes bundled with excellent documentation describing the various configuration options available. Like *Zend Accelerator*, *PHPA* keeps up to date with PHP releases very well, and even has a download for the CVS release of PHP for the more adventurous among us. Also like *Zend Accelerator*, *PHPA* has an optional web front-end for administration and monitoring, and it comes with a selection of features that *Zend Accelerator's* equivalent doesn't, including a "Detailed file info" view. The focus here is quite different – functionality rather than favourable looks – and this makes it more useful than Zend's version. With regards to performance, *PHPA* is sometimes almost neck-and-neck with *Zend Accelerator* which is remarkable considering it was mostly developed by just one person – Nick Lindridge. Incidentally, Nick wrote an excellent PDF document on how *PHPA* works, which provides a good insight into PHP acceleration on the whole – read the box *How PHP optimisation works* for more information. *PHPA* fell down on the Basic Tests benchmark (where it placed behind *Afterburner*) and also the Smarty test. Given that *PHPA* is actually mentioned and recommended on the *Smarty* site, I think it's quite likely the two development teams will work together to improve performance (if they aren't doing so already), so that just leaves the Basic Tests.

The contents of these tests are discussed later on, in the section *What's in the tests?*, but suffice to say that the Basic Tests are very small scripts performing single operations (counter scripts, XML parsing, etc), so it is possible that *PHPA* is better at handling larger scripts than smaller. In the short time between the tests being performed and this being written, a newer version (*1.3.3*) has been released, so it is safe to say that *PHPA* is still being improved. Overall, *PHPA* is a great product to use, let down only by under-performing in the popular *Smarty* template engine. Through community help, it has been extended to work on four platforms other than Linux (BSDi, FreeBSD, OpenBSD, and Solaris), so if you're using these platform in conjunction with your Linux server, *PHPA* makes a great choice to unify on a single product. I've also heard on the grapevine that a Windows version of *PHPA* is "now a real possibility", which will be welcome news to those of you who also have Windows somewhere in your server mix.

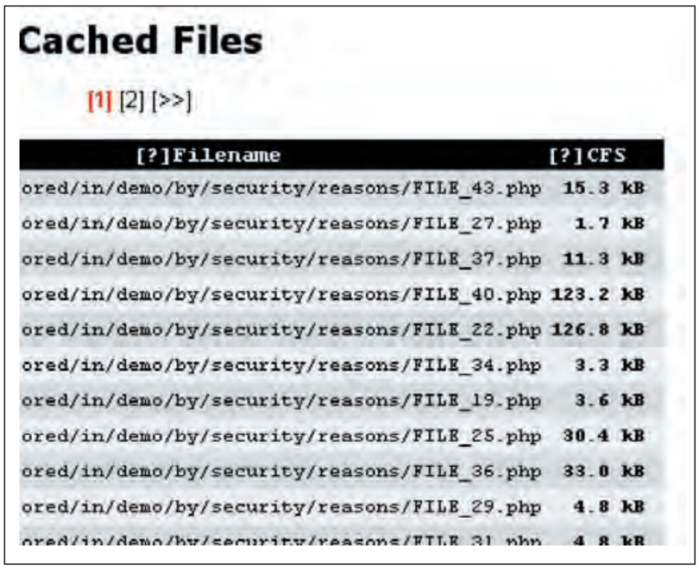

Cached Files

[1] [2] [>>]

```
            [?]Filename                              [?]CFS
ored/in/demo/by/security/reasons/FILE_43.php  15.3 kB
ored/in/demo/by/security/reasons/FILE_27.php   1.7 kB
ored/in/demo/by/security/reasons/FILE_37.php  11.3 kB
ored/in/demo/by/security/reasons/FILE_40.php 123.2 kB
ored/in/demo/by/security/reasons/FILE_22.php 126.8 kB
ored/in/demo/by/security/reasons/FILE_34.php   3.3 kB
ored/in/demo/by/security/reasons/FILE_19.php   3.6 kB
ored/in/demo/by/security/reasons/FILE_25.php  30.4 kB
ored/in/demo/by/security/reasons/FILE_36.php  33.0 kB
ored/in/demo/by/security/reasons/FILE_29.php   4.8 kB
ored/in/demo/by/security/reasons/FILE_31.php   4.8 kB
```

The *ionCube* management interface needs a bit of work, but beneath its ugly exterior lies smart and compact functionality.

LINUX Format VERDICT

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

Excellent performance and the best documentation and support available. Lags behind Zend by a smidge.

LINUX Format **RATING** 9/10

How does code optimisation work?

With ionCube's *PHPA* being the fastest free PHP accelerator out there, Nick Lindridge wrote an essay on how the PHPA optimisation works. Read the full text at www.php-accelerator.co.uk/PHPA_Article.pdf. Part of it is reproduced here with permission.

***PHPA* performs a number of peephole optimisations on the compiled code to reduce its size and improve performance. Due to the way that compiled code is often generated, such code can be inefficient unless optimisations are made prior to code generation, or if the compiler is particularly dumb. For example, the compiled code for if statements without an else clause contains an unnecessary jump to the very next instruction at the end of the if block, and code may contain jump instructions that jump to other jump instructions. The handling of double quoted strings is also inefficient, with the language parser tokenising each word and separator within the string, and code generation producing individual string concatenation instructions. Other optimisations include turning post-increment/post-decrement into pre-increment/pre-decrement where possible, elimination of unconditional jumps to other unconditional jumps or return statements, reordering of some code constructs to eliminate jumps, and removing unreachable code.**

What's in the tests?

A total of seven tests were conducted to get the results you see. The first four, shown collectively as the Basic Tests, involve file reading, database access, XML parsing, and a mixed test. The mixed test simply uses the PHP **include()** function to run the other tests. The fifth test loads the index page of a default *Smarty* install, with no configuration options set. The sixth test loads the forum index page from a default install of *phpBB*, with several forums on there to better simulate a live site. The seventh and final test loads the main page from *phpMyAdmin*, which had been configured to load data from the local *MySQL* server.

All scripts were served from the same computer, a 500MHz x86 PC with 256MB RAM and running a fresh install of Debian 2.2, with *Apache 1.3.26* and PHP 4.1.2. No optimisation or tweaking has been performed; everything used the "factory default" settings where possible. The testing was performed on a separate computer using *ApacheBench* with the parameters **-n 1000 -c 10** (run 1000 requests, 10 at a time). Each graph represents the number of requests served per second by *Apache* in each test.

Detailed Results

The form book

Test	Unaccelerated	Zend Accelerator	APC	Bware Afterburner	ionCube PHPA
File manipulation	107.49	122.37	113.99	114.12	111.7
MySQL querying	13.66	16.18	13.7	15.74	13.72
XML parsing	81.72	106.23	85.42	96.85	95.92
All four at once	12.17	15.01	12.16	14.51	12.65
Smarty	3.89	9.22	4.71	4.73	6.78
phpBB	7.44	17.18	10.7	11.92	17.13
phpMyAdmin	6.75	17.11	10.53	11.31	16.85

Comparison Table

Feature by feature

Feature	Zend Accelerator	APC	Bware Afterburner	ionCube PHPA
Installation	Excellent	Good	Fair	Good
Documentation	Good	Fair	Poor	Excellent
Configuration Options	Good	Fair	Good	Excellent
Web front-end	x	x		
Works with latest PHP	x	x	x	
Performance	Excellent	Poor	Fair	Good
Optimises Code	X	X		
Support	Excellent	Fair	Poor	Excellent
Price	Fair	Excellent	Excellent	Excellent
Licence	Zend Licence	QPL	LGPL	No licence(!)

Conclusion

On the basis of pure speed alone, *Zend Accelerator* leads the pack – often by a fair margin. However, while it does come with a good range of features and an excellent front-end, it's dragged away from the top slot by expensive licence costs and a lack of included source code. With the Zend team having such a close relationship with the developers of PHP, I can certainly imagine that *Zend Accelerator* will keep on improving. As many companies still feel better about products they have paid for, Zend has a good market for its product – hardware costs are as high as ever, and ISPs are more than eager to squeeze as much as they can out of software before upgrading their hardware, even if that means paying a lot for it. For them, the only choice really is *Zend Accelerator*, especially when you remember how much more eager managers are to buy things that come with corporate support. Pretty much everyone else, though, should perhaps rule out such an expensive option, and take a close look at the competition. That leaves us with *APC*, *Bware*, and *PHPA*. *APC* gives acceleration a fairly good shot, but the

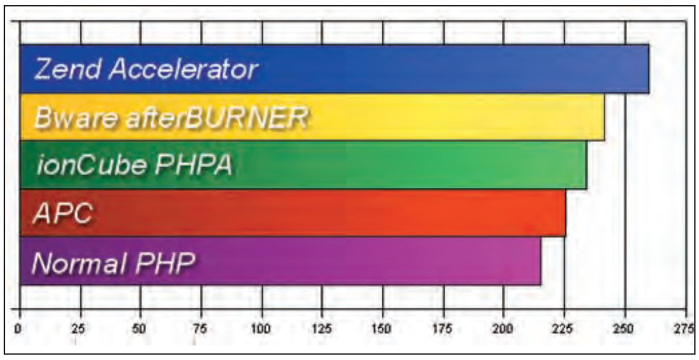

The Basic Test score is the cumulative number of requests per second served across the four basic tests. Test 2 involved opening a database connection and performing a query, so this placed a large strain on the database server as opposed to PHP, and this accounts for the lesser difference between the best performer and un-accelerated PHP. *Afterburner*, *PHPA*, and *APC* all come fairly close together, with *Zend Accelerator* breaking away from the pack. It should be noted that these tests really are basic, so it's impressive to see *Zend Accelerator* get such a large lead.

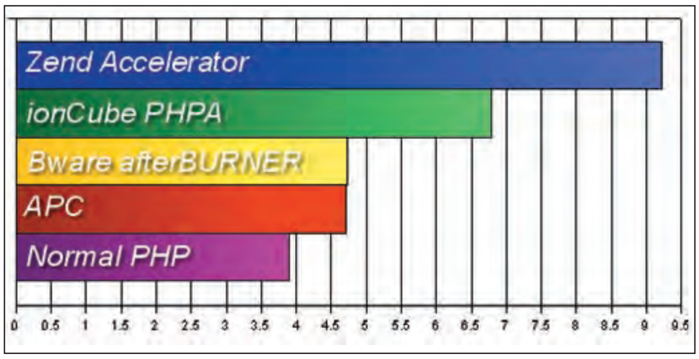

It's at this point that *Zend Accelerator* really stretches its legs – although it's worthwhile noting that even *Zend Accelerator* only managed to serve 10 requests per second. However, even 10 requests per second is two-times faster than un-accelerated PHP, which is a vast difference. *PHPA* makes a good attempt here and is definitely much faster than both *Afterburner* and *APC*, but it's still not a patch on *Zend Accelerator*. While the difference between 7 and requests per second and 9 per second may not seem like a lot, it certainly counts when faced with hardware pushed to its limits.

numbers speak for themselves: *APC* is out-paced by all the other accelerators on the market, and, other than adding cache control functions to PHP, has little to recommend it. *APC* is old, and shows it – consider it as a last resort only. Bware performs much better than *APC* in the tests, even managing to come second in the Basic Tests. However, in the other three tests, *Afterburner's* performance languishes near the bottom with *APC*. Like *APC*, it comes with few extras to help it stand out, and its cross-platform nature, whilst nice to know, won't merit a second thought for most Linux admins.

As such, I can conclusively say that ionCube's *PHP Accelerator* is the best acceleration program on the market right now. The documentation is chatty and easy to understand; the range of features is a breath of fresh air in this category; and the speed, while pipped by Zend, it is certainly a massive improvement over un-accelerated PHP. Users coming to this as their first accelerator will find the web front-end a great help, and the price is, of course, unbeatable. Being open-source software, the support is what we've all grown used to – post a message on the forum, and more likely than not Nick (the author) will answer you.

Things to watch out for

Don't get bitten by avoidable mistakes

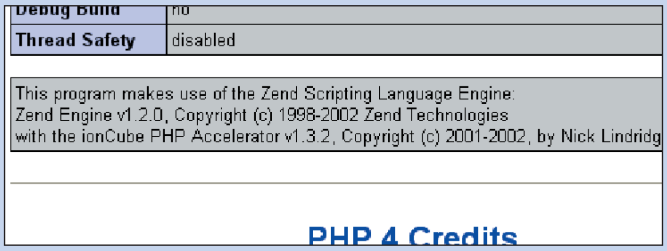

Debug Build	no
Thread Safety	disabled

This program makes use of the Zend Scripting Language Engine:
Zend Engine v1.2.0, Copyright (c) 1998-2002 Zend Technologies
with the ionCube PHP Accelerator v1.3.2, Copyright (c) 2001-2002, by Nick Lindridg

PHP 4 Credits

Check to make sure your accelerator has been installed correctly using phpinfo().

While I recommend that everyone uses a PHP accelerator (even *APC* is better than nothing!), there are a few points that you need to be wary of.

- **Not all programs work with the latest release of PHP (*e.g.*, *APC*, *Afterburner*)**
- **Not all programs work with older releases of PHP (*PHPA*, for example, was built and tested on releases >= PHP 4.0.6)**
- **Not all programs work on Web servers other than *Apache1.3*.**
- **Read the installation documentation thoroughly – some programs (*e.g.*, *Afterburner*) require specific configuration options to enabled caching, whereas others (*e.g.*, *Zend Accelerator*) have "optimisation level" configuration options.**
- **If you don't change your scripts regularly, some accelerators (*e.g.*, *APC*) allow you to disable automatic change checking. This means that it no longer has to check to see if a script has been updated, which provides an extra, minor speed boost.**
- **Load up phpinfo() once you have installed your caching program, just to make sure it has been installed correctly. All the programs above add an extra line near the top of the phpinfo() page saying they have been installed successfully.**
- **Don't be afraid to carry out some benchmarks of your own. As you can see, the performance gain offered by each product varies according to the use to which you put it, so it's good to see how each program performs with your particular setup.**
- **If you're running PHP as a CGI, then strongly consider switching to an *Apache* module – you'll see a much bigger speed-up doing this than you would by enabling a caching solution. Better yet, do both!**

I've been reliably informed that Zend will be releasing a new version of the *Zend Accelerator* at the International PHP Conference in November (it will probably be released by the time this is printed!), and it's possible that we may see another leap forward in PHP acceleration technology. Of course, I also know that ionCube are putting a lot of their time towards R&D into new acceleration techniques as well as better performance for larger sites. Until the next big releases of the two acceleration leaders, I don't doubt that the undisputed king of PHP acceleration is Nick Lindridge, with the ionCube *PHP Accelerator*. LXF

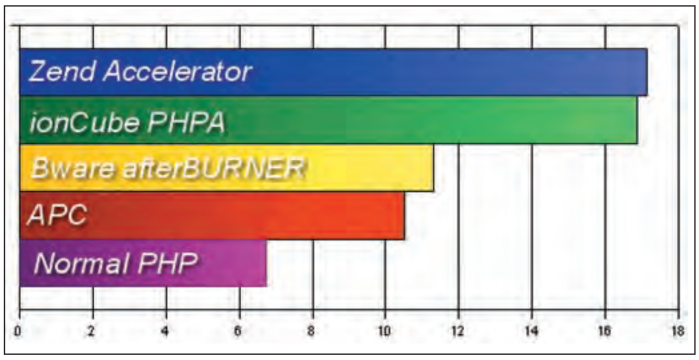

***phpBB* is a very popular bulletin board system written using PHP, and it's here that ionCube's *PHP Accelerator* finally catches up with *Zend Accelerator*.**
Whilst *Zend Accelerator* pips *PHPA* to the top, it's certainly not by a long way, and both *Afterburner* and *APC* come way down the chart in comparison.
Given the immense popularity of *phpBB* and systems ike it, it's good to see both *Zend Accelerator* and ionCube's *PHP Accelerator* almost tripling the processing speed of PHP.

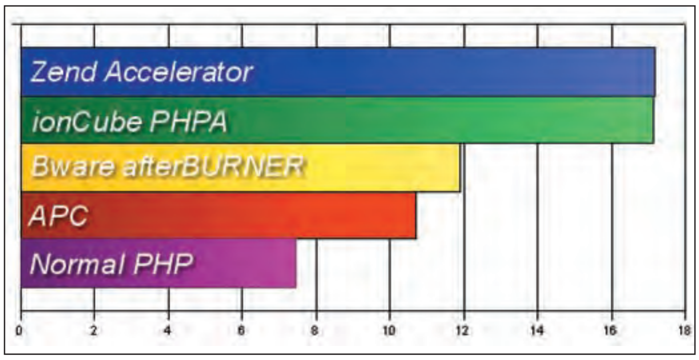

The last test was running *phpMyAdmin*, the popular *MySQL* administration front-end by Ratschiller *et al*. Again we see both *Zend Accelerator* and *PHPA* holding the top positions, with *Zend Accelerator* edging out ionCube's *PHP Accelerator* by a fraction more than with the *phpBB* test.
***Bware* again squeaks into third place by a hair, leaving *APC* last.**
***phpMyAdmin* relies on database connectivity a lot more than *phpBB*, and, combined with *ZA's* test results on the *MySQL* basic test, this may signal *ZA's* superior *MySQL*-related acceleration.**

SCANNING with Linux

Richard Drummond unravels the mysteries of SANE to help you get your scanner working in Linux, and offers some tips on scanning quality.

What is it that's stopping most users adopting Linux wholeheartedly and removing Windows from their systems? The chances are most desktop users dual-boot, not because Linux lacks some crucial application, but because they have some piece of hardware that they think is either unsupported or too difficult to get to work with Linux. Top on this list of troublesome hardware are printers and scanners which, considering how vital they are today's digital office, can prove a real set-back on the desktop.

Scanner support in Linux is provided by a piece of software called *SANE* which stands – optimistically, some might say – for *Scanner Access Now Easy*. *SANE* is a framework, an API for allowing applications to talk to image acquisition devices such as scanners. The job of controlling a particular device or class of device is done by an appropriate driver, known as a *SANE* back-end. The beauty of *SANE* is that it lets any application – or *SANE* front-end – talk to any supported device without worrying about the specifics of any particular device. In this respect, it's a lot like *TWAIN*, the system found in the Windows and Mac worlds, although *SANE* is much more flexible.

The ultimate question for any user is, 'Is my scanner supported by *SANE*?'. Alas, there's no quick answer here. *SANE* does provide back-ends for an impressive array of devices, including popular models from Agfa, Epson, HP and Mustek. But you'll have to visit the *SANE* website and check whether your device is supported in the current release. If you're scanner isn't listed and you're feeling adventuresome, also look at the page listing third-party drivers which haven't yet been included in the official version of *SANE*. You should read any info regarding your scanner's *SANE* driver thoroughly before proceeding. This should tell you any kernel or other requirements it has, and what features are supported.

SANE is supplied with all the major distros these days, and if you're lucky, the kernel and the version of *SANE* installed on your system will already support your scanner. Load up your favourite *SANE* front-end and start scanning. If you're not so lucky, you may need to install a newer or patched version of *SANE*, or rebuild your kernel.

> "Top on the list of troublesome hardware are printers and scanners – a set-back on the desktop for today's digital office"

Driving in SANE

SANE back-ends are user-space drivers: they talk to your scanner at the application level; you still need a kernel driver to handle the hardware-level communication between your system and the scanner. For SCSI scanners, this begins with a driver for your SCSI adaptor; for USB scanners, you need a working USB subsystem in your kernel. Fortunately, most distros will ship with a kernel with all the necessary drivers supplied as modules. Your distro's hardware detection system may even load them automatically.

A quick way to check whether you have the appropriate kernel support for your USB or SCSI scanner is to use the *SANE* command **sane-find-scanner** (assuming you have *SANE* installed). This probes the USB and SCSI buses and reports any scanners that are found. Note, however, that just because a scanner is listed here, it doesn't necessarily mean it is supported by *SANE*; it just means that the kernel can access it. If your scanner is not reported, then you may need to re-build your kernel with the necessary drivers compiled in or built as modules. We haven't got the space to go into much detail here, but we can give you some pointers as to what kernel modules you may need to get your scanner working. Look at the table 'Kernel drivers' for a summary. For more general help with building kernels see the Kernel HOWTO at www.tldp.org.

SCSI scanners should present no problems at the kernel level as they don't need any special kernel drivers, just a supported adaptor and the SCSI generic driver. Simply rebuild your kernel with the appropriate modules and ensure they are loaded at run-time.

USB scanners are a different story, because a USB scanner could actually be a SCSI scanner that has been given a USB interface. This is the case with the HP 5300C (a re-badged Avision scanner) and similar devices, which are supported by the *hpusbscsi* driver, and the Microtek X6USB, which is supported by the *microtek* driver. In both cases, theses devices will appear to the system as SCSI scanners, and hence you'll also need the SCSI generic driver. Other USB scanners require the generic USB *scanner* driver. To support any USB device, you'll also need appropriate host-controller support for your hardware, that is the *OHCI* or *UHCI* drivers, the *USB core* driver and so on. If you have a *USB hotplug* daemon running, this should take care of loading the appropriate drivers at run-time; if not, add the modules to your /etc/modules file.

Parallel port scanners are less well supported by *SANE* and by the kernel. Like USB scanners, some parallel scanners may be SCSI devices in disguise. However, the official kernel doesn't include SCSI-over-parallel drivers, so you'll need to patch your kernel source to be able to build these (see www.torque.net/parport/ppscsi.html). Other devices may require support for user-space parallel drivers (this option can be found in the character devices page of xconfig), essential if you want to give non-root user access to the scanner. Whichever type of device you have, you'll obviously need a parallel port driver in your kernel.

Finally, IEEE1394 (Firewire) devices are now beginning to be supported by *SANE*. These require a kernel with IEEE1394 support and the *SBP2* (Serial Bus Protocol) driver. This provides a SCSI interface, so again you'll need the SCSI generic driver.

Scanner access

The final piece of the puzzle as far as kernel issues go is to ensure that your system has the correct device files installed, so that *SANE* can make use of the kernel drivers you have just built. If your system uses DevFS, then the drivers will automatically create their own device nodes as required. Otherwise, you may need to install them manually.

Devices which use the SCSI generic interface, whether they are real or emulated SCSI devices, require a device file of the form /dev/sgX, where X depends on the ordering of SCSI devices attached to your system. (Have a look in /proc/scsi/scsi. If your scanner is the fourth device, it will be accessed via /dev/sg4.). An appropriate range of /dev/sgX files should already be present on your system, although on some systems these might be ordered by letter rather than number (that is, you'll have /dev/sgd rather than /dev/sg4). To create them yourself, use a command of the following form:

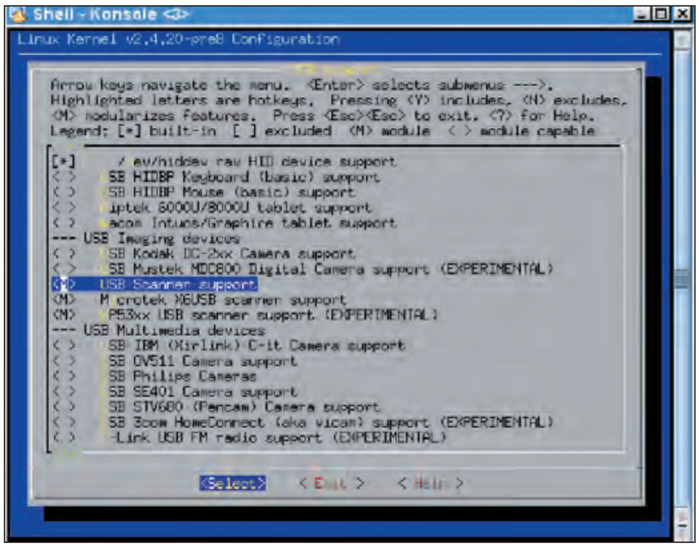

```
Linux Kernel v2.4.20-pre8 Configuration

 Arrow keys navigate the menu.  <Enter> selects submenus --->.
 Highlighted letters are hotkeys.  Pressing <Y> includes, <N> excludes,
 <M> modularizes features.  Press <Esc><Esc> to exit, <?> for Help.
 Legend: [*] built-in  [ ] excluded  <M> module  < > module capable

[*]     /dev/hiddev raw HID device support
< >   USB HIDBP Keyboard (basic) support
< >   USB HIDBP Mouse (basic) support
< >   Aiptek 6000U/8000U tablet support
< >   Wacom Intuos/Graphire tablet support
--- USB Imaging devices
< >   USB Kodak DC-2xx Camera support
< >   USB Mustek MDC800 Digital Camera support (EXPERIMENTAL)
<M>   USB Scanner support
<M>   Microtek X6USB scanner support
<M>   HP53xx USB scanner support (EXPERIMENTAL)
--- USB Multimedia devices
< >   USB IBM (Xirlink) C-it Camera support
< >   USB OV511 Camera support
< >   USB Philips Cameras
< >   USB SE401 Camera support
< >   USB STV680 (Pencam) Camera support
< >   USB 3com HomeConnect (aka vicam) support (EXPERIMENTAL)
< >   D-Link USB FM radio support (EXPERIMENTAL)

          <Select>    < Exit >    < Help >
```

Depending on your scanner model and your distro, you may need to rebuild your kernel for scanner support.

```
mknod /dev/sgX c 21 X
```

It is common practice to create a link from the device /dev/scanner to the actual SCSI generic device, to make the scanner quicker to find. *E.g.*:

```
ln -s /dev/sg4 /dev/scanner
```

If, on the other hand, your scanner uses the generic USB scanner module, you need to create a /dev/usbscanner file. Do this with

```
mknod /dev/usbscanner c 180 48
```

Don't create a link from /dev/scanner to /dev/usbscanner, though, because this will just confuse *SANE*. If you have more than one USB scanner, then you can make nodes of the form /dev/usbscannerX, with X replaced with the number of the scanner. Minor numbers 48 to 63 are reserved for USB scanners (you can have up to 16, if you want), so 48 is the first scanner, 49 the second, and so on.

Last, but not least, you need to ensure that your scanner's device file has permissions set so that it is accessible by non-root users. The

»

Kernel drivers

Which driver do I need for my scanner?

DRIVER	CONFIG OPTION	DEVICE NODE	NOTES
sg	CONFIG_CHR_DEV_SG	/dev/sgX	SCSI generic driver. Required for all SCSI and emulated SCSI scanners.
scanner	CONFIG_USB_SCANNER	/dev/usbscanner	Generic USB scanner.
hpusbscsi	CONFIG_USB_HPUSBSCSI	-	Supports USB Avision scanners such as the HP 5300C. Requires sg.
microtek	CONFIG_USB_MICROTEK	-	Supports Microtek X6USB and similar. Requires sg.
ppdev	CONFIG_PPDEV		/dev/parport Support for user-space parallel drivers. May be required for some parallel-port scanners.
sbp2	CONFIG_IEEE1394_SBP2	-	IEEE1394 serial bus protocol. Supports Firewire scanners. Requires sg.

«

typical way to do this would be to set its ownership to the group 'scanner' and add any user who is permitted to use the scanner to this group. *E.g.*:

```
chgrp scanner /dev/sg4
chmod g+rw /dev/sg4
```

Note that you need both read and write access for proper scanner operation.

Using SANE

If you don't already have *SANE* installed, if you need a newer version to support your scanner, or if you want to add a third-party driver, then you'll probably want to build *SANE* yourself – but, fortunately, this is quite straight forward. The *SANE* package is supplied as two tarballs, called *sane-backends* and *sane-frontends*. The former contains the *SANE* drivers, library and tools and is all you really need, unless you really want the *xscanimage* and *xcam* applications supplied in the frontends package (there are better *SANE* front-ends available).

To build the backends package, it's a simple case of unpacking the tarball and performing the usual **configure /make** incantation. If you are adding a third-party driver, you should do that first, and most add-on drivers will supply instructions on how to do this or even supply a script to do the job for you. *SANE* doesn't have any peculiar build requirements, although some parallel drivers may need the *libieee1284* package (see http://cyberelk.net/tim/libieee1284/) and some USB drivers may need *libusb* (see http://libusb.sf.net/). You should perform the **./configure -help** command to see what options apply to you before building. You should at

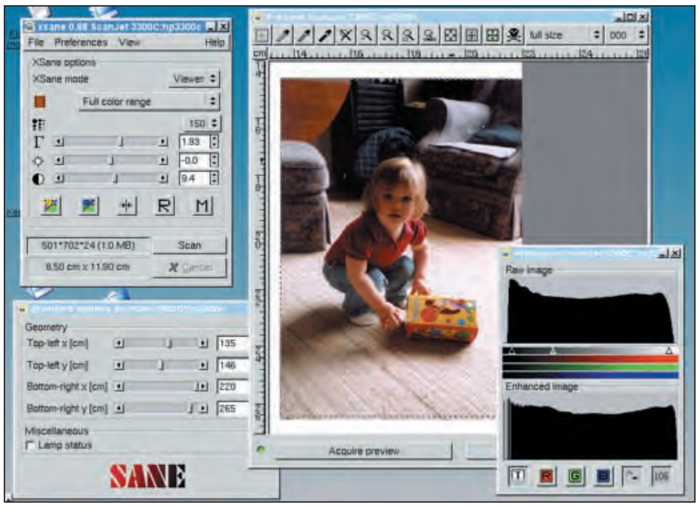

XSane **is currently the most powerful** ***SANE*** **front-end, but the competition are catching up fast.**

SCANNING TIPS!

Find your scanner sweet spots!

Many people don't realise that scanners, like many other devices that interface to an analogue world, have a sweet spot. Due to the way scanners, particularly colour ones, work, there are always some areas of the scanning bed that give better reproduction than others, and some where system noise, for want of a better term, interferes with the results.

This can make a profound difference, especially if you realise that the 'noisy' areas are most often around the edges of the scanning area – where most people like to line up their photos or whatever against a clean edge.

Sweet spots can vary depending on the resolution you are scanning at, and whether you are using greyscale or colour mode. If you want to be really picky, you could prepare a sweet-spot 'map' for different resolutions and colour depths.

So, how can we easily make a sweet spot 'map' in the first place? What you need to do is scan a pure surface (and it helps to clean the glass first. White is usually best. It is important that you scan a clean, white surface. A piece of paper gives a rough idea, but good quality, thick uncoated card is better. To eliminate the possibility of imperfections in the card giving false readings, you can turn it round and do another. You can then manipulate the image to include only the common areas of interference.

When you have the scan, you may already be able to see some coloured areas. Load the image into ***GIMP*** **and use the Levels or Curves tool to maximise the contrast, which will reveal the effect more clearly. When scanning large objects which have to go over noisy areas of the scanning bed, it's possible to use the map you created to correct the image (as a 'Divide' layer in** ***GIMP*** **for say), but it can take a lot of work to actually improve an image.**

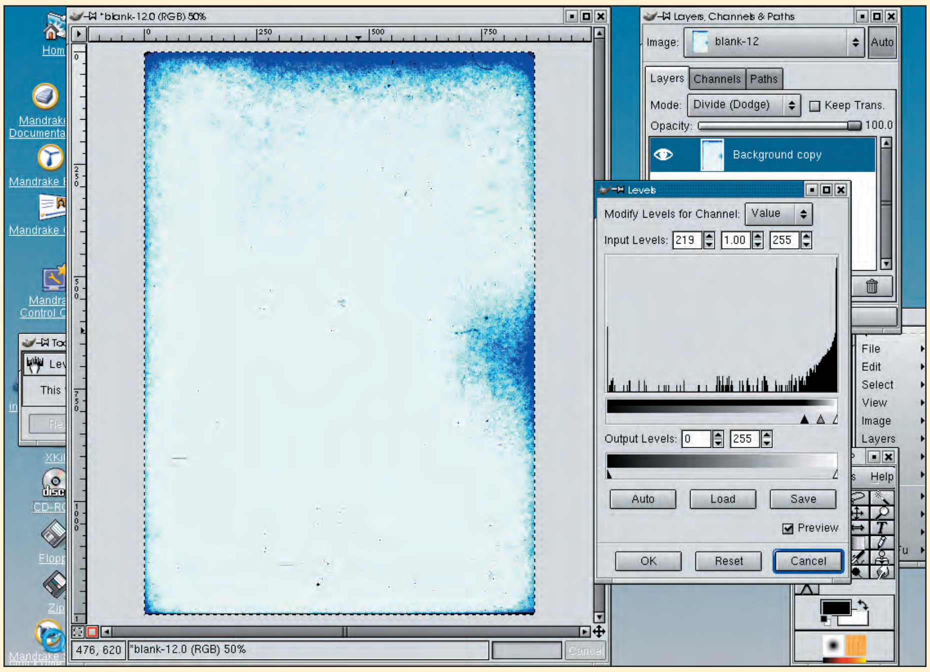

This test clearly shows areas of concern on an A4 flatbed. The upper edge corresponds to the open end of the lid, so a possible cause here is light leaking in. The dead zone on the side could be due to a number of factors. Note that the blue cast to most of these areas indicate that the blue head is less susceptible to the problem.

least supply the **--prefix** option to specify the install location, which will typically be /usr or /usr/local. If you have an existing *SANE* install, it's safest to remove that before installing a new version.

To check that your scanner is driver has been successfully built and installed, execute the command

```
scanimage -L
```

This will try all the installed back-ends and report any devices found. In my case, for example, this finds my USB ScanJet 3300C and returns with the line

```
device `niash:hp3300c' is a
 Hewlett-Packard ScanJet 3300C
 flatbed scanner
```

Here **niash:hp3300c** is *SANE's* device name for my scanner. *SANE* uses a simple naming scheme for devices. The first part, in this case **niash:** is the name of the sane back-end. The second part is driver-specific and will usually be the path to the device file via which the scanner can be accessed, such as /dev/usbscanner0 or /dev/sg2. When using *SANE* applications, if you don't specify a default device to use, *SANE* will search for one, which significantly slows down application start up.

Now you know how to access your device, you can test out the driver by supplying the **-T** switch to scanimage. This will perform a number of tests and report whether the driver passes or fails. For example:

```
scanimage -T -d niash:hp3300c
```

If this fails, consult the *SANE* FAQ and your *SANE* driver's manual pages. The latter should contain additional information about the driver and what options its supports. (Most *SANE* backends can use a file in /etc/sane.d/ called <back-end>.conf for storing default settings.) If all goes well, you're ready to use a full-blown *SANE* front-end and do some proper scanning.

SANE front-ends

The *SANE* package includes several *SANE* front-ends, applications which make use of *SANE* drivers for image acquisition. These include the command line *scanimage*, which we have been using, and the GUI-based *xscanimage* and *xcam*. The first of these is useful for testing and batch purposes, but not up to everyday scanning, while *xscanimage* and *xscan* are too basic to be really useful. Thankfully, there are several good *SANE* application available.

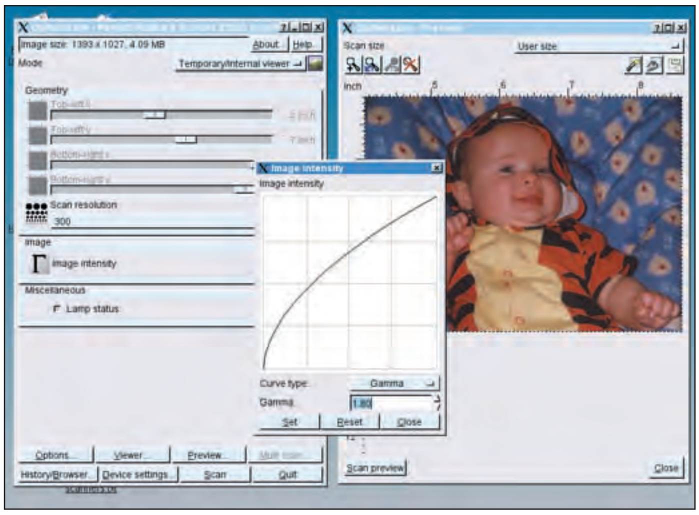
Image size: 1393 x 1027, 4.09 MB
Mode: Temporary/internal viewer
Scan resolution: 300
Curve type: Gamma
Gamma: 1.80
Scan size: User size

***QuiteInsane* is a new *SANE* front-end that looks promising, although much work needs to be done yet.**

»

Correcting wonky scans

SCANNING TIP!

Sometimes when scanning pictures, almost inevitably if you are keeping away from the edges, you end up with a crooked scan. *GIMP* is here to help you though! We have exaggerated a wonky scan here in our example

Double click on the Transform tool in the *GIMP* toolbox to get up the tool setting dialog box. Change the "Tool Paradigm" setting to Corrective.
This means that the tool is set up to correct the transformation effect you set.

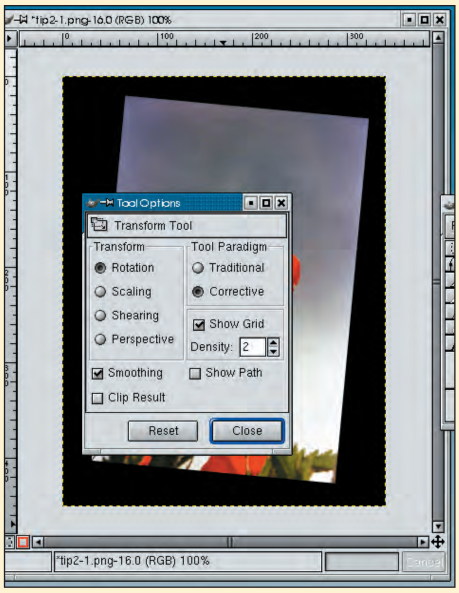
*tip2-1.png-16.0 (RGB) 100%

Tool Options – Transform Tool

Transform: ● Rotation ○ Scaling ○ Shearing ○ Perspective

Tool Paradigm: ○ Traditional ● Corrective

☑ Show Grid

Density: 2

☑ Smoothing ☐ Show Path

☐ Clip Result

Now use the Transform tool and rotate the grid to square up with your scanned image. The Transform function will correct this angle, making your scan straight again. The same technique can be used to straighten parts of a photograph that were wonky to begin with.

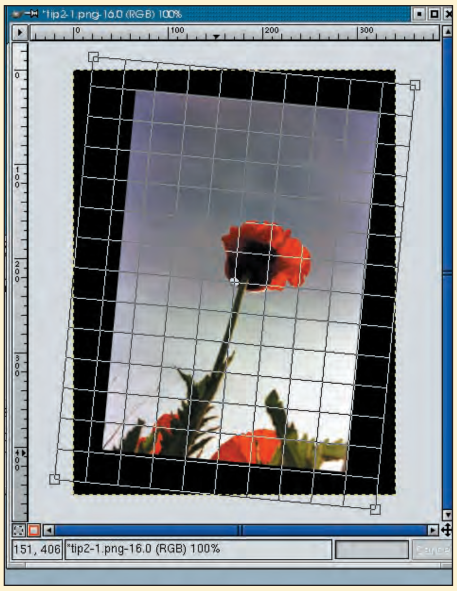
*tip2-1.png-16.0 (RGB) 100%

151, 406 *tip2-1.png-16.0 (RGB) 100%

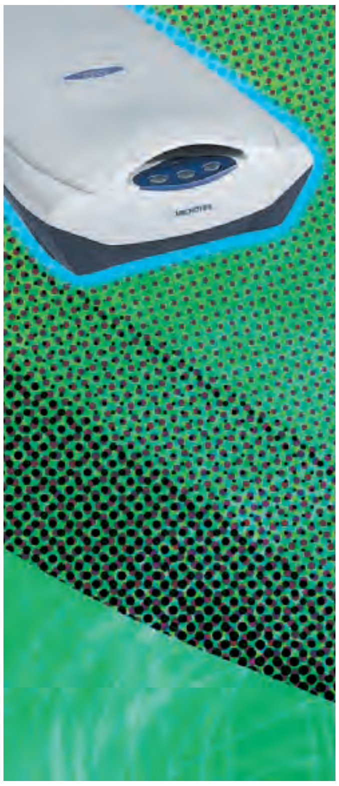

XSane

Web: www.xsane.org

Xsane is currently the most functional *SANE* front-end, and the one most people will prefer to use. It provides a friendly, *GTK*-based GUI, and the current version will build against either *GTK-1.2* or *GTK-2.0*. It works either as a stand-alone application or as *GIMP* plug-in and supports *JOCR* for character recognition. *XSane* provides an excellent range of features, including good pre-scanning image-enhancement tools. You can easily crop, rotate and zoom in the preview window to control the geometry of your scan, and there's a histogram function for the quick and accurate setting of image intensity. *XSane* saves images in a variety of formats, print images, or send them via fax or email. It's built in image viewer is rather primitive, however, and has only basic post-scan image-processing tools, only two filters – blur and de-speckle – and no undo function.

QuiteInsane

http://quiteinsane.sourceforge.net/

QuiteInsane is a fairly new *SANE* front-end and still in an alpha state. It has a pleasant *Qt*-based GUI which will build against either *Qt2.x* or *Qt3.x*. It's preview functions are less easy to use than *Xsane's*, and it's noticeably lacking a histogram tool. Instead, there are various functions for setting the intensity/value curve, including a classic gamma function. It doesn't update the preview on the fly, so it's difficult to set this appropriately, however. *QuiteInsane* makes up for these shortcomings by supplying a good range of tools and filters for enhancing your image post scanning. This includes colour and intensity controls, sharpen and blur filters and effects such as posterize. Stability can still be a problem.

QuiteInsane integrates well with _JOCR_ for OCR, and lets you view scanned text and interpreted text side-by-side.

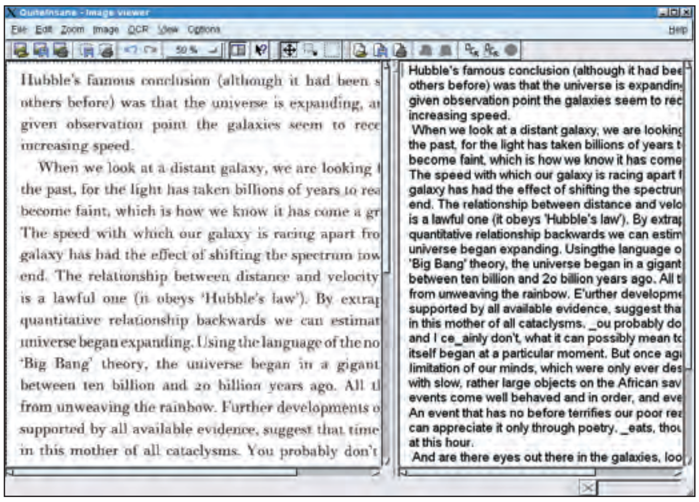

QuiteInsane - Image viewer
File Edit Zoom Image OCR View Options Help
50 %

Hubble's famous conclusion (although it had been s
others before) was that the universe is expanding, a
given observation point the galaxies seem to rece
increasing speed.
When we look at a distant galaxy, we are looking t
the past, for the light has taken billions of years to rea
become faint, which is how we know it has come a gr
The speed with which our galaxy is racing apart fro
galaxy has had the effect of shifting the spectrum tow
end. The relationship between distance and velocity
is a lawful one (it obeys 'Hubble's law'). By extrap
quantitative relationship backwards we can estimat
universe began expanding. Using the language of the no
'Big Bang' theory, the universe began in a gigant
between ten billion and 20 billion years ago. All tl
from unweaving the rainbow. Further developments o
supported by all available evidence, suggest that time
in this mother of all cataclysms. You probably don't

Hubble's famous conclusion (although it had bee
others before) was that the universe is expandin
given observation point the galaxies seem to rec
increasing speed.
When we look at a distant galaxy, we are looking
the past, for the light has taken billions of years t
become faint, which is how we know it has come
The speed with which our galaxy is racing apart f
galaxy has had the effect of shifting the spectrun
end. The relationship between distance and velo
is a lawful one (it obeys 'Hubble's law'). By extraj
quantitative relationship backwards we can estim
universe began expanding. Usingthe language o
'Big Bang' theory, the universe began in a gigant
between ten billion and 2o billion years ago. All t
from unweaving the rainbow. E'urther developme
supported by all available evidence, suggest tha
in this mother of all cataclysms. _ou probably do
and I ce_ainly don't, what it can possibly mean to
itself began at a particular moment. But once agi
limitation of our minds, which were only ever des
with slow, rather large objects on the African sav
events come well behaved and in order, and eve
An event that has no before terrifies our poor rea
can appreciate it only through poetry. _eats, thou
at this hour.
And are there eyes out there in the galaxies, loo

Kooka

Web: www.kde.org/kooka

Kooka is the *SANE* front-end featured in KDE 3.x, although versions for KDE2.x do exist. It is built on the KDE *libkscan* library, which can provide scanner services to all KDE applications, and is used latest version of *KOffice*. *Kooka* presents a very configurable user interface and great integration with other KDE apps, but I've never had much success with it myself. In terms of scan control features it looks about on a par with *QuiteInsane*, but as I've never been able to get it to do a preview scan, I can't really comment. Its a promising-looking project, though, so stay tuned.

Automatic reading?

OCR or Optical Character Recognition is any process for converting the printed page into text in the computer's memory. Not surprisingly this is difficult to do, involving complex image-processing techniques, and early OCR packages gained a poor reputation. In recent years, OCR software has improved immeasurably, and processional OCR software is often used commercially for archiving documents.

In the Free Software world, there's at least one useful OCR program, and that's *JOCR*. (Confusingly, this application used to be known as *GOCR* – and the program executable is still called **gocr** – but there's another, separate project called *GOCR*.) *JOCR* is a command-line tool, although a simple *GTK* front-end is also included in the package and all the popular *SANE* front-ends can integrate with *JOCR*. To use *JOCR* from the command line you simply need to supply a image file for it to decode.

```
Gocr scan.pnm -o scan.txt
```

It only works with PNM, PGM, or PPM images, but with other image formats you can use tools in the NetPBM package to convert them first. For JPEG images, the *djpeg* tool from the *libjpeg* package is also useful. For example:

```
djpeg -gray myscan.jpg | gocr -o out.txt
```

SANE drivers

Where to find information

SANE home page
www.mostang.com
SANE back-ends
www.mostang.com/sane/sane-backends.html
Third-party back-ends
www.mostang.com/sane/sane-backends-external.html
Official SANE FAQ
www.xs4all.nl/~ljm/SANE-faq.html
USB scanner resources
www.buzzard.org.uk/jonathan/scanners-usb.html
Parallel scanner resources
www.buzzard.org.uk/jonathan/scanners.html
HP OfficeJet project
http://hpoj.sourceforge.net/
Epson scanner drivers
http://khk.net/sane/

A useful feature in recent versions of *JOCR*, if you need to scan a lot of pages which use the same font, is the ability to create a dictionary mapping glyphs it cannot recognise to the appropriate characters or strings. You can run **gocr** in an interactive mode – with the switch **-m 130** – and it will prompt you to enter values for any characters it fails to recognise, with the option of storing those values in the dictionary.

JOCR is a work-in-progress and results can be variable. It is rather slow, uses a lot of memory, and has problems with complex page layouts, serif fonts, italics and ligatures. It's definitely a project worth supporting, however.

Not just for scanners

Some other devices that work with *SANE*

SANE provides back-ends for a wide range of image acquisition devices, but it's not just about flatbed scanners, you know. The *v4l SANE* back-end supports any video capture device with a kernel *Video4Linux* driver, including TV cards and web-cams. Here, the *SANE* device name specifies the device file to use, for example, v4l:/dev/video0. If that's not enough, the *gphoto2* driver lets you access any digital camera supported by the *Gphoto2* library (see www.gphoto.org). With this driver, you must specify the parameters of the device to use in the /etc/sane/gphoto2.conf file.

What on Earth are... MANPAGES?

Learning how to learn: **Marco Fioretti** goes into the most venerable Linux help system to explain how it works, and why it is still a Good Thing.™

Whenever I ask how some program works the answer is "see the man (page)". Do I have to visit the program author, or leave a message on his pager?

No. You are just being told that before wasting everybody's time you should have read the online *manual* page(s) of that program with the Unix *man* command. Practically, you have to type **man** followed by the command name inside a terminal.

What would I find in a man page?

An explanation of what the command does, and of all the options with which it can be invoked.

From a terminal? In this century? Why not some multi-neural-quantum system? Maybe marketing-oriented and customer-focused?

The man documentation system appeared more than thirty years ago, way before HTML browsers were invented. Many fundamental Unix commands have not changed since then, so they just had their original man page updated. Typing **man** at the prompt may be much faster than starting *Mozilla*...

This is all well and good if I *already* know the command name!

Enter the *apropos* command. It extracts from a database a one line description of all the commands whose man page contains the word given as argument. The same result can be obtained with **man -k** (**k** is for keyword here).

What if I just want to know what some command is for?

Use the *whatis* command, or just **man -f**: they both return a one line description of the app's purpose:

```
$ whatis grep
grep (1) - print lines matching a pattern
```

Do man pages exist only for Unix commands?

No, also for a lot of programming docs. There are man pages for almost every C lib routine or system call, and for many Perl modules and similar objects.

What are those numbers that are sometimes associated with the command name?

The man pages are grouped in sections, and each section has its own directory. Traditionally, they are:

1 User-level commands
2 System calls
3 Library functions
4 Devices and device drivers
5 File formats and conversions
6 Games
7 Various miscellaneous stuff - macro packages etc.
8 System maintenance and operation commands
9 Kernel Routines
(l) Local
(n) New
(o) Old

Typing **man <number> intro** will output a short description of the kind of material is dealt with in the corresponding section.

With this in mind, it is easy to understand that **man 1 grep** looks for the grep man page only in section 1, and **grep (1)** means "the grep command, as described in section 1 of the manual". Specifying the number makes the search faster, and sometimes is just needed to avoid ambiguity:

```
whatis socket
socket (2) - create an endpoint for communication
Socket (3pm) - load the C socket.h defines and
 structure manipulators
socket (7) - Linux socket interface
socket (n) - Open a TCP network connection
```

Four different meanings of socket, with their own docs.

How is a man page structured?

The actual layout can vary, but the default template hasn't changed much since the beginning. It includes the following sections, usually in this order.

NAME – The name of the command, and the one line description already mentioned. Example: "**ls** - list directory contents".

SYNOPSYS – "a compendium... giving a view of the whole" (*Webster's College Dictionary*); more on this later.

DESCRIPTION – A more detailed description (what else?) of the purpose and capabilities of the program.

OPTIONS – A complete list of all the command line options, also called *switches*, that can be written after the program name when it is launched. Very often, options have both a short and a long form: the first is faster to *write*, the second, being self-explanatory, is easier to *read* . *E.g.* (again from the *ls* man page):

```
-a, --all (do not hide entries starting with .)
```

As you can see, the short form is perfect when entering commands in a hurry, the other is best suited for complex shell scripts which will be used for a long time and/or by many people.

ENVIRONMENT VARIABLES – A list of all the external variables influencing the program behaviour.

DIAGNOSTICS – When programs finish they return a numeric value (**exit status**). The shell or script which invoked them can then decide the next action on the basis of that value. The meanings of the exits states available is explained here.

BUGS, REPORTING BUGS – These two sections report (hopefully) all the known bugs of the corresponding program version, and the email address to which new bug reports should be sent.

FILES – The configuration files read by the program when it starts, and/or the ones where it logs its actions.

SEE ALSO – Links to other commands used by this one, similar commands, or other documentation.

Finally, almost all man pages also include a **COPYRIGHT** and/or an **AUTHORS** section.

MANUAL
BUGS
OPTIONS
FILES
AUTHOR
NAME
SYNOPSIS
DESCRIPTION
DIAGNOSTI

« This synopsys thing looks a random mix of assorted parentheses, dashes and random letters. Is that supposed to be a human language?

No, not really, although some system administrators speak almost in the same way. The synopsys shows the several ways in which the command can be run, listing all its options and arguments.

Hold on: what is the difference between options and arguments?

Usually, the options are the different *things* that a command can *do*: they roughly correspond to the several buttons or menu entries of a GUI program. The arguments are the *objects* (files, users or hardware peripherals) on which the commands *acts*.

If the synopsys is just a list, why didn't they write a plain table?

Because it is essential to know which options can be combined, (and in which order), and which not. *E.g.*, the *tar* command can create or read archives, but certainly not do both things in the same run. A good synopsys contains all these rules, and their syntax is relatively simple. Let's look at the synopsys of *GNU tar*, version *1.13.25* (heavily trimmed for simplicity) :

```
tar c|t|x [-v] [ -f, --file F ][ -Z, --compress]
filename1 [ filename2, ... filenameN ]
directory1 [ directory2, ...directoryN ]
```

The synopsys doesn't says what each parameter is for, because that is explained later. Another reason for this structure is that people who already know the command but forgot the name of same options need read only the synopsys.

The vertical bar (|) indicates mutually exclusive actions: in the example, *tar* can be invoked with the **c** (**Create**) *or* the **t** (**lisT**) *or* the **x** (**eXtract**) option . Square brackets delimit facultative arguments (**-v** generates **verbose** output if present), while commas separate equivalent versions of the same option (**-Z**, **--compress**). The program arguments usually come last: the example above shows that *tar* can act on any combination of one or more files or directories.

This is making me sick! Why do I have to read a PhD thesis when all I really need is a three-line HOWTO?

This is the hardest part to get used to, isn't it? Yes, man pages explain how something *works*, not how it *can be used*: that is left as an exercise for the reader.

Jonathan Lasser, author of *Think Unix*, says the first Unix thing to learn is "how to teach yourself. Most computer people think that everyone is born knowing how to interpret documentation, but it's not the case: it is a learned skill."

This is a *Good Thing*™: command line programs with many options, combined with pipes and scripting, mean that even non programmers can easily automate almost any task. Ten minutes spent reading a man page can save you many hours in the future.

I like eye candy: isn't there any GUI interface?

Predictably, there are several choices. *xman* has been around for a long time; some other common options are listed in the box (*above right*). Of course, KDE and GNOME have their own viewers, or the ability to read nicely formatted man pages directly in the standard file manager. Entering a **man://some_command** URL in *Konqueror*, for example, will display in nice colours the corresponding man page.

My browser is up all the time anyway, can't I read man pages from it?

Why not? man pages can be converted to plain HTML automatically with *man2html*, included in many distros. The output will preserve all the bold or italics text, and generate links to other man pages mentioned in the text. A fairly complete list of man page converters, including some working in CGI mode, is available at www.w3.org/Tools/Man_faq_filters.html

Did somebody already do it? Can I just read them somewhere online?

Yes, there are searchable interfaces are at **http://www2.linuxpakistan.net/man.php** and **http://linux.ctyme.com/**. Remember, however, that you may find a version of the command for another Unix.

Emacs brews coffee and irons my shirts:will it also do man pages?

Why, yes! In *Emacs*, **Esc X man** (or **manual-entry**) will prompt for the man page to load, and display it with all the correct faces. Since **manual-entry** can accept the output of another elisp command, and **current-word** returns the word the point is on, a macro based on "manual-entry (current-word)" will try to display the man page of the selected word.

Emacs is evil, what about vi/Vim?

The pager actually used by man can be specified in the **MANPAGER** environment variable:

```
col -b | view -c 'set ft=man nomod nolist' -
```

will use *vi/Vim* syntax higlighting and nav commands.

Couldn't I just just display the man file with more, less, or cat?

To begin with, manual files are compressed to save space. In the second place, they also include some markup to indent the text, display section headers in bold face, and so on. To read a man page, one has to decompress it, submit it to the proper formatter (*nroff* or similar), and pipe the output to a pager like *less*. The man command automates this procedure.

Can I save the plain text into a file?

Of course. There are several ways to do it: you could simply pipe the output of **man** to *col* (try **whatis col** to see what it does...), and save the output:

```
$ man grep | col -b > grep_man.txt
```

The man command above may even be substituted by *groff* and *zcat* (or *cat* if the source is not compressed) running on the source file:

```
$ zcat grep.1.gz | groff -t -e -mandoc -Tascii
```

How can I print the man pages with bold face and all?

The procedure is very similar to the one we just described. All you need to do is to tell ***groff*** to generate PostScript output instead of ASCII (replacing **-Tascii** with **-Tps**), and pipe the output to the printer (after checking that it can print PostScript, of course!):

```
zcat manfile.1 | groff -man -Tps | lpr
```

Man page viewers

More sophisticated man interfaces

Pinfo (http://zeus.polsl.gliwice.pl/~pborys/): allows browsing of both *man* and *info* pages with a text interface similar to *lynx*.

TkMan (http://tkman.sourceforge.net) can also browse both kind of pages (warning: some version force the user to patch *Tk8.0*). *TkMan* can also hyperlink pages among them, search text (plainly or with regular expressions) and do much more.

ManEDit (http://wolfpack.twu.net/ManEdit/) displays man pages, but was also designed specifically to edit them with all bells and whistles (drag and drop templates, preview, etc...).

Man just displays gibberish, or nothing at all!

Sometimes, **man somecommand** might just quit saying "No manual entry, or format its output incorrectly: in the last case, check first the **LOCALE** settings of your terminal (in Red Hat 8.0, for example, which supports UNICODE by default, you could set the **LANG** variable back to **en_US** or **C** for a correct man page display). Other important environment variables are:

MANPATH – A colon separated list of the directories actually containing the man pages

MANWIDTH – The width of the man output

MANPAGER (or, alternatively, **PAGER**) **–** As already explained, if the built-in way used to display the pages doesn't satisfy you, change the value of this variable

LANG – In an ideal world, selecting a certain language during the installation would make *all* the documentation available in that language. In this case, the value of the **LANG** environment variable would tell to the man command to look first for the version in that language, showing the English one only if the first is missing

The man system just takes care of itself, right?

Wrong. The correct behaviour of the *whatis* and *man* commands is based on a database created with the **makewhatis(8)** program. This database, and its index, should be rebuilt every time there is a change in the man pages. Strictly speaking this means any time a program is added, patched, updated or removed, not to mention distribution upgrades. Running **makewhatis** via *cron* is probably the best solution, if you update your system often.

I have a computer with very little disk space. How do I avoid to install documentation?

Some programs have all their documentation in a separate add-on package, so not installing them solves the problem for good. The main package managers have an option to not install any file marked as documentation (**--excludedocs** for *RPM*). Of course, if the package author didn't do his or her job, that option will accomplish nothing. In such a case, or when installing from a tar archive, you have to remove all the documentation by hand, unless the Makefile has an option to not install any documentation.

I wrote a program: how do I generate a man page for it?

A programmer documenting his code? Allow me to recover from shock! In the meantime, please accept our eternal gratitude, and grab the man page HOWTO at: **www.tldp.org/HOWTO/mini/Man-Page.html**

What about style?

A poorly written man page will cause almost as many curses and complaints from users as a missing one. Some basic tips to write a good, useful man page are:

- Include some examples, and their *actual* output.
- Test them first, and specify on which Unix/Linux version they worked.
- Remember the difference between users and developers: a man page must explain usage. Of course, do also document the internals of your program, but in a separate document, and point to it from the man page.
- Document also *all* the settings in the configuration file (maybe in a separate man page, see **man man.config**).
- A man page is an official document: be polite, clear, and don't use words or expressions that will be understood only by native English speakers.
- Humour is fine, but in small doses. And run that spell checker, please!
- Check if the source causes any warnings from *groff*, and how the result looks in the graphical viewers mentioned in the Man page viewers box, *above left*.
- Format the **NAME** section as explained in the HOWTO, to allow correct indexing
- Include a valid email address or website for reporting bugs.

You mentioned here and there some so-called "info pages." What are they?

Info pages (or, more exactly, "nodes") are another Unix documentation system. From the end user point of view, the main difference is that they are hyperlinked to each other, so one can browse them very much like webpages.

Of course, trust Unix to complicate things by having at least two unrelated ways to do anything...

The original reason why the *info* system was created is that, while a man page is the condensed instruction manual of some Unix command, an *info* document is a (hypertext) book, which may be about just any possible argument. In practice, the two systems became interlaced, and started to point to each other. Many man pages, especially those for GNU project commands, give a very terse explanation and then just point to the *info* node of the corresponding command for the full documentation.

So, there is a complete Unix online help, after all!

Unix has online help all right, but it is sadly (very) far from being complete. As a home assignment, try what follows on your favourite distribution:

create a list of all your executables (Hint: run **find** on all directories in your **$PATH**)

run **man <command> > /dev/null** on each element of that list count how many times it just says "No manual entry for <command>"

On an almost standard Red Hat 7.3, this little test gave 1191 missing man pages on 2270 binaries. Granted, many programs today have their docs somewhere else in the system, but a man page at least pointing to that somewhere would be nice...

This is great! I can study how a command works on my home Linux box, and show it off tomorrow in our Solaris lab!

Alas, no. The man system (and the knowledge you gain from it) is not completely portable. The first obvious trap is the fact that different versions and/or flavours of Unix carry different versions or implementations of the same command. The *find* command in Solaris 2.7 supports the **-name** option, to find only the files whose names match the given pattern, but the GNU version also offers the **-iname** variant, which is case insensitive. Be warned, especially when moving scripts!

Other pitfalls are in the *man* command itself. Some options are only present on certain platforms (**man -t** doesn't work on HP-UX), or have different names, and the same applies to environment variables. Even some source man pages may not be readable on all platforms.

What is the most useful things I can do about man pages?

First of all use them to learn. Right after that, always report (both to the program and distribution maintainers) when they are missing or outdated! Even better, write some documentation yourself, and send it to the author. If nothing else, he'll have a template to start from.

There is also a project for this, the Missing Man Pages Project (or m2p2), www.netmeister.org/misc/m2p2/index.html. It "aims to provide the man-pages...where the author did not provide an adequate one". Please contribute to it: there is no fun in yelling "read the man pages" if they are not there... LXF

Tutorials »

Our experts offer help and opinions on a whole host of Linux applications

Your guide to getting things done!

Whether you are just starting out in Linux, or an experienced veteran, there's always more to learn. Every issue of *Linux Format* is packed full of practical advice, and nowhere is it more concentrated than in our tutorials. pages.

Here you'll find expert guides to all sorts of things, from Basic Linux usage to understanding and deploying network solutions, from simple script coding to the complexities of Perl regular expressions, Java server apps and more. We aim to bring a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact us, by post, by email (linuxformat@futurenet.co.uk) or log on to our website and post your suggestions in our special forums? (www.linuxformat.co.uk). Hope to hear from you soon!

Nick Veitch EDITOR

THIS MONTH...

System programmimg

First we take a byte at moving data to and from the filesystem, then talk stat structures with the inode table and show how to design a security tool **p74**

OpenOffice.org

Installing a dictionary and getting the most out of *Writer* - the *OpenOffice.org* word processor – all the power-user tips that you need, and how to live with the proprietary docs that you are sent. **p78**

Doxygen »

Creating code is all well and good, but if you haven't a clue how it works next time you look at it, you've only done half a job. Jono Bacon looks at comments and documentation tools **p82**

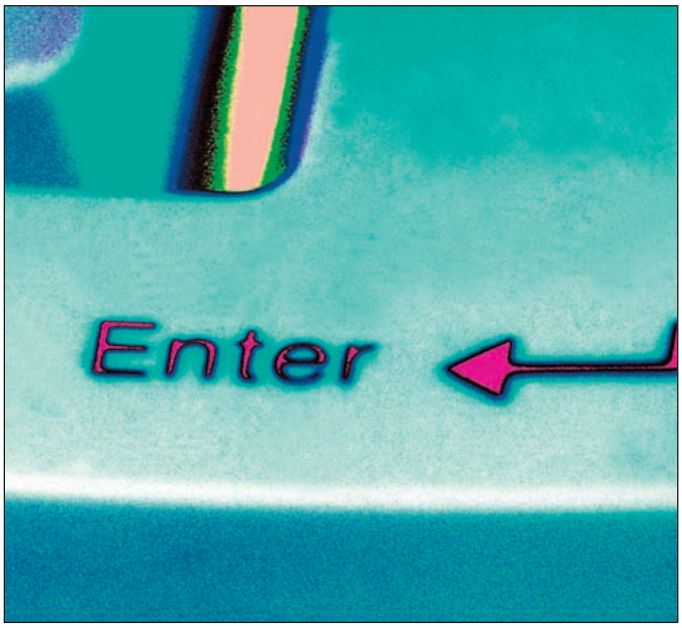

Perl

After the apocalypse: a quick tour of Parrot – the byte-code interpreter destined for Perl 6 **p86**

PHP

Harnessing the power of XML, and XSLT, for your PHP-driven website – and how to avoid common XML errors **p88**

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.

TIP OF THE MONTH! Killing processes

Even on the most stable of operating systems, sometimes an application or a process will run away. Either because of system problems (lack of memory perhaps?) or bad coding, it can be difficult to get through to these apps when you no longer want them.

This is where the *kill* command comes into it's own. *kill* simply sends a process signal to shutdown the runaway application, and should work in most circumstances. All you have to provide is either the process name (*i.e.* the name of the executable, usually) or it's Process ID number. The latter is important when multiple instances are running.

If you don't know the process ID (and why should you) a program like *top* can usually reveal it, or you can type:

```
ps -A
```

at the command line to get a list of all running processes. Then a simple:

```
kill 21798
```

where the number is the PID should solve your problems. There are other friendlier ways to use *kill* too – check out the man pages, they are well worth reading.

If you do want to *kill* all instances of something, then you want to be using, er, *killall*:

```
killall mozilla
```

will kill all *Mozilla* processes – handy if you aren't sure which is causing the problem.

And if all this seems like too much effort, *XKill* is a really handy way of killing runaway X processes. Run *XKill* and the cursor turns into a skull and crossbones. Now just click on the window associated with the task that you want to be rid off and hey presto, it's gone!

That should sort out your runaways. See you next month for more tips!

UNDERSTANDING THE FILESYSTEM

Talking files

Chris Brown takes a byte at moving data to and from the filesystem, and shows how to design a security tool.

Last month, we looked at the notion of system calls in general, and inflicted upon ourselves a couple of specific calls, **fork()** and **exec()**, which are at the heart of the Linux programming model. Congratulations to those hardy souls who made it to the end of last month's tutorial and have come back for more. This month, we're going to look at how programs can interact with the file system. Whilst there's quite a lot of detail here, the conceptual hill is not so steep as for **fork()** and **exec()**.

Accessing a file's data

There are lots of fancy ways a program can use to read and write data in the filesystem, ranging from formatted, text-oriented output functions like **fprintf()** to SQL queries on relational databases. Here, though, we're down at the bedrock – the underlying system calls that support file I/O operations – which simply move data between a file and a buffer in your program. Now, there's a strong tradition in the Unix world of beginning with a program which writes "hello world" to the screen, and usually this is done by calling the formatted output function **fprintf()** from the standard library. Here, though, we're using a low-level **write()** system call. (The line numbers are for reference, they are not part of the file.)

```
1  main()
2  {
3      write(1, “hello world\n”, 12);
4  }
```

The first argument to the **write()** call is a file descriptor which specifies the output stream we want to write to. Here, we're using descriptor **1**, which is one of 3 descriptors which are automatically available to any program. Descriptor **1** is called the standard output: by default output written to this descriptor appears on your screen. The second argument to **write()** specifies the data to be written and the third argument is the number of bytes to write. (**write()** is not smart enough to figure out the length of the string for itself.)

There are two other standard descriptors available to processes; these are **0** (standard input) and **2** (standard error). By default, reading from descriptor **0** reads from the keyboard, and writing to descriptor **2** also writes to the screen, as shown *above right*. If you want to read and write files, instead of using pre-defined descriptors like standard output, you have to open them yourself. This example makes a copy of the file /etc/passwd in a file *thecopy*:

```
1  #include <fcntl.h>
2
3  #define BUFSIZE 100000
4
5  int main()
6  {
7      char buffer[BUFSIZE];
8      int  count, fdin, fdout;
9
10     fdin  = open(“/etc/passwd”, O_RDONLY);
11     fdout = open(“thecopy”, O_RDWR | O_CREAT, 0644);
12     count = read(fdin, buffer, BUFSIZE);
13     write(fdout, buffer, count);
14     close(fdin);
15     close(fdout);
16 }
```

Line 7 defines a data buffer of 100000 bytes which will be used to park the data in on its way from one file to the other. Line 10 opens the input file for reading only; the **open()** call returns a small integer value called a file descriptor, which we store in **fdin**. Later, in line 12, we reference this descriptor in a **read()** call. Line 11 opens the output file. The second argument says that we want to open it for writing, and to create it if it doesn't exist. The third argument specifies the mode (access permissions) to give to the file if it gets created. (The octal value **0644** corresponds to the permissions **rw-r--r--**; you may be familiar with this representation of file permissions if you have used the *chmod* command.) Since the name of the output file is a relative pathname (it doesn't start with a "/"), the file will be created in the current directory.

Line 12 reads the input file into the buffer, up to a maximum of **BUFSIZE** bytes. The **read()** call returns the number of bytes it actually read and placed into the buffer; assuming that the file is shorter than 100000 bytes, this will be the total size of the file. In line 13, however many bytes were read in to the buffer are written back to the output file. Lines 14 and 15 close the file descriptors. We are just being tidy minded here; Linux will implicitly close any open descriptors when a process terminates anyway.

For simplicity, this example assumes that the input file is not bigger than 100000 bytes. To copy arbitrarily long files, we would need to wrap lines 12 and 13 inside a loop. There are a couple of other deliberate simplifications in this code. First, we've hard-wired the names of the input and output files into the code. More likely, we'd take these from command line arguments. Second, we have not attempted to handle (or even detect) errors.

Accessing a file's attributes

In addition to the data held in a file, the Linux filesystem maintains other attributes for each file in a data structure called an inode. There is one inode for each file and they are stored on disk in a table called (would you believe) the inode table. If you have ever run the command **ls -l** you have already seen most of these attributes. Here's an example to remind you; I'm going to assume that you already know what the various fields in this line of output mean:

```
$ ls -l tutorial2
-rw-r--r--    1 chris    users  4329 Aug 30 13:54 tutorial2
```

Here, we're going to access the same info from within a program, using the **stat()** system call. Let's dive straight in with an example.

```
1  #include <sys/types.h>
2  #include <sys/stat.h>
3
4  int main()
5  {
6      struct stat info;
7      stat(“tutorial2”, &info);
8     printf(“size = %d, owner = %d\n”,
9              info.st_size, info.st_uid);
10  }
```

There are several things in need of explanation here. The central idea is that **stat()** returns the file attributes in a structure called a stat structure. The members of this structure correspond, roughly, to the fields in the output from **ls -l**. The structure is defined in the file /usr/include/sys/stat.h, but I wouldn't recommend looking in there (at least, not after a heavy meal) because it gets quite messy. The details fall more readily to hand in the manual page – **man stat** lists the members of the structure:

```
struct stat {
    dev_t        st_dev;      /* device */
    ino_t        st_ino;      /* inode */
    mode_t       st_mode;     /* protection */
    nlink_t      st_nlink;    /* number of hard links */
    uid_t        st_uid;      /* user ID of owner */
    gid_t        st_gid;      /* group ID of owner */
    dev_t        st_rdev;     /* device type */
    off_t        st_size;     /* total size, in bytes */
    blksize_t    st_blksize;  /* blocksize for filesystem I/O */
    blkcnt_t     st_blocks;   /* number of blocks allocated */
    time_t       st_atime;    /* time of last access */
    time_t       st_mtime;    /* time of last modification */
    time_t       st_ctime;    /* time of last change */
};
```

We see, for example, one of the structure members is called **st_size**. It contains the size of the file in bytes. Its data type is **off_t**. This is defined in the file /usr/include/sys/types.h; like most of the types encountered in this structure, it's simply a 32-bit integer.

Lines 1 and 2 of our listing tell the C compiler to process the contents of the two relevant header files, so it knows about the type definitions and what a stat structure looks like. Line 6 declares a stat structure called info, whose address is passed to the **stat()** call in line 7. The first argument to **stat()** is the name of the file. We've hard-wired a fixed name in here for simplicity. Because it's a relative pathname, stat will look for the file in the current directory. We could have used an absolute path name, such as /home/chris/tutorial2, just as well.

When **stat()** returns, the attributes of the file will have been placed in our stat structure, info. In lines 8 and 9, we print a couple of these out. If we run this program, we'll see output like this:

```
$ stat_demo
size = 4329, owner = 1000
```

We see that the size (**4329**) agrees with that reported by *ls*. The owner is shown not as a text login name but as a numeric user ID (**1000**) because that's what is stored in the inode. If we want to manually translate from the numeric UID to the login name, we must look in the password file, like this maybe:

```
$ grep 1000 /etc/passwd
chris:x:1000:100:Chris Brown,,,:/home/chris:/bin/bash
```

We see that UID **1000** is **chris**, which agrees with the *ls* output. When we do an **ls -l**, ls does a similar lookup (by calling the function **getpwuid()**, which is known as a resolver) to map the numeric ID back to the login name.

Error handling

Wouldn't it be wonderful if we could assume that every operation our programs tried to perform would succeed, that every file we expected to be there was there, and that every piece of data we encountered was correctly formatted? In reality, developers know that much of their code is there to handle errors and special cases.

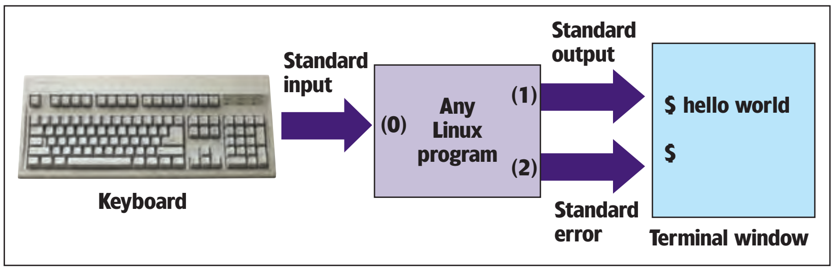

Default connections for a program's I/O streams.

At the system call level, errors are notified in two stages. First, most system calls return **-1** to indicate failure. Second, the specific error condition is returned in a global variable called **errno**. There is a long list of symbolic constants defined in /usr/include/asm/errno.h for the various error codes that can be returned. For example, the constant **ENOENT** is returned if a call tries to access a non-existent file, and has the value **2**. You can test for specific error conditions:

```
  if (errno == ENOENT) {
      /* Handle the error */
  }
```

For really cheap-and-cheerful error handling, you can call **perror()**. This is a library routine that prints out a (rather cryptic) error message for the current value of **errno**. Here's a modified version of our stat demo called *stat_error* which illustrates the use of **errno** and **perror()**:

```
#include <sys/types.h>
#include <sys/stat.h>
#include <errno.h>

int main()
{
    struct stat info;
    if (stat(“xyzzy”, &info) == -1) {
        printf(“errno = %d\n”, errno);
        perror(“xyzzy”);
        exit(1);
    }
    printf(“size = %d, owner = %d\n”,
             info.st_size, info.st_uid);
}
```

Here's the output when we run it. Note error message from **perror()**.

```
$ stat_error
errno = 2
xyzzy: No such file or directory
$
```

Changing file attributes

The **stat()** call gives read-only access to a file's attributes. You can't change the attributes by modifying the stat structure and writing it back out again. However, there are system calls for changing certain specific attributes, such as the access mode and the ownership. In fact, if you know how to do these two things from the command line, you pretty much know how to do them from a C program. The system calls, **chmod()** and **chown()**, even have the same names as the equivalent command line tools.

The **chown()** and **chmod()** calls are simple enough not to warrant a full program to explain them. Here're a couple of examples:

```
  chown(“tester”, 1000, -1);
  chmod(“tester”, 0644);
```

The second argument to **chown** is the numeric UID of the new user. The third argument lets us change the group ownership » too, **-1** means we want to leave it alone. Only processes running as root (user ID **0**) can successfully execute this call.

There is an important class of programs in Linux called filters. A filter is a program which reads a single input stream, transforms it in some way, and writes a single output stream. Filters are at the heart of the command line "tool-building" philosophy which Linux inherited from Unix, in which solutions are composed by using a set of general-purpose tools in combination.

Programs like *cat*, *sort*, *grep*, and *head* are examples of filters. Filters follow an important convention in the way they handle command line arguments. If they receive arguments (other than options beginning with **-**), they treat them as file names, and will open and process the contents of each file in turn. If they do not receive any command line arguments, they will read and process their standard input. This means that a filter can be used to process either a file, or data coming down a pipe from some other process:

```
$ head -5 tutorial2
```

shows you the first five lines of the file tutorial2, whereas

```
$ ls -l | head -5
```

shows you the first five lines of the output from **ls -l**.

Accessing command line arguments

We already know enough about doing I/O to write a filter, but to do the job properly we need to know a bit about how a program can retrieve its command line arguments. When the **main()** function of a program is called, it receives two parameters which give it access to its command line. Traditionally they are called **argc** and **argv**. **argc** is an integer and tells you how many arguments you got, including the name of the command itself. **argv** is an array of strings, containing the actual arguments. (Before the C language purists amongst you set quill to parchment to point out that **argv** is actually a pointer-to-pointer-to-char, let me confess that the notion of an "array of strings" is a slightly loose description. But it will do.) So for example if a program is invoked like this:

```
$ someprog  -z  banana  plum
```

then **argc** will be **4** and the array elements **arv[0]** through **argv[3]** will correspond to the strings **someprog**, **-z**, **banana** and **plum**.

To illustrate, here's a little program that prints its arguments. This is similar to the Linux *echo* command.

```
main(int argc, char *argv[])
{
    int i;
    for (i=1; i<argc; i++)
        printf("%s ", argv[i]);
    printf("\n");
}
```

It simply prints the strings **argv[1]** through **argv[argc-1]**.

Our own version of head

We can now write our own version of the *head* command. Call it *five*, as it prints the first five lines of its input, and discards the rest:

```
1  #include <fcntl.h>
2
3  main(int argc, char *argv[])
4  {
5      char buf[1];
6      int fd, count = 0;
7
8      if (argc > 1)
9          fd = open(argv[1], O_RDONLY);
10     else
11         fd = 0;  /* Use standard input */
12
13     while (read(fd, buf, 1) > 0) {
14         if (count < 5) write(1, buf, 1);
15         if (buf[0] == '\n') count++;
16     }
17 }
```

Line 5 declares a one character buffer to park the characters in on their way through the filter. Processing the characters one at a time like this is potentially very inefficient but it a simple scheme which will serve our present purpose. The logic in lines 8 through 11 implements the standard filter behaviour of "read from a file if you've got one, else read your standard input." The loop at lines 13 through 16 reads in the characters one at a time until the **read()** returns zero, indicating it has reached the end of the input. Line 14 copies the character to standard output if it's within the first five lines. Line 15 checks to see if the character was a newline character and increments the line count if so. Notice that the program doggedly consumes the whole of its standard input, not just the first five lines.

We can use this command in a pipeline like so:

```
$ sort /etc/passwd | five > firstfive
```

The *five* command is not aware that its standard input is coming from a pipe and its standard output has been redirected to a file.

Checking for changes in file status

Our masterwork for this month brings together many of the features we've just talked about. It's the bare bones (well, more like just the marrow really) of a security tool which monitors the file system for changes. A fully-developed tool of this type, such as *tripwire*, can be invaluable to detect tampering with the file system. Ours is rather trivial, in that it monitors only a single file.

There are actually two programs. The first, which we'll call *snap*, because it takes a sort of snapshot, captures a file's attributes using **stat()**, and squirrels them away in a file. It also stores a simple checksum of the file's contents, so we'll be able to detect if the file contents have changed. The second program, which we'll call *check*, and which is intended to be run at some later time, stats the file again and compares the result with the copy that *snap* saved, printing out any differences it finds. Here's the code for *snap.c*:

```
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <stdio.h>

main(int argc, char *argv[])
{
    struct stat info;
    int fd, sum=0;
    char buf[1];

    /* Check that the user supplied a filename */
    if (argc != 2) {
        fprintf(stderr,"usage: snap filename\n");
        exit(1);
    }

    /* Get the file's attributes */
    stat(argv[1], &info);

    /* Compute a checksum for the file */
    fd = open(argv[1], O_RDONLY);
23     while (read(fd, buf, 1) > 0)
24         sum += buf[0];
25     close(fd);
26
27     /* Save the checksum in the stat structure */
28     info.st_blocks = sum;
29
30     /* Write it out to a file */
31     fd = open(“snap.out”, O_RDWR | O_CREAT, 0600);
32     write(fd, &info, sizeof info);
33     close(fd);
34  }
```

The program reads the name of the file to take the snapshot of from the command line. Lines 12-16 check the user remembered to actually supply a filename. If not, it prints out an error message and exits. Line 19 captures the file's attributes, and the loop at lines 23-24 reads the file a byte at a time and adds them all up to yield a simple checksum. The idea here is to be able to detect changes in the file's contents. In reality we'd need to use a cryptographically secure checksum (such as MD5) for this; our simple scheme is too easy for a wily hacker to foil by making “compensatory” changes to the file which result in the checksum staying the same.

For convenience, at line 28, we press one of the less interesting members of the stat structure into service to store this checksum in. Finally, lines 31 to 33 write the stat structure out to a file *snap.out*.

On to the second program, check.c:

```
1  #include <sys/types.h>
2  #include <sys/stat.h>
3  #include <fcntl.h>
4  #include <stdio.h>
5
6  void compare_stat_structures(struct stat old, struct stat new)
7  {
8     if (old.st_uid != new.st_uid)
9        printf(“owner changed; was %d, now %d\n”, old.st_uid,
 new.st_uid);
10    if (old.st_mode != new.st_mode)
11       printf(“mode changed; was %o, now %o\n”,
 old.st_mode, new.st_mode);
12    if (old.st_size != new.st_size)
13       printf(“size changed; was %d, now %d\n”, old.st_size,
 new.st_size);
14    if (old.st_blocks != new.st_blocks)
15       printf(“checksum changed\n”);
16 }
17
18 main(int argc, char *argv[])
19 {
20    struct stat original, current;
21    int fd, sum=0;
22    char buf[1];
23
24    /* Check that the user supplied a filename */
25    if (argc != 2) {
26       fprintf(stderr,”usage: check filename\n”);
27       exit(1);
28    }
29
30    /* Get the file’s current attributes */
31    stat(argv[1], &current);
32
33    /* Compute the current checksum */
34    fd = open(argv[1], O_RDONLY);
35    while (read(fd, buf, 1) > 0)
36       sum += buf[0];
37    close(fd);
38    current.st_blocks = sum;
39
40    /* Get the original stat structure */
41    fd = open(“snap.out”, O_RDONLY);
42    read(fd, &original, sizeof original);
43    close(fd);
44
45    compare_stat_structures(original, current);
46 }
47
```

Lines 6-16 define a function called **compare_stat_structures**. It's not as complicated as it looks. It takes two stat structures (the old one and the new one) and compares some of the fields of the structures, printing out any differences it finds. For brevity, we compare only the file's ownership, access mode, and size, plus the checksum. Clearly the function could be extended to compare other attributes. In the **main()** function, lines 24-28 check as before that a filename argument was supplied to the program. Line 31 captures the current stat structure, and lines 33-38 compute the current checksum. This code is the same as in the *snap* program. Finally, lines 40-43 read the original stat structure back from the snap.out file, and pass the original and current stat structures to compare_stat_structures for comparison.

Here's an example of the command line dialog when we run these programs. I did these as root, mainly so we could see an example of the file ownership being changed.

```
1  # snap tester
2  # check tester
3  # chown root tester
4  # check tester
5  owner changed; was 1000, now 0
6  # chmod 644 tester
7  # check tester
8  owner changed; was 1000, now 0
9  mode changed; was 100777, now 100644
10 # echo hello >> tester
11 # check tester
12 owner changed; was 1000, now 0
13 mode changed; was 100777, now 100644
14 size changed; was 318, now 324
15 checksum changed
16 #
```

At lines 1-2 we run *snap* on a file called *tester*, then run *check*. Not surprisingly, there are no differences, so there's no output. At line 3 we change the file's ownership and re-run check; the change is duly reported at line 5. In line 6 we change the access mode and try again; both changes are now reported (lines 8,9). Finally at line 10 we append some text to the file and try once more. As expected, lines 14,15 report the change in the file size and the checksum.

There are any number of ways this program could be extended. A little error checking would not go amiss, for a start. Storing the stat structure in a fixed file (*snap.out*) is really dumb; it would be better to store it in a database with the filename as the key. It could be extended to operate on an entire dir, or descend recursively through a subtree of the file system. It could read the list of the files and directories to be monitored from a config file

NEXT MONTH

Pipes and other inter-process communication mechanisms, and six Really Useful Things to do with signals.

WORD PROCESSING

OpenOffice.org for power users

Open Office.org's (OOo) suite of tools is one of the best things to happen to the Open Source movement in a long time. This month Neil Lucock puts away the chisel and the marble slab and tries to write without using a hammer.

On the CD and DVD

Before we get down to any real work with OOo, I'm going to assume that you have used word processors before, perhaps MS *Word*, *Abiword* or *KWord*. What I'm going to do is look at the word processor part of the suite, show how to set up a dictionary, add more fonts and show you what you need to know to get the most from the best free word processor available for Linux. And I'm writing this article with it...

Spelling

When you launch *OOo*, you are presented with the word processor. (You can change this, if you want to). The first thing I needed to do was set it up to make learning *OOo* easier. If you look under Tools on the menu or click the icon with 'ABC' and a tick underneath it to the left of the working area, you can ask *OOo* to either spell check a word or to autocorrect. Autocorrect is a useful feature, if you always type "hte" instead of "the" it will save you some effort. However, I am going to type *OOo* a lot in this article, so I need to prevent it "correcting" the two initial capitals. One way is to go to Tools>Autocorrect/Autoformat and choose Options from the tabbed dialogue. The second tick box down allows you to prevent it from correcting two initial capitals. That's not the most efficient way of working, we are denying ourselves a useful tool by doing that. If we choose "Exceptions" from the tabbed menu, the lower part of the dialogue box (helpfully called "Words with TWo initial capitals") allows us make a custom rule. I typed *OOo* into the box, told it to add it and now *OOo* just underlines it rather than alters it. The spell checker can be told to ignore all instances of "OOo" in the text, but the next time we spell check it, it will ask you for confirmation that the word is correct. Either live with it or add it to the list of correct words.

Installing the dictionary. Once you have unzipped the file into the user/wordbook directory, go to Tools>options and edit the language modules.

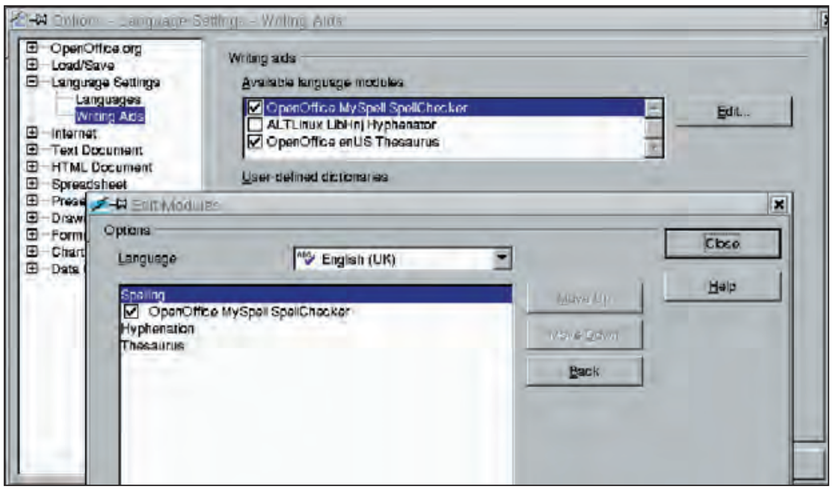

Available language modules

☑ OpenOffice MySpell SpellChecker
☐ ALTLinux LibHnj Hyphenator
☑ OpenOffice enUS Thesaurus

Options

Language: English (UK)

Spelling
☑ OpenOffice MySpell SpellChecker
Hyphenation
Thesaurus

Installing a dictionary

The next problem is that OOo comes with a US English dictionary and while US English is comprehensible to the British, I would prefer to write in my own variant of the language (despite mad spellings, crazy pronunciation and other irritations that make it so hard to learn). I might want to write in French, so this procedure applies equally to installing a dictionary for a second language. First go to the *OOo* website and download the dictionary (en_GB.zip) from the link. Uncompress all the files inside the zip file into the user/wordbook directory in *OOo*. Note that there are two "wordbook" directories in OOo, there's also one in share. Ignore it. The zip file has a few files inside it. You must ensure the original file called "dictionary.1st" is overwritten, as the new version has the lines:

```
DICT en US en_US
DICT en GB en_GB
```

That second line tells it to make the GB English dictionary available. You now need to go to Tools>Options and look at Language Settings. There are two options. Choose "Languages". If there is a tick by "English UK" then it will allow you to use it. It probably won't, so go to Writing Aids and edit the available language modules. All you have to do is choose "English UK" from the options in the drop-down box. When you return to Languages you should find that the GB English option now has a tick by it. Choose it and enjoy correct British spelling. It's not the most obvious method of installing a dictionary, but it does the job.

Now we have got *OOo* using the correct dictionary, we can start looking at the Word processor in a bit more detail. Make no mistake, the only light thing about *OOo* is the price, everything you need is there. It even fills in long words for you (tools>spellcheck>autospellcheck). Type a word and it automatically offers you the rest of the word the second time you start typing it. Press the **Enter** key to accept the suggestion (it doesn't always get it right) or continue typing to finish a different word. It's not a unique feature, but it's nice to have.

What it all does

We need to look around the interface and figure out what all the icons do. When we look at the other parts of the suite, (in the articles after this one) we'll just explain the different icons and sub-menus that *OOo* displays. Most icons will be familiar, but a few have unexpected behaviours. Some icons have small green arrows to indicate that there is a sub-menu, with others you have to click and hold to get a menu. For example, the text colour

Autopilot

Easy flying

OOo contains an "autopilot" feature. Basically, it's a wizard to help you with common tasks. You select File>AutoPilot and choose (for example) Fax. The process is easy and it gives you loads of options. You can insert information from a database, put in graphics (such as a logo), and a footer message (such as your name and product details). *OOo* allows you to save the fax you made as a template so that you can reuse the design. (You need something like *Efax* or *Hylafax* to allow *OOo* to send it, it treats it as another printer). Autopilot allows you to make letters, memos, agendas and a few other useful things.

icon: click it once and the selected text changes to the last selected colour. Hold it and it shows you all the colour options available. Some of the menus tear off and stay open, a useful property that makes it easier to use a group of tools.

At the top of the program is the usual menu bar. You can get at all *OOo's* functions from here. Note that the menus are context-sensitive. If there's nothing selected, the Edit menu will not give you the options of "Cut" or "Copy". Until you have copied or cut something, you can not see the "Paste" option under Edit.

You might want to go to Tools>Options and put in your user data and set a few other options. If you want *OOo* to start as a drawing, you do it under the Options>Load/Save>General. You can tell *OOo* to autosave your work, but you are using Linux, remember? – this program also runs on Windows. That's why it's there.)

Top line icons

Underneath the menu bar (labelled Top Line in the picture) there is a URL entry field. If you click it open, it shows you the files it has opened previously. You can also type in the filename and path to open a new file. You have a Stop button next to it. I suspect this is left over from *StarOffice 5.2*, it used to include a web browser in the interface. You can stop a file (or web site) loading if you feel the need. The Edit File button has odd behaviours. If you click it, it offers to save your changes, then gets rid of the second tool bar at the top and the side tool bar. Clicking again reverts to how it was before you clicked it. The New file icon opens a new document really quickly. The first time you click it you sit there wondering where all your text went. It's still there, just close the new empty document to get it back. If you click and hold, you get a menu that allows you to choose the type of document to open. The next icons are Open, Save, Print and the Cut, Copy and Paste group. The Undo and Redo icons follow. The Navigator allows you to get around a large document (amongst other functions) If you place a bookmark (from the Insert menu), name it and then go to somewhere else in the document, you can then click the Navigator button and choose the bookmark you just made. Your cursor jumps to the place where you made the bookmark.

The Stylist applies styles to paragraphs in your document. *OOo* includes a few ready made styles to use but does not display them in the drop-down box. Click Format>styles>catalog and you can use one of the supplied styles or define a new one. You might want all your headings to use an Ariel font at 24 points in red. Just select a paragraph (*i.e.* a heading), and apply your style to it. The icon next to it is for editing hyperlinks. Finally, there is the icon to start the Gallery. *OOo* includes many of the backgrounds, 3D models and sounds from *StarOffice*. Clicking the icon brings up the gallery browser window, allowing you to choose from a range of graphics to insert. Once you've had a look, click the small arrow at the bottom to hide it. The push-pin icon fixes it in its own window, leaving it open while you write. The icon is a toggle switch, clicking it a second time shuts the gallery down, presumably saving a bit of memory.

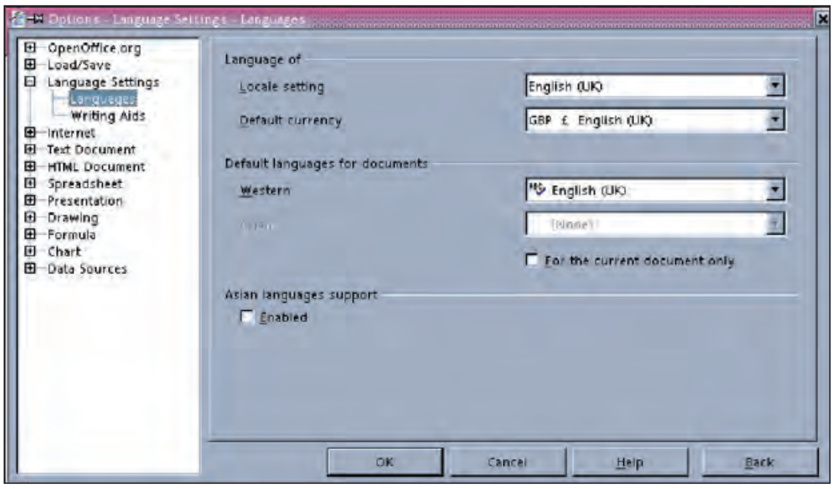

Options - Language Settings - Languages

Language of

Locale setting: English (UK)

Default currency: GBP £ English (UK)

Default languages for documents

Western: English (UK)

(None)

□ For the current document only

Asian languages support

□ Enabled

This is the second step to install the dictionary. Note that there is now a tick in the "default languages for documents" box. This is the desired end result.

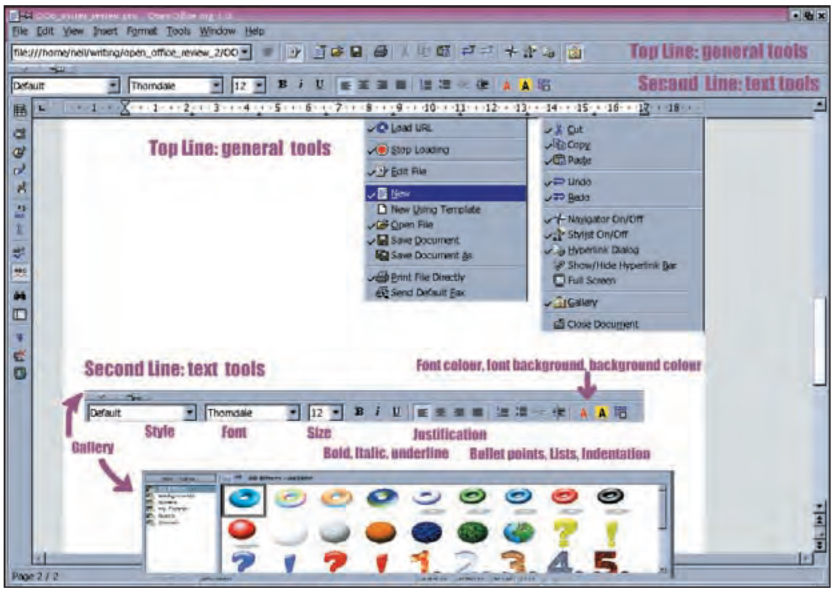

Top and Second lines of icons with The Gallery pasted underneath. It normally appears under the Second Line.

Second line icons

Most of the contents of the second line are obvious to anyone. The Style options is the drop down box on the left. *OOo* does not come with a lot of fonts, so we might want to make some of the other fonts on our machine available. First find out where they live. If you open "Find files" and look for *ttf you should find a few. KDE keeps some in usr/share/fonts/ttf

Open a terminal and become root. Navigate to your OpenOffice directory and find the file called *spadmin*. If you attempt to open this from a *Konqueror* window by clicking on it, it just starts the *OOo* program. If you type

```
./spadmin
```

in a terminal window it brings up a dialogue box. It's really for printer admin. You use the Fonts button to add some new ones. It copies the fonts unless you tell it to "create soft links only" which »

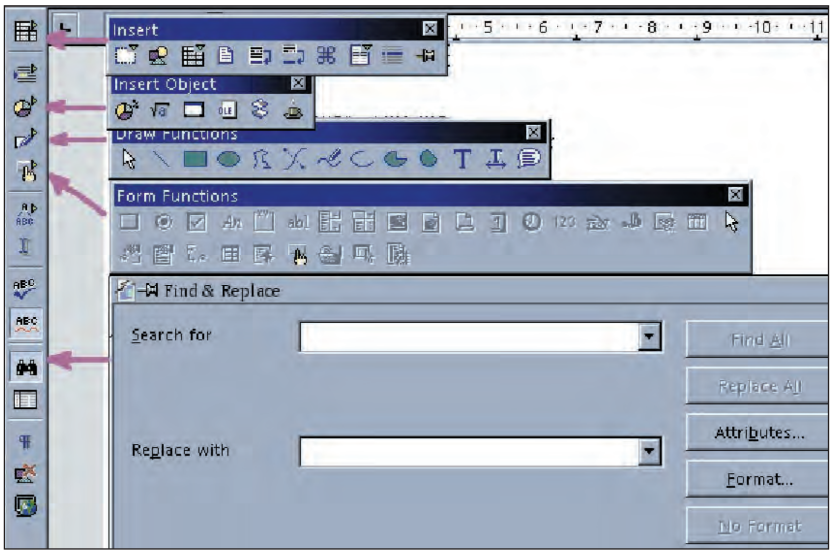

The sub-menus on the side tool bar.

saves some disk space. Do not just drop copies of your other fonts into the *OOo* directory, you'll break the program. (Guess how I found this out?) If you do break it, there's a Setup script in the OOo directory. This allows the program to mend itself

There are two bullet point icons, one for bullet points, the other for a numbered list. *OOo* has a really excellent range of configuration options for bullet points. You set how you want them to look under Format>Numbering/Bullets. This makes work done in the presentations part of the suite very impressive, but the word processor benefits from having it available too. The button at the far end of the row is for putting a background colour in a paragraph.

The vertical sidebar

The vertical sidebar at the left contains many sub-menus. The top icon duplicates a lot of the functions on the Insert menu. It allows you to put a multi-column frame in your page then paste some text and some pictures into it. Paste foreign characters or symbols in from here too. Once you use the tool, the icon for that task stays on the sidebar. This is handy if you want to reuse that function (just click it again) but can be confusing if you are trying to learn the menus.

The second one down allows you to paste various fields including the date, time, page numbers, a macro or a hidden message into the text. You can see the hidden messages by clicking View>fields on the menu bar. It's the sort of thing you will either never use or one of those tools you start to use and wondered how you ever managed without it.

Working down the sidebar the next icon has a pie chart icon. It inserts objects such as charts, applets and the *OOo* Math function creator. Below that we have the drawing functions tool bar. You can instantly start drawing something within the word processor. I intend to cover drawing in *OOo* in a later article.

OOo makes webpages and the next icon down allows you to make forms to use on your webpage. You can insert Radio buttons or images, tied into the macro scripts in *OOo*. The Help files have more on this.

The "Direct Cursor" is activated by the **I** with a drop-shadow. It places a small blue triangle at the left margin. If you move it down a few lines below the last line you typed and start typing, the text appears where you typed, leaving a few lines empty. It also works to accurately place text if you click in the middle of the page and start typing. It puts in carriage returns for you. I've already covered spell checking, the two icons allow you to start a spell check or turn on and off the automatic check.

The binoculars brings up the Find & Replace box. This has more functions than other word processors. If you want to find all instances of the word "OOo" and make them appear in a different font and colour, it's very easy. Enter the text you want to find or replace in the top box. "Attributes" refers to the text already in the document, so if you want to get rid of the bizarre font or colour you used on some words, but leave the ones you changed to Courier alone, it can do it. Type in what you want to replace it with. "Format" refers to the replacement text, so if you want to replace "OOo" with "Open Office.org" in purple Ariel at 14 points you can. The "No Format" button removes the Attributes and Format qualities you have defined. Remember that if you don't put something in the "Replace with" box, you'll replace your searched items with nothing. **Ctrl-Z** will undo it.

I intend to cover the next icon, data sources, in a later article. The printer's paragraph mark turns non-printing characters on and off. Word processors use these to format the display on-screen. Occasionally you do something that messes up your carefully formatted page and you can not seem to get rid of the mistake. Tell it to show all the hidden formatting characters, then delete whatever is causing you grief.

OOo likes memory, if you need to include a few large pictures in your work, it might slow the machine right down. You can use the next icon down to make it display placeholders rather than pictures.

The icon at the bottom of the sidebar shows you what your

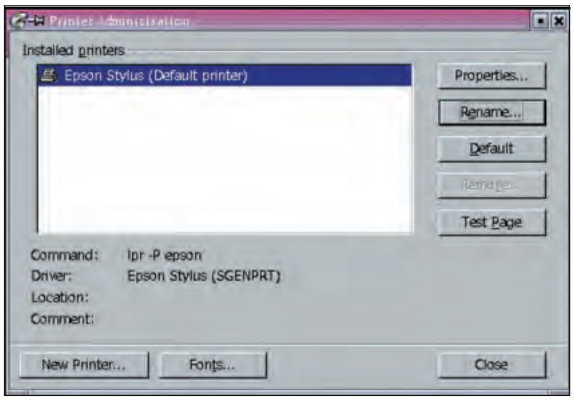

Epson Stylus (Default printer)

Command: lpr -P epson
Driver: Epson Stylus (SGENPRT)
Location:
Comment:

To add new fonts type ./spadmin and you get here. The font button is at the bottom.

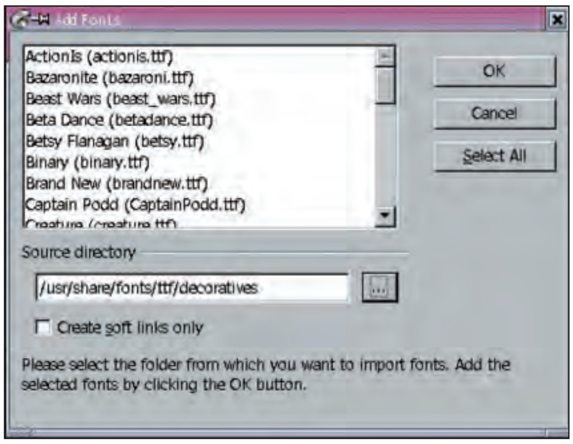

ActionIs (actionis.ttf)
Bazaronite (bazaroni.ttf)
Beast Wars (beast_wars.ttf)
Beta Dance (betadance.ttf)
Betsy Flanagan (betsy.ttf)
Binary (binary.ttf)
Brand New (brandnew.ttf)
Captain Podd (CaptainPodd.ttf)
Creature (creature.ttf)

Source directory
/usr/share/fonts/ttf/decoratives

□ Create soft links only

Please select the folder from which you want to import fonts. Add the selected fonts by clicking the OK button.

Adding more fonts. You next click "Select All", then "OK".

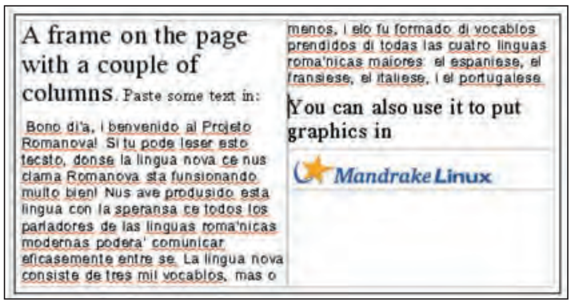

A frame on the page with a couple of columns. Paste some text in:

Bono di'a, i benvenido al Projeto Romanova! Si tu pode leser esto tecsto, donse la lingua nova ce nus clama Romanova sta funsionando multo bien! Nus ave produsido esta lingua con la speransa ce todos los parladores de las linguas roma'nicas modernas podera' comunicar eficasemente entre se. La lingua nova consiste de tres mil vocablos, mas o menos, i elo fu formado di vocablos prendidos di todas las cuatro linguas roma'nicas maiores: el espaniese, el fransiese, el italiese, i el portugalese.

You can also use it to put graphics in

Mandrake Linux

Multi-columned frame inserted into the page. The top icon on the vertical tool bar allows this.

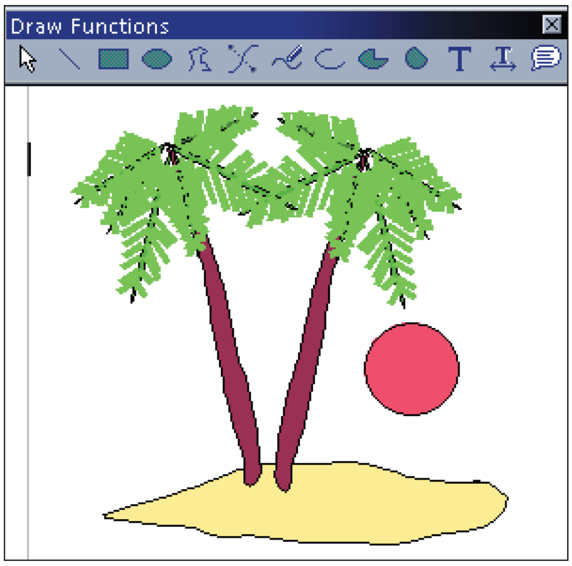

A quick sketch done directly in *Writer* with the drawing tools.

doc will look like on-line. *OOo* will happily make webpages, so this might be considered a web preview.

Whenever you are working there is a right-click menu. Once you know what's on it you'll find that it's much quicker to use than moving the mouse to the edge of the screen to click on something. See the screenshot for details of the contents.

Helpful

It is often a valid criticism that while Open Source software is well written, stable and free, you don't always get the most out of the coder's efforts. The people who write code are not always motivated to write Help files or able to explain in layman's terms how their application works. If you don't know what the program can do or can't figure out how to do something, the developers have wasted their time. *OOo* has an excellent and comprehensive Help system. If you use autocorrect it occasionally puts up a discreet light bulb picture in the lower right corner to tell you that it has done something, but unlike some word processors, it's never intrusive and if you ignore it it vanishes after a short time.

The Help files are astounding. You can read them by subject, there's an index with a search function and bookmarks so you can find that useful page again easily. It's the best Help system I've encountered in a free program, the only criticism I could make is that there could be a few interactive tutorials included to help real beginners (the program includes macros, so it could be done in *OOo* by someone talented) but even without them, the materials provided are well written and always provided satisfactory answers.

MS Word Import

Living with the proprietary world

I almost forgot to mention one of the key reasons to use *OOo*. Tell people that Linux is free, that the office suite is free, and they'll be unimpressed. Cost is not the only factor in choosing an office suite. How well something works has to be taken into consideration. Yes, *OOo* might be free, stable and capable, but the office manager still needs to know how to read MS *Word* documents that other people send. Most importantly, can you read the company's existing documents with the new system? *OOo* isn't the only Linux office suite to open MS *Word* .docs reliably, but it is the best free one. I quickly tested it on all the MS *Word* docs I could find and it opened them faultlessly. Even one I did at work with a map included.

The only minor problem was when someone had used a font that I don't have in Linux, it substituted another but it did not look nice (but it didn't look nice in Windows either). It also opened the MS *Office's* "Contemporary fax.dot" file (the template, see the screen shot) too, so you can still use fax templates that you have a licence for. However, be careful, in some of the new Microsoft end-user licences they forbid you to use their work on other operating systems. Read licences carefully before accepting them.

OOo runs on Windows, so try a version to see how good it is. Once you have it, think hard if you really have to buy the latest MS *Office*. And once you can open your *Word*, *Excel* and *PowerPoint* documents, your business doesn't need Windows to be productive. What are you waiting for? Save yourself some money, get Linux installed!

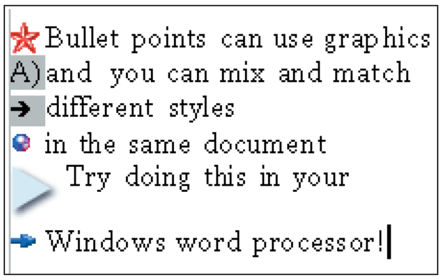

Bullet points can use graphics

A) and you can mix and match

different styles

in the same document

Try doing this in your

Windows word processor!

The variety of bullet points, all in the same document. Easy!

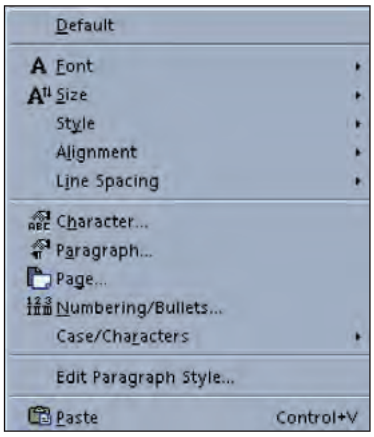

Right: **The right mouse-click menu. It's quicker than moving the mouse but not as fast as the keyboard shortcuts.**

OOo is big and having looked at it in-depth, I was very impressed. It did everything I could think of (and more) and it worked reliably. It is packed full of features and although I'm wary of saying it is "better" than MS *Word*, there are certainly a lot of tools in *OOo Writer* are not in *Word*. (Compare *Word's* "Find and Replace" with *OOo's*.) There were a few quirks. If you open a menu but don't select anything it leaves the menu open. You have to click somewhere else to close it. When it starts sometimes you have to wait for a few seconds after the document opens before you can edit it, but it's certainly no worse than other programs of a similar size. Get it and learn it and you'll reach the same conclusions as I did. It's a killer office app! LXF

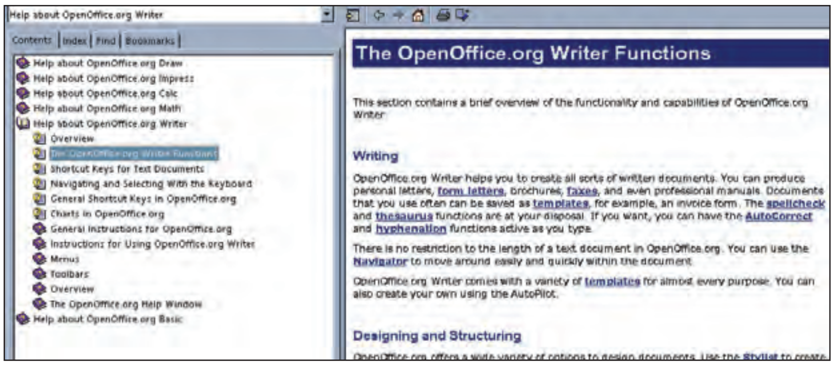

Help about OpenOffice.org Writer

Contents | Index | Find | Bookmarks

- Help about OpenOffice.org Draw
- Help about OpenOffice.org Impress
- Help about OpenOffice.org Calc
- Help about OpenOffice.org Math
- Help about OpenOffice.org Writer
  - Overview
  - The OpenOffice.org Writer Functions
  - Shortcut Keys for Text Documents
  - Navigating and Selecting With the Keyboard
  - General Shortcut Keys in OpenOffice.org
  - Charts in OpenOffice.org
  - General Instructions for OpenOffice.org
  - Instructions for Using OpenOffice.org Writer
  - Menus
  - Toolbars
  - Overview
  - The OpenOffice.org Help Window
- Help about OpenOffice.org Basic

The OpenOffice.org Writer Functions

This section contains a brief overview of the functionality and capabilities of OpenOffice.org Writer

Writing

OpenOffice.org Writer helps you to create all sorts of written documents. You can produce personal letters, form letters, brochures, faxes, and even professional manuals. Documents that you use often can be saved as templates, for example, an invoice form. The spellcheck and thesaurus functions are at your disposal. If you want, you can have the AutoCorrect and hyphenation functions active as you type.

There is no restriction to the length of a text document in OpenOffice.org. You can use the Navigator to move around easily and quickly within the document.

OpenOffice.org Writer comes with a variety of templates for almost every purpose. You can also create your own using the AutoPilot.

Designing and Structuring

OpenOffice.org offers a wide variety of options to design documents. Use the Stylist to create

The impressive OOo Help system. It's easy to use, complete and intelligent.

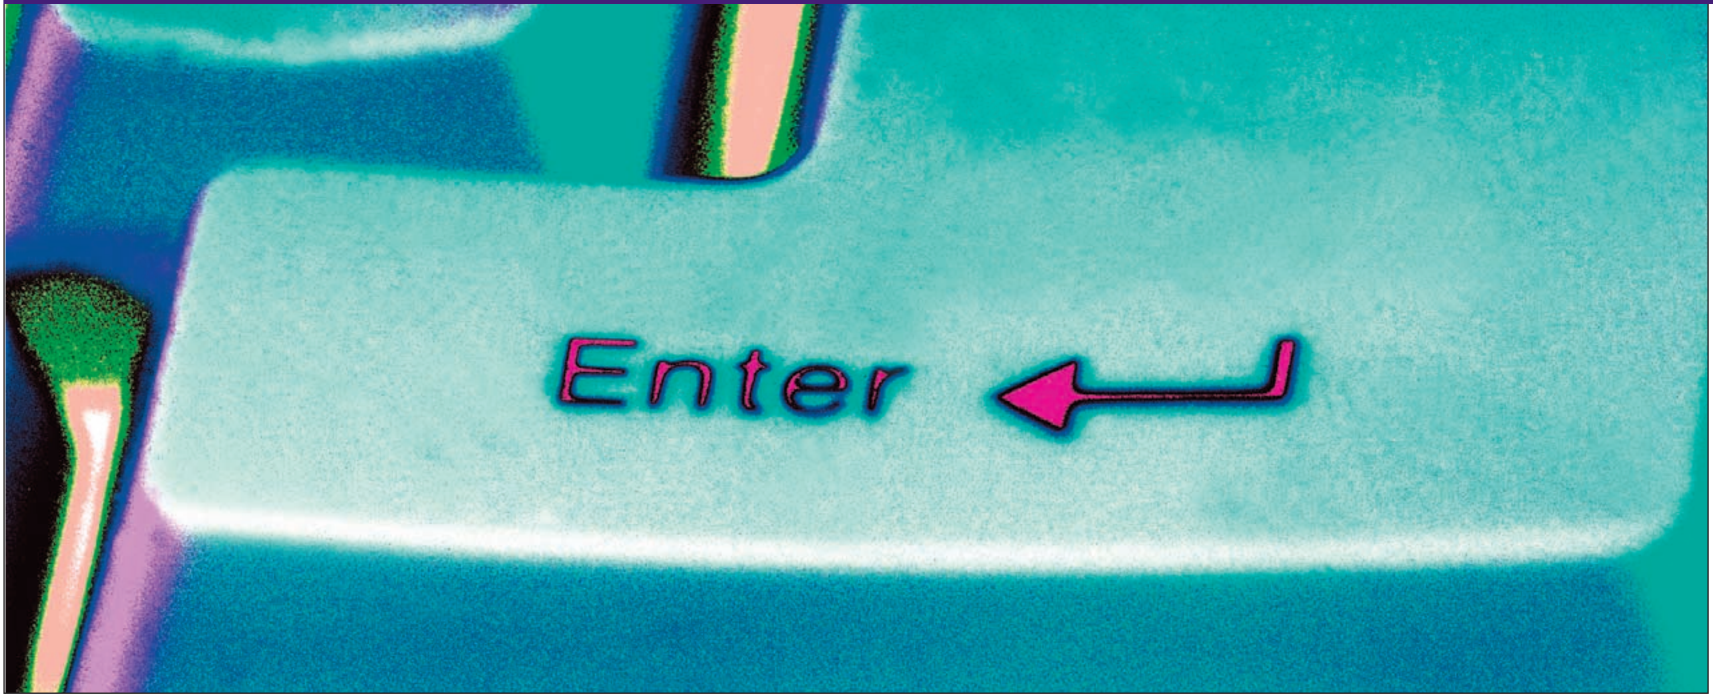

CODE COMMENTING

Fair comment

Wishing to be understood, Jono Bacon gets help from Doxygen in providing well commented code.

Code. The media that many of us spend hours looking at, days improving, and weeks optimising. At its most basic level, code can be easy to read and understand, but as a project developer, the ease of reading the code can readily fade – and you will be presented with a mess of spaghetti code that is both difficult to understand and off-putting.

This shameful scenario has spurred on the development of a range of code documentation processors. These processors are used to format and create online, printed and other documentation from your source code, and can also interpret special symbols within your code to adjust the documentation and formatting.

There are a range of these documentation processors, and some of them are discussed in the *Other Documentation Processors* box.

The comment concept

Before we take a look at one of these snazzy documentation processors – *Doxygen* – we will see what programmers used to use before the advent of these newer tools to document their code comments.

Comments are an important part of your programming routine. Although it is often taught to lay good comments in your code, and to document your bizarre yet amazingly powerful algorithms, many programmers rarely leave enough comments in their code.

When using comments in your code, it is important to remember why you are using them – to help the reader (which is often yourself) understand what is going on. Bearing this in mind, here are a few guidelines for comment common sense:

- Ensure that your comments are easy to read and legible. Comments incorporating bad grammar will not do the reader any favours
- Keep your comments short and to the point. Long, drawn out pages of comments that are not specifically in context, will just bore the reader instead of informing them.
- Keep the technical jargon to a sensible level. If you are documenting a function that is discussing how a particular process works, do not expect the reader to be familiar with all elements of the process and hence use the jargon to discuss it. By all means don't avoid the jargon, just don't overdo it.

Clean commenting

Although many would think comments are a simple process to grasp, it is amazing how comments can hinder understanding of a program, process or method. Much of this lies in how the comment is presented to the reader. Take for example two comment styles for the same function:

Example1:

```
/* This method will do something very snazzy */

MyMethod()
{
}
```

Example 2:

```
/* ————————————————————————————
  ///              myMethod              ///
  ————————————————————————————
  |                                        |
  |      This method will do something     |
  |              very snazzy               |
  |                                        |
  ———————————————————————————— */

MyMethod()
```

```
{
}
```

The two examples show two very different ways of documenting a method. The first is very short and efficient and is in my view the preferred type, particularly given that the information in the comment is very short. The second example is more verbose, and although it may grab your attention, imagine 30 of those comments in a file – it would bloat the file and make it much longer than required – and harder to find what you are looking for.

When documenting methods and functions, it is also advisable to say what the function should do, what it expects, and what it should output. This type of commenting always ensures that the reader knows the extent of the method functionality. For example

```
/* FUNCTION: to show the details of a staff member
   REQUIRES: a two digit integer (eg. 24)
   OUTPUTS:  displays the staff member (text)   */

void showStaffMember( int id )
{
}
```

This comment clearly states the information a programmer would need to know to use the method **showStaffMember()**. The information is detailed but easy to understand and states the details in a concise manner.

Using Doxygen

We will now take a look at one of the most popular documentation processors available and one that has a good level of features available. *Doxygen* is Free Software, is maintained by a team of dedicated developers, and can be used to document C, C++, Java, IDL and even PHP. One of the nicest elements of *Doxygen* is just how nice the output is visually – it looks very nice.

To get started with *Doxygen*, you first need to install it. There are a variety of ways of installing *Doxygen* on Linux:

For RPM based systems (Mandrake, Red Hat, SuSE, etc.):

The RPM's for *Doxygen* are located on Dimitri Papadopoulos' site at **http://perso.club-internet.fr/dpo/rpm/**. You can also get RPM files for your particular distribution at **www.rpmfind.net**. You can then install the package with:

```
rpm –I packagename.rpm
```

For Debian systems:

You can use your good friend *apt* to get *Doxygen* for you with:

```
apt-get install doxygen
```

To compile:

You can compile the source code by first getting the latest source .tgz package from **www.stack.nl/~dimitri/doxygen/download.html** and using the following commands to compile it:

```
gunzip package.tgz
tar xvf package.tar
./configure
make
```

Getting started

To get started with *Doxygen*, we first need to create some code that we will use as an example for documenting. We will create two C++ files for this purpose:

File1, myclass.h:

```

#include <iostream.h>

class MyClass
{
        public:
                MyClass();
                ~MyClass();

                int testOne( int a );
                void testTwo();
};

```

File2, myclass.cpp:

```

#include "myclass.h"

MyClass::MyClass()
{
}

MyClass::~MyClass()
{
}

int testOne( int a )
{
}

void testTwo()
{
}
```

We will also create a directory in our project directory to hold the documentation that is generated. We will call this doc:

```
mkdir doc
```

Now we have set up some code to be documented, we will now look at configuring *Doxygen* to our needs.

Configuring Doxygen

Doxygen works in a similar fashion to the other build tools you will use on Linux such as *Automake* and *Autoconf*. *Doxygen* requires a configuration file that is read by the *Doxygen* parser to determine what gets documented and how it is formatted. We can generate a new file with the command:

```
doxygen -g <config file name>
```

We can call the configuration file anything we want (a file called 'Doxygen' is produced if we do not specify a filename). We will keep it simple and call it doxygen.conf:

```
doxygen -g doxygen.conf
```

If you look in the directory you will see the new file has been created. We can now load the file into our favourite text editor to start tweaking it.

You will see that the format of the file is similar to many other Linux configuration files. The format uses the system:

```
FEATURE = SETTING
```

With this system you can specify how the feature works by supplying text on the right side of the **=** sign. An example is if you turn on **EXTRACT_ALL**:

```
EXTRACT_ALL = YES
```

To begin with we will set our project name (**PROJECT_NAME**) to **"Doxygen Test"**. You will need to put this in double quotes as » there are spaces in the name. We will also set **PROJECT_NUMBER** to **1.0**. The **OUTPUT_DIRECTORY** line specifies an absolute (*e.g.*. /home/jono/doxytest/doc) or relative (*e.g.* ./doc) directory in which our generated documentation will be stored (note that this is relative to where we execute the command that generates the documentation later).

We can now save the changes to our configuration file and use it to generate some documentation. This can be done with:

```
doxygen doxygen.conf
```

You will then see some messages flash down the screen as *Doxygen* parses and builds the documentation for you. You can then go to the directory that you specified in **OUTPUT_DIRECTORY** and you will see two directories created – html and latex. If you go into the HTML directory and load the index.html file you can view your documentation. It is a good idea to read the documentation in a CSS capable browser such as *Mozilla* or *Konqueror*.

When you look at the generated documentation, you may notice how sparse it is of the information on your methods. The reason for this is that by default we have not specified what to document and as there are no special *Doxygen* symbols in the code, very little is generated.

We can turn it on by setting **EXTRACT_ALL** to **YES**. This setting will document everything it sees in the code. This is often useful for seeing how *Doxygen* works without writing lots of special comments. Run the generation command again to generate your documentation:

```
doxygen doxygen.conf
```

Taking a look at what is created

We will now take a quick peek at what we have been generated by *Doxygen* on our simple class. When you load the main index.html page, you will see the main documentation topics at the top of the page. These include:

MAIN PAGE – This takes you back to the main root of the documentation.

COMPOUND LIST – This shows a list of the data structures in the program.

FILE LIST – This is a list of the files that have been identified by *Doxygen*.

COMPOUND MEMBERS – This is a list of the members that have been parsed in the classes.

FILE MEMBERS – This shows a list of members and the files they are in.

You will also notice that everything is completely cross-referenced, so you can easily jump between these different sections of the documentation.

Expanding our documentation

Now we have got a basic set of documentation generated, we can look at some more of the options in the configuration file to how they can extend our documentation:

SOURCE_BROWSER

Set this option to **YES** to have the code that is being analysed included in your documentation via reference. This option is particularly useful for visualising code, and is recommended.

INLINE_SOURCES

You can use this option to have your source code included directly within the documentation instead of referencing it. This can be useful unless there is a vast amount of code that may bloat the documentation.

REFERENCED_BY_RELATED

With this option all methods that are documented will be shown by the methods that reference them.

REFERENCES_RELATION

With this option, all methods that are documented are listed within the function.

Detailed docs with symbols

So far we have focussed on the documentation that can be generated with pre-existing code. While this is without doubt a useful tool, this is only half of the picture with *Doxygen*.

Doxygen will allow you to insert special comments in your code that can not only be used as regular code comments, but also used by *Doxygen* to add to and format our generated documentation.

Adding comments is a simple matter and you merely add a few extra symbols to the code. There are essentially three ways of leaving a comment:

Adding a second * after the opening /* (this is also how *JavaDoc* does it):

```
/**  This is a comment */
```

Adding a ! after the opening /* :

```
/*!  This is a comment */
```

Adding an extra / after a single line C++ style comment:

```
///  This is a comment
```

Adding an extra **!** after a single line C++ style comment:

```
//!  This is a comment
```

You can of course have multiple line comments such as:

```
/** This is a
  * very nice and
  * concise comment */
```

When writing a comment, there are essentially two different parts of the comment block – a brief and a detailed section. These can be differentiated with **\brief** command within the comment:

```
/** \brief This is a brief section.
  *
* This is a more detailed section. */
```

When you use these comment blocks in your code, you will need to place them *before* the item that you want to document. *E.g.*:

```
/** \brief A method that adds two numbers together
  *
  * this method will take the two parameters
  * a and b and add them together
  * and return the value as an integer            */

int addNumber( int a, int b)
{
}
```

There are also some additional commands you can use inside your comments for methods:

\param – This is a parameter that is passed to the method.

\return – This specifies what is returned and in what type.

\sa – This is *See Also*. You can use this to suggest other methods that the reader may find useful.

Let us now take a look at our previous code and add some of these new *Doxygen* comments to it:

```
#include <iostream.h>

/** \brief An example class to use for Doxygen
```

```
         *
         * This class does nothing useful aside from providing
         * us some code that we can document.              */

        class MyClass
        {
                public:
                        /// Constructor
                        MyClass();

                        /// Deconstructor.
                        ~MyClass();

                        /** \brief A test method that
                         * takes a parameter.
                         *
                         * This method does nothing
                         * useful but let's us use
                         * Doxygen to document a
                         * method that take a
                         * parameter.
                         *
                         * \param a is an integer
                         * \return an integer is returned
                         * \sa testTwo()              */

                        int testOne( int a );

                        /** \brief A test function that
                         * does not take a parameter.
                         *
                         * This method does not take
                         * a parameter is and is used
                         * only to show how Doxygen
                         *  works.
                         *
                         * \return Returns nothing
                         * \sa testOne( int a )       */

                        void testTwo();
        };
```

When you have added these comments to the myclass.h source file, you can then generate the documentation again with:

```
        doxygen doxygen.conf
```

You can see how the additional information is cleanly laid out in the various sections.

Tweaking the generated output

Doxygen has some features in it's configuration file that are useful for tweaking the output that it generates. This is particularly useful for HTML output as it can really customise the look and feel of the generated pages.

Most of the configuration is done using the configuration file we generated earlier. There is a section in it specifically for setting up your HTML output, and here are some of the options worth noting:

GENERATE_HTML – Set this to **YES** if you want to have some HTML output generated.

HTML_OUTPUT – This specifies the name of the directory in which the HTML output will be stored. Remember that this is used in conjunction with the **OUTPUT_DIR** that was set earlier.

HTML_HEADER and HTML_FOOTER – These can be used to specify a header and footer that can be used to sandwich the documented information in a customised format.

HTML_STYLESHEET – Use this to specify a stylesheet that can tweak the appearance of the output. Particularly useful for fine-tuning font presentation on the page.

Latex generation

Many people only use *Doxygen* for the generation of HTML documentation. Although this is fine, there are some other output formats that are useful; one of the most useful being *LaTeX*. *LaTeX* is a professional quality printing language that is often used for technical manuals.

Doxygen will generate *LaTeX* support by default, but will not actually create the output for you. This is no mean feat (assuming you have the *LaTeX* software installed on your system) and can be done by typing **make** in the 'latex' directory that is created when you generate your documentation.

It is also worth noting that there are a variety of settings in the configuration file for setting up *LaTeX* support.

Summary

Commenting and documenting our projects is a job that is loved by some and hated by many. In these pages we have looked at the common methods and pointers for commenting code, and we have looked at using a tool such as *Doxygen* to help us with this often unenviable task. *Doxygen* is yet another tool that has been developed, tested and extended to satisfy the requirements of many people.

It is a truly useful piece of software and is increasingly used in free, commercial and educational software projects. More specifically, *Doxygen* is used by (amongst others) the game *Pingus*; the libraries *KDE-DB* and *Gandalf* – not to mention the *GNU Standard C++ Library*; *Rosegarden* sound software; the *GNUnet* p2p infrastructure; the *AbiWord* word processor; the *Orca* CMS tool and the HP Officejet Linux driver; as well as the GPL'd desktop OS, AtheOS.

We all need documentation at some point. Document as you code, and pick a good tool for the job.

Other documentation processors

Different languages, different systems

There are a variety of documentation processors available for various languages and systems. Here are some of the best.

DOXYGEN – *Doxygen* is rapidly becoming one of the most popular documentation processors available. *Doxygen* is free software with a full feature-set that can generate documentation from standard source code and special inserted symbols. Documentation can be formatted in HTML, Man Pages, *LaTeX* and other formats.
www.doxygen.org

JAVADOC – *JavaDoc* is a documentation processor for not surprisingly...Java. Automatic generation of documentation can be done in various formats and methods/classes can be configured to be documented.
http://java.sun.com/j2se/javadoc/

KDOC – *KDOC* is a C++ and IDL documentation tool that was developed specifically for documenting the KDE libraries. *KDOC* works like many of the other processors by using symbols to document the code, and by looking at the code syntax to document the API.
www.ph.unimelb.edu.au/~ssk/kde/kdoc

ADOC – *aDoc* is a documentation processor written in *awk* and shell scripting code. The processor can be used to generate *LaTeX* pages that are similar to the structure of man pages. *aDoc* seems to support almost every variety of programming language.
www.sect.mce.hw.ac.uk/~peteri/adoc

CODOC – *CODOC* is a C documentation processor. It will allow you to create a template file where your formatting instructions can be processed.
http://codoc.sourceforge.net/

BYTECODE INTERPRETER

Parrot soup

Charlie Stross investigates the plumage of Parrot, the bytecode interpreter behind the forthcoming Perl 6

Perl 5 came out in 1994, or thereabouts. Perl 6 isn't out yet, but work on the language began in 2000 – it's having a long gestation. Some concrete material has, however, emerged. First, a lengthy public consultation period took place; then Larry Wall and his helpers began to emit a series of *Apocalypses* – design documents covering different aspects of the Perl 6 functional requirements.

A lot has changed since Perl 5's heyday. Newer, cleaner languages have come along: Python, in particular, has gained a considerable following for its elegant modularity, and nobody can be unaware of the existence of Java and C# (although the latter is primarily a Microsoft toy, and the former is aimed more at application development than Perl was, initially). Meanwhile, Perl 5's weaknesses have become apparent. It's not impossible to do large-scale application development in Perl 5, but it *is* hard to maintain consistency between the work of novices and Perl 5 experts, and good development practice is hard to establish.

Perl 6 is therefore a work in progress that is intended to address the weaknesses of Perl 5 – modularity, maintainability, extensibility, and suitability for large-scale application development. Perl 6 will be mostly backward-compatible: the target is for 90% of Perl 5 programs to run without changes, and for 90% of the remainder to work with only minor changes. But it's going to be a whole new language, written from scratch, and the first part of it to emerge in the form of code is already available: *Parrot*.

Pretty Polly

Parrot takes its name from an April Fool's joke by Simon Cozens. Simon announced a merger between Perl and Python, orchestrated by Larry Wall and Guido von Rossum, who were going to collaborate on a new language called Parrot. As it happens, there's a core of truth in the joke. Both Perl and Python are hybrid languages that contain both a compiler and an interpreter. The compiler generates a parse tree which the interpreter then executes. Unlike a purely interpretative language, this permits the benefits of compilation (notably various optimisations), while keeping the benefits of interpretation (easy debugging). Perl 5's structure is a bit of a mess internally, so one of the goals of the design exercise was to separate the compiler from the bytecode interpreter.

At this point, Cozens and friends had a brain-wave: if they were going to build a bytecode interpreter (like the Java Virtual Machine) why not do it properly? Why not write one that could run languages other than Perl – like Python? Or Ruby, Tcl, and anything else anyone felt like writing a compiler for? And so, to *Parrot*: *Parrot* is a bytecode interpreter optimised for running high-level languages. It'll be the back-end of Perl – but there will be other compilers for it as well. Indeed, given the Perl language's innate suitability to pattern-matching and compilation tasks, Perl 6 will be a toolbox for designing and building application domain-specific languages. And it all hinges on *Parrot*.

Parrot is implemented in C – for maximum portability. (C is about the only language that is available everywhere, and sufficiently fast.) It's capable of interpreting several million opcodes per second on a typical PC. However, it also has a JIT (Just In Time) subsystem which permits compiled bytecode to be translated to native machine code and executed directly (only on Intel x86 or Alpha CPUs, running Linux or BSD, right now).

Unlike the Java virtual machine, Parrot is a register machine, that recognizes four basic data types: **INTVAL** (an integer type guaranteed wide enough to hold a pointer), **FLOATVAL** (architecture-independent floating point), **STRING**, and **PMC** (a low-level abstract scalar data type: **Parrot Magic Cookie**). Each data type has 32 associated registers, for example **I0 .. I31** (**INTEGER** registers **0** to **31**). The core *Parrot* assembly language is simplistic but has some interesting opcodes you don't normally see on real microprocessors – such as string operators. Despite being a register machine, *Parrot* has a stack and permits you to push and pop frames of registers on a stack – segregated by register type (so that, for example, you can push all the **PMC** and **FLOATVAL** registers without touching the **INTVAL** and **STRING** registers). There's also a scratchpad mechanism for keeping track of the name, type, and attributes of variables (which might be needed by a higher level language for checking constraints), and exception handling mechanisms.

It's no surprise to find that *Parrot* provides support for handling objects and compiled bytecode libraries (modules) and shared libraries. More interestingly, there's low-level support for obtaining locks on objects and registers and for spawning new interpreters and for resynchronising threads: *Parrot* is designed to support multi-threading. You can find all the gory details in the (mercifully short and painless) *parrot_assembly.pod* file that comes in the docs subdirectory of the *Parrot* distribution.

One point to note: *Parrot's* data types aren't polymorphic. On the other hand, they don't support the intricacies of variable type behaviour you'd expect in Perl or Python. Policy on type conversion needs to be specified by the compiler – or established by adding new data types. *Parrot's* mechanism to do this is to define a new object class with appropriate methods, then pass around pointers to objects in this class using the **PMC** registers and opcodes. Tools are provided that allow you to generate a template for your new class, which then needs filling in and compiling as an extension to *Parrot*. In this way, *Parrot* can support scalar types for Python and Perl that have incompatible behaviours.

So what's it good for?

You can grab *Parrot* (current release: 0.0.7) from CPAN and compile it using the GNU utilities – if you're unused to this, extract the contents of the tar.gz archive and read the file README inside it. Once you've built *Parrot*, you can run the regression tests to verify that it functions as advertised, then read the documentation in the docs subdirectory. Most of the really interesting stuff is buried in a subdirectory called languages. This is where the compilers live – programs written in *Parrot* assembly language that compile other languages down to *Parrot* code. For a squeaky-clean new system, *Parrot* has already spawned a lot of compilers.

Parrot is an assembly language, but an abstract one – optimised for the job of writing compilers and interpreting the code they spit out, rather than optimised for the job of running on physical silicon. A whole bunch of small mini-languages have been implemented in *Parrot* so far, for testing purposes. First in the list is probably the *Parrot* compilation system itself – the *Parrot* compiler will compile any *Parrot* program written with fully qualified opcode names. There are also a couple of toy compilers, *Cola* and *Jako*; *Cola* is a smallish C-like or Java-like language that's under development, while *Jako* is another very simple language with just enough complexity to implement **while** loops. At a higher level, small compilers for Forth, Scheme, and BASIC (admittedly an old-fashioned line-numbered BASIC compiler) are provided – if you want proof of concept of *Parrot's* suitability for compilation, the startling number of compilers that have landed on top of what is, after all, only release 0.07 is remarkable! And finally, there are the beginnings of the prototype implementation of Perl 6 itself.

There are a couple of different components to Perl 6. One is the regex compiler. Regular expressions are a mini-language in and of themselves; the proposed Perl 6 extended regular expressions are actually rather more powerful than the basic pattern-matching expression language, and Perl 6 *Apocalypse 5* adds a whole slew of new facilities, such as named rules and facilities for defining grammars. To implement regular expressions in Perl 6, a separate regex compiler and mini-language is provided, and the core of this is provided in parrot/languages/regex – a regular expression compiler. There's also the separate Perl 6 compiler itself – or the sandbox in which it is to be developed – and a miniperl compiler (which builds a minimal Perl core without most of the language extensions we use, but sufficient to help bootstrap the Perl build process).

It needs to be stressed that as of this release, none of these languages are much use. But as Benjamin Franklin commented, "how much use is a baby?" Parrot is going to grow up, and when it does, expect these mini-languages to be joined by full-blown implementations of Perl 6, Python, and possibly Ruby, Java, and Tcl as well.

Parrot says "hello, world!"

Here's a simple piece of Parrot code (lifted shamelessly from the manual):

```
set I1, 10
set N1, 3.1415
set S1 "Hello, Parrot"
```

All we're doing here is inserting values into some registers – the integer **10** into the **I1** (integer 1) register, the float **3.1415** into **N1** (number 1) register, and the string **"Hello, Parrot"** into the **S1** string register.

We can print like this:

```
print "Register S1 contains: "
print S1
print "\n"
```

And we can loop like this:

```
          set I1, 1
          set I2, 10
REDO:     print "Hello "
     print I1
     print "\n"
     inc I1
     le I1, I2, REDO
DONE:     end
```

Which should (your columnist's rusty assembly language skills notwithstanding) print **"Hello <n>\n"** repeatedly for values of **<n>** from 1 to 10. And it's a testimony to Parrot's easy learning code that indeed it worked first time. It's functionally equivalent to the following Perl code:

```
 $a = 1;
 $b = 10;
 do {
    print "Hello $a\n";
    $a++;
 } while ($a <= $b);
```

Compare this to the fun you'd have with a real microprocessor's assembly language – which as a general rule don't come with explicit string data types, so you'd have to define **"Hello $a\n"** as an array of **chars** and loop through it for each **print** statement! Parrot is not only a great platform for Perl 6, it's an easy and rewarding tutorial environment for learning basic assembly language programming.

Parrot futures

The evolving interpreter

This tutorial is based on Parrot 0.0.7, which was the current release in August 2002. Parrot is moving fast, and 0.0.8 came out on September 2nd, followed rapidly by 0.0.8.1, after most of this article was written.

Parrot is moving fast; there are a number of significant changes in 0.0.8.1. A number of new language front-ends have found their way into the Parrot distribution, including the foundations of the planned Perl 6 regular expression compiler (which compiles down to Parrot bytecode). The regexp compiler isn't complete; in particular, several major Perl 5 regexp features such as character classes, word boundaries, back references, and look ahead/look behind assertions aren't implemented yet, and some Perl 6 features (mostly regexp operations on arrays and hashes, and weird stuff like two-dimensional regular expressions) aren't there either. The Perl 6 compiler is coming along, and the beginnings of some new compilers have appeared – Python, Ruby, and (rather surprisingly) Befunge-93. (What next, INTERCAL?)

The documentation has been updated a little, and a load of new pmc classes have appeared, including the basics of the boolean and keyed data types.

Also a bit boggling are some new tools that have shown up in the language's subdirectory – specifically, parsers that are designed to translate BNF pseudocode into grammatically correct Perl 6, and BNF grammars for a number of languages. The importance of this toolkit can't be understated if Perl 6 is to reach its goal of becoming a sort of language for creating application domain-specific sublanguages.

XML AND XSLT

Practical PHP programming

This month, **Paul Hudson** demonstrates how to put the power of XML to use in your sites.

Although greatly touted several years ago, it has only been the current drive towards Web services that encouraged people to take on XML. Because it is human-readable, XML is a great way to allow programs and companies to inter-operate smoothly, and therefore is perfect for the Web. Unsurprisingly, a lot of work has been put into PHP support for XML, and its related technology, XSLT. Note that this article does not strive to be an XML tutorial, although I'll try to make things as easy as possible for non-XML users.

Introduction to XML using PHP

XML, being human-readable, needs to be parsed by PHP. That is, PHP needs to convert the text into something it can understand and allow you to work with. PHP allows you to choose two methods of parsing – event-based (using the *Expat* library), and tree-based (**DOM XML**). At the time of writing the **DOM XML** functions in PHP are experimental, so for this article we'll be covering PHP XML parsing using the event-based parser *Expat*. *Expat*, being an event-based parser, processes your XML document and reports events through callback functions you specify. To give you an idea how it works, consider the following XML:

```
<lie>
ASP is better than PHP
</lie>
```

As you can see, the XML contains only one element, **<lie>**, which is the root element. When this element is encountered by *Expat*, it calls the function you specified for the start of an element, passing in the parameter **lie**. Then it calls the function you specified for handling character data, passing the parameter **ASP is better than PHP**. Finally, it reads **</lie>**, and calls the function you provided for closing elements, passing in **lie** again. When PHP calls your function for starting element, it also passes in an array of attributes that parameter has. This is shown below:

```
<lie type="statistic">
ASP is 20% more popular than PHP
</lie>
```

In that example, PHP would pass an array to your start element function containing one element, **type**, with the value **statistic**. Now you have an idea of how an event-based parser works, we can take a look at how PHP implements *Expat*.

Firstly, the function **xml_parser_create()** is used to create an instance of the *Expat* parser. **xml_parser_create()** returns a resource handle to your XML parser, which you can then use in other XML functions. There's a corresponding **xml_parser_free()** function which takes this resource handle as its only parameter and frees the memory assigned to the XML handle. Be a good programmer: free your XML handles as soon as you're done with them! Once you have a handle to an XML parser, you can assign your callback functions. **xml_set_element_handler()** and **xml_set_character_data_handler()** are the two key PHP functions here. The former allows you to assign the functions *Expat* should call when it finds a starting element (*e.g.* **<BOLD>**) and when it finds an end element (*e.g.* **</BOLD>**). **xml_set_element_handler()** takes three parameters: the handle to your parser, the name of the function to call when it finds a starting element, and the name of the function to call for end elements. **xml_set_character_data_handler()** is much the same, and configures PHP to use its second supplied parameter as the function to handle character data. So, it takes two parameters: the handle to your XML parser, and the function to use for character data. The parameters to these functions must be defined in a very specific manner. This is shown in the code block below:

```
function startElement($parser, $name, $attrs) { }
function endElement($parser, $name) { }
function charData($parser, $data) { }
xml_set_element_handler($parser, "startElement", "endElement");
xml_set_character_data_handler($parser, "charData");
```

The first three lines are empty function definitions for the start element handler, the end element handler, and the character data handler. While you can change the name of the functions, the parameters they accept must remain as shown there. The final two lines demonstrate configuring PHP to use **startElement**, **endElement**, and **charData** as callback functions.

Get Syndicating!

Motivational resources

Here is a short list of six popular RSS files available online, for you to try out:

Linux Planet – www.linuxplanet.com/rss

Linux Weekly News – www.lwn.net/headlines/rss

Motivational quote of the day – www.quotationspage.com/data/mqotd.rss

Dictionary.com Word of the Day – www.dictionary.com/wordoftheday/wotd.rss

PHP Builder articles – www.phpbuilder.com/rss_feed.php?type=articles&limit=20

***LXF* itself! – www.linuxformat.co.uk/backend.php**

Note that a list of RSS files is up online at the Nautilus site – check out http://nautilus.eazel.com/rss/. Be warned, though, some of the RSS files are more complicated than our simple code can handle. Ampersand (&) is reserved in HTML and XML, as are <, >, and “. As a result, these are encoded in XML as &amp;, &lt;, &gt;, and &quot; – you may need to unscramble these in order for certain newsfeeds to work.

Before you crack on parsing the XML itself, it's very common to want to configure the *Expat* parser using the function **xml_parser_set_option()**. The most common option to set is case folding, which is basically a fancy term for uppercasing. All values passed in as the **$name** and **$attrs** parameter in your **startElement** and **endElement** handlers are automatically converted to uppercase if case folding is enabled, excluding the values of elements in the attributes list. Case-folding is enabled by default – you can use the following to disable the feature:

```
xml_parser_set_option($parser,
 XML_OPTION_CASE_FOLDING, false);
```

Parsing XML

Parsing XML requires two steps in its simplest form – reading in the file, and handing the contents over to *Expat*. Reading in the file to parse is easy enough (refer back to *LXF31*), however the actual parsing itself needs explaining. Here is the code to use for parsing an XML file – it won't work yet, because there's no XML file. Save it in your PHP scripts directory for now, as *lxfxml1.php*.

```
$file = '/wom/bat/test.xml';
if (!($fp = fopen($file, “r”))) die(“Can't open file!”);

$parsesuccess = true;

while($data = fread($fp, 512)) {
  if(!xml_parse($parser, $data, feof($fp))) {
    $parsesuccess = false;
    break;
  }
}

fclose($fp);

if(!$parsesuccess) {
    die(sprintf(“Fatal XML error: %s at line %d”,
xml_error_string(xml_get_error_code($parser)),
xml_get_current_line_number($parser)));
} else {
   echo “Parsing completed successfully!”;
}
```

The variable **$parsesuccess** will be used to determine whether any errors were encountered while parsing the XML, and by default we say there are no errors. We then enter into a **while** loop. The criteria for the loop is **$data = fread($fp, 512)**. PHP evaluates this as “**fread()** 512 bytes from the file at **$fp**, and assign it to **$data**”. If there are no more bytes to read from **$fp**, **$data** will be set to false, and the loop will end. Otherwise we call the main parsing function for **XML: xml_parse()**. **xml_parse()** takes three parameters: the XML parser to use, the data to send to be parsed, and a boolean value, “is this the last of the data?”. The first parameter for us is filled by the return value of **xml_parser_create()**, and the second parameter was just set in the criteria for the **while** loop. The final (optional) parameter is fulfilled in the code by a function call, **feof()**, which takes one parameter itself (a file handle), and returns true if the file handle pointer is at the end of the file, or false if otherwise – perfect for the third parameter of **xml_parse()**!

If **xml_parse** returns **false** (it failed to parse the data), we set **$parsesuccess** to false and break out of the **while** loop. Once the loop completes, we **fclose()** our file pointer, and check the value of **$parsesuccess**. If our parsing failed, we print an error message using values from the XML parser. **xml_get_error_code()** gets the last error number from your parser, and **xml_error_string** converts that code into something meaningful. **xml_get_current_line_number()** returns the line number of the current line being parsed, which, if there was an error, was the line on which the error occurred.

Errors with XML

And how to avoid them

XML might seem easy to those coming from HTML, but don't rush in too quickly – it's a little more complicated than HTML. These simple guidelines should give you the basic information you need to be able to use XML easily with PHP.

An XML document can be said to be either well-formed, valid, both, or neither. A valid document is automatically well-formed, but a well-formed document is not necessarily valid. The definition of valid means that the XML document matches a Document Type Definition (DTD) – a document which defines how the XML should be grammatically. “Well-formed” means that a doc fits the basic XML syntax rules. The XML syntax rules are surprisingly simple, but are key to being able to write XML hassle-free:

- **All elements must have a closing tag. <title> must have </title> somewhere after it. You can also use <title/> if you do not want to have any character data in there – the slash in the opening element auto-closes the element.**
- **XML elements are case sensitive. <body> must have a matching </body> – </BODY> will not do.**
- **XML elements must be properly nested. <item><title>foo</title><item> is acceptable. <item><title>foo</item></title> is not.**
- **All XML documents must have a root tag. This is the element which contains all others. In our example, we used <news>**
- **Attributes must always have quotes. This means <item type=fun> is right out :)**

If you follow those, you shouldn't experience any problems at all using XML. Remember that PHP allows you to output error string reported by *Expat* – this is normally very helpful, especially when combined with the line number.

So, now we have all we need to create a full XML parsing script – you will find a complete example on your coverdisc.

You will notice we reference a file '/wom/bat/test.xml'. Just in case your XML is rusty, here's a simple XML example file. Save it somewhere on your computer that can be accessed by PHP, and point the **$file** variable to it to have your script parse it.

```
<?xml version=”1.0”?>
<news>
<item type=”fun”>
  <title>Monkeys type Shakespeare</title>
  <body>Blah blah blah</body>
</item>
<item type=”science”>
  <title>Man walks on Pluto</title>
  <body>Foo foo foo</body>
</item>
</news>
```

Now that we have an XML file and a PHP script to parse it, load the script into your web browser – you won't see much (yet), but it should at least say “Parsing completed successfully!”

Working with XML data

At this point, our script can read and parse XML. However, we are not yet able to do anything useful with the data PHP finds because the callback functions are empty! Let's expand these handlers to make some sense of the XML file being read. As you can see, **startElement** takes three parameters – **$parser**, **$name**, and **$attrs**. **$parser** sends us a reference to the XML parser that found the element. Using that we can call other functions, like **xml_get_current_line_number()**. **$name** is the element name as read in by *Expat* – note that this may be modified if you have case folding enabled. In our above example, this will be strings like **item**, **title**, and **body**. Finally, we get an array of the attributes in the element stored in **$attrs**, where the array keys are attribute

»

names and the array values are the attribute values. Let's take a look at a better **startElement** handler to add some formatting (note: this assumes case folding is turned **on** – the default setting):

```
function startElement($parser, $name, $attrs) {
  $lineno = xml_get_current_line_number($parser);
  switch ($name) {
    case "ITEM":
      echo "News item found on line $lineno - type is
'{$attrs['TYPE']}'<BR>";
      break;

    case "TITLE":
      echo "News title found on line $lineno<BR>";
      break;

    case "BODY":
      echo "News body found on line $lineno<BR>";
      break;
  }
}
```

Improving the **endElement** function is almost the same, with the difference that ending elements don't have any attributes, and so there is no **$attrs** parameter.

Finally we come to the **charData** function. This works much the same as in **endElement**, except that it takes a **$data** parameter filled with character data. XML is often very specific about character data, so you may find your **$data** variable includes linebreaks or other whitespace at the beginning and/or end – use **trim**() to remove this kind of thing. A better **charData()** function might be this:

```
function charData($parser, $data) {
  $lineno = xml_get_current_line_number($parser);
  $data = trim($data);
  echo "Character data found on line $lineno. Data found was
 $data<BR>";
}
```

Take a look at the output from the script – if you see a lot of whitespace, try adding the following line to the **charData()** function, just after the **trim()** line:

```
if ($data == '') return;
```

Moving on to XSLT

Raw data should be stored in XML – no layout or styling info should be included, just pure information. The layout of your pages should be written in HTML or preferably XHTML – this allows you all the design possiblities modern web browsers offer. Your styling (colours, fonts, etc) should be written using CSS, linked to from your HTML/XHTML. Finally, we have one other technology: eXtensible Stylesheet Language Transformations, or just XSLT. XSLT is used to transform XML data into HTML or other formats for reading by clients. I say "or other formats" because with XSLT you can transform your XML data however you want – you can make it HTML, XML, or even WML (for WAP devices). With the aid of the *Sablotron* library, PHP has full support for XSLT processing. Simply tell PHP what XML file to read from, what XSLT stylesheet to apply to that file, and PHP will generate the appropriate output. XSLT works simply by pattern matching elements you choose, and allowing you to read values from matching markup. For the purpose of this article, we'll be parsing a popular XML file format called RSS, which stands for RDF Site Summary; RDF itself stands for Resource Description Framework. RSS is used most commonly to describe news headlines on sites – a site putting an RSS file online is offering people the chance to download that RSS file and syndicate their news. This is what we're going to try to do here, except, just to keep focused on the matter at hand, we'll be using an offline RSS file. First, grab the RSS. I'll be using the headlines from Linux Weekly News, available from www.lwn.net/headlines/rss, but you're welcome (encouraged!) to try other sources – remember that you may need to view the source code to get the correct content. Here's an (edited for length) copy of what was there at the time of writing:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<!DOCTYPE rss PUBLIC "-//Netscape Communications//DTD
 RSS 0.91//EN"
"http://my.netscape.com/publish/formats/rss-0.91.dtd">
<rss version="0.91">
<channel>
    <title>LWN.net</title>
    <link>http://lwn.net/</link>
    <description>Complete and concise coverage of events in
 the Linux
community</description>
    <language>en-us</language>

    <item>
      <title>LWN.net Weekly Edition for September 12,
 2002</title>
      <link>http://lwn.net/Articles/9185/</link>
    </item>
    <item>
      <title>A couple of KDE security advisories</title>
      <link>http://lwn.net/Articles/9572/</link>
    </item>
    <item>
      <title>Stable kernel prepatch 2.4.20-pre6</title>
      <link>http://lwn.net/Articles/9533/</link>
    </item>
</channel>
</rss>
```

As you can see, it's just XML – nothing particularly fancy. The **<channel>** element contains what we're looking for – the **<title>** of the site we're syndicating from, the **<link>** to that site, etc. Most importantly, it contains news **<item>**s, each with the title of the headline, and a URL to the full story. Save a copy of the LWN RSS file somewhere you can access it using PHP. Now we have XML, let's look at the absolute simplest XSL required to transform the headlines into something legible. As I said earlier, PHP combines the XSL with the XML to produce HTML which is then

Learning more about XML and XSL

Useful links

If you want to learn more about XML, you have a good selection of books and websites to take a look through. Perhaps the most pertinent book, just recently released, is *XML and PHP*, by Vikram Vaswani (New Riders). I have only had time to flick through the book so far, but it certainly seems to cover a great deal. Learn more at www.xmlphp.com

For those of you feeling more adventurous who are interested in a hard-core XML book, I recommend *Mastering XML*, by Navarro, White, and Burman (Sybex). To be honest, *I* still haven't managed to get all the way through this monster – it's very complete, taking you from beginners' XML through XSLT, and on to various examples of popular XML languages. Highly recommended, and the RRP is just £30. If you're not willing to fork out some cash for a book, look no further than the XML FAQ, kept online at www.ucc.ie/xml – the FAQ covers a lot of what you've read here, but often in more detail and at a generally slower pace. You might find it of help if you didn't quite understand a particular concept presented here.

Configuring PHP for XML

Compiling PHP modules

If you missed issue 30 where I gave out my list of recommended configuration options for PHP, you may not have XML/XSLT support compiled into your PHP module. If you find the code example above or on your coverdisc don't work, try re-configuring PHP with the following extra options:

```
--with-xml --enable-xslt --with xslt-sablot
```

Before re-configuring, be sure that the *Expat* library is installed somewhere your compiler can find it. If you compile PHP as a module for a fairly recent release of *Apache*, PHP will automatically use the bundled *Expat* library from *Apache*.

Also, you will need to have the *Sablotron* library installed somewhere your compiler can find it. Visit www.gingerall.com to download the *Sablotron* library *before* you try to re-configure PHP.

sent to the browser, so you will notice there is HTML inside the XSL. Save this file somewhere PHP can get it, as with the RSS file.

```
<?xml version="1.0" encoding="utf-8" ?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns="http://my.netscape.com/rdf/simple/0.9/">
<xsl:output method="html" indent="no" encoding="utf-8"/>

<xsl:template match="/">
  <html>
    <head>
      <title>XSLT</title>
    </head>
    <body>

    <xsl:for-each select="/rss/channel/item">
      News Item: <xsl:value-of select="title"/><BR/>
    </xsl:for-each>

    </body>
  </html>
</xsl:template>
</xsl:stylesheet>
```

Line one defines the document as being XML – you might not have been expecting to see it inside XSL. Well, luckily for us, XSL is an XML language itself – that is, it follows all the same syntax rules as XML, which it very easy to pick up once used to XML itself. You can spot XSL-specific markup because it starts with **xsl:**. Note that because XSL is XML, and your XSL may contain HTML, and HTML is *not* XML-compliant, you technically need to write your HTML as XHTML. Sadly, I don't have room to explain XHTML to you, but, putting it very simply, just write your HTML while following the XML syntax rules given in the box *Errors with XML*. One line two, we define the start of the XSL stylesheet and also say where copies of XML grammar definitions can be download. **<xsl:output>** lets you define how you want your XSL-processed results formatted – there's a list of all the options and accompany explanations at **www.w3.org/TR/xslt#output**. Template rules, defined using **<xsl:template>**, identify XML nodes using the **match** attribute, and, if they find a matching element, the contents of the template will be parsed. In our template, we match **/**, which matches the root node of an XML document. If *Sablotron* finds a root element, it processes the contents of the template (**<html><head>**, etc), otherwise it jumps to the corresponding **</xsl:template>** and continues from there. XSL pattern matching certainly isn't regex, but it is very powerful for its purpose. *E.g.*, **section[@class="appendix"]//content** matches any **<content>** element with a **<section>** ancestor element that has a class attribute with the value appendix. When *Sablotron* matches a root node, it continues parsing, and that means reading **<html><head>.....<body>**. As these aren't XSL instructions, *Sablotron* considers it output. **<xsl:for-each>** is just like the **foreach()** construct in PHP or Perl, which is basically just a **for** loop for arrays. If you consider the news **<item>**s in our RSS file as an array, **for-each** cycles through the array of these elements. You pick the XML elements that you want to iterate through using the **select** attribute of **for-each**. In the example, we use **/rss/channel/item**, which translates to **<item> elements with the parent <channel> and the grandparent <rss>**. Each tag picked up that matches the select attribute is made available for further processing, using the instructions between the **<xsl:for-each>** and the corresponding **</xsl:for-each>**. In our example, every time an /rss/channel/item is found, *Sablotron* will parse the line starting "News Item". This includes yet another piece of XSL, **<xsl:value-of>**, which we use to extract a text value – we select the element **title**, which is the **<title>** child element of the current /rss/channel/item in our **for-each** loop. If you look back at the RSS above, you will see we have three **<item>**s, each with their own **<title>** child element. Once we match the root element using our template, we **for-each** through each of these **<item>**s, picking out its **<title>** as we go. Once all the **/rss/channel/item** elements have been iterated through, processing continues at **</body>**, and the remainder of the stylesheet is just closing tags.

Now we can see the code to combine our RSS with our XSL.

```
<?php
$xsltproc = xslt_create();
echo xslt_process($xsltproc, '/path/to/lwn.rss', '/path/to/thexsl.xsl');
xslt_free($xsltproc)
?>
```

Save that file in your PHP scripts directory, and load it up into your web browser. You should see the list of LWN headlines shown in your browser, parsed neatly. You will probably want to alter the XSL to add more HTML formatting, but that's the easy part. For now, try changing the RSS file to another one: your output will change accordingly, with no reprogramming necessary. I'll leave it as an exercise to you to change your XSL to handle the channel description info. Here's a hint, though – you can have as many **<xsl:template>**s as you need, each matching various parts of a doc.

The PHP code itself is only three lines, which shows how easy PHP makes XSL transformations. **xslt_create()** and **xslt_free()** work just like **xml_parser_create()** and **xml_parser_free()** respectively, except they return an XSL parser as opposed to an XML parser. The only other line in there is **xslt_process**, which takes a minimum of three parameters – the XSLT parser to use (returned by **xslt_create()**), the XML input file, and the XSL input file. **xslt_process()** sends the processed data back as its return value. As with XML parsing, PHP gives you two functions, **xslt_error()** and **xslt_errno()**, that allow you to echo out info about any errors that occurred during **xslt_process()**.

Conclusion

If you were new to XML before reading this article, you should now be able to see a little of how useful it can be, especially so on the Internet. PHP does a great job of seamlessly integrating both XML and XSL into the language, and, with our XSL example being just three lines long, it couldn't be much easier! LXF

About Paul Hudson

Paul Hudson is a London-based web developer specialising in PHP and Perl. He can be emailed at hudzilla@php.net

NEXT MONTH

Next month we're going to explore the possibilities of using PHP to automate your site design through the use of templates. We'll be taking a look at how to create your own template system, and also what off-the-shelf solutions have already been made. If there is something you want to see covered here, drop us an email and we will try our best to incorporate it into a future instalment.

Answers

If you are really stuck and the HOWTOs yield no good result, why not write in? Our resident experts will answer even your most complicated problems!

Experts this month

Whatever your question is, we can find an expert to answer it – from installation and modem woes to network administrations, we can find the answer for you – just fire off a letter or email and it'll all be taken care of.

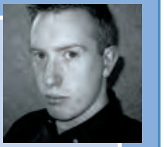

LXF answers guy **David Coulson** is a networking and security guru with plenty of sysadmin experience to boot.

NIck Veitch is the editor of the magazine, and answers your easy questions! Or indeed anything to do with *Grub, Lilo, netatalk, vi*...

3D woes

Q First of all, your mag is great. A little thin, though fun/educating. I have used Linux on and off since 1992. I first used Slackware 1.0 beta and then tried all the different distros. I now use a combo of Mandrake 8.2, Slackware 8.1 and Sorcerer. I love Linux for all it is, but I can't seem to shake that WinXP installation. There's still one more frontier as far as I'm concerned, and that is 3D.

I have a Hercules 3D card based on the nVIDIA TNT2 chipset and with a good 32MB of RAM. My Win32 system uses that smoothly and let's me run the newest games in a fair resolution. I recently tried the *Wolfenstein* binaries from idsoftware: full of anticipation I started playing, but to my big disappointment the game was bleaker than a pale hand of ice.

No fog, no lighting, no deep 32 bit textures, no speed, no picturesque resolutions, no FUN!

I run the *3.3.6 XFree86* with the "3D-support". I know my card isn't 100% configured as that XFree86 version is too old. But I run with an identical chipset to the one I support, so it can't be all driver-issues I am facing.

The *3dmesa* libraries are either very incomplete or they have a setup-config-function which I am totally oblivious of. Can you please do a tutorial/piece on 3D graphics at the hardware/driver-level. I regard myself as a pretty adequate Linux user, but this stuff has really left my head twisting... But then again, that may be a direct consequence of the lack of sleep.
Sverre Eldøy

A You may wish to check, using **glxinfo**, that *OpenGL* is enabled correctly on your *XFree86* display. Should *OpenGL* not be working, you may want to just upgrade to *XFree86 4.1.0* and setup the **glx** extension and try again, as *XFree86 3.3.6* is rather old and *Mesa* has improved significantly since then.

It does sound quite like *Wolfenstein* decided that it couldn't use *Mesa*, isn't working happily on your system, and turned many graphics capabilities off. *Mesa* does not have quite as many features as 3D libraries on Windows, but it is better than it sounds like you are experiencing.

Xine libs

Q I've installed Red Hat 7.3.. but *Xine* can't play DVDs on my laptop (TOSHIBA 1800 series Celeron 1100).I thought something wrong with my Red Hat Linux..so I uninstalled Red Hat Linux 7.3 and then I decided to install Mandrake Linux 8.2.. but I still have some problem as Red Hat Linux can't play my DVD.. (I've even already downloaded *libdvdcss, libdvdread, libdvdnav*).

And I tried to open *Xine* from terminal command line ..this is what I've got:

```
This is xine (X11 gui) - a free video
 player v0.9.8
(c) 2000, 2001 by G. Bartsch and
 the xine project team.
Built with xine library 0.9.8 [Thu 14
 Mar 2002 10:35:25]-[gcc version
 2.96 20000731 (Red Hat Linux 7.2
 2.96-107)]-[Linux 2.4.9-21smp i686].
Found xine library version: 0.9.8
 (0.9.8).
Display is not using Xinerama.
main: probing <Xv> video output
 plugin
video_out_xv: Xv extension is
 present but I couldn't
find a usable yuv12 port.
              Looks like your graphics
 hardware driver doesn't support Xv?!
load_plugins: video output plugin
/usr/lib/xine/
plugins/xineplug_vo_out_xv.so: i
nit_video_out_plugin failed.
main: probing <SyncFB> video
 output plugin
video_out_syncfb: aborting. (unable
 to open device "/dev/syncfb")
load_plugins: video output plugin
/usr/lib/xine/
plugins/xineplug_vo_out_syncfb.s
o: init_video_out_plugin failed.
```

The Linux *OpenGL* library, *Mesa*, requires support from *XFree86* in order to provide the best performance.

```
main: probing <XShm> video output
 plugin
video_out_xshm: video mode depth
 is 16 (16 bpp),
TrueColor, not swapped,
       red: 0000f800, green:
 000007e0, blue: 0000001f
yuv2rgb: using MMXEXT for
 colorspace transform
load_plugins: video output plugin
 XShm successfully
loaded.
main: probing audio drivers...
main: trying to autoload 'oss' audio
 driver:
audio_oss_out: Opening audio
 device
audio_oss_out: using device
 >/dev/dsp<
audio_oss_out: using
 SNDCTL_DSP_GETODELAY
Segmentation fault
```

Can you please help me out..
Mirza Kurima

A segmentation fault error means that there was a bug in either the program, or a library, which you are running. It is possible to debug a segmentation fault, using *strace* and *gdb*, although you will need to build the libraries and *Xine* from source, with debugging symbols.

Alternatively, you could try to use a different DVD player, such as *ogle* or *mplayer*, or download the Red Hat RPMs of the libraries you are using with xine, along with the RPMs of xine. If you use Mandrake RPMs with Red Hat, or Red Hat RPMs with Mandrake, you may have problems.

Video adaptor

Q I'm trying to play vcd/dvd via *MPlayer* and *Xine*.

I have *MPlayer* working, but have to use -vo X11 instead of xv (which I'm told is better?).

I did an xvinfo and got the following result:

```
markr@krynn:~/.xine> xvinfo
X-Video Extension version 2.2
screen #0
no adaptors present
```

I'm having difficulty getting any information on why it thinks I have no adaptors, does anyone have any advise or links to more information where I can read how to get an adaptor configured/installed?
Mark, from the LXF Forums

A *xvinfo* is not saying that you have no video adaptors, only that you have no X Video adaptors. An XVideo adaptor is independent of your video card, although it does require the video card to support it. You've not said which video card you have, nor which version of *XFree86* you are running, so you may want to look at upgrading your release of *XFree86* if www.xfree86.org suggests that your video card will function as an XVideo adaptor. If your video card is not supported, then there isn't too much you can do, other than swap it out for something else.

Linux DVD players, including *Xine*, give much better performance when using the XVideo extension, but not all video cards are supported.

MS Bootstrap

Q I found in your *Linux Format* July 2002 a question about installing Mandrake 8.2 over Mandrake 7.2 on Sony Vaio, signed by Colin Thomas. I have had similar problem on a normal PC machine.

I am not Linux guru but I solved problem by installing Windows again and Linux new version after that. Everything worked fine.

I think that this may help Colin Thomas in solving problem.
Vedran Vucic

A It's interesting that installing Windows helped Linux to work. It's possible that the drivers in Windows gave certain components in the system the desired kick to help get them up and running in Linux. Of course, this is quite a rare occurrence,

»

A QUICK REFERENCE TO: SSH

***SSH* is a secure alternative to *telnet*, although it does have a great many more features, which makes it particularly useful. From a usage point of view, *SSH* is more like *rsh* than *telnet*, although as it supports a variety of authentication methods, and can do nifty things like copying files from one host to another.**

***SSH* requires a *SSH* server on the machine you are *sshing* to, as well as a client. The most common *SSH* server is *OpenSSH*, although there are a great number of *SSH* clients available for various operating systems. For Windows users there is *puTTY* and *SecureCRT*, and for platforms without specific *SSH* clients, there is the Java-based *Mindterm* client.**

However, for Linux users, the 'ssh' client from *OpenSSH* is a great choice, as it supports all of the *SSH* protocols and is freely distributable and usable. The original *SSH* distribution is now non-free for commercial users, so is best avoided if you think that you may be breaching their licence agreements.

The simplest thing to do with *SSH* is to shell into a server, which is done with:

```
ssh host.domain.tld
```

If we don't supply a username, it will use the username of the user which executed ssh. We can supply a username, either with a -l switch, or by prefixing the hostname with user@

```
ssh user@host.domain.tld
ssh host.domain.tld -l user
```

To copy files from one host to another, we use a secure version of *cp*, known as *scp*. If we wanted to copy 'file' from here, to a remote machine, we just do:

```
scp file host.domain.tld:
```

which will put 'file' into our home directory on the remote system, and call it 'file'. We can also supply an alternative directory and filename, just like regular 'cp'

```
scp docs/file.txt
 host.domain.tld:/tmp/ourfile.txt
```

***scp* will equally work the other way around, so we can copy files from remote machines to our local system.**

```
scp me@host.domain.tld:log.txt
 ~/docs/changelog.txt
```

Each ssh command, including *scp*, will request your password, which can become quite annoying. Instead, we can authenticate using a RSA or DSA key. Creating a key is done with ssh-keygen, and to create a DSA key pair we just do:

```
ssh-keygen -d
```

When it asks for a passphrase, we can simply press enter twice, and it will not ask for it. If security is paramount, one can supply a passphrase, and use ssh-agent to request it once and it will then authenticate without the passphrase for the rest of the session.

Once we have a DSA key pair, we need to copy the public key from our .ssh directory, to our remote host:

```
scp ~/.ssh/id_dsa.pub
 host.domain.tld:
```

On the remote host, we need to add the contents of id_dsa.pub to ~/.ssh/authorised_keys2:

```
cat id_dsa.pub >>
 ~/.ssh/authorised_keys2
```

When we next *ssh* into the system, it will not ask us for a password, which makes the use of *scp* far more transparent.

FREQUENTLY ASKED QUESTIONS: STORAGE

FAQ I installed Linux, and it asked me to make lots of different partitions. Can I just install it in one?

Linux can be installed into as few, or as many, partitions you desire. Usually, one has at least two partitions, one being for regular storage, and another for swap space, as swapping to a partition is nicer than swapping to a file on a filesystem.

Having more than one partition allows you to distribute your Linux installation over more than one disk, or allow you to reinstall without destroying the contents of /home. It can also be very useful if you are trying out distros, as you could share /home between more than one.

FAQ Which partitions should I create?

The most common directories which have their own partitions are /, /home, /var, /usr and /boot. /home stores user specific files, so isn't usually destroyed when you reinstall. /usr contains all of the programs and libs you use, and is always distribution specific. While one could use the / partition for /usr, it's usually more desirable to have a small / partition, and a large /usr partition. /var stores log and spool files, so the filesystem there is generally optimised to store a large number of small files, where as most others will be storing a small number of much larger files. /boot is only needed if your system has issues booting from beyond 1024 cylinders, and in such a case, is created at the very beginning of the disk, so that it does not go beyond the 1024 cylinder boundary.

FAQ Can I resize partitions?

The *parted* utility allows partitions to be resized, and you can then resize the filesystem using a utility such as *resize2fs*. However, if you think you will be resizing partitions frequently, or needing to create new ones, using the *Logical Volume Manager*, *LVM*, is a much better choice, as it permits the creation of new Logical Volumes, which are similar to partitions, without having to reboot or restart anything.

FAQ Can I use LVM rather than regular partitions?

Indeed, although your kernel will need to know how to deal with *LVM* when you first boot up, so keeping / as a regular partition is often a wise choice. Having the rest of the filesystems as LVM logical volumes allows much greater flexibility, and rather than having to limit your partitions to the size of the disk, it will allow you to span more than one disk with a single logical volume, making much better use of the storage space available.

FAQ What about RAID? Can I setup an array with lots of disks?

Linux supports hardware RAID using many controllers, from companies including Intel, Mylex and 3ware. Alternativly, if you want a cheaper option, one can use software RAID, and create an array which requires the kernel to handle the synchronisation. While this is slower than using hardware RAID, if you have a lightly loaded system, but require redundancy, then it is a good choice. For a simple RAID-1 mirror, having two disks mirrored using software RAID can improve read speed and, of course, gives great redundancy should one of the drives bite the dust.

FAQ Does Linux support journaled filesystems?

There is a wide range of journaled filesystems for Linux, including *ext3*, *ReiserFS*, *JFS* and *XFS*. Both *XFS* and

« and there is generally a method by which one can enable hardware without resorting to installing Windows.

You didn't say what your specific problems were, although it would be intriguing to hear of others who have found that installing Windows, or indeed any other OS, has helped get Linux up and running properly.

Fluxconf

Q I use your version of *Fluxbox* most of the time but cannot get *FluxConf* to install on SuSE 7.3. It gives an error wanting <gtk-config> or <pkg-config conf> and suggests looking at www.gtk.org. However a search of this site returns zero find. I also loaded *GTK+* from the same disc (Development) but this did not help; perhaps I put it in the wrong place. Suggestions?
Prof. Joe Lamb

A *gtk-config* is found in the *gtk+-devel* package, which also includes the headers required to compile *GTK+* apps. You did not say if you used RPMs or sources to install *GTK+*, but unless you are familiar with how *GCC* locates headers, and the use of /etc/ld.so.conf, then you will want to install the SuSE distribution of *GTK+*, and the appropriate development package, after removing the *GTK+* package you installed earlier.

If you installed the source distribution of *GTK+*, you will find *gtk-config* within /usr/local/bin, which may or may not be in your **$PATH** variable.

Install prefixes

Q I'm trying to install *Ogle* from *LFX31D* and I am having some problems. I installed *a52dec-.7.3* and *ogle-0.8.4* and they seemed to configure, make and install OK. When I tried to configure *ogle_gui-0.8.4* I get a config error saying "You must install ogle and ogle_gui in the same prefix". The INSTALL text suggested the most common thing to do was to add --with-dvdcontrol=prefix to the ./configure command. Their sample was /usr/local for the prefix. I got the same error when I tried this. What is a prefix and what prefix does *ogle* use by default?

I'm running Xandros beta 3 and it is an incredible release. I can't wait to see your review of this product. I love *Linux Format* and look forward to each issue. Thanks in advance
Art Atkinson

A The prefix option sets the directory into which the files will be installed when you do a **make install**. Usually, /usr/local is used when you compile programs from the source code, so as not to get muddled up with files installed as packages in /usr. The general default for ./configure is /usr/local, although it can differ depending upon what you are installing. You may want to try installing *ogle* with;

```
./configure —prefix=/usr/local
```

to ensure that it installs into /usr/local.

Configure scripts log to 'config.log', so checking that file should indicate exactly what the problem is, and why *ogle_gui* does not like the prefix you are supplying. If all else fails, checking the ogle site, or contacting their mailing list may be your best option, as if you are supplying the same prefixes, or none at all, to both *ogle* and *ogle_gui*, then it is likely a problem with *ogle*, rather than something you are doing.

PCMCIA

Q I am running SuSE 8.0 on a Toshiba S1640CDT laptop, and I am trying to get my PCMCIA ethernet card running. I'd

JFS come from large companies, which have ported the filesystem from their own OS, so it is a more tried and tested system, even if not on Linux. *ext3* is simply a journaled version of *ext2*, so if you should need to recover your system and don't have an *ext3* kernel available, one can simply mount it *ext2* and fix it, then remount it *ext3* later. *ReiserFS* is a filesystem written for Linux from the ground up, and is very popular with people with large arrays, storing hundreds of thousands of files, as it handles lots of small files much better than *ext3*.

FAQ Are there any other filesystems which can improve my system?

tmpfs is particularly useful, as it is a dynamic version of *ramfs*. Rather than mounting /tmp as a regular filesystem, if we mount it with tmpfs, our system will allocate RAM for it on an as-needed basis. If /tmp doesn't have anything in it, it won't use any memory, so we're not wasting tens of MB of expensive system RAM on nothing. As *tmpfs* uses RAM for storage, it is significantly quicker than a disk-based filesystem, making it ideal for temporary file storage, although when the system is rebooted, the contents of /tmp will disappear.

be grateful for any help you can give.

The card is a Linksys etherfast 10/100 PCMPC200 v.2. The weird thing is, I bought another card, a sitecom cardbus 10/100, which completely refused to work with the 2.4 kernel, although I tweaked everything in the whole world.

After finally giving up on it, I installed the linksys, which worked immediately as eth0. One ifconfig -a later, I had a complete home LAN.

However, after a reinstall (don't ask!) I cannot get the card recognised. No beeps, nothing. I have attached:

```
/etc/pcmcia/config
/etc/pcmcia/config.add_kernel
/etc/pcmcia/network
/etc/pcmcia/network.opts
/etc/pcmcia/config.opts
/etc/pcmcia/config.add_kernel.opts
```

and printouts from *cardctl*:

```
toshiba:~/Desktop # cardctl ident
Socket 0:
  no product info available
Socket 1:
  no product info available
  manfid: 0x13d1, 0xff02
toshiba:~/Desktop # cardctl status
Socket 0:
  no card
Socket 1:
  3.3V CardBus card
  function 0: [ready]
Socket 0:
  not configured
Socket 1:
  Vcc 3.3V  Vpp1 3.3V  Vpp2 3.3V
  interface type is "cardbus"
  irq 11 [exclusive] [level]
  function 0:
```

***cardmgr* gives a null response unless I look at ps ax, following a restart of PCMCIA, which gives:**

```
 2238 ?    S      0:00
/sbin/cardmgr -k -m
 /lib/modules/2.4.18-4GB/ -n pcmcia
```

ifconfig -a gives the following (I assume sit0 is the firewall. It doesn't seem to be the card, as allocating a network address does not allow connection to the network):

```
toshiba:~/Desktop # ifconfig -a
lo       Link encap:Local Loopback
          inet addr:127.0.0.1
 Mask:255.0.0.0
          inet6 addr: ::1/128
 Scope:Host
          UP LOOPBACK RUNNING
 MTU:16436  Metric:1
          RX packets:94 errors:0
 dropped:0 overruns:0 frame:0
          TX packets:94 errors:0
 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:7048 (6.8 Kb)  TX
 bytes:7048 (6.8 Kb)
sit0      Link encap:IPv6-in-IPv4
          NOARP  MTU:1480  Metric:1
          RX packets:0 errors:0
 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0
 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:0 (0.0 b)  TX
 bytes:0 (0.0 b)
```

Please help, as having spent 100s of hours trying to learn Linux, & tons of money buying the distros, books and compatible hardware, I don't want to be forced to drop Linux at this stage. I know the thing works, but I can't reproduce how it did so.
Paul Hawkes

A When you insert a PCMCIA card, the system should make a couple of beeps. If the last beep is high pitched , then the card will generally have installed correctly. If it is low pitched, then there is a problem. If you get no beeps at all, then it would suggest that something isn't very happy with your PCMCIA set up.

Many 2.4 users find they have more success using the *pcmcia-cs* distro of modules, rather than the kernel PCMCIA system. You did not indicate which you were using, either now, or before you reinstalled, so if you are currently using the kernel modules, you might want to remove them and build the *pcmcia-cs* package from source against your current kernel.

By looking at /var/log/messages, or /var/log/syslog, you can spot what *cardmgr* is doing and which module it is trying to load when you insert the card. The sit0 interface is used to create IPv6 tunnels over existing IPv4 networks, so isn't anything to do with your PCMCIA config, or its networking.

DNS resolution

Q **I need to set up my Linux box to look at another IP on my network as its name server, I can do this in SuSE using *YAST*, however I would rather know which file this is being written to, so that I can edit it myself, anyone know where this is stored? Thanks**
Paul Norris

A The file you are looking for is /etc/resolv.conf, and you'll need to edit the **nameserver** lines within this file. As always, the man pages are useful, so **man 5 resolv.conf** will give you all the information you need.

External monitor

Q **I have Mandrake 8.2 with Kernel 2.4.18 and KDE 2.x on my IBM A30p laptop with a Gericom CT1997PF monitor.**

I cannot find a way of telling Mandrake to run my monitor at the correct resolution. The laptop has a 1600x1200 TFT screen and X runs fine. However, when I boot with my external monitor attached, I get 640x480 (or possibly 800x600).

I've run *XFSetup*, and Mandrake's own resolution programs, but the screen resolution stays the same. Mandrake's utility tells me to run KDEINIT: KWIN after I change the resolution, but then before I can do anything, it reboots the PC!. Why and how do I run KDEINIT?

I tried inserting a line manually in XF86Config as mentioned in a previous issue, but no change. I have even deleted XF86Config, and X still runs at the same low res.

It appears that X does not know about my CRT monitor, and does not have a name for it. If I look in XF86Config, there is only one monitor section, that for my TFT screen. How do I find out the video display device names in order for me to manually insert a Monitor section into XF86Config to enable the high resolutions?

I know the monitor works at these resolutions as I can boot up at 1600x1200 without the monitor attached, and then connect it while X is running. Then the monitor works at 1600x1200!

I am a new user of Linux so I know very little. I tried to fathom out the *insmod*-type utils, but I became lost and still could not find out the name Linux had assigned to my monitor.

I have spent months on this on a daily basis. Do I need to throw out

»

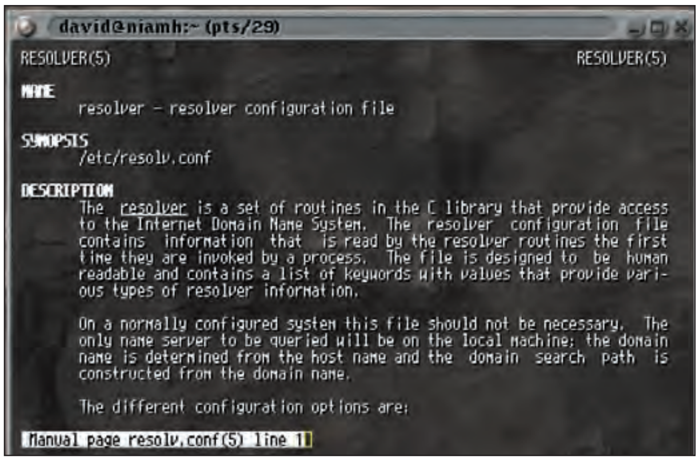

```
RESOLVER(5)                                                    RESOLVER(5)

NAME
       resolver - resolver configuration file

SYNOPSIS
       /etc/resolv.conf

DESCRIPTION
       The  resolver is a set of routines in the C library that provide access
       to the Internet Domain Name System.  The  resolver  configuration  file
       contains  information  that  is read by the resolver routines the first
       time they are invoked by a process.  The file is designed to  be  human
       readable and contains a list of keywords with values that provide vari-
       ous types of resolver information.

       On a normally configured system this file should not be necessary.  The
       only name server to be queried will be on the local machine; the domain
       name is determined from the host name and the  domain  search  path  is
       constructed from the domain name.

       The different configuration options are:

 Manual page resolv.conf(5) line 1
```

The *resolver* library uses /etc/resolv.conf to find out which DNS servers it can query.

« Mandrake and move to another Linux? I hope not at it is otherwise a very easy to use system. PS, what's the difference between XF86Config and XF86Config-4 ?
Owen Widdicombe

Starting with your last question first, XF86Config-4 and XF86Config are simply alternative names for XFree86's config file. XFree86 will check a selection of locations for it's config – it will first look for XF86Config-4, and then for XF86Config in each of these. Usually, you would have either one or the other, although many distros symlink XF86Config-4 to XF86Config. If you're getting confused, make sure you only have one of them.

XFree86 logs to /var/log/XFree86.0.log, so when you start X up with the monitor attached, copy this file out to somewhere safe and look through it to see what the system is complaining about. Assuming that you can't change the resolution using **Ctrl**-**Alt**-Numpad**+**/**-**, then XFree86 probably thinks that the frequencies you have supplied are inappropriate for your monitor, and used what it thought was best. By checking the log file, you should be able to edit XF86Config to stop this happening.

The **kdeinit: kwin** output you are seeing is from KDE when it starts up, and launches the *kwin* window manager. It is not actually telling you to run this, only supplying some log output so you know what is going on.

UDF filesystem

Q I have a computer with SuSE 7.3 installed at home. It has a DVD drive and a separate CDROM R/W unit and everything works well. I could even install SuSE 7.3 from the DVD rather than 7 CDROMs provided.

You recently had an article in *LXF30* August 2002 about the EPIA mini-ITX motherboard which I found most interesting. I was unable to find a local supplier but obtained one via the distributor you helpfully provided. In addition I purchased a suitable case and, because of the diminutive nature of this item, a combined CDROM R/W and DVD read drive. I have assembled everything and overall I am very pleased with the setup and particularly the reduced decibels. I can even hear the hard drive kicking in now.

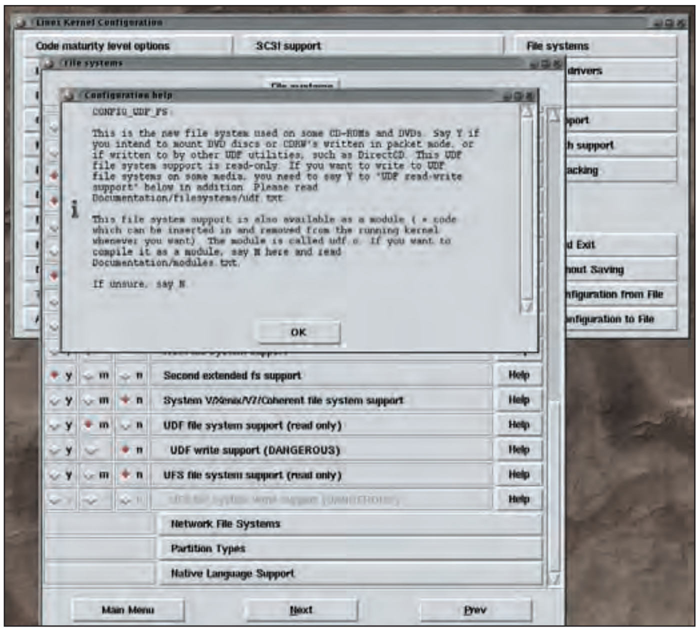

```
CONFIG_UDF_FS

This is the new file system used on some CD-ROMs and DVDs. Say Y if
you intend to mount DVD discs or CDRW's written in packet mode, or
if written to by other UDF utilities, such as DirectCD. This UDF
file system support is read-only. If you want to write to UDF
file systems on some media, you need to say Y to "UDF read-write
support" below in addition. Please read
Documentation/filesystems/udf.txt

This file system support is also available as a module ( = code
which can be inserted in and removed from the running kernel
whenever you want). The module is called udf.o. If you want to
compile it as a module, say N here and read
Documentation/modules.txt.

If unsure, say N
```

```
● y  ○ m  ○ n   Second extended fs support                      Help
○ y  ○ m  ● n   System V/Xenix/V7/Coherent file system support  Help
○ y  ● m  ○ n   UDF file system support (read only)             Help
○ y  ○    ● n   UDF write support (DANGEROUS)                   Help
○ y  ○ m  ● n   UFS file system support (read only)             Help
```

DVDs use the UDF filesystem, but not all kernels have it available.

As you might have guessed I do have a problem which is as follows. I have not been able as yet to read anything on the DVD drive but I have no problem at all with the CD R/W drive. I am using an LG Model RW/DVD GCC-4120B Rev 2.02 which appears in the BIOS setup screens as HL-DT-ST RW/DVD GCC-4120B.

I have been using Linux for quite a while and prefer it to Windows. However I do not always find it easy to resolve difficulties when they occur. There is of course a wealth of docs to hand for Linux, but the weight of this sometimes makes it difficult to know where to start.

Any help or pointers you can throw my way would be much appreciated. I suspect an article, in your excellent mag, on the mounting and functioning of CDROM, DVD and the combined drives would not go amiss.

Keep up the good work with the mag. I think the balance is pretty good for the readership as a whole, although personally I would prefer more articles for beginners or inexperienced users.
Philip Riddle.

A You've not said exactly why you can't read from DVDs, although we suspect it's an issue with filesystems. CDs are generally burnt using the ISO9660 filesystem, whereas DVDs use UDF. Many kernels come pre-compiled with ISO9660 support, although UDF may or may not be included. You can check which filesystems your kernel supports by looking at /proc/filesystems, and seeing if **udf** is listed. If it is not there, you can try to **modprobe udf**, but you may have to recompile your kernel with the appropriate support.

You can then mount a DVD with;

```
mount -t udf /dev/cdrom /mnt/dvd
```

You may wish to setup /mnt/dvd in /etc/fstab, so that instead of using /mnt/cdrom with iso9660, you have a separate mount point which uses udf.

DHCP

Q I've got SuSE 8 Pro running and have eth0 as an external interface to the Net. Eth1 is my internal network serving a couple of PCs. I installed DHCPD and configured the dhcpd.conf file for my network. However when I start *dhcpd* I get a message like below:

```
Internet Software Consortium DHCP
 Server V3.0.1rc6
Copyright 1995-2001 Internet
 Software Consortium.
All rights reserved.
For info, please visit
 http://www.isc.org/products/DHCP
Wrote 0 leases to leases file.
No subnet declaration for start
 (0.0.0.0).
** Ignoring requests on start. If this is
 not what you want, please write a
 subnet declaration in your dhcpd.conf
 file for the network segment to
 which interface start is attached. **
```

I've checked out some docs on the 'Net but can't find anything where a server has two interfaces. I tried it inside the network with a Linux PC with two NICs and it works fine – it seems to default to eth0 though.

Does anyone know what the command is to make it go to eth1?
From the LXF Forums.

A You will need to tell the *dhcpd* daemon to only bind to eth1, and not eth0. *dhcp* can take a list of interfaces as a command line option, so you will need to look in /etc/init.d/dhcpd and edit it appropriately, so that rather than running;

```
dhcpd
```

it runs

```
dhcpd eth1
```

Many distributions have a **DHCPD_INTERFACE** variable in this file, so you may be able to simply include 'eth1' in there, rather than editing the rest of the *init.d* script.

Your error would also suggest that you have one of your ethernet devices up, but without an IP address. Without more configuration information, it's impossible to say whether or not this in error, but if eth1 does not have an IP address in the correct range, then *dhcpd* will not function correctly. 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 complied 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).

MAGAZINE FEATURING:
Sun's move into the Linux server market – with LX50 review, The Liberty Alliance, Systems programming, using OpenOffice.org, Homebase

November 2002

Product code:
LXFB0033(cd)
LXFDB0033(dvd)

DVD HIGHLIGHTS:
Debian 3.0, mFighter, OpenOffice.org (bugfix release), DVDRip, Cyrus IMAP Server

CDs HIGHLIGHTS:
Kylix 3, GNUCash, BXPro, KDevelop, Opera, Vega Strike, Parted, AnjutaIDE, GTransferManager

MAGAZINE FEATURING:
Building better databases, 'Trusted Computing' – beware Palladium, USB 2.0, firewall roundup, Amiga emulation, Gentoo review

October 2002

Product code:
LXFB0032(cd)
LXFDB0032(dvd)

DVD HIGHLIGHTS:
Knoppix, Drip, Squeak, extra FlightGear maps, Ogg Vorbis 1.0, Knoppix

CDs HIGHLIGHTS:
FlightGear (runs from disc), Aglaophone, UAE, Clam Antivirus, Perl 5.8, Quanta Plus, Netclipboard, Mah-Jong, HTML-Mason, WebSuck, Epsutil

MAGAZINE FEATURING:
Linux goes to Hollywood, HTML Editors roundup, Internet security special, Ruby scripting language

September 2002

Product code:
LXFB0031(cd)
LXFDB0031(dvd)

DVD HIGHLIGHTS:
Slackware 8.1, Cinelerra, Ogle & Zine, Gnome2, MultiCD, Phobia III, Grip, Zinf, TkVoice

CDs HIGHLIGHTS:
Slackware 8.1, Fluxbox, Kallers, Torcs, SaveMyModem

MAGAZINE FEATURING:
Ultimate Office: Every current office solution on test, plus Quantum computing, Red Hat 7.3, PHP tips, Eden motherboard test

August 2002

Product code:
LXFB0030(cd)
LXFDB0030(dvd)

DVD HIGHLIGHTS:
Intel's C++ & Fortean complilers, CPAN, Gnome2, Gentool Linux, Boson, Mozilla, Normalize, EarCandy

CDs HIGHLIGHTS:
Highlights: Intel's C++ & Fortean compilers, Gnome2, Transcode, Album Cover Grabber, Office Suites

MAGAZINE FEATURING:
Customise your kernel. WineX latest, Inside IPv6, Astronomy applications roundup

July 2002

Product code:
LXFB0029(cd)
LXFDB0029(dvd)

DVD HIGHLIGHTS:
Highlights: OpenOffice.org 1.0, Mozilla 1.0rc3, Netscape 7.0, Opera 6.0, Mac on Linux

CDs HIGHLIGHTS:
Highlights: OpenOffice.org 1.0, Evolution, Omnis Studio 3.01, Clam AntiVirus, Python 2.21

MAGAZINE FEATURING:
Mozilla special issue, Video editing, Networking with Macs, KDE3 uncovered, Crossover Office

June 2002

Product code:
LXFB0028(cd)
LXFDB0028(dvd)

DVD HIGHLIGHTS:
KDE3, Mozilla 1.0, Gentoo, Beehive, PixiePlus, Funk You, NetComics, Fluxbox

CDs HIGHLIGHTS:
Highlights: KDE3, Mozilla 1.0, Linux from Scratch, Tomsrtbt, GKrellm, evoBB

Other back issues also available, please call for further information, giving the issue number and/or date required

To order, call the back issues hotline on 0870 4448645

or +44 1458 271178 for overseas orders

Remember to quote the issue code shown above and have your credit card/switch card details ready.

Prices (inc post and packing): CD – UK£5.99, US & Eur £6.99, rest of world £7.99
DVD – UK £6.99, US&Eur £7.99, rest of world £8.99.

LINUX FORMAT
www.linuxformat.co.uk

Coverdisc

Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format* CD, and this month it's playtime...

On the CD

Wherever you see this logo it means there's related stuff on the CD

Essential info

On page 107 we have grouped together essential info on the different types of packages on your coverdiscs – along with instructions for installing source packages.

Important notice

Before you even put the CD or DVD in your drive, please make sure you read, understand and agree to the following:

The *Linux Format* CD is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker on ANY new software. While every care is taken in the selection, testing and installation of CD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system which may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all important files. Please read individual licences for usage terms.

READ ME FIRST

The nights are drawing in and the weather is getting colder, which traditionally means more time spent computing. To make the loss of summer easier to bear, we have brought you another varied collection of software to play with, with the emphasis very definitely on play with one of this month's highlights. *Unreal Tournament 2003* hasn't been released yet, but the demo is available and on the CDs and DVD. There have been several "live" CD Linux distributions lately. DVD users got the excellent Knoppix a couple of months ago, now CD users needn't feel left out because the first CD this month contains DemoLinux.. Then there is the usual, wide ranging selection of programs. The second CD has everything from small system utilities to major office suites. See the listings, either on these pages or on the discs for a complete breakdown of what's available.

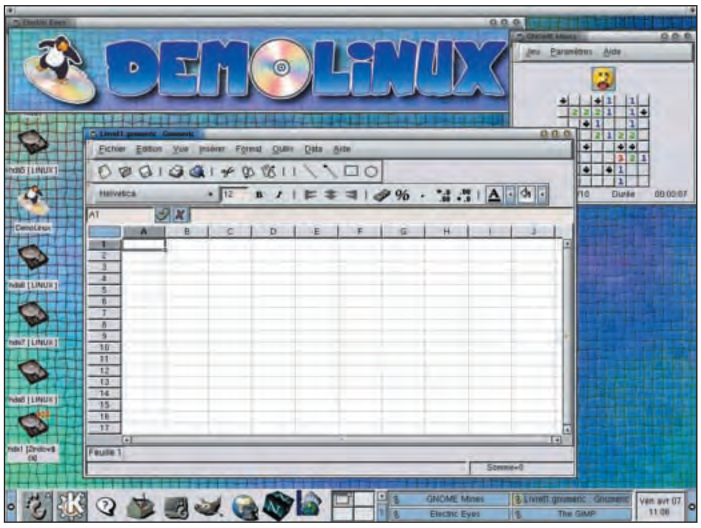

Run Linux direct from the CD. Show your friends or use it as a rescue disc.

Demo Linux

The first CD contains DemoLinux. This is a "live" Linux CD. That is, it boots into a working Linux environment rather than requiring installation to your hard disk. DemoLinux is covered in detail in the Mini Distros *RoundUp* this month, so we'll just take a quick look at getting it running here. Being a 'live' distro, you can run it simply by booting from the CD. If you have difficulties booting from this, or any other CD, see the *Booting CDs* box. Unfortunately, this only applies to CD users. We were unable to transfer the boot process to DVD in a reliable manner, so DVD users get an ISO image that has to be burned to CD-R first. See the instructions at the end of this section for how exactly to do that. Once booted, follow the on-screen prompts and you'll have a fairly substantial Linux system running without touching your hard drive.

Games/UT2003

Unreal Tournament 2003 is a lightning-fast, no-holds-barred sport of the future, where warriors face off like the Gladiators of ancient Rome to determine the ultimate combatant. Super human reflexes, awesome weapons and the skill to use them effectively are what separate the good from the bad, the living from the dead, the victor from vanquished. Well, that's what the website says, but don't take Epic's word for it, try the demo for yourself. Before installing it, make sure you have a suitable graphics card and drivers. This basically means an nVIDIA card with their closed source drivers. See the *nVIDIA only?* box for details on this.

You can install *UT2003* as root or as a user. In either case, installation is simply a matter of typing

```
sh /mnt/cdrom/Games/
UT2003/UT2003-Demo-Linux.sh.bin
```

If you install as root, the game is installed into /usr/local/games/ut2003_demo, with a link to start it at /usr/local/bin/ut2003_demo. If you install as a user it tries to install the game to ~/ut2003_demo/ and create the startup link at ~/ut2003_demo. Naturally, this fails. The solution is to make a minor change to the name of the install directory when the installer asks you where it should be installed. Once installed, you can start it from the KDE or GNOME menu, or by typing:

NVIDIA only?

Proprietary texture compression

_UT2003_ requires an nVIDIA card, GeForce 2 or better. _UT2k3_ uses Savage Texture compression, known as S3TC. The only drivers to support this are the closed source drivers from www.NVIDIA.com/view.asp?IO=linux_display_1.0-3123, and older cards like the TNT2 don't have the necessary hardware support.

Users of Radeon cards can use the Xi Graphics drivers from www.xig.com. These are commercial drivers with S3TC support, they should work with _UT2k3_.

The reason for this unfortunate state of affairs is that S3TC is proprietary so cannot be added to open source code. That's why nVIDIA's drivers can support it. Although free-of-charge, they are closed source.

```
ut2003_demo
```

if installed as root, or:

```
~/ut2003_demo\
```

if installed as a user.

If you need to remove the *UT2003* demo from your system, run the uninstall program in the ut2003_demo directory, as the same user as when you installed it. Although once you've played it, the only time you're going to want to remove it is when you get the full version.

Desktop/**Cardfile**

Linux is endowed with some very powerful relational database servers, such as *MySQL* and *PostgreSQL*. However, for some tasks don't want immense power, you just want a simple tool to do a simple job. Using *MySQL* with a web front end to keep a few addresses could easily fall into the overkill category, which is where a program like *cardfile* comes in.

As you've probably already guessed from the name, *cardfile* is a flat file database, like a set of index cards. It won't handle relational queries or complex transactions, but that's not what it's for. It will however, keep an index of your CD or video collection and let you find what you need as quickly as possible, which is what a database is all about really.

Desktop/**Rute**

Every so often a package comes along that doesn't really fit into any of the categories into which the coverdiscs are divided. Rute is one of them, but it is far too useful to leave off the CD just because we can't think of an obvious home for it.

Rute is another of those recursive acronyms that are so popular (I wonder how long it will be before we see HURD-style "mutually recursive acronym pairs" applied to Linux programs). Rute stands for Rute Users Tutorial and Exposition, which largely tells you all you need to know about it. Rute is a comprehensive tutorial and reference, for everyone from new users to system administrators.

Unlike most other packages on the CD, this one doesn't even need to be installed. There are two versions on the CD, one each in HTML and PDF. Read them direct from the CD or copy your preferred format to your hard drive for instant access to a wealth of reference material.

***UT2003* – test your skills with wanton acts of personal violence.**

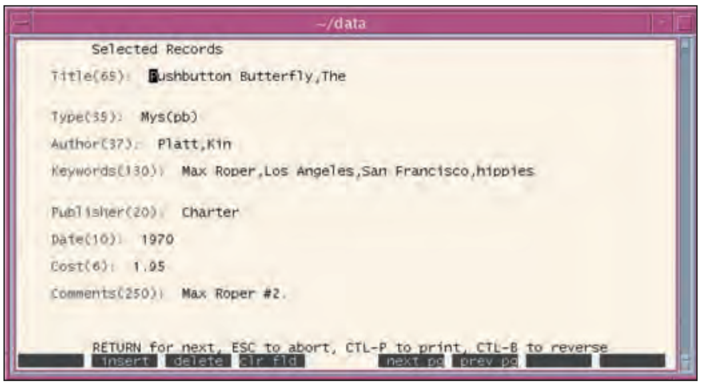

```
~/data
     Selected Records
  Title(65):  Pushbutton Butterfly,The

  Type(35):  Mys(pb)
  Author(37):  Platt,Kin
  Keywords(130):  Max Roper,Los Angeles,San Francisco,hippies

  Publisher(20):  Charter
  Date(10):  1970
  Cost(6):  1.95
  Comments(250):  Max Roper #2.


     RETURN for next, ESC to abort, CTL-P to print, CTL-B to reverse
         insert   delete   clr fld            next pg   prev pg
```

***Cardfile* – keep track of your data without learning SQL.**

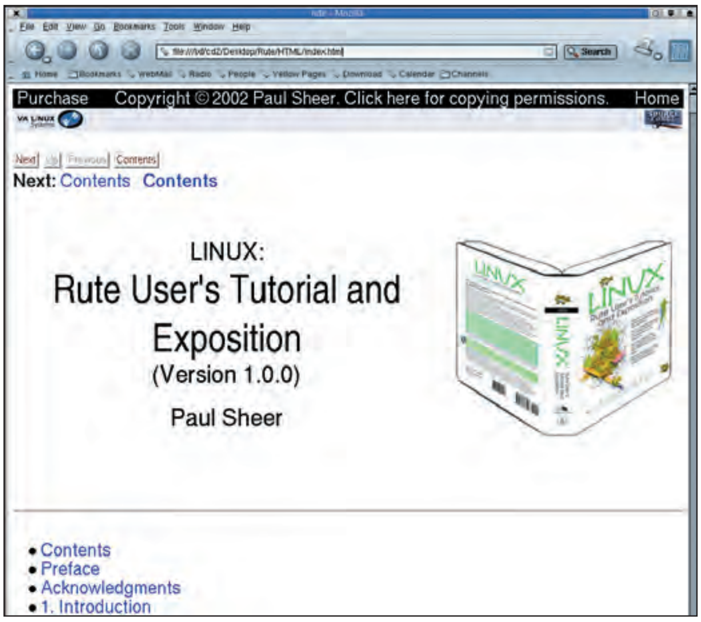
File Edit View Go Bookmarks Tools Window Help
Search
Home Bookmarks WebMail Radio People Yellow Pages Download Calendar Channels

Purchase Copyright © 2002 Paul Sheer. Click here for copying permissions. Home

Next: Contents Contents

LINUX:
Rute User's Tutorial and Exposition
(Version 1.0.0)
Paul Sheer

- Contents
- Preface
- Acknowledgments
- 1. Introduction

Rute is a comprehensive tutorial and reference on all things Linux.

Internet/**Mozilla**

It took years for *Mozilla* to reach *1.0* release status, but no sooner had it been released than work started on version *1.1*. Now that is finished (in fact, *1.2* has been available as an alpha release for a few weeks). It's a big package, which is why we weren't able to bring it to you last month, we ran out of room. If you want to install from the binary archive, make sure you read the README file first, actually that advice applies when installing any software. Alternatively, there are binary and source RPMs for those of you using RPM-compatible systems, and a source tarball of course.

Internet/**MozillaExporter**

If you are one of those unfortunate souls who prefer Linux but have to use "another operating system" at work, you can improve matters by installing *Mozilla* for Windows. Once you've installed it, you have to start configuring it all over again, or do you? *Mozilla Exporter* is a small program to convert a Linux *Mozilla* configuration to one suitable for loading by the Windows version. Run this at home, save the converted configuration to a floppy disk (or email it to yourself) and you can recreate your setup at work in a matter of seconds.

Development/**Kylix3update**

Murphy struck again last month. We received an update to *Kylix 3* the day after the coverdiscs were mastered. It's

»

not a major patch, but it does fix a couple of locale-related issues, so we made sure it was on this month's discs. To install it, you must be logged in as the same user as when you installed *Kylix3*. Then type

```
tar xzf /mnt/cdrom/
Development/Kylix3update/
k3_open_update_2.tar.gz
          cd k3_open_update_2
          ./install.sh
```

System/ DariksBootAndNuke

Let me start by warning you that this is a very dangerous program, ***read the documentation before trying it***. Deleting files from a hard disk doesn't actually remove them, they are easy to recover. Even overwriting them with null or random data doesn't stop recovery of the original data with suitable equipment. This is a boot disk that will completely erase all data on your hard drives, way beyond recovery. Why would you want to do this? You may be selling an old computer or, as happened to me recently, need to return a faulty drive under warranty. You don't really want to give people your password files, copies of emails, browser cookies with web site logins and any number of other confidential files you could have on the disk.

This package contains a floppy disk image. Once written to disk you boot from the disk and it will remove all trace of your data from the drive. This involves a great deal of writing and overwriting to the disk, it's the sort of thing you need to leave to run overnight, or longer.

OpenOffice.org

You'll find a copy of the latest *OpenOffice.org* in the Magazine directory, to go with the tutorial this month. The Office directory contains some useful additions to the package. *OOoDictionaries* contains spelling and hyphenation dictionaries in many languages. To help you get them installed and working we have *OOoDictionaryInstaller*. This is an automated dictionary installer for *OpenOffice.org*, taking care of the otherwise fiddly process for you.

To make the suite even easier to use, we have included *OpenOffice.org Quickstarter*. This is a KDE panel applet to start any of the various modules quickly.

It also has a history of recent documents and adds various other functions to make the whole package quicker to use. If you have sufficient memory, it will even keep a hidden

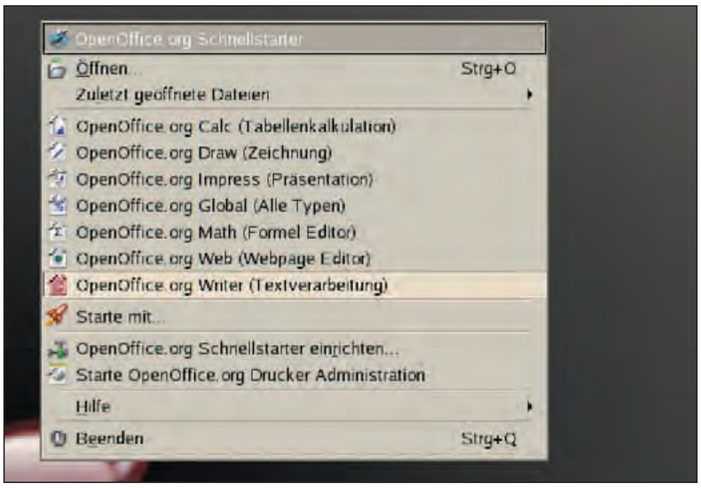
OpenOffice.org Schnellstarter
Öffnen...
Strg+O
Zuletzt geöffnete Dateien
OpenOffice.org Calc (Tabellenkalkulation)
OpenOffice.org Draw (Zeichnung)
OpenOffice.org Impress (Präsentation)
OpenOffice.org Global (Alle Typen)
OpenOffice.org Math (Formel Editor)
OpenOffice.org Web (Webpage Editor)
OpenOffice.org Writer (Textverarbeitung)
Starte mit...
OpenOffice.org Schnellstarter einrichten...
Starte OpenOffice.org Drucker Administration
Hilfe
Beenden
Strg+Q

Instants to the various *OpenOffice.org* modules from the KDE panel, with *OpenOffice.org Quickstarter.*

CD CONTENTS AT A GLANCE

Disc A

Distros

Demo Linux	Bootable disc giving a complete demonstration Linux distro. Also very useful as a rescue disc.

Disc B

Magazine

HotPicks	All the programs covered in this month's HotPicks section.
Java	The files to go with this month's Java tutorial
Kylix	Example files from the Kylix tutorials
MailServers	The mail servers covered in last month's round up
MiniDistros	The Mini distros from this month's article.
OpenOffice.org	Bug-fixed update to OpenOffice.org
Perl	Example scripts from this month's Perl tutorial
PHP	PHP source code and example scripts
PHP-Accelerators	PHP accelerators covered in this issue.
Scanners	The scanner software featured this month
WhatOnEarth	The man page software covered this month

Desktop

Allin1	Fluxbox dock applet to monitor CPU load
BitfluxEditor	Browser-based WYSIWYG XML editor
Cardfile	Simple flatfile UNIX database with a curses interface
CTWMThemesSystem	Adds basic themes capability to the ctwm window manager
EgoFileManager	GTK+2 file manager
EvilFinder	Explains why any person, object or other term is truly evil
GKrellMLaunch	GKrellM plugin for one-click access to applications
GNOMEClipboardManager	Manage your clipboard selections
GNOMEFileManager	File manager for GNOME 2
GtkDiskFree	Shows how much free space is available on your filesystems
iRPM	Simplifies RPM installation
Konserve	Small backup application for KDE 3
MultiGnomeTerminal	Enhanced version of gnome-terminal
QuantaPlus	Web editor for KDE supporting HTML and more
RPMWizard	Installation program for RPM packages
Rute	The definitive guide for new Linux users
SqlGui	Graphical MySQL and PostgreSQL frontend for KDE
Syscriptor	Displays information about your hardware
Tksql	Edit the tables of PostgreSQL databases
Webmin	Web-based interface for system administration
XMLTV	Process TV (tvguide) listings and manage your TV viewing
Xslay	Destroys all windows not known to the window manager

Development

BASHDebugger	Patched BASH that enables better debugging support
BusyBox	Tiny versions of many common UNIX utilities
cnikJEditor	Edit text, HTML pages, Java classes and C++ code
Daemon	Turns other processes into daemons
DrJava	Integrated Java development environment
EnterpriseGantt	Gantt chart library
JpGraph	OO Graph drawing library for PHP 4.0.2 and above
Kylix3update	An update to Kylix 3 from last month's cover discs.
libslack	General utilities library
Tcl_Tk	Portable scripting environment for Unix, Windows and Mac
TinyCCompiler	Small, fast, unlimited and safe C compiler
Zoinks	Programmer's editor and development environment

Games

GLAsteroids	3D clone of the arcade classic 'Asteroids'
UT2003	Lightning-fast, no-holds-barred sport of the future
UT2003area	Tools for quering Unreal Tournament 2003 servers
Wakkabox	Block shuffling puzzle game using SDL
Xbreaky	xbreaky is a breakout game for X.

Graphics

Axiom	Photo and news gallery written in PHP
Camsource	Daemon to grab image frames from a video4linux device
iCam2	Video4linux/imlib2 based webcam application
Netpbm	220+ programs to convert, analyse and process image files
RivaTV	drivers for nVidia graphics boards with video-in
TIFF2PNG	Convert images from TIFF to PNG format
XPaint	Image editor supporting most standard paint program options

Internet

Batemail	Strips potentially nasty attachments from email messages
BidMonkey	Automatically place a last minute bid on an eBay auction
Chk4mail	Lists the number of read/unread messages in all mailboxes
Emphetamine	Extensible download manager for the GNOME desktop
Eqlplus	Modem combining strategy based on IP masquerading and eql

Booting CDs

Troubleshooting a broken BIOS

Many of our cover CDs and DVDs are bootable. This doesn't present a problem for the vast majority of readers, but a few do have difficulties. There are two main ways of making a CD bootable (DVDs are just big CDs as far as this is concerned), the first is to embed the image of a bootable floppy on the disc. The system boots from that as though it were booting from floppy and this "disk" contains the information needed to transfer the boot process to the CD and run whatever software the boot needs to run, such as auto-starting a distribution installer. Any PC that is able to boot from CDROM should be able to boot from one of these discs.

The other method uses ISOLINUX. This gives more control to the creator of the discs and is generally considered a better way to do it. However, a small number of PC BIOSes will not boot from an ISOLINUX disc. This was the case with the recent Slackware coverdiscs. This can often be fixed by downloading and installing a BIOS update from your motherboard manufacturer's website.

If this doesn't work, or your system just cannot boot from floppy (this usually only applies to older hardware) the solution is to use a boot disk. This performs exactly the same task as the embedded fake floppy disk on the CD, but it's a little less convenient that booting from the CD directly.

Most bootable CDs contain the image of the floppy disk needed for this on the CD. It's usually in a directory in the root of the CD, in the case of DemoLinux the directory is called "dosutils". There are two clues that you've found the right directory, the presence of a file 1474560 of bytes, the size of a 1.44MB floppy disk and, usually, a file called *rawrite.exe*. The latter is a Windows program to write the image file to a disk and usually has an accompanying documentation file.

Why a Windows program? Well, you don't want to have to install Linux from the CD just to create a disk so you can install Linux from the CD... To create the disk in Windows, type

```
cd D:\dosutils
rawrite -f boot.img -d A
```

in a DOS box. To do the same in Linux, type

```
dd if=/mnt/cdrom/dosutils/boot.img
 of=/dev/fd0
```

using the correct path to the image file of course.

Then reboot and your system should boot from the floppy disk and transfer control to the CDROM.

instance of *OpenOffice.org* in memory, to avoid the normally lengthy startup time each time you want to use it.

Sound/**NoteEditor**

Traditional music editing software is less common than tracker-type sequencer programs, but we have one here for you. *Note Editor* is an editor for music notation that supports an unlimited number of staffs and up to nine voices per staff.

It can import files from MIDI keyboards and the TSE3 sequencer engine and output in a variety of formats used by other music programs.

If you are more interested in a sequencer, take a look at *Anthem*, also in the Sound directory. This is an advanced MIDI sequencer that allows you to record, edit and playback music using a "sophisticated and acclaimed object oriented song technology". I think what they're trying to tell us is that it is both powerful and easy to use.

FreeGuide	TV guide program that grabs listings from the Internet
Galeon	GNOME Web browser based on the Mozilla rendering engine
HTTrack	Offline web browser utility
Khttrack	Offline browser utility with a KDE Wizard Interface
Mozilla	Version 1.1 of Mozilla
MozillaExporter	Convert Mozilla configurations from Linux to Windows format
Muttprint	Formats the printing of Mutt and other mail clients
RabbIT	Speed up surfing over slow links like modems
SpamProbe	Spam detection program

Mobile

FakeAP	Generates thousands of counterfeit 802.11b access points
RoadMap	Displays street maps from the US Census Bureau
Rocks	Protects sockets-based applications from network failures
Wellenreiter	Discovery and auditing of 802.11b wireless networks
WireKismet	Small GTK frontend for Kismet that was written for the iPaq

Office

ConsultantCommunicator	Manage multiple projects, clients or tasks
KOffice	Integrated office suite based on the KDE libraries
OOoDictionaries	Spelling and hyphenation dictionaries for OpenOffice.org
OOoDictionaryInstaller	Automated dictionary installer for OpenOffice.org
OpenOffice.orgQuickstarter	Quickly start the different OpenOffice.org modules

Server

Apache	The world's most popular HTTP server V2
Boxhog	Warn when someone is using too much of a server's resources
Ecartis	Mailing list server, the successor to Listar
EcartisWebGUI	Simple but useful Web GUI for the Ecartis listserv
GeekLog	Weblog software for performance, privacy, and security
GtkFtpd	FTP server with a personal, convivial interface
mod_accounting	Apache module to record traffic statistics
mod_ifexists	Apache conditional configuration support
MyPhpTracker	Web tracking system
MySQL	Widely used and fast SQL database server
PHP-Nuke	Web portal and online community system
ProMA	Web interface for managing users in a ProFTPd/MySQL setup
SecurityFilter	Security and authentication for Java Web applications
SherlockHolmes	System for gathering, indexing and searching textual data

Sound

Anthem	MIDI sequencer to record, edit and playback music
Festival	Multi-lingual speech synthesis system
GNOMEALSAMixer	Sound mixer for GNOME and ALSA
GQmpeg	Frontend to various audio players
KDESpeaker	Text-to-speech output for Konqueror and the Kate editor
NoteEditor	Music notation editor for an unlimited number of staffs
TerminatorX	Realtime audio synthesizer
TSE3	Powerful open source sequencer engine written in C++

System

BuffaloSnooper	Trace library calls (somewhat like SnoopDOS)
CheckInstall	Creates Slackware, RPM or Debian packages from tarballs
DariksBootAndNuke	Self-contained boot floppy to securely wipe hard disks
DebSync	Synchronize a bunch of Debian GNU/Linux machines
FirewallBuilder	GUI and policy compilers for various firewall platforms
FTimes	System baselining and evidence collection tool
lm_sensors	Monitor system temperatures, voltages and fans
OpenwallLinuxKernelPatch	Collection of security "hardening" features for the kernel
RPMrebuild	Build an RPM file from a previously installed package
Speakimage	Rescue CD based on Linux 2.4.19 and speakup for the blind
TapeOrientedBackup	General driver for making and maintaining backups

Essentials

Allegro	Multi-platform game library
Alsa	The Advanced Linux Sound Architecture project
Avifile	Library to read and write compressed AVI files
GTK	GNU's GIMP Tool Kit graphics library
Kernel	The latest kernel source and patches.
lesstif	LGPL'd implemention of the OSF/Motif standard GUI toolkit
libmcrypt	Encryption library
libSMTP	Mail transfer libs
libstdC++3	OO grafted onto C, apparently quite popular for Linux coding
libxml	XML C library developed for the Gnome project
Mesa	3-D graphics library that uses the OpenGL API
ncurses	Text-based interface creation library
OggVorbis	Open, professional audio encoding and streaming technology
SBdriver	Drive that soundblaster card!
SDL	Groovy games library
SVGAlib	Yet another display alternative

Coverdisc

Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format* DVD, and it's a bit of a desktop special...

Here we are with another collection of extra software, over and above what is included on the two CDs. Unlike the first CD, the DVD won't boot into DemoLinux, you'll have to burn the image file from Magazine/MiniDistros to a CD-R first. Otherwise, you get exactly what's on the CDs and then a lot more. Here's a selection of the extras.

Desktop/ KDE3.1beta2

This is not the final release of KDE 3.1 by any means, but it is getting close enough to try this taster. This is the third public release of KDE 3.1, there was one alpha release before the betas, so the nastiest bugs should be gone. However, it is probably not stable enough to run on a system you need for daily use. If that's the case, why would you want to run it at all? There are several reasons; you may want to develop software that uses new features of KDE 3.1, by running the beta you may be able to give useful feedback to the developers to help eliminate bugs, or you may just like having the latest toys to play with.

There are a couple of ways around this potential instability (apart from not trying it at all and waiting for the release version). You may have more than one Linux distribution installed, especially if you like trying out all the alternatives on the *Linux Format* DVDs, so you could install 3.1 on one of them. Another alternative is to install it alongside KDE 3.0, enabling you to choose between the two, and remove the beta if it proves to be not yet ready for you.

Running two copies of KDE naturally requires that they are installed in different places. How you do this depends on whether you are installing from tarballs or RPMs. When installing from source, use the **--prefix** option to install the KDE 3.1 tree in a different location. For example, if you have 3.0 installed in the default location of /usr/local/kde, use

```
./configure --prefix=/opt/kde3
```

to keep 3.1 completely separate. You could also use something like /usr/local/kde31 to keep them side by side.

The situation with the RPMs could be more complex. The *rpm* command does allow you to alter where a package, or part of a package, is installed with the **--relocate** option. This only works if the packages are relocatable, and none of the KDE ones are. The beta RPMs put everything in /opt/kde3, so you'll need to check where your current KDE 3 setup is installed, with

```
rpm -ql kdelibs | more
```

Assuming you don't currently have KDE3 installed in /opt/kde3, you can install the 3.1 beta with

```
rpm -ihv /mnt/cdrom/Desktop/
KDE3.1beta2/Mandrake/*.rpm
```

or

```
rpm -ihv /mnt/cdrom/Desktop/
KDE3.1beta2/SuSE/*.rpm
```

depending on which of the distros you use. There were no Red Hat RPMs available from KDE at the time the DVD was compiled. If you now restart your computer, you'll find yourself back in KDE 3.0. To use 3.1, boot into console mode and execute the following commands, you could put them in a script to make things easier

```
export KDEDIR="/opt/kde3"
export PATH="$KDEDIR/bin:$PATH"
export KDEHOME=$HOME/.kde31
export
 LD_LIBRARY_PATH="$KDEDIR/lib"
```

Then you can launch 3.1 with

```
kde3
```

If you need to switch back to 3.0 during a session, log out of KDE, log out of the console and log back in then run

```
kde
```

Logging out and back in resets the command and library paths, **kde**, as opposed to **kde3** starts the standard KDE 3.0 system. It is important to log out and back in when switching, otherwise you could end up with KDE loading libraries from the wrong source.

Don't forget, this is a beta, which means two things. It is likely to contain bugs and it has been released so we can help development by providing useful bug reports. If you have a problem, don't just delete it and wait for the final release, report the bug. You don't have to be a coder to help with open source development.

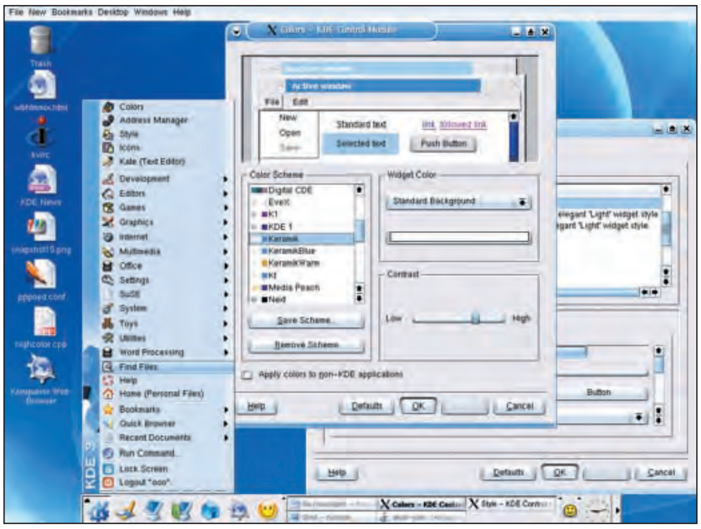

Konqueror **gets tabbed browsing in KDE 3.1.**

On the DVD

Wherever you see this logo it means there's related stuff on the DVD

Important notice

Before you even put the DVD in your drive, please make sure you read, understand and agree to the following:

The *Linux Format* DVD is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker on ANY new software. While every care is taken in the selection, testing and installation of DVD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system which may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all important files. Please read individual licences for usage terms.

Distros/SmoothWall2

SmoothWall has been a popular product for some time. In case you haven't come across it before, it turns an old PC into a powerful, yet easy to configure firewall and router. The first version stopped tantalisingly short of the magic 1.0 release number, at 0.9.9. Now SmoothWall 2 has been released as a beta, this is the second beta release, in fact. Hardware requirements are incredibly modest. An old 486 will do perfectly, I've run the previous version on an old P100. You'll need a monitor, keyboard and CDROM drive on the machine for the duration of the installation only. After that, all configuration and control is done remotely, via a browser interface. The software is supplied as an ISO image, burn this to CD and boot from that CD.

SmoothWall takes over the whole hard drive. I'll repeat that because it's important, SmoothWall *reformats the **entire** drive*, you'll lose anything currently on there. After that it's simply a matter of following the prompts. If you connect via an ADSL or cable modem that uses ethernet, you'll need two ethernet cards in the SmoothWall box, one to connect to the Internet and one for your local network. There is also support for USB ADSL modems from Alcatel (the BT frog) and Fujitsu. Keep and eye on the SmoothWall web site for security patches, there has been one released so far for this version, it's on the DVD. The patch is applied once SmoothWall is running, using the browser interface.

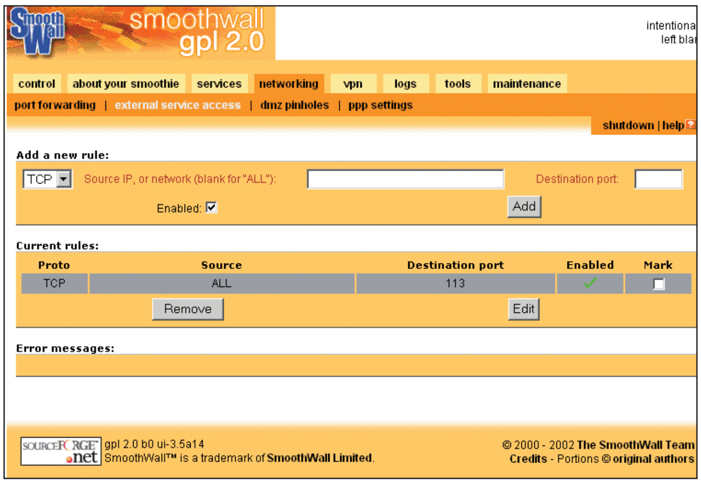

Add a new rule:

TCP Source IP, or network (blank for "ALL"): Destination port:

Enabled: ☑ Add

Current rules:

Proto	Source	Destination port	Enabled	Mark
TCP	ALL	113	✓	□

Remove Edit

Error messages:

gpl 2.0 b0 ui-3.5a14
SmoothWall™ is a trademark of **SmoothWall Limited**.

© 2000 - 2002 **The SmoothWall Team**
Credits - Portions © **original authors**

Put that old PC to use, protecting your system instead of gathering dust. SmoothWall turns even an old 486 into a powerful firewall/router.

Desktop/**GCompris**

The best time to convert computer users to Linux is when they're young, before they get set in their ways. So here's a package for children between the ages of three and ten. *GCompris* is a suite of educational programs to teach children to use a computer from an early age, moving on to teach things like arithmetic and spelling through play.

This is a modular program, designed around the concept of boards for different games. Full developer documentation is included so that others may add new games. Whether you'd want your children using your computer, putting toast in the DVD drive when you're not looking, is another matter. Although it does give you an excuse to buy a new one so the kids can have the old one.

Graphics/**MoviX**

Linux distributions range from the general purpose, like Debian or Red Hat, to the single-minded, like SmoothWall. Here is another example of the latter. It is in the Graphics directory because that's all it does. MoviX is a mini-distribution designed to play movies directly from a bootable CD. Usage is very simple, although you can make it easier by pasting the **mkisofs** command line from the README into a one line shell script. All you do is unpack the archive, copy the movies you want to play into its directory and run *mkisofs*. Burn the resulting image file to a CD and you have a disc that will play the movies you've copied to it when you boot from it.

MoviX uses *Mplayer* to play the movies, so it can handle a wide range of file formats. Uses of this program include creating your own movie collection, ready to watch on any computer, to creating demonstration CDs showing animated presentations or slideshows.

Catch 'em young. *GCompris* is an educational suite to get children comfortable with computers (and Linux) from an early age.

Office/**Scribus**

Desktop publishing is not too well catered for on Linux, but *Scribus* is a good attempt at creating a DTP package with professional features. It may be a while before *Scribus* is up to producing a magazine like *Linux Format*, but it is already more than capable of producing documents like newsletters and stationery. *Scribus* is capable of producing output suitable for professional imagesetting machines, meaning it can generate files to be sent to a commercial printing company, as well as printing directly to a desktop printer. Installation follows the usual

```
./configure
make
make install
```

routine, but there are a couple of other steps to take care of too, relating to the program's use of PostScript fonts. Read the accompanying README file (although you do that anyway, don't you?) before installation.

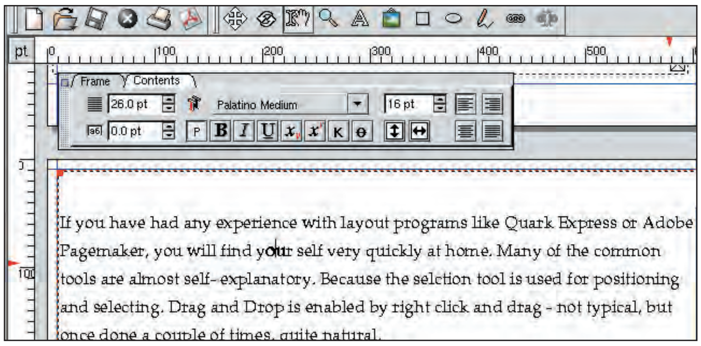

26.0 pt Palatino Medium 16 pt
0.0 pt

If you have had any experience with layout programs like Quark Express or Adobe Pagemaker, you will find your self very quickly at home. Many of the common tools are almost self-explanatory. Because the selction tool is used for positioning and selecting. Drag and Drop is enabled by right click and drag - not typical, but once done a couple of times, quite natural.

***Scribus* bring easy to use DTP with quality output to Linux.**

Desktop/ **GNUPrivacyGuard**

Information is power, they say, and an increasing proportion of information is travelling over insecure networks, like the Internet. It is impossible to be sure that no one is able to intercept your emails, but you can make sure no one is able to understand them. The GNU Privacy Guard, usually shortened to GPG, is the open source alternative to PGP (Pretty Good Privacy). It provides a highly secure, yet simple to use, means of encrypting your emails, or any other files, so that they may only be read by the intended recipient. It's so secure that, unless you add yourself to the recipient list, you can't even decode the files yourself after processing with GPG.

It does more than encrypt files. You can also add a digital signature to a message. any attempt to tamper with the message, even changing one letter, will make the signature invalid and let the recipient know that the message they received was not sent by you.

If this all sounds a bit cloak and dagger and you are thinking you have no need for such security, that you have nothing to hide in your communications, ask yourself how you would feel if your bank sent your next statement on a postcard, for anyone to read. LXF

» DVD CONTENTS AT A GLANCE

Desktop

2fax	Converts ASCII files into fax (TIFF) formatted files
AFT	Almost Free Text is a document preparation system
Atool	Manages file archives of various types (tar, tgz, zip etc)
CdrLin	Java-based GUI for cdrecord, mkisofs, cdda2wav etc.
CheckBookTracker	Fully-functional money management program
DatabaseSchemaDesigner	Creates and maintains schemas for SQL databases
DI	Disk information utility
Dropline	GNOME Desktop that has been tweaked for Slackware Linux
DRT	Design recovery tool for interactive graphical applications
eXchaNGeR	Browser and editor framework, written in Java
GCompris	Complete educational suite for children 3 to 10 years old
GNUPrivacyGuard	GNU's tool for secure communication and data storage
GSpeakers	GNOME loudspeaker design program
GTK+XFce	Easy-to-use and easy-to-configure environment for X11
GxCFG	Substitute for xdpyinfo, xset, xgamma, and xvidtune
InputReplayModule	Record and replay the input of a process
Kbrite	Applet to control the brightness of a Sony VAIO LCD screen
KDE3.1beta2	This is the third pre-release of the forthcoming KDE 3.1.
Kemistry	Collection of chemistry applications for KDE
Koha	Library and collection management system
KPopup	GUI to send and receive Microsoft(tm) Winpopup messages
Kwatch	Graphical KDE 2.0 user interface for watching log files
LEDapplet	Shows the status of the Caps-, Num- and Scroll-Lock keys
LinuxSMS	Perl script to send SMS to GSM phones
Mad.Thought	Web-based PHP diary/journal
MySQLNavigator	MySQL database server GUI client program
PalmZeteticStrip	An interface to the Palm database (PDB) files for Strip
PFAedit	PostScript font editor
PgAccess	Tcl/Tk frontend for PostgreSQL
Pip	Wrapper for programs that won't use stdin and stdout
PostgresFrontEnd	Connect to and work with a Postgres database server
QuikDic	Dictionary that will look up words over the Internet
ROX-Session	Simple and easy-to-configure session manager
SchemeScribe	Text processor best suited to writing technical documents
sed	The GNU Stream Editor
Shed	Hex editor with a friendly pico-style interface
Snowdrop	Steganographic text document and C code watermarking tool
TeXmacs	Free scientific text editor
VMWare	Run 'virtual machines' inside Linux
Wine	Implementation of the Windows 3.x and Win32 APIs
XGlobe	Displays a globe on your X desktop
Xnee	Records, distributes and replays X protocol data
XORSuite	OTP encryption suite

Development

AllegroSpriteEditor	Creates sprites for use in video games
Anthill	Bug tracking database system written in PHP
AutoconfMacroArchive	Documented and tested macros using the autoconf tool
Autospec	Creates RPM spec files automatically from a tarball
CheapThreads	Portable C routines for implementing sychronous threads
CVSweb	WWW interface to explore a CVS repository
Dachshund	Software modeling tool for the GNOME environment
EditX	Basic text editor construction kit for X
ezV24	Programming interface for Linux serial ports
Jikes	Fast, simple Javasource code to byte code compiler
KCachegrind	Visualizes traces generated by profiling
libsndfile	C library for reading and writing sound files
LinCVS	Graphical frontend for the CVS-client
PiccoloToolkit	Create robust, full-featured graphical applications in Java
popt	Library exists for parsing command line options
Prima	Extensible Perl toolkit for multi-platform GUI development
TOAD	Create graphical user interfaces with C++
wxWindows	GUI framework for easy cross-platform programming
XmlConfig	Instantiate and configure Java objects from an XML file

Distros

BanSheeLinux_R	Two-floppy rescue system
CoolLinuxCD	Bootable CD with a live distribution based on RedHat 7.3
Devil-Linux	Special Linux distribution for firewalls and routers
ServerdiskDisketteDistro	Floppy disk distribution that includes FTP and HTTP servers
SmoothWall2	Turn an old PC into a firewall/router

Games

CrimsonFields	Tactical war game in the tradition of Battle Isle
Crossfire	Multi-player graphical arcade and adventure game
D20ListView	Manage D&D 3rd Ed. combat
Epiphany	Multiplatform clone of the game Boulderdash
Freeciv	Multiuser reimplementation of Civilization
InfernalContractorII	Raise mayhem in this fast-moving office simulation
JumpNBump	Funny multiplayer game
PachiElMarciano	Platform game inspired by Manic Miner and Jet Set Willy
Q3Log	Generates HTML stats pages from Quake 3 Arena log files
Reaper	Graphics-intensive 3D space game
SINS	Snake game controlled by only one key

Graphics

DMIGallery	Editor-oriented image gallery application
Drip	DVD to divx conversion tool
GRASSGIS	Geographic Resources Analysis Support System
ImageMagick	Automated and interactive manipulation of images
Kodisein	Design, display and animate scenes containing 3D objects
MonAlbum	MonAlbum is a PHP and MySQL based photo album.
MoviX	Boot a PC from CD and play all the movie files on it
MPlayer	Movie and animation player
VideoDiskRecorder	Digital sat-reciever program using DVB technology

Mobile

gnuPod	Use your iPod under Linux
Keyring	Store passwords and data in a resonably secure fashion
SyncPOD	Syncs a local directory with your iPod

Office

AccWhizz	Multilingual accounting application for small businesses
Cacti	Create graphs and populate them with data in MySQL
FavorinTime	Client-server group calendar using the KDE desktop
Roomba	Property Management System and Central Reservation System
Scribus	Easy-to-use tool for simple Desktop Publishing
Zeiberbude	Point of sales program for cyber-cafes

Server

AlarmPinger	Monitors various IP devices by simple ICMP echo requests
ApacheTomcat	Java Servlet and JavaServer Pages for Apache
BobsVirtualCDServer	Remote operated network virtual CD server.
ChronoStats	Fast, easy log parser for Half-Life: Counter-Strike
DarkPortal	Web portal system
DBaliasd	Rewrites aliases of MySQL database names with real values
ffRADIUSLibrary	C routines implementing RFCs 2865 and 2866
htCheck	Link checker derived from ht://Dig
InsightLogger	GUI-based Web server log file analysis tool
JBoss	Enterprise JavaBeans application server
Lcd4web	Write text to an LCD using a Web interface
Logrep	Extract and present information from various logfiles
MiddleMan	Fast HTTP/HTTPS proxy server
MIMEDefang	MIME email scanner to protect Windows clients from viruses
MySQLImport	Import files into MySQL databases via the Web
NetVisModule	Simulate, analyze and visualize social networks
NiftyPHPUtilities	Collection of small utilities written in PHP
PHP	Scripting language, especially suited for Web development
PostGIS	Adds support for geographic objects to PostgreSQL
PostgreSQL	Robust, next-generation, Object-Relational DBMS (ORDBMS)
PostgreSQLMaintenance	Runs routine queries and vacuums on PostgreSQL servers
Protoscope	Shows HTTP communication while testing web applications
RHLinuxDomainController	Fulfils the role of a Windows NT domain controller
SPIP	Free publication system for the Internet
Squid	High performance Web proxy cache
WebCP	Control Panel for Web hosting companies and resellers
WebShell	Remote UNIX shell that works via HTTP
WorldPayJunior	Track transactions made via WorldPay Select Junior

Essentials

gcc-3	The best compiler in the world?
LDP	A complete mirror of the Linux Documentation Project

Please note that contents of Internet, Sound and System directories can be found listed on the DVD.

User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally natter about all things Linux. You can find lots more information online at: www.lug.org.uk

1 Hampshire
URL www.hants.lug.org.uk
Contact Hugo Mills

2 Bristol & Bath
URL www.bristol.lug.org.uk

3 Scottish
URL www.scottish.lug.org.uk

4 Oxford
URL www.oxford.lug.org.uk
Contact Alasdair G Kergon

5 Kent
URL www.kent.lug.org.uk
Contact John Mills

6 Brighton
URL www.brighton.lug.org.uk
Contact Johnathan Swan

7 Worcestershire
URL www.worcs.lug.org.uk
Email info@thirdeyedevelopment.com

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.wolves.lug.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://gllug.linux.co.uk/
Contact John Southern

20 Surrey
URL www.surrey.lug.org.uk
Contact Jay Bennie

21 Cambridge
URL www.cam-lug.org

22 Devon & Cornwall
URL www.dclug.org.uk
Contact Simon Waters

23 Falkirk
URL www.falkirk.lug.org.uk

24 Manchester
URL www.manlug.mcc.ac.uk
Contact John Heaton, Owen Le Blanc

25 Hertfordshire
URL www.herts.lug.org.uk
Contact Nicolas Pike

26 West Yorkshire
URL www.wylug.lug.org.uk
Contact Jim Jackson

27 Sheffield
URL www.sheflug.co.uk
Contact Richard Ibbotson

28 Staffordshire
URL www.staffslug.org.uk

29 North East
URL www.shofar.uklinux.net/NELUG

30 London
URL www.lonix.org.uk

31 Thames Valley
URL www.sclug.org.uk

32 Liverpool OpenSource
URL http://linux.liv.ac.uk/_liv_linux_ug/
Contact Simon Hood

33 Deal Amiga Club
Email superhighwayman@hotmail.com
Contact John Worthington

34 Chesterfield
Email spirelug@yahoo.co.uk
Contact Robin Needham

35 South Derbyshire
URL www.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 ben@ilovephilosophy.com

39 Cheshire
URL www.sc.lug.org.uk
Contact Anthony Prime – enquiry@sc.lug.org.uk

40 North Wales
URL www.northwales.lug.org.uk
Contact Jonathan Cole

41 Midlands
URL http://midlandslug.port5.com/
Contact Pete Thompson

42 Cumbria
URL www.cumbria.lug.org.uk
Contact Jamie Dainton

43 Dorset
URL www.dorset.lug.org.uk
Contact John and Mat

44 Shropshire
URL www.shropshire.lug.org.uk
Email shropshire@lug.org.uk

45 South West
URL www.southwestlug.uklinux.net
Email southwest@lug.org.uk

46 South Wales
URL www.sw.lug.org.uk
Contact Tim Bonnell

47 North London *.– see 87*
URL www.kemputing.net/lug/anlug-aims.html

48 Malvern
URL www.malvern.lug.org.uk
Contact Greg Wright

49 Huddersfield
URL www.hud.lug.org.uk
Contact Adam Brookes

50 Nottingham
URL www.nottingham.lug.org.uk
Contact Godfrey Nix

51 St Albans & Luton
URL www.lust.lug.org.uk
Contact Michael Culverhouse – mike@easily.co.uk

52 Wrexham
Contact Paul Kersey-Smith
Email paul@pkls.fsnet.co.uk

53 Preston & Lancs
URL www.preston.lug.org.uk
Contact Phil Robinson

54 Derry
URL www.derry.lug.org.uk

55 Isle of Wight
URL www.iow.lug.org.uk
Contact David Groom – info@iow.lug.org.uk

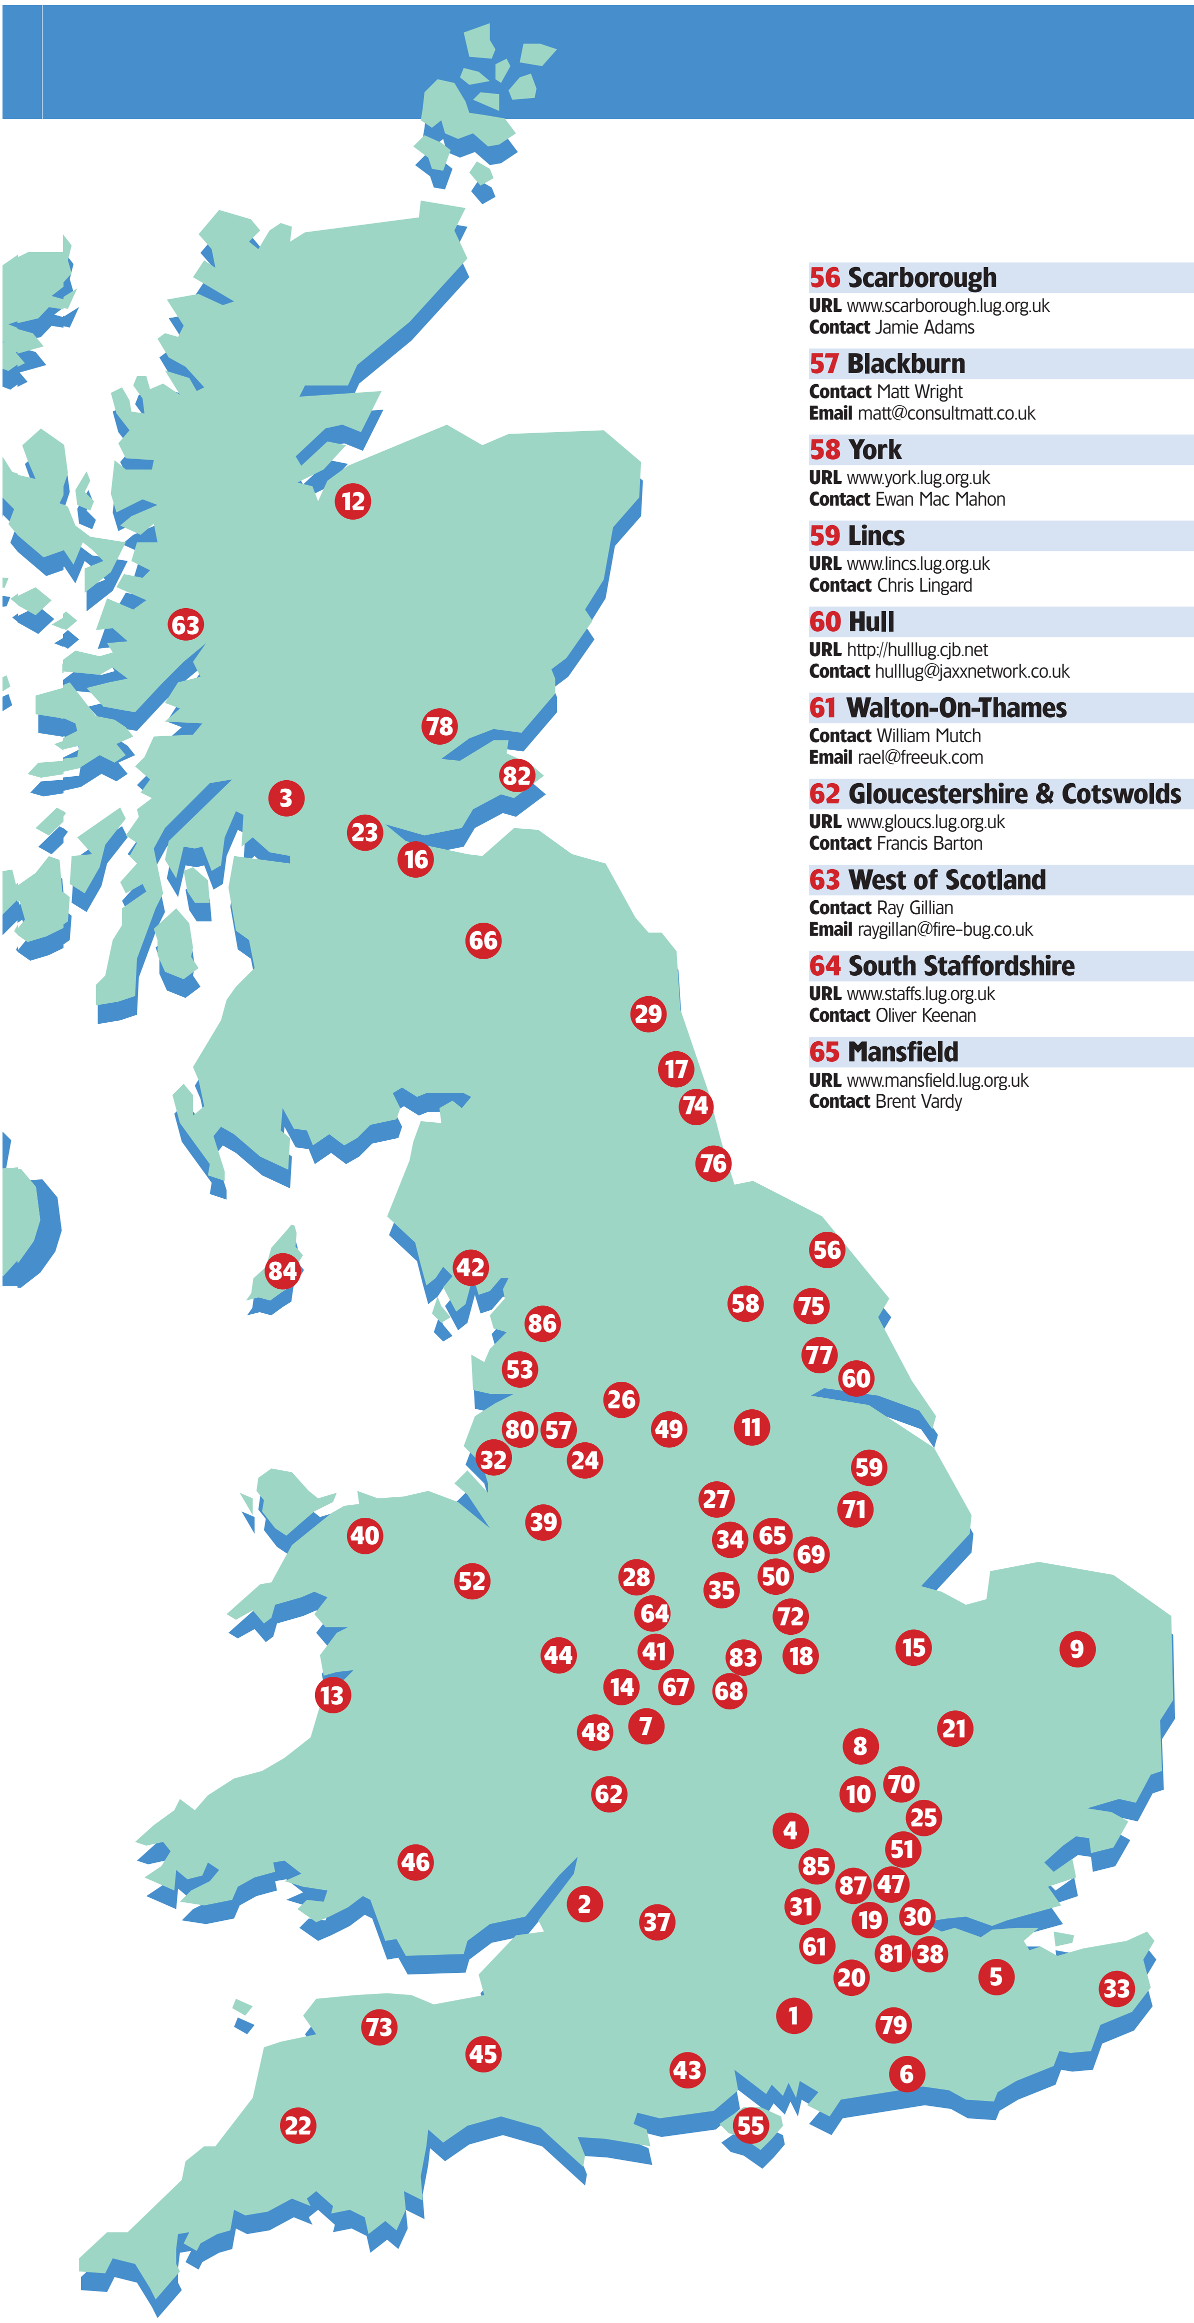

56 Scarborough
URL www.scarborough.lug.org.uk
Contact Jamie Adams

57 Blackburn
Contact Matt Wright
Email matt@consultmatt.co.uk

58 York
URL www.york.lug.org.uk
Contact Ewan Mac Mahon

59 Lincs
URL www.lincs.lug.org.uk
Contact Chris Lingard

60 Hull
URL http://hulllug.cjb.net
Contact hulllug@jaxxnetwork.co.uk

61 Walton-On-Thames
Contact William Mutch
Email rael@freeuk.com

62 Gloucestershire & Cotswolds
URL www.gloucs.lug.org.uk
Contact Francis Barton

63 West of Scotland
Contact Ray Gillian
Email raygillan@fire-bug.co.uk

64 South Staffordshire
URL www.staffs.lug.org.uk
Contact Oliver Keenan

65 Mansfield
URL www.mansfield.lug.org.uk
Contact Brent Vardy

66 Borders
URL www.linux.bordernet.co.uk
Contact Welby McRoberts

67 South Birmingham
URL www.sb.lug.org.uk
Contact Tim Williams

68 Coventry
Contact Darren Austin
Email info@coventry.lug.org.uk

69 Newark & Lincoln
URL www.newlinc.lug.org.uk

70 Bedfordshire
URL www.beds.lug.org.uk
Contact Neil Darlow

71 Lincoln
URL www.lincoln.lug.org.uk
Contact Jon Shamash

72 Loughborough
URL www.loughborough.lug.org.uk
Contact Martin Hamilton

73 Exeter University
Contact Nicholas Murison
Email N.J.Murison@exeter.ac.uk

74 Sunderland
Contact Thomas Croucher
Email thomas.croucher@sunderland.ac.uk

75 East Yorkshire
Contact Daniel Gallacher
Email sharkonline@whatemail.com

76 Cleveland Open Source Group
Contact Haniff Din
Email openlug@digitalmedia.co.uk

77 Beverley
Contact Vladimir Lukyanov
Email vladimir_lukyanov@hotmail.com

78 Dundee & Tayside
URL www.dundee.lug.org.uk
Contact Duncan Gauld

79 Sussex
URL www.phpworld.co.uk/~swlug
Contact Gareth Ablett

80 Wigan & St Helens
Contact Paul F Johnson
Email paulf.johnson@ukonline.co.uk

81 Brixton
URL www.communitytechnology.org.uk/~linuxhome
Contact R M Sanchez

82 St.Andrews, Fife
Contact Stuart Anderson
Email stuart@nx14.com

83 Nuneaton
URL www.nuneaton.lug.org.uk
Contact S Prosser

84 Isle of Man
REVISED DETAILS
URL www.iom.lug.org.uk
Email helix@manx.net

85 Aylesbury
URL www.aylesbury.lug.org.uk
Email drbond@educational-computing.co.uk

86 Lancashire
NEW DETAILS
URL www.lancasterlug.org.uk
Contact Jonathan Hill

87 North London *– see 47*
NEW DETAILS
Contact Jason Hodson
Email jason@voyagercomputers.co.uk

LUG OF THE MONTH!

Belfast

Users of Linux, and GNU software in general, are welcome to join the Belfast Linux Users Group (BLUG). People from all over Northern Ireland gather at our monthly meetings, where we have techical talks; installation and fixing sessions; and socialise over a cup of coffee. We post details of our meetings at our website.

If you have a problem with your system, want more information on a topic or just want to share your thoughts, please join our mailing list.

We have an IRC channel, **#blug**, on the OFTC network, as one of our members runs a local node. We use the mailing list as our main communication channel however.

BLUG also has strong links with the Belfast Perl Mongers group, the BelfastWAN group, and other local user groups. As such, we have a healthy mix of Unix gurus, interested novices, and people from other fields of expertise.

We are a friendly bunch of people, and would invite any interested person, whether they run Linux or not, to come along to one of our meetings and talk to us – or to join our mailing list and introduce themselves to the group.

More information is available on our website.

www.belfastlinux.org

Worldwide Linux User Groups

Free Software users across the globe

Africa

EGYPT
URL www.linux-egypt.org
Contact Hesham Bahram

GAUTENG, SOUTH AFRICA
URL www.glug.org.za
Email glugmin@revolution.org.za

Australia

ADELAIDE
URL www.linuxsa.org.au
Email mtippet@anu.edu.au

ALICE SPRINGS
URL www.aslug.org.au

MELBOURNE, VICTORIA
URL www.luv.asn.au
Contact luv-committee@luv.asn.au.

PERTH
URL http://plug.linux.org.au/

Europe

AUVERGNE
URL www.linux-arverne.org
Email Cyril.Hansen@wanadoo.fr

COSTA DEL SOL (English speaking)
URL www.fuengirola.lug.org.uk

DENMARK
Alssund www.alslug.dk
Esbjerg www.eslug.dk
Fyns www.flug.dk
Midt-og Vestjylland www.mvjlug.dk
Nordjylland www.njlug.dk
Skåne Sjælland www.sslug.dk
Trekantsområdet www.tlug.dk
Vest-fyn www.haarby-net.dk/vflug
Århus www.aalug.dk

EIRE
URL www.linux.ie
Email root@linux.ie

URL www.dilu.org
Email glossary@dilu.org

GOTHENBURG
http://nain.oso.chalmers.se/LUGG/

UK: Don't forget the distro-specific lists:
URL www.lug.org.uk/maillist.html

India

URL www.linux-india.org
Email newsmaster@linux-india.org

TRIVANDRUM
URL www.river-valley.com/tux
Email anil@river-valley.com

North America

ALASKA
URL www.aklug.org
Email deem@wdm.com

ALBERTA
URL http://calgary.linux.ca/

BATON ROUGE
URL www.brlug.net
Email dpuryear@usa.net

BAY AREA
URL www.balug.org
Email aftyde@balug.org

CLARKSVILLE, TN
URL www.cllug.org
Email tux@cllug.org

DENVER
URL http://clue.denver.co.us/

FLORIDA
URL www.flux.org

LOS ANGELES
URL www.lalugs.org
Email dank@alumni.caltech.edu

NORTH COLORADO
Email nclug@nclug.org

OTTAWA CANADA
URL www.oclug.on.ca

TAMPA
URL www.suncoastlug.org
Email president@suncoastlug.org

UHACC Normal, IL
URL www.uhacc.org
Email lug@uhacc.org

VIRGINIA TECH
URL www.vtluug.org
Email nega@vt.edu

South America

BUENOS AIRES
Email dcoletti@impost.com.ar

CHILE
URL www.linux-chile.org

MONTEVIDEO
URL www.linux.org.uy

PARAGUY/ ASUNCION
Email rolgiati@conexion.com.py

SAO PAULO
URL http://gul.ime.usp.br/
Email gul@ime.usp.br

Spreading the word

Jono Bacon continues his advocacy column with a look at what GNU/Linux can do to benefit schools.

Last month we looked at how Linux could be targeted at differing types of people/organisations in various ways. This month we will continue to identify our audience and look at their needs.

The true test of advocacy is not in telling someone how great something is – any fool can do that, but it is in telling how great something is for *them*. To do this successfully we need to understand the audience.

Let us take a school, for example. Think to yourself what the issues surrounding IT may be at a school. Well, we all know that schools are increasingly strapped for cash, so cost is an issue. Educational software would be another benefit; ease of administration is another (many mathematics teachers have to double as IT staff); and also some administration software is another.

If you have difficulty thinking of some issues, call a local school or college and ask them; I am sure someone can help you.

The key now is to link together your research of Linux/Free Software with these issues. This is applied best by writing the Linux benefits on one side of a piece of paper and the requirements on the other side, and linking them with lines. We now have a clear action plan of what needs to be communicated to the school.

With this knowledge of what the key links are between the requirements and benefits, we can now research these specific areas and gather information that can be presented to a school representative/teacher. This is an important step, and it should be noted that your information should be as correct as possible; try to avoid mailing list posts that Google picks up as fact – stick to the facts related to the connections you made earlier.

With these connections under our belt, and the research that backs them up, we will begin working on the best presentation of this material to the audience next month.

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*

Art Editor Julian Jefferson *julian.jefferson@futurenet.co.uk*

Production Editor Richard Smedley *richard.smedley@futurenet.co.uk*

Editorial Contributors
Jono Bacon, Neil Bothwick, Chris Brown, David Cartwright, Andy Chanelle, David Coulson, Chris Denton, Richard Drummond, Hoyt Duff, Robert Fenwick, Marco Fioretti, Paul Hudson, Jon Kent, Neil Lucock, Charlie Stross

COVER CD PRODUCTION
CD Editor Neil Bothwick
Senior CD Editor Kate Hadley

ART CONTRIBUTORS
Photography James Wilson, Gavin Roberts
Illustration Paul Bateman, Chris Winn, Shane Collinge

ADVERTISING SALES
Portfolio advertising manager Matt Stanley 01225 442244
matt.stanley@futurenet.co.uk
Portfolio deputy ad manager Diane Clydesdale
Business development manager Damian Hughes
Account managers Zoe Drew, Giles Crosthwaite Scott,
Ad sales executives Jo Doyle, Stephen Hall

MARKETING AND PROMOTIONS
Marketing Manager Fiona Tully
Subscriptions Executiver Nick Skardon

PRODUCTION
Production Co-ordinator Diane Ross
Group Production Manager Clare Tovey
Advertising Production Co-ordinator Jo Crosby

MANAGEMENT
Publisher Dave Taylor
Group Advertising Manager Simon Moss
Group Publisher John Weir
Circulation Director Sue Hartley
Managing Director Colin Morrison

DISTRIBUTION AND CIRCULATION
Circulation Manager Jamie Malley
jamie.malley@futurenet.co.uk

Distributed by Seymour Distribution, 86 Newman Street,
London W1T 3EX
Tel +44 (0)207 3968000
Overseas Distribution by Future Publishing Ltd.
Tel 01225 442244.

Overseas Licences
Simon Wear *simon.wear@futurenet.co.uk*
Tel +44(0)1225 442244 **Fax** +44 (0)1225 732361

Contact Details
Linux Format, 30 Monmouth Street, Bath BA1 2BW
Tel +44 (0)1225 442244 **Email** *linuxformat@futurenet.co.uk*

Subscriptions and Mail Order
Phone +44 (0)1458 271178. See page 98
Email *linuxformat.subs@futurenet.co.uk*

Disclaimer All tips in this magazine are used at your own risk. We accept no liability for any loss of data or damage to your computer, peripherals or software through the use of any tips or advice.

Printed in the UK by Midway Clark (Holt) and © Future Publishing Ltd 2002

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

the future network
Media with passion
www.thefuturenetwork.plc.uk

NEXT MONTH

Issue 35 on sale Friday 29 November

CHRISTMAS IS SORTED!

Our Xmas issue is looking like a right little cracker. Amongst other things, we'll be taking a look at some great geeky stocking-fillers and gadgets, and we'll also have a very special present for you on our Coverdiscs!!!

HAVE YOUR SAY!

Of course, this is also the time of year you get to give other people presents. And what could be more thoughtful than a *Linux Format 2002 Award*! Will **Smoothwall** get more prizes? Do those **GNOME** chaps deserve a gong this year? Can anything beat **OpenOffice.org** – only you know the answer.

PLUS:

Don't miss our huge roundup of development IDEs, Reviews of Mandrake 9, Webmin, Unreal Tournament (maybe), NetOp remote control and more!

DON'T MISS YOUR COPY OF LINUX FORMAT!
Get it delivered to your door every month – subscribe on page 98

The exact contents of future issues are subject to change

Contacts

EDITORIAL

Editor
Nick Veitch
nick.veitch@futurenet.co.uk

Art editor
Julian Jefferson
julian.jefferson@futurenet.co.uk

Production editor
Richard Smedley
richard.smedley@futurenet.co.uk

Contributors
Warren Brown, Greg Ungerer

Photography
Superstock, M-Systems, PhotoDisc

Publisher
Dave Taylor

ADVERTISING

Portfolio ad manager
Matt Stanley
matt.stanley@futurenet.co.uk

Portfolio deputy ad manager
Diane Clydesdale
diane.clydesdale@futurenet.co.uk

Business development manager
Damian Hughes
damian.hughes@futurenet.co.uk

Account managers
Zoe Drew,
Giles Crosthwaite Scott

Ad sales executives
Jo Doyle, Stephen Hall

LinuxPro is a part of
LINUX FORMAT
magazine

Published by
the future network
Media with passion

30 Monmouth Street,
Bath BA1 2BW
Tel +44(0) 1225 442244

Welcome

Twenty-four pages of real-world Linux for IT professionals

Welcome once again to *Linux Pro*. This time we kick off with a more detailed look at the current *lingua franca* of web services, SOAP. You'll find web services are easier to implement and potentially more useful than you might have previously expected.

Laziness is a virtue, for without the urge to limit effort, we would never come up with ideas like remote system administration. If you spend more time travelling to your servers than servicing them, you should take a read of our cover feature this month, which gives a brief introduction to some of the tools available and what you should know before you start saving yourself some time.

In a further embedded episode this issue, we'll be looking at how developers can free themselves from bulky, power-hungry and noisy disk drives by designing around Flash memory and EPROMs. This is how the big boys do it, and it's probably a lot easier than you think. Our thanks to SnapGear for sharing their thoughts on this very interesting topic.

Thanks again for all your feedback over the past few months, but do keep you comments and suggestions coming.

Nick Veitch Editor
nick.veitch@futurenet.co.uk

Contents

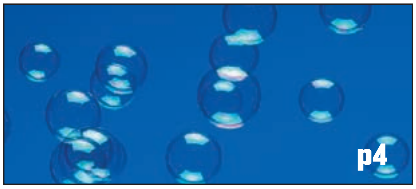

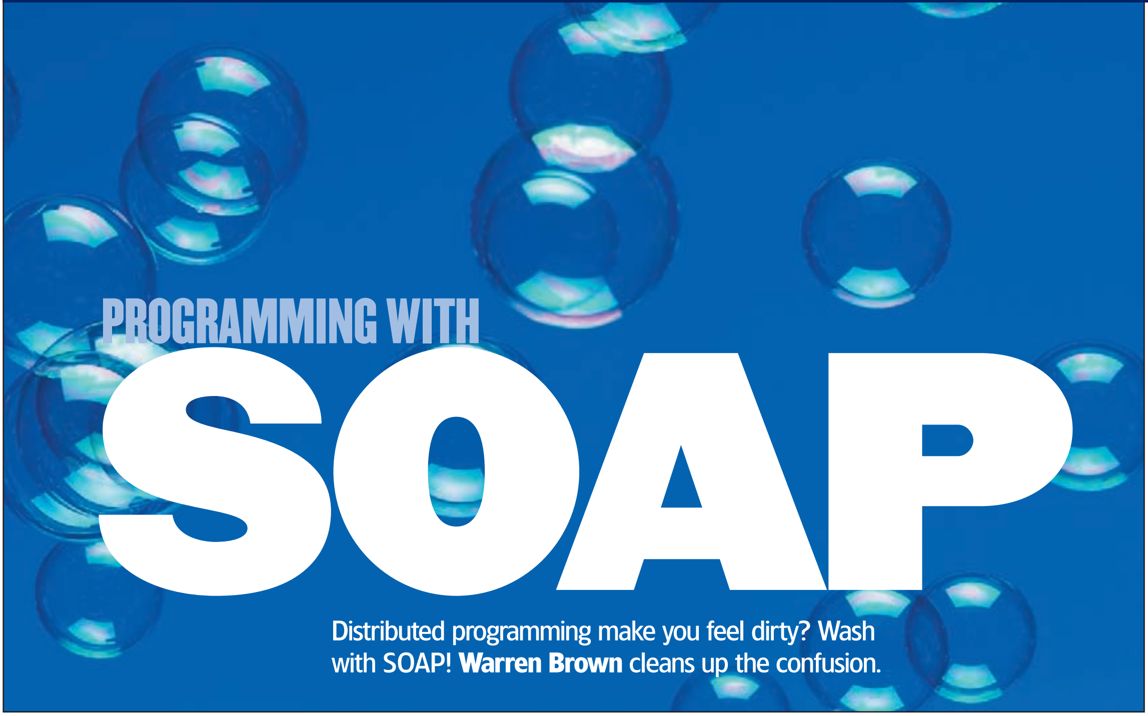

PROGRAMMING WITH SOAP

Distributed programming make you feel dirty? Wash with SOAP! **Warren Brown** cleans up the confusion.

SOAP is a lightweight protocol for exchange of information in a decentralised, distributed environment – well, that's the definition of SOAP from www.w3.org. Say it out loud in pubs and clubs, and see what response you get. (Save yourself by saying you're with *Candid Camera*!)

Basically, you can write a program 'client' that can get information from another program 'server' somewhere else on the network. "Ah! network programming" I hear you say. Similar, but with SOAP and other technologies such as Java's (RMI) and Corba you can actually use objects or call functions that reside on the server, and use the information they return.

For example, you can have a server that if you send it a postscript file, will return the document as a text file. Now try doing that with ftp!

So, why not use CORBA or Java?

The nice thing about SOAP is that it is firewall friendly, quite a nice thing if you don't like talking to firewall engineers, begging them to open ports through the firewall, compiling lengthy documents on how safe it is to open that particular port. SOAP can act as a CGI script or can be told to use a particular port. The nice thing about having it run as a CGI script, is that port 80 (http) is usually open on firewalls to allow web traffic through. Therefore your SOAP service is allowed through as well – pretty cool, I hear you say.

(There is, however, a quite reasonable case to be made that firewall avoidance is a bad thing. Indeed many in the SOAP community would object quite strongly to the notion that SOAP uses http merely to sneak through the already open port 80. Current debate around the reuse of http as a substrate for other applications-level protocols is summarised in rfc3205 **http://ietf.org/rfc/rfc3205.txt**. Of course with so many applications now working over http, it is becoming less sound to base your firewall/security policy on port numbers, anyway.)

So now you can have your SOAP service that drops you to a root shell on the server that you want to infiltrate! (Only kidding Mr Firewall engineer!)

The key word with soap is *service*. A service is, well, a service. You can have a service that translates Dutch to English, Google has a search service, Amazon.com has a book finder service. Just go to http://xmethods.com and check out some of the services you can use.

Let's get dirty with soap

To write programs with soap, you will need to get a few things:

gSOAP from
www.cs.fsu.edu/~engelen/soap.html
Java Development Kit (*i.e.* *jdk1.3*) from
http://java.sun.com
Apache Xerces 2 >= 2.0.1 jar files from
http://xml.apache.org (The .bin tar.gz)
And possibly Perl

To install *gSOAP* do the following:
Unzip the *gSOAP* package into a directory of your choice.

Go into the gsoap directory which is usually called soapcpp-linux, and copy *soapcpp2* to one of the bin directories *i.e.* /bin /usr/bin /usr/local/bin. *soapcpp2* is the soap preprocessor that does all the cool things to your files.

Next we need to compile the *stdsoap* files.

```
g++ -c stdsoap2.c
cp stdsoap2.o /usr/lib/stdsoap2_c.a
g++ -c stdsoap2.cpp
cp stdsoap2.o /usr/lib/stdsoap2_cc.a
```

That will compile the source code we need for our programs, and stick it in /usr/lib. You can put it wherever you want, however I find /usr/lib to be the most convenient.

The first one is for C programming and the second is for C++ programming.
Now do the following:

```
cp stdsoap2.h /usr/include
```

The header files for the *stdsoap* source code. You'll need it later.

Now for the *Apache Xerces* part. Install the JDK – mine installed itself into /usr/java.

Untar the *Apache Xerces* files into the lib directory for Java. *i.e.* /usr/java/lib It should create a directory called xerces Mine is /usr/java/lib/xerces.

If it doesn't create a directory then just place the jar files into /usr/java/lib/xerces.

That's all you have to do for that part. Now go back to the soapcpp-linux directory you installed previously.

WSDLCPP installation

WSDLCPP is really cool. It takes .wsdl files and converts them into .c and .h files. Now if you go to http://xmethods.com, you can download some .wsdl files and try it out. It helps if you install it first!

Go into the wsdlcpp directory of *gSOAP i.e.* soapcpp-linux-version/wsdlcpp. Inside you will find some java files. Typically *NodeSearch.java*, *NsNodeSearch.java*, *NsNode.java*, and *wsdlcpp.java*.

Make sure your classpath contains path information to these java files. You can try the following bash shell commands if you have trouble.

```
for x in `ls /usr/java/lib/*.jar;
 ls /usr/java/lib/xerces/*.jar`;
 do export CLASSPATH=$x:$CLASSPATH;
 done
```

Type the following to compile the *wsdlcpp* compiler

```
javac *.java
jar -cvf wsdlcpp.jar *.class
cp wsdlcpp.jar /usr/java/lib
```

And that's *gSOAP* installed.

Taming the slippery beast

Now that we've gone through the hard work of installing *SOAP*, it seems only fair to give it a go!

The simplest thing to do then is to write a 'client' that uses *SOAP*. It will be a bit difficult to test seeing as we don't have a server, but that's where **http://xmethods.com** comes to the rescue, they have tons of servers that we can use.

So point your browser at **http://xmethods.com**, click through to the *full list*, and choose Lotto. You never know, it could give you the winning numbers. From the Lotto web page right click on **http://reto.checkit.ch/Scripts/Lotto.dll/wsdl/IgetNumbers** and choose 'save as'. Click 'save' to save it to a directory of your choice.

I saved mine in /tmp/stuff and it was called *IgetNumbers.xml*. In actual fact it is a .wsdl file which we can use to create a client from.

Now we need to work on the file we downloaded, so find the file you downloaded and follow the yellow brick road!

```
java wsdlcpp IgetNumbers.xml
```

This should create two files, *IgetNumbers.h* and *IgetNumbers.c*.

Open up *IgetNumbers.c* and see what it contains!

It should contain the following:

```
#include “soapH.h”
#include
 “soapIgetNumbersservice.nsmap”
main()
{
        struct soap soap;
        soap_init(&soap);


        if (soap_call_tns__getLottoNumbers
 ( &soap, “http://reto.checkit.ch/Scripts/
Lotto.dll/soap/IgetNumbers”,
 “urn:getNumbersIntf-IgetNumbers#
getLottoNumbers”,/* xsd__int Min,
 xsd__int Max, xsd__int Numbers,
 xsd__int count,  struct
 tns__getLottoNumbersResponse
 {xsd__string _return_;  } *out*/))
                soap_print_fault(&soap,stderr);


}
```

As you can see the *wsdlcpp* program has done quite a lot of work of making this easy for us.

We just need to create the *SOAP* serialisers and we are almost finished with our client. Type the following:

```
soapcpp2 IgetNumbers.h
```

As you can see, this produces quite a lot of files, some of which we don’t really need at this time. Reopen *IgetNumbers.c* and change it to the following.

```
#include “soapH.h”
#include
 “soapIgetNumbersservice.nsmap”
#include <iostream.h>

int main()
{
        struct soap soap;
        soap_init(&soap);
        // If a proxy server gets in the way
        // uncomment the following
        /* soap.proxy_host = :insert proxy
 host here”;
           soap.proxy_port = insert proxy port
 here;         */
        struct
 tns__getLottoNumbersResponse r;
        xsd__int Min=1;
        xsd__int Max=49;
        xsd__int Numbers=6;
        xsd__int count=1;

        if (soap_call_tns__getLottoNumbers
 ( &soap, “http://reto.checkit.ch/Scripts/
Lotto.dll/soap/IgetNumbers”,
 “urn:getNumbersIntf-IgetNumbers#
getLottoNumbers”,Min, Max, Numbers,
 count,  &r )==SOAP_OK)
                cout << “Lucky numbers this
 week are “ << r._return_ << endl;
        else
                soap_print_fault(&soap,stderr);
        return 0;
}
```

As you can see from the previous example. **struct tns__getLottoNumbersResponse**, **xsd__int Min**, **xsd__int Max**, **xsd__int Numbers** and **xsd_int count** were all commented out. All you have to do is declare them in your code, and use them in the calling function. Modify the return statement to check for **SOAP_OK**, which means that everything went well, and print out the result, which in this case is 6 numbers from 1 to 49.

To compile the code into a client just do the following:

```
g++ -o client IgetNumbers.c
 soapClient.cpp soapC.cpp
 /usr/lib/stdsoap2_cc.a -I.
./client
Lucky numbers this week are
 4,27,47,34,9,44
```

If you have a proxy server in the way, then you will need to uncomment the proxy code above.

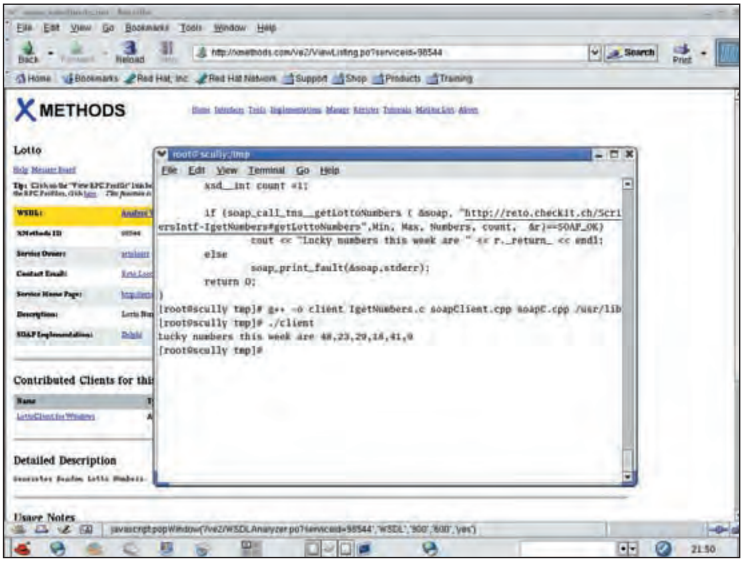
XMETHODS
Lotto
root@scully:/tmp
File Edit View Terminal Go Help
Contributed Clients for this
Detailed Description
Usage Notes
http://xmethods.com/ve2/ViewListing.po?serviceid=98544
21:50

```
        xsd__int count =1;

        if (soap_call_tns__getLottoNumbers ( &soap, "http://reto.checkit.ch/Scri
ersIntf-IgetNumbers#getLottoNumbers",Min, Max, Numbers, count,  &r)==SOAP_OK)
                cout << "Lucky numbers this week are " << r._return_ << endl;
        else
                soap_print_fault(&soap,stderr);
        return 0;
}
[root@scully tmp]# g++ -o client IgetNumbers.c soapClient.cpp soapC.cpp /usr/lib
[root@scully tmp]# ./client
Lucky numbers this week are 48,23,29,18,41,8
[root@scully tmp]#
```

The finished Lotto client in action.

There you have it, a fully working albeit simple soap client. I don't know about you, but whoa, I'm impressed.

Thankfully this soapie ends well

OK, so we can write clients. (Or do we just help out?) What about writing a server?

Well Watson me old matie, that's where deductive reasoning, and a good mastery of *jeet-kune do* comes in. Well maybe not!

To write a server we basically do things in reverse.

First of all we need a header file, and an idea for a service of course. Let's have a service that gives fortunes. Here's what out header file (fortune.h) looks like.

```
typedef char* xsd__string;
int bc__fortune(void *_, xsd__string
 *fortune);
```

gSOAP need the **void** *_ because we aren't sending anything to the client. The **typedef** helps *gSOAP* to map **char** to an **xsd__string** properly.

The **bc** in front of the function is for namespace issues, and it makes *soapcpp2* write a .wsdl file.

Now do the following

```
soapcpp2 fortune.h
```

As before, you will get a lot of files. Impress your friends and tell them you wrote them yourself.

Two important files are *bc.nsmap* and *bc.wsdl*. The *bc.wsdl* file is what you publish for other people to create clients from.

Let's have a look inside the *bc.wsdl* file.

```
<?xml version="1.0" encoding="UTF-8"?>
<definitions name="%{Service}%"
 xmlns="http://schemas.xmlsoap.org/wsdl/"
 xmlns:SOAP-ENV="http://
schemas.xmlsoap.org/soap/envelope/"
 xmlns:SOAP-ENC="http://
schemas.xmlsoap.org/soap/encoding/"
 xmlns:SOAP="http://
schemas.xmlsoap.org/wsdl/soap/"
 xmlns:WSDL="http://
schemas.xmlsoap.org/wsdl/"
 xmlns:xsd="http://
www.w3.org/2001/XMLSchema"
 targetNamespace=
"%{URL}%/%{Service}%.wsdl"
 xmlns:tns="%{URL}%/%{Service}%.wsdl"
 xmlns:bc="%{URL}%/bc.xsd">
...
```

Those tags **%{...}%** in the code need to be changed. To change them you can use Perl (Don't you just love Perl!)

»

I'm going to put my server in http://localhost/cgi-bin/fortune.cgi. Here are the changes I need to make

```
perl -p -i -e ‘s/%{Service}%/fortune/g’ *
perl -p -i -e
 ‘s/%{URL}%/http:\/\/localhost\/cgi-bin/g’ *
perl -p -i -e ‘s/%{URI}%/urn:bc_fortune/g’ *
```

That will set up our nsmap and wsdl file quite nicely. The documentation for *gSOAP* says that you can do the substitution with commands in the header file (chapter 12), but I couldn't get it to work and Perl does it fine for me.

Now we need to write the code to implement the server. Follow me men (sorry, people)!

fortune.cc

```
#include <sys/wait.h>
#include <unistd.h>
#include <stdlib.h>
#include “soapH.h”
#include “bc.nsmap”
int main()
{
	soap_serve(soap_new());
}

int bc__fortune(struct soap *soap,void *_,
 xsd__string *fortune)
{
	FILE *bc;
	int result;
	*fortune = (char*)soap_malloc(soap,
 8192);
	bc = popen(“/usr/games/fortune”,”r”);
	if(!bc)
	{
		perror(“popen”);
		return SOAP_FAULT;
	}
	while((fgets(*fortune,8192,bc)));
	result = pclose(bc);
	if(!WIFEXITED(result))
		return SOAP_FAULT;

	return SOAP_OK;
}
```

As you can see from the code the main part is pretty small. When a connection comes in, the **soap_server** routine grabs it and calls the correct function, in this case there is only one function to be called *i.e.* **bc__fortune**.

Our function prototype and actual function name differ slightly. We need to include the soap structure, this is a bit like **this** in C++ programming.

The function's main task is to capture the fortune program's output and return it to the client, we do this using the **popen** command. You can find out more about this command from a very good book called *Linux Application development* (Michael K. Johnson and Erik W. Troan, Addison-Wesley Pub Co; ISBN: 0201308215). If the program fails at all we tell the client by sending a **SOAP_FAULT**, but if everything went well, then we send **SOAP_OK** and our fortune output.

Let's compile the code and see what it does:

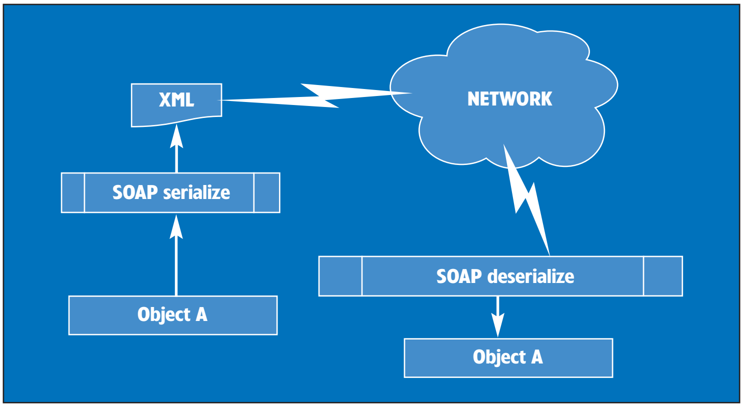

That distributed environment in action.

```
g++ -o fortune.cgi fortune.cc soapC.cpp
 soapServer.cpp /usr/lib/stdsoap2_cc.a -I.
cp fortune.cgi /var/www/cgi-bin
```

On my machine running Red Hat 7.3 the cgi-bin is located at /var/www/cgi-bin.

We don't need to do anything else! No message looping or **ORB->run** commands. We are using **Apache** to do our dirty work for us.

Me server, you client

With no further ado, a client must be present to run our cool new server. Let's write one, because we know how.

Make another directory, call it fortune_client and copy the *bc.wsdl* file into it. Go to that directory and type the following. (I bet you know what it is?)

```
java wsdlcpp bc.wsdl
soapcpp2 bc.h
```

If all goes well, then you should have a pile of files in your new directory. Open up the *bc.c* file, you should see the following:

```
#include “soapH.h”
#include “soapfortune.nsmap”
main()
{
        struct soap soap;
        soap_init(&soap);


        if (soap_call_tns__fortune ( &soap,
 “http://localhost/cgi-bin/fortune.cgi”,
 “urn:bc_fortune#fortune”,/* struct
 tns__fortuneResponse {xsd__string
 _fortune;  } *out*/))
                soap_print_fault(&soap,stderr);


}
```

Now we just need to modify it a bit. I prepared one earlier, so let's take a look at it.

```
#include “soapH.h”
#include “soapfortune.nsmap”
#include <iostream.h>

int main()
{
        struct soap soap;
        soap_init(&soap);
        struct tns__fortuneResponse r;
        if (soap_call_tns__fortune ( &soap,
“http://localhost/cgi-bin/fortune.cgi”,
“urn:bc_fortune#fortune”,&r)==SOAP_OK)
                cout << “Fortune for today is:\n”
<< r._fortune << endl;
        else
                soap_print_fault(&soap,stderr);
        soap_end(&soap);
        return 0;
}
```

As previously I've changed one or two things.

```
g++ -o client bc.c soapC.cpp
 soapClient.cpp /usr/lib/stdsoap2_cc.a -I.
./client
Fortune for today is:
The only difference between a rut and a
 grave is their dimensions.
```

You can also write *SOAP* programs that use specific ports, just check the documentation for the minor changes you need to make.

Read the docs

That about wraps it up for this simple tutorial on *SOAP*. For more information regarding this remarkable protocol. Go to the *gSOAP* web page, get the documentation (**www.cs.fsu.edu/~engelen/soapdoc2.html**) and join the mailing list.

Just remember when in doubt ‘Use the source, Luke!’ ■

SOAP Message Embedded in HTTP Request Example

```
HTTP/1.1 200 OK
Content-Type: text/xml; charset=”utf-8”
Content-Length: nnnn

<SOAP-ENV:Envelope
  xmlns:SOAP-ENV=”http://schemas.xmlsoap.org/soap/envelope/”
  SOAP-
ENV:encodingStyle=”http://schemas.xmlsoap.org/soap/encoding/”/>
   <SOAP-ENV:Body>
       <m:GetStockPrice xmlns:m=”http://stocks.com”>
           <Price>42.34</Price>
       </m:GetStockPrice>
   </SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```