LXF42 Linux Format July 2003 ● XINIT SHARQ ● ALL ABOUT X ● WINE ● STEGANOGRAPHY ● DISTRIBUTED COMPUTING ● MYSQL 4.0.4 IN DEPTH ● VISUAL SLICKEDIT ● DHCP ● BLENDER ●

■ 7 PAGES OF EXPERT ADVICE ■ DUAL-XEON SERVER ON TEST ■ STEGANOGRAPHY SECRETS

# LINUX FORMAT

www.linuxformat.co.uk

**PROBLEMS?**
Not any more! Advice from the experts on: ● Hardware ● Debian ● Hotplugging ● Digital cameras ● Filesystems ● Ogg ● and more
**Page 90**

# ALL ABOUT X

**INSIDE XFREE86:** ● What is X? ● New features ● Future of X ● X alternatives ● Drivers and fonts ● Developer spats **p44**

## USE WINE TO RUN WINDOWS SOFTWARE

Use MS Office, Internet Explorer and even Battlefield 1942 from linux with WINE **p52**

LINUX PRO

FROM THE MAKERS OF LINUX FORMAT

JULY 2003

ONLY AVAILABLE WITH LINUX MAGAZINE

Grow your enterprise through open source

Novell: "freeing users from their proprietary shackles" – the new Netware strategy p4

Oracle: "We see the growth in Linux business as being fundamental to the Oracle economy" – unbreakable Linux p14

PLUS

LICENSING
Software licences explored and explained – are you using the right one?

FLEX AND BISON
Exploring [illegible] variables and language structure

STORAGE
The NFS network filesystem is far from obsolete – how you can use it and what the future holds

SYSTEM MONITORING
NetSaint is dead, long live Nagios, the new king of open source monitoring

## HOWTO:

| | |
|---|---|
| Choose the perfect distro | **p68** |
| Configure DHCP | **p70** |
| Manipulate files with PHP | **p76** |
| Use Python packages | **p80** |
| Light scenes in Blender 3D | **p84** |

## SHARQ ATTACK!

Xinit's powerhouse dual-Xeon server put to the test **p30**

## HIDDEN IN PLAIN VIEW

The secrets of steganography **p58**

## DISTRIBUTED COMPUTING

How projects like Seti@Home and Chessbrain challenge the notion of free clock cycles **p62**

## MYSQL 4.0.12 IN DEPTH

Bigger, better, faster, more **p28**

DVD issue also available
Printed in the UK
LXF42 JULY 2003
£5.99
the future network
Media with passion
07
9 771470 423019
(£5.99 outside UK and Republic of Ireland)

THE UK'S BEST-SELLING LINUX MAGAZINE

# The freedom dividend

Open Source software is great. One of the greatest benefits of having an open development model is that everyone can have opinions, and make changes if they so desire. In many ways this is ably demonstrated by the recent spat involving the core developers of XFree86. As a project, XFree exemplifies the open advantage – with it's roots in a time when a morass of commercial software each had their own take on the X11 'standard', XFree86 emerged as the reference platform for implementation of X11R6. It has been so successful that very few commercial competitors still exists, and those that do are generally aimed at more specific markets. You can find out a lot more about X, and why developer infighting can actually be a good thing, by reading our cover feature this issue.

What about data that you don't want to be free? The best place to hide a tree is in a forest. It seems the best place to hide secret text may well be in an image, or perhaps even in some spam email. For more on the intriguing world of cryptography, see our mini-feature on page 58.

There are many reasons why people may want or need to run proprietary software from other platforms. In the case of people migrating from Windows, they usually end up with dual boot systems so they can still use that particular finance package, game or network application. But there are alternatives. WINE has come on in leaps and bounds since the last time we covered it, and coupled with commercial versions such as WineX and *Crossover Office*, more Windows apps than ever will now run happily within Linux. Catch up with the state of play on page 52.

In spite of the searing heat inside *LXF Towers*, we have also toiled and sweated into the small hours to bring you reviews of the hottest new hardware and software (definitely check out the books section and the Sharq server) as well as dozens of pages of tutorials. If you have an idea for something you'd like to see here, just mail me or visit the forums at **www.linuxformat.co.uk**.

Nick

**Nick Veitch** EDITOR

Everything you ever wanted to know about X... and then some... p44

Which WINE? Our connoisseurs guide starts here p52

Is that really spam, or is someone trying to pass you secret documents? Steganography decrypted p58

## AIMS OF THE MAGAZINE

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

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

## MEET *LINUX FORMAT'S* TEAM OF WRITERS...

**Andrew Channelle**
Fast becoming the 'install guy' of the team, Andy now shares his distro experiences in a new Beginners' tutorial.

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

**Richard Drummond**
Indiana is a flat and boring place, so Rich has no distractions from writing for us!

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

**Paul Hudson**
Former web designer, Paul is discovering that journalism isn't all about long lunches. Well, not all the time.

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

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

**Mike Saunders**
*LXF* veteran, Mike is top of the popular software pickers again as he hosts our Hot Picks section.

**Patrick O'Brien**
Published by O'Reilly and IBM, the author of the new *Py* module has a *deep* affection for Python scripting.

**Amias Channer**
BBLUG stalwart, Perl devotee and developer, our local University calls him when its computer systems throw a wobbly.

## CONTACT US

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

Subscriptions/back issues:
**subs@futurenet.co.uk**

Technical help/Ask the Experts:
**lxf.answers@futurenet.co.uk**

Disc problems:
**lxf.support@futurenet.co.uk**

General enquiries:
**linuxformat@futurenet.co.uk**

Website: **www.linuxformat.co.uk**
Or send your letters to:
***LINUX Format*, Future Publishing,**
**30 Monmouth Street, Bath, BA1 2BW**
**Phone:** 01225 442244
**Fax:** 01225 732295

More contact info on p114

LXF42 July 2003

# Contents

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

■ 7 PAGES OF EXPERT ADVICE ■ DUAL-XEON SERVER ON TEST ■ STEGANOGRAPHY SECRETS

PROBLEMS? Not any more! Advice from the experts on: ● Hardware ● Debian ● Hotplugging ● Digital cameras ● Filesystems ● Ogg ● and more Page 90

LINUX FORMAT www.linuxformat.co.uk

ALL ABOUT X

INSIDE XFREE86: ● What is X? ● New features ● Future of X ● X alternatives ● Drivers and fonts ● Developer spats p44

USE WINE TO RUN WINDOWS SOFTWARE Use MS Office, Internet Explorer and even Battlefield 1942 from Linux with WINE p52

HOWTO: Choose the perfect distro p68 Configure DHCP p70 Manipulate files with PHP p76 Use Python packages p80 Light scenes in Blender 3D p84

MYSQL 4.0.12 IN DEPTH Bigger, better, faster, more p28

SHARQ ATTACK! Xinit's powerhouse dual-Xeon server put to the test p30

HIDDEN IN PLAIN VIEW The secrets of steganography p58

DISTRIBUTED COMPUTING How projects like Seti@Home and Chessbrain challenge the notion of free clock cycles p62

DVD issue also available Printed in the UK LXF42 JULY 2003 £5.99 $15.99 outside UK and Republic of Ireland

THE UK'S BEST-SELLING LINUX MAGAZINE

LINUXPRO
Grow your enterprise through open source
PLUS LICENSING FLEX AND BISON STORAGE SYSTEM MONITORING

## COVER HIGHLIGHTS

## ALL ABOUT X 44

Can twenty-year-old technology still cut it at the heart of the Linux OS?

**28 MySQL 4.0.12 review**
Love or hate it, it's the most popular Open Source DB

**30 Sharq attack!**
If 1U isn't enough, try this penny-pinching 2U

**52 Use WINE to run Windows apps**
Use MS Office, Internet Explorer and play games!

**58 Hidden in plain view**
Steganography and the art of concealment

**62 Distributed Computing**
Our *What On Earth...* feature has everything you need

30 Xinit "supports Linux without question" – and it's a bargain!

> 'One in every 199 UK emails was classified as spam in Jan 2002. By June 2002, it was one in 36; and now it stands below one in five!'

**40** Why we all need spam-filtering apps

> 'MPlayer means Linux users don't have to reboot into Windows or mess with WINE just to watch certain video files.'

**38** It won't help 'DVD Jon' dodge jail though

## LINUXPRO

**Novell** Grow business by adopting the Open Source ethos **Flex and Bison** Compiler writing **Oracle** Unbreakable Linux **Licence to code** Which licence means what? **Storage** NFS **System monitoring** Nagios

LINUXPRO
Grow your enterprise through open source
PLUS LICENSING FLEX AND BISON STORAGE SYSTEM MONITORING

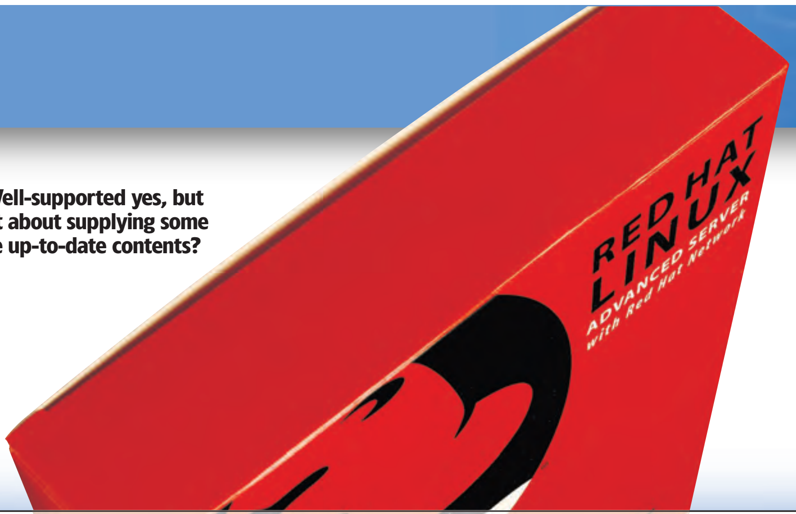

18 Well-supported yes, but what about supplying some more up-to-date contents?

**06 News**
Analysis of the latest Linux stories

**12 Mailserver**
Share your opinions with everyone!

**18 Red Hat Enterprise Linux**
RH for the enterprise

**20 GCC 3.3 review**
Compiler takes a solid step forward

**21 Java 2 Std Edition 1.4.2 review**
New functionality and speed boost

**22 Seapine TestTrack Pro 5.1.1 review**
Got bugs? Track 'em down!

**25 VisualSlickEdit 8.0 review**
Pro alternative to Open Source IDEs

**32 Neoteris Access 3000 review**
The pinnacle of 'set and forget'

**34 Book Reviews**
Treeware both terrific and terrible

**38 HotPicks**
The best and latest Linux apps

**68 Beginners' tutorial**
Part 1: Choosing a Linux distro

**70 DHCP tutorial**
Bread-and-butter installing tips

**72 GIMP tutorial**
Wash the colour back in!

**76 PHP tutorial**
Writing PHP extensions

**80 Python tutorial**
PyCrust is dead! Long live Py!

**84 Blender tutorial**
Lighting your scenes

**88 Runtime Revolution tutorial**
Final part of our mini-series

**90 Answers**
Our experts answer your queries

**110 User Groups**
Local support and Linux advocacy

**114 Next month**
August in the crystal ball...

25

52

32

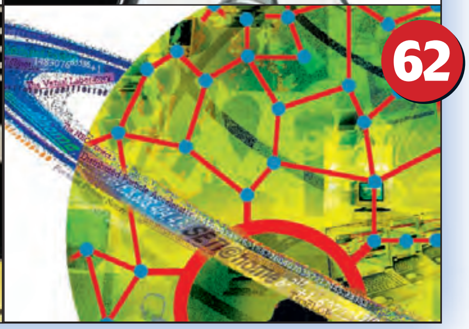
62

## COVERDISCS

A DVD or 2 CDs packed full of the latest Linux goodies 100

### CDS A AND B

**Morphix** modular Live CD distro; **GuardDog** user-friendly firewall; **NVidia XDrivers** latest and fastest drivers; **HDDtemp** keep your PC interior cool; **WINE** wave goodbye to dual boot; **GCC 3.3** compiler at the heart of Linux gets upgrade; **SETI** "...pray that there's intelligent life somewhere up in space 'cause there's not much down here on Earth..."; **TORCS** car racing; **Linux Cookbook** 1,500 tips and tricks

### DVD

**Opera 7.1 beta** browser; **GAIM** Instant Messaging Magic! **ABCDE** burn your discs; **KDE-Sounds/ Aqua Fusion** new visual and audio KDE themes; **Pingus** Lemmings-esque game

On the CD and DVD

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

# Newsdesk

● IBM touts business and home Linux ● Linux on iPod ● Athenyx – the X11 alternative ● Opera 7.1.1 Lycoris GamePak launched ● Thai govt subsidises Linux hardware ● Ardour audio studio ● Win4Lin

IBM VS SCO CONTINUED...

## Who is SCO's next target?

After launching its stunning $1 billion lawsuit against IBM (see *LXF40*), SCO seems to have set its sights on alienating the entire Linux community, and a large proportion of its own customer-base. Despite earlier assurances that the billion dollar lawsuit against IBM would not impact on Linux users, SCO's Darl McBride has gone on the offensive, now claiming that significant portions of the company's UNIX code has illegally found its way into Linux. In a letter to commercial SCO Linux users, McBride said that there was "evidence that portions of UNIX System V software code have been copied into Linux." He accused "unrelated and unknown software developers" of deliberately trying to obfuscate their original source by "copying and modifying" SCO code.

"We believe that Linux infringes on our UNIX intellectual property and other rights. We intend to aggressively protect and enforce these rights," McBride wrote, citing recent efforts to close down file sharing services as a model for future actions. "Similar to analogous efforts underway in the music industry, we are prepared to take all actions necessary to stop the ongoing violation of our intellectual property or other rights." Alongside the letter sent to customers, McBride also announced the closure of SCO's Linux operations and a plan to reveal the 'stolen' code, though this would be done under strict non-disclosure agreements (NDA).

The company then went further, claiming in a press release the Linux was in fact an unauthorised derivative of UNIX and that "legal liability for the use of Linux may extend to commercial users."

Vice President Chris Sontag claimed there were intellectual property issues with Linux. "When SCO's own UNIX software code is being illegally copied into Linux, we believe we have an obligation to educate commercial users of the potential liability that could rest with them."

### Whose code is it?

However, SCO's claims were substantially undermined by Novell, from whom SCO originally bought the 'contractual' rights to UNIX in 1995.

In a letter to McBride, Novell CEO Jack Messman expressed his support for Linux and the Linux community, and pointed out that Novell, not SCO, still own the intellectual property rights associated with UNIX System V: "We believe it unlikely that SCO can

### Links

**Open Source Initiative response by Eric Raymond: www.opensource.org/sco-vs-ibm.html**
**Latest developments from SCO: www.sco.com/scosource/**
**Updates from *Linux Format*: www.linuxformat.co.uk**

BUSINESS WANTS ENTERPRISE LINUX

# Peoplesoft respond to customers' demands

**Constant requests from customers** have convinced PeopleSoft to port its entire range of enterprise applications to Linux. The 170 products should be available by the end of the year. Spokesman David Sayed said many of the company's customers had used Linux for roles such as file and web serving and other infrastructure roles, but now there was a definite desire to integrate the OS into more general enterprise tasks.

The distribution of choice for PeopleSoft will be *Red Hat Advanced Server*, while a hardware tie-in with IBM will see both operations pushing PeopleSoft's range of Linux applications on xSeries servers, *DB2 Universal Database* and *Websphere Application Server*.

PeopleSoft CTO Rick Bergquist said Linux was now "ready for prime-time" and was more than capable of running mission-critical applications. "We're teaming with IBM to deliver PeopleSoft applications in a Linux environment, which gives our customers greater choice and is another milestone in our commitment to open standards."

demonstrate it has any ownership interest whatsoever in those copyrights".

## Action reaction

Open Source advocate Bruce Perens said that discontinuing their Linux distro was a non-event as they had already 'killed their business.' "No one will touch them," he said. "Linux customers can easily go elsewhere, and Unix customers are going to Linux, since what they see coming out of SCO is just nuts." Highlighting a recent report which gave recent versions of Linux a 'better grade than low-end UNIX in most categories,' Perens said this was a last-ditch attempt by SCO to "wring every penny out of the business on their way out.

"So, let's talk about copying. Isn't it more likely that parts of the publicly-available GNU/Linux system were copied into the secret SCO source code, than the other way around? It's also interesting that SCO claims a lot of their damage is around the Itanium chip. IBM didn't do that part of the Linux kernel. The leading developers of that were at HP, and the folks who did it wrote a book about it. It's called *'IA64 Linux Kernel'* by David Mosberger and Stephane Erdman (foreword by Bruce Perens). In that book, David and Stephane explain what they wrote."

## Gates in a frame-up?

It seems a knee-jerk reaction – blame everything on Bill Gates! – but hot on the heels of the SCO's decision came a press release from Microsoft and SCO announcing that MS had bought a UNIX license.

Microsoft lawyer Brad Smith said acquiring the license was consistent with "Microsoft's ongoing commitment to respecting intellectual property and the IT community's healthy exchange of IP through licensing. This helps to ensure IP compliance across Microsoft solutions and supports our efforts around existing products like services for Unix that further Unix interoperability."

Despite claiming that no Linux vendor would be under the radar for its legal action, SCO won't be leaving the UnitedLinux Consortium. In response, SCO's main UL partner, SuSE, said the whole thing was "curious" and said despite requests for clarification, SCO had so far declined. "We are not aware, nor has SCO made any attempt to make us aware, of any specific unauthorised code in any SuSE Linux product. As a matter of policy, we have diligent processes for ensuring that appropriate licensing arrangements are in place for all code used in our products."

"Linux customers can go elsewhere, and Unix customers are going to Linux overwhelmingly since what they see coming out of SCO is just nuts" – Bruce Perens

Red Hat President Matthew Szulik said his company and its community development partners would take valid intellectual property rights seriously. "We respect and take effort to maintain the legal and technical integrity of valid intellectual property, including patents, copyright and trademark. When the integrity of the Red Hat brand is publicly called into question, we will defend the use of Red Hat Enterprise Linux by our customers." Users fearing fallout from the action could contact Red Hat's legal department for advice.

Bradley Kuhn, executive director of the Free Software Foundation said SCO appeared to be using scare tactics to drum up license fees, but that the action would backfire. Kuhn told Australia's SME.com that the FSF had spent nearly two decades "carefully and arduously collected copyright assignments on each contribution to the GPL'ed programs on which we hold copyright

"We carry out due diligence to ask contributors if they have any reason to believe that trade secrets, patents, or other copyright claims cover their work before they submit it to us. We then collect a copyright assignment from the contributor to ensure that we hold proper title to the software."

He added that those using FSF copyrighted programmes "know as much as one can know that the software has been examined carefully, that its authors certify that the work is their own, and that the authors have no knowledge of other claims conflicting with its licensing under GPL or LGPL."

Former SCO developer Christoph Hellwig said the possibility of 'retrofitting' UnixWare code into Linux was remote as the kernel internals were so different. "It might be more interesting to look for stolen Linux code in Unixware, I'd suggest with the support for a very well known Linux filesystem in the Linux compat addon product for UnixWare," he wrote. The case almost certainly continues...

**Jono Bacon**

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

COMMENT

## Safety in numbers

**"Personally I think this SCO vs IBM case will be forgotten when it has finished. Let us not forget that IBM are a huge company, and have probably got beefcake lawyers that are going to be truly thorough at every step of the way. SCO seem to me to be less likely to be successful in court due to the fact that they are (a) smaller (b) have less money and (c) their case is not 100% solid.**

**There are a number of issues that could make life difficult for SCO. Firstly, they seem to be trying to attack a number of Linux companies, issuing statements to them and basically stirring up the FUD. Aside from the technical claims that I see as not entirely clear, they also have the challenge of dealing with themselves. If IBM has misappropriated UNIX code into Linux, SCO itself is partially to blame as they distributed Linux code as well. Could this could invalidate their claim?**

**To me it seems quite clear how (in my personal opinion) SCO has approached this case. They used to be a Linux company, shouting about Linux and how great it is, then they saw the UNIX opportunity, took it, took on the many existing UNIX contracts and now put Linux on a back burner. They then saw Linux as a principle threat to their new source of revenue. Although a practical business decision, they have essentially turned their backs on their history (and the Linux community) and are just following the easier money. Although I find this disgraceful from a personal view, who knows; maybe IBM will buy SCO and then open-source the original UNIX code? Time will only tell..."**

LINUX CONSOLIDATION

# IBM goes for big and little iron

Big Blue is bolstering its Linux operations with launches at both ends of the IT sector. The company's mainframe business gains a new z Series server that improves on performance of its flagship z900 by a factor of three. Moreover, the eServer z990 introduces for the first time the ability to scale from one to 32 processors as demand dictates without the need for any downtime.

William Zeitler, IBM senior vice president, said the z990, code named T-Rex, was not just the most powerful server in the company's history, it also raised the bar for system builders. "This technology is critical to customers who rely on it to manage unpredictable, on demand environments where performance, security and reliability are so critical."

Like its predecessor, the z990 has the capacity to run hundreds of virtual Linux servers from a single machine giving 'data centre performance in a server the size of a fridge.'

At the other end of the market, the ThinkCentre range of small footprint desktop PCs replaces the aging NetVista range. Most interesting of all is the M50 which boasts SuSE and RedHat certification.
www.ibm.com

**The ThinkCentre M50 is IBM's first desktop to be Linux certified.**

## Linux Web Watch/

Xbox and Linux in car multimedia

**Just don't watch *The Matrix* on Xbox while you're driving!**

Linux on iPod

News

**Linux on iPod – there's more to life than music you know...**

handhelds.org

Welcome to handhelds.org!

What is handhelds.org?

News

**www.handhelds.org – serving the Linux/iPaq community since 1999.**

**Don't throw out your old PDA. Give it an OS transplant!**

# In praise of the curious

## Some people just ain't happy with what they've been given. Thank heavens!

Progress, George Bernard Shaw wrote, always depends on the actions of the unreasonable, those men and woman who persist in trying to adapt to them. This is true of many Linux hackers who won't settle for being told what a certain device is supposed to do.

Take, for instance, Jon Snowdon who took a modded Xbox, the Linux/Xbox distro and Freevo software to make a personal multimedia player for his car. The device's 40GB disk is stuffed with music, films and emulated games to while away those long traffic jams. He documented his efforts at www.target-earth.net/xbox/.

Bernard Leach had a smaller goal: to bring his favourite OS to his favourite music player. His efforts to Linux on an iPod has spawned its own project page on sourceforge (http://ipodlinux.sourceforge.net/index.shtml) and has so far got as far as adding Firewire support to the uClinux port that forms the basis of the distro.

It's not all-new stuff though. www.handhelds.org has been around for a while and is still dedicated to putting Linux on anything that will fit in your pocket. www.handhelds.org has a wealth of material available, especially for iPaq users with designs on adding Linux to their PDA.

Another old hand is PsiLinux (http://linux-7110.sourceforge.net/) which, as a project, recently managed to bring Linux to Psion's netBook machines. If you have an old Psion PDA lying around, the site also hosts distro for the Revo and 5MX among others.

## NEWSBYTES

■ Linux Distro vendor **Gentoo** has started a games division and their first release is self-booting CD of *America's Army*, the FPS developed by the US Military. Gentoo Games' stated aim is to bring cutting edge gaming to Linux. See **www.gentoo.org.**

■ When ***Mozilla's*** developers announced their plans to split the project into separate browser and mail elements (LXF, issue 41) fears were expressed about the fate of *Composer*, one of the few WYSIWYG HTML editors available on Linux. However, the project has been rescued by Dan Glazman, a *Mozilla* contributor and Netscape employee who said he didn't want to see such potential go to waste.

■ Talking of ***Mozilla***, its name-change problems continue with both *Firebird* and *Minotaur* being dropped as names due to trademark issues. The browser and mail client are currently living under the aliases *Mozilla Browser* and *Mozilla Mail Client*.

■ UK PC builder **Hi-Grade's** latest Xperian range of desktop PCs ship with Windows XP as standard, but a 'hidden' Linux distribution offers instant access to DVD/CD and MP3 playback, without having to 'boot' the machine. Linux was apparently chosen over a BIOS solution for its ability to access files on NTFS file systems.

■ **Lycoris** has finally launched its GamePak for Desktop/LX. The package is based on TransGaming's WineX implementation. More at **www.lycoris.com**

■ ***Denim*** is a new web designer which, literally, takes your scrawl efforts and turns them into a navigable website. Running on Linux, Windows and MacOS-X with appropriate JDK versions, *Denim* (and its *Silk* development system) has been designed for collaborative web development and can be downloaded from this link below: **http://guir.berkeley.edu/projects/denim/**

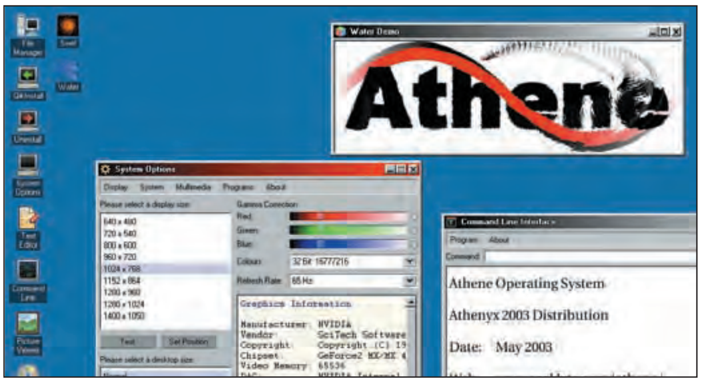

System Options

Display System Multimedia Programs About

Please select a display size:

640 x 480
720 x 540
800 x 600
960 x 720
1024 x 768
1152 x 864
1280 x 960
1280 x 1024
1400 x 1050

Gamma Correction: Red, Green, Blue

Colours: 32 Bit 16777216

Refresh Rate: 85 Hz

```
Graphics Information

Manufacturer:  NVIDIA
Vendor:        SciTech Software
Copyright:     Copyright (C) 19
Chipset:       GeForce2 MX/MX 4
Video Memory:  65536
```

Command Line Interface

Athene Operating System

Athenyx 2003 Distribution

Date: May 2003

**Athenyx can make systems look like Windows – your users'll never know!**

### NEW DISTRO

# Athene offers alternative to X11

**After the recent debates over the** merits or otherwise of X11, it is little wonder that the spotlight has been thrown on a number of potential successors. The most high-profile alternative so far comes from Rocklyte Systems and is built around the SciTech SNAP Graphics system and the Linux kernel. Rocklyte claims that Athenyx 2003 is 'one of the highest performing operating systems available today.'

Headline features include support for over 100 graphics chipsets (thanks to the SciTech core), TrueType font support, plug and play for most multimedia hardware and an XML-based scripting language for user interface development.

As well as being a full OS, a cut-down version of the Athene interface can be installed on top of any Linux distro and, in the near future, will feature an X11 compatibility mode. There is also a Windows edition available. The OS ships with a pair of interfaces, aping Windows and the classic Amiga UI. SciTech itself has a beta release of the latest Linux version of *SNAP Graphics* for download from **www.scitechsoft.com**. SNAP Graphics is a replacement for the graphics drivers within X and provides true plug-and-play support for most common chipsets. It is currently certified for latest releases of SuSE, MDK, RH, Slackware, Debian and TurboLinux. **www.rocklyte.com/athene.**

### NO PLAYING CATCH-UP!

# Opera 7.11 unites Linux and Windows versions

**Demonstrating its continued** commitment to Linux, Norway's OperaSoft has unleashed the latest Linux edition of its flagship browser in tandem with the Windows release. In addition to the roll call of features that Opera users have come to love, like tabbed windows, mouse gestures and a very skinnable interface, the company claims version 7.11 includes a number of unique features which have the potential to really alter the way you browse.

Major additions include *The Wand*, to simplify password logins; *FastForward*, which tries to anticipate where you want to go next; and a slideshow which takes a batch of photos from a website and displays them, in full screen if necessary, in a presentation. Most attention, however, will be focused on *M2*, Opera's integrated mail client.

**Opera's integrated mail client M2 proves to be very useable.**

## Hoyt Duff

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

### COMMENT

# 733t THEFT?

“Warez is nothing new, but the growth in Linux warez is alarming. With most Linux software provided free (as in beer), it's surprising that the OS sense of community (free as in speech) is being slowly eroded by those who trade in the exchange of non-free (as in beer or speech) apps; look at **alt.binaries.warez.linux** if you don't believe me. As of this writing, I see offered: *VMWare ESX server, Cold Fusion, Crossover Plugin, IBM DB2* and *Tivoli Administrator, Maya, SuSE OpenExchange* and *WineX* among others as well as license key generators. (It's telling of their mindset that some pirates also post software that can be downloaded legit for no charge.)

Speaking with Jeremy White of CodeWeavers about the blatant pirating of their product, he found it personally disappointing that someone in the community who found value in paying for the app would then have so little regard for his company's efforts as to do such a thing. I find it totally outrageous.

Why does this happen? Eloquent arguments abound, but some feel a moral imperative to steal just because they can or because they feel entitled. They twist the "free as in speech" mindset to their own crooked way and adopt a rebellious, Robin-Hood-like rationalisation to assuage their guilt and justify their resentment of those who "have".

That's just wrong. One can rationalise it any way it makes one feel comfortable, but in the final analysis it's just common theft. Stealing is wrong. Software license fraud is wrong. Using warez is wrong. The improper use of copyrighted property is wrong. We may assign differing degrees of wrongness to it, but it's *all* wrong no matter how one weasels around the fact. Wrong.”

GOVERNMENT-SUBSIDISED LINUX HARDWARE

# Laptops for the masses

The Thai government has contracted Hewlett Packard to build a low-cost Linux laptop with the intention of improving computer literacy among the population. The £285 devices are based around a 1GHz Celeron chip, 128MB RAM and a 20GB disk, but feature no removable media. A similarly specified desktop model is being built by a trio of domestic system manufacturers, Belta, SVOA and Computec and is priced at £159.

The cost of the devices has been heavily subsidised by the Thai government, and arrangements have also been made with the country's banking sector to offer preferential loans to buyers. The first consignment of 100,000 machines is expected to sell out quickly leading to the second release phase, which will see a further one million machine roll off production lines. Both machines will be preloaded with the Thai Language Edition of Linux.

Given the prevalence of counterfeit goods trading in Asia, so the decision not to fit media like CD or DVD burners to subsidised machines isn't surprising.

## Embedded Linux News

● **Sharp** is following up its successful Zaurus PDA with a pair of clamshell designs featuring upgraded hardware and an improved screen. The usual sales pattern – release machines in Japan, gauge reaction and then release elsewhere – is being adhered to, so don't expect to get your hands on a non-import version until next year. Both new machines run 400MHz Xscale processors from Intel with either 64 or 128MB RAM. Pricing has yet to be announced.

● **Sharp** has also launched a stand alone IC card reader (right) built on Linux. The 'Touch and Go' system is designed for entry systems or fare collection. The big advantage is that the system can be installed in areas where network infrastructure is not available, as the device is able to operate without a connection to a central server.

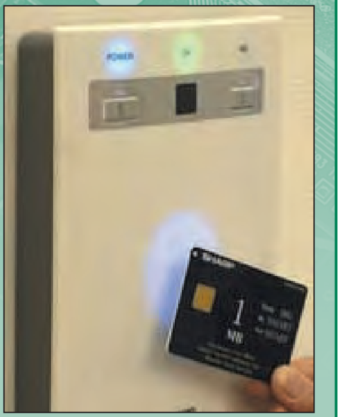

● **www.linuxdevices.com** has recently updated its list of cool Linux gadgets. The latest collection includes everything from bog standard PDAs to more esoteric fare such as Wakamaru, a 3.3 foot tall yellow robot designed to provide 'companionship or caretaker' functions. There is also a very useful list of pre-built devices suitable for personal embedded Linux projects. No talking toasters though.

● **NEC** is the latest mobile phone builder to announce its support for Linux as an OS for G3 phones. NEC is said to be considering using X as its phone windowing system instead of relying on a tailored solution such as Qt/embedded.

TUX TUNES!

## Ardour heads for official Linux music studio release

**One of the great hopes for Linux** as a genuinely useful music studio tool has entered feature freeze in preparation for its first "official" release after three years of work. Ardour has ambitions to be a professional quality multichannel, multi-track digital audio workstation for Linux and uses established technologies such as ALSA, JACK and LADSPA plugins. Developer Paul Davis says the application will support Steinberg's VST plugins as soon as the company allow redistribution of VST header files, though he's not holding his breath.

Ardour's featureset is certainly impressive: 16- and 24-bit recording, support for pro-level audio interfaces, non-destructive, non-linear editing, time compression and stretch and 'per track' speed variation, meaning individual tracks can be non-destructively time-shifted, and a fully featured mixing environment.

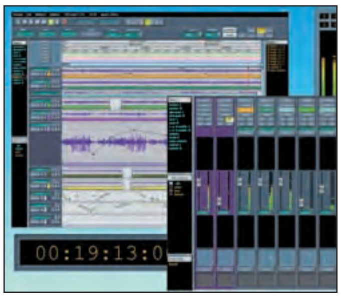

**Ardour's user interface mimics a traditional mixing environment.**

**http://ardour.sourceforge.net**

## NEWSBYTES

eCosCentric

Home
- Company
- eCos
- Services
- Support
- Testing
- Resources
- DevZone

News
FAQ
Contact
Search

eC
development ar

Staffed by the c
software engine
systems and co
technologies in
by professional

Unlike many oth

■ **eCos** has released the first version of its RealTime Operating System (RTOS) since splitting from Red Hat. ECos 2.0 is available on all major architectures and boasts support for 74 publicly available evaluation and development platforms. John Dellaway said the developers had work hard on making installation of eCos a lot simpler than previous versions. "All the tools needed to evaluate eCos are installed directly from the CD-ROM and need no further configuration." See **www.ecoscentric.com**

■ **Sendmail** has solicited support from both Hewlett Packard and Intel for an attempt to bring email to 'deskless' workers. Sendmail says that half of all workers in the average enterprise are not tied to a desk and organisations find it difficult to organise email access for this significant part of the workforce. The solution, which competes with MS Exchange Server and Lotus Notes from IBM, will be powered by the Carrier Grade Edition of SuSE Linux's Enterprise Server.

■ **Sun's** Scott McNealy told a conference audience that his company should have jumped on the Linux bandwagon two years ago and unambiguously said that he thought the future of Linux lay on the desktop. Not the server space...

■ **Citrix,** the company which developed the original Windows Terminal Server, has been moving its attention to Linux recently and has a pair of projects ready for exposure. First is the Citrix Secure Gateway, which provides secure access to the company's application hosting servers, while the second is a new version of its ICS client which enables users to run Microsoft Office on entry-level Linux PCs.

■ **Fujitsu** is increasing its Linux operations after revealing that Linux servers may account for 25 per cent of its business within three years. The group recently inked deals with both RedHat and SuSE to bring Advance Server and Enterprise Server respectively to Intel 32- and 64-bit system as well as RISC and Sparc based servers.

■ **Hewlett Packard** has become a preferred supplier of Red Hat's complete range of enterprise level servers. HP also announced a revenue of £1.3 billion on Linux operations for last year.

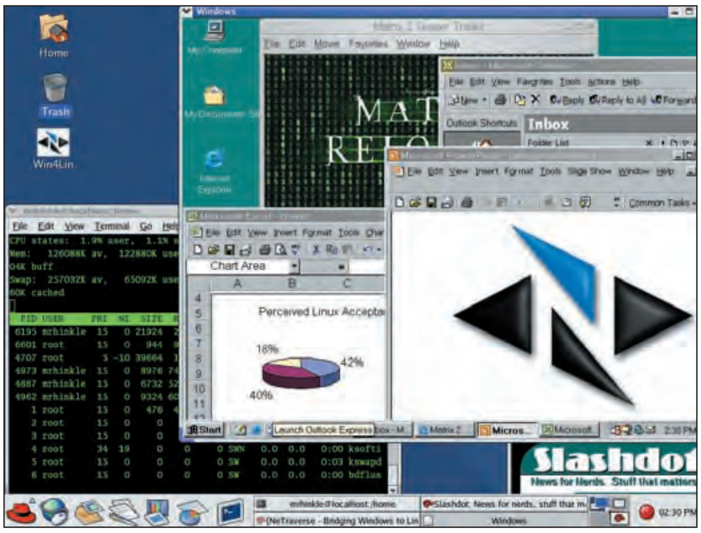

```
CPU states:  1.9% user,  1.1% s
Mem:   126088K av,  122880K use
04K buff
Swap:  257032K av,   65092K use
60K cached

  PID USER     PRI  NI  SIZE  R
 6195 mrhinkle  15   0 21924  2
 6601 root      15   0   944  9
 4707 root       5 -10 39664  1
 4973 mrhinkle  15   0  8976 74
 4887 mrhinkle  15   0  6732 52
 4962 mrhinkle  15   0  9324 60
    1 root      15   0   476  4
    2 root      15   0     0
    3 root      15   0     0
    4 root      34  19     0     0     0 SWN   0.0  0.0   0:00 ksofti
    5 root      15   0     0     0     0 SW    0.0  0.0   0:03 kswapd
    6 root      15   0     0     0     0 SW    0.0  0.0   0:00 bdflus
```

**NeTraverse claims *Win4Lin 5.0* offers better than native speed in many cases.**

### WINDOWS VIRTUALIZATION

# Win4Lin's high-five

**NeTraverse has announced an** upgrade to its popular virtualisation software, *Win4Lin*. The marketing focus of this release is providing an upgrade path for users looking for an alternative to the latest versions of Windows, while retaining compatibility with their legacy applications and documents.

James P. Curtin, NeTraverse President and CEO, said the new features in *Win4Lin 5.0* would encourage migration to Linux where legacy support has been a significant barrier. "This marks the expansion of *Win4Lin* beyond core application support to deliver value to a wider community not previously addressed," he said. "This release is critical to an increasing number of users bound by that one legacy dependency keeping them from moving to Linux."

The latest version includes support for secure, authenticated connections with Novell servers, basic DirectX features – with more robust support promised in the near future – and enhanced sound support. Networking is also improved with a more complete *Winsock 2* implementation and the introduction of a virtual network card.

The company also claim installation, which involves the loading of a custom kernel image, has been redesigned to be fully graphical.

### VIRUS SCANNING

# Protection against the growing Linux viral threat

**BitDefender, an anti-virus company** well respected in the Windows world, has released a new version of its software for Linux users. The company claims this is the most advanced multiplatform solution available, with tools for on-demand scanning of command line and shell scripts, individual files or whole filesystems. *BitDefender 6.5* not only protects against UNIX worms such as Morris and Scalper, but can also be used to stop Windows viruses reaching your network. Basic requirements are a Pentium 166, 32MB RAM and glibc 2 or above. Download it for free from **www.bitdefender.com**.

**Virus protection is a sensible precaution, Linux or not.**

## David Cartwright

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

### COMMENT

# Win? D'oh!

**"Anyone who wants to run a proper Web server under Win2K has to go out and spend hundreds of quid on *Windows 2000 Server* (or *Advanced Server*), because the Web server you get with *W2K Professional* is cut-down and can't run multiple virtual servers on a single machine. It's a significant cost – hence the myriad of Linux machines running *Apache*.**

**But with *Windows Server 2003*, Microsoft has announced a secret weapon in its fight against OSS – the *Web Edition*. Here we have a version of *Windows Server* that's aimed solely at websites. It costs a little over £200 (*ie* it's affordable) and it's a very cut-down version of its larger siblings (maximum RAM 2GB, a maximum of two CPUs in the server, no clustering, no terminal services *etc*; but on the face of it, it's just about right for the average front-end Web server.**

***WS2003 Web Edition* is certain to tempt some people away from their Linux machines. So the Linux distributors need to redouble their efforts to give the public at large a reason not to – the "Linux is free, *Windows Server* is expensive" argument just went out of the window. So distro makers need to ensure that the installers are kept simple even though the underlying functionality grows (they're doing an great job already, they just need to retain focus). Also, the Linux community needs to address the apps side of the equation too – I wouldn't have put many of the websites I've worked on on Linux if Chili!Soft hadn't made a nice ASP engine for it. Let's get solidly behind projects like Mono (www.go-mono.com), whose aim is to produce an open version of the .NET development platform for *nix systems, and show the world that there's more to Linux than just the price tag."**

# 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**

## Time tinkering

In the reply to Carl Maycock's letter (*LXF40* where he had returned to Windows because he was spending more time configuring Linux than using it), you asked us to give our experiences with problems concerning Linux. Here's mine:

About a couple of months ago I decided to upgrade my system, so I made sure I went for companies who supported Linux:

- Nvidia Geforce 2 Graphics card
- AMD Athlon XP processor
- ECS Elite motherboard (L7VTA) with Via's KT400 chipset
- Soundblaster PCI soundcard.

I thought this little lot would be very simple to set up under Linux, how wrong I was!

I think the main problem is the KT400 chipset, as Linux doesn't fully support it yet. The Soundblaster card loops the samples when using Play and freezes the browser when surfing the Net. I've tried both the OSS Free and Alsa drivers, so I've removed it and I'm just using the internal sound for now. Given that Soundblaster cards are standard, I'm really surprised by this.

The Nvidia card works fine with the Xfree driver, but there is no AGP or GLX, pretty pointless with a fast Graphics card noted for its 3D! I've tried to install the official Nvidia drivers but they just freeze the system. Yup, I've spent hours on the Net to find out how to get them working, but there's no info on the Soundblaster card with the KT400 chipset, and the Nvidia Linux forum is stuffed full of people with problems with Nvidia cards and the KT400.

None of the workarounds work with my system, though needless to say my shiny new system works perfectly (relative term) with MS Windows, so the problem is obviously Linux. As you can imagine, I'm pretty disappointed with this – I'm no hacker so I can only wait and hope Via and Nvidia can work with the community and in the meantime tinker with Linux but do most of my important stuff under Windows.

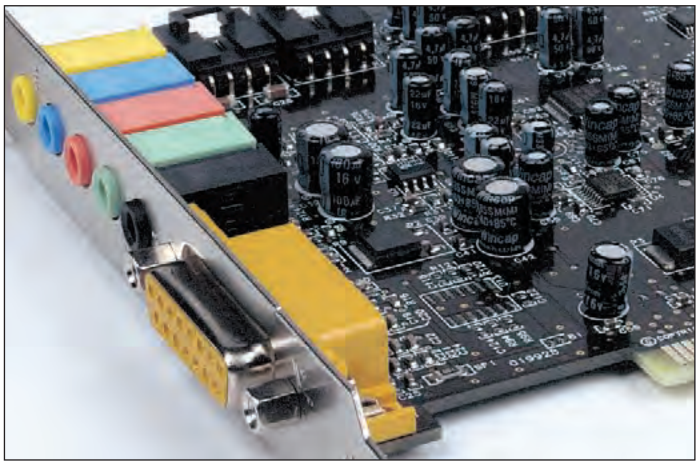

**Difficulty getting hardware to work for which you have the latest Linux drivers? The problem could be elsewhere, like the motherboard.**

## ★Letter of the month

This month's winner receives a copy of **Wiley's *Linux Process Manager* ISBN 0-4708-4771-9**

The Linux Process Manager
The internals of scheduling, interrupts and signals
John O'Gorman
WILEY

## Free as in Freedom

I've just experienced the full joys of Open Source and it's great. I've been OSS for years now, I had an account on a Linux box run by friends at Glasgow University in the mid 1990s, have had a dual boot PC running Linux for at least 3 years and have used and promoted Linux and *OpenOffice.org* at the College where I run the computer systems. But up until now I have only used Open Source software as freeware, I didn't care that I could download the source, all I cared about was not having to pay.

Not any more! Before I describe what changed, I will point out that I'm never going to be a great programmer – I only work one day a week running computer systems. The other days of the week, I'm a minister working for the church, but I know my way around Perl for scripting purposes and have written web pages in PHP.

I was in the process of moving my College away from *Lotus Organiser* to *WebCalendar,* a web-based organiser written in Perl, that would allow our faculty to use their calendar even when they were away from the College on business, which happens quite a bit. As one or two of my users were getting to grips with *WebCalendar*, they commented that they found it really frustrating that they couldn't specify an end date for an event but had to calculate the number of days to repeat it. So, I had a look at the web page, saw how a simple javascript function could calculate the number of days given a start date and an end date and wrote such a function. I then wanted to put my function into *WebCalendar*. I had a look at the Perl code and realised I couldn't make head nor tail of it.

So I emailed my idea and code to the writer of *WebCalendar*, who emailed me back telling me which Perl function was responsible for the code I was interested in. I was able to add my javascript and give my users the function they wanted. But even better the writer of *WebCalendar*, Maorong Zou, was able to recode my javascript as perl and the function will be included with the next version.

Isn't Open Source software and the attitude of Open Source programmers great? Try doing that with any closed source product!
**David Petticrew,** *via email*

We're glad you took the time to give feedback to the developers. Sadly, too many people just treat GPL software as freeware, and when it doesn't do exactly what they want, they just moan about it and try to find something else to solve their problem rather than assisting development of the original program by giving feedback. For the most part developers are pleasant human beings who actually enjoy getting feedback (the polite stuff anyway) and will welcome ideas to improve their software.

**As far as I am concerned, I love the concept of GPL and Open Source, and Linux may be ready for the Business Desktop, but it has a long way to go before it's ready for the Home computer user whose hardware is a bit more cutting edge.**
**R Sutcliffe,** *via email*

Thanks for taking the time to share your experiences. As regards the Soundblaster, the Linux drivers and the Creative ones work fine, so the problem is unlikely to be there. I think the big problem you have is not with the individual components, but with the motherboard. Via is pretty good on Linux support these days, but often lags behind Windows development. The likely problem with the graphics card may well be the AGP support.

The thing to remember here is that it is the hardware manufacturers letting you down. They put lots of time and effort into Windows drivers, not so much into Linux support.

**While there may be many individuals within LUGs who have experience of scientific Linux use, most groups try to cater for a broad range of users.**

## Return to *NIX

**I started scientific programming on Deuce in 1960, and went through a series of computers until I got a Whitechapel Unix Workstation in mid 1980s. I then moved to a Sun, and have been with that ever since. I therefore have some 20 years running Unix, including superuser system administration.**

**Being now semi-retired, I have been forced to Gates PC systems for exchanging Word documents etc. I do need to keep my Unix-based software up and running, and am very aware that Microsoft's inefficient software is strangling the potential of a PC's performance. My Suns are now dying, and I naively think that I can move all my Unix software (10,000 lines Fortran, plus scores of scripts piping data between programs, output devices and files etc) to a PC running Linux.**

**I have been reading *Linux Format*, a very helpful and lucid publication, and am aware that systems have developed somewhat since my early Solaris Unix, but can generally understand what is going on. I would be grateful for some general advice on the best routes to take.**

**1 What modes of combining Windows XP Pro and a Linux distro on a single PC do you recommend?**

**2 What Linux 'flavour/s' do you recommend for use for self-developed Fortran and C software, together with Office XP Pro compatibility. I obviously use email, and a limited amount of browsing, but have no games interest. I would prefer fairly user-friendly software, and/or software recognisable or quickly assimilated by a former Sun Unix superuser. I would need a decent Fortran compiler, and a Fortran callable line drawing capability. Obviously I would like a system to which I can easily port all my shell scripts etc.**

**3 What books do you recommend as manuals to consult when keeping up with general system administration, and are there any good places in the North West where I can go and browse Linux books before choosing ones I fancy?**

**4 Are there any other warnings I should heed before I jump in?**

**5 I have tried several Linux User Groups but my emails are usually returned address unknown, or I get dead links etc. Can you recommend a lively active Linux User Group that focuses on scientific programming?**

**6 As I am about to purchase new hardware, any Linux specific pointers to watch for would be appreciated. At present I am tempted by an Evesham top end machine. My ideal system would be Linux based, but with an option to load Microsoft-based software on top when required for compatibility, and also to use MS software where this is unavoidable. It may be that I need a dual boot system, perhaps using 2 hard discs. What I certainly need is the ability for Linux/Fortran based programmes to motor at maximum speed in one window, while a minimal amount of CPU keeps up with my report typing in another. I am reluctant to fully divorce myself from Microsoft at this stage. I am about to purchase a new high level PC, and really do not have desk top space for 2 PCs immediately to hand.**

**I would be most grateful for your advice so that I can make the most from changing back to a Unix environment. I am really looking forward to being able to type grep again, and watch the**

»

### Posting to the forum
**The LXF online community**

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

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

**GIve a hoot – contribute!**

**instantaneous response.**
**Roger Gibson,** *via email*

Welcome back to the world of *nix programming! To answer some of the questions you raised:

**1** You could mix both on one partition. If you are using XP's NTFS filesystem, only Mandrake will currently be able to resize it, or you will need a tool like Partition Magic/ Acronis Partition Expert. The easier option is probably to install a new HD dedicated to Linux.

**2** In terms of Fortran development, the choice of distribution shouldn't matter much. They all include the g77 compiler, though I would suggest you also try Intel's Fortran Compiler (**www.intel.com/software/products/compilers/flin/**)

**3,4,5** There are a great number of books available, with new releases all the time. I would certainly recommend O'Reilly's Unix Power Tools. As for sysadmin, there are plenty of specific books and some more general. Flick through the pages of our book reviews section. Also ask around at your local LUG. The nearest one to you that I know is pretty active is South Cheshire LUG (www.sclug.org.uk)

**6** The things to watch for with hardware are that it is well supported by Linux. Particularly the graphics card and the chipset used by the motherboard itself can sometimes cause problems. Find out specifically what hardware any potential system uses, then check it out either at the distribution specific sites, or just do a Google search. Most hardware is OK these days (built-in Winmodems rarely work though), but very new stuff may be harder to get working. Your LUG will be able to give you some help here.

You can also seek advice on the Linux Format forums on our website – **www.linuxformat.co.uk** where there are lots of helpful people with a huge amount of experience in installing Linux.

## SuSE struggling

**I attempted to have a look at SuSE 8.2 last weekend, as my understanding is that it is aimed towards the home user. Do the guys at SuSE not realise that downloading from their site is damn near impossible for a complete beginner like me? Having decided to migrate from Windows (after a few months of reading your magazine) I expected an operating system directed towards the home user to be a simple case of 'put this little CD in the drive and press the button, well done, good lad, here's a treat for ya', but having not been able to get to grips with the way SuSE host it on their website, I never got to the first step. Seems a bit strange to hide something from your targeted users; do they not realise that people won't migrate if it's made more difficult than it has to be?**
**Garrett Kennedy,** *via email*

We think the idea is that if you are a beginner and need support etc, you should perhaps buy a boxed version.

05/28/2003: City of Munich Replaces Windows with Linux
05/23/2003: Test SuSE Linux Free of Charge
05/22/2003: HP Achieves First Linux Clustered Oracle Applications Standard Benchmark
05/22/2003: More Performance for Lotus Domino 5 and 6: UnitedLinux Extension Pack for Lotus Domino
05/16/2003: D.H. Brown Study: SuSE leads in Linux improvements

**Struggling alone is not the best introduction to Linux, be it SuSE or any other flavour. Why not contact your local LUG for help and support?**

## EFI handicapped?

**Putting the BIOS on the HDD**

**On reading *What On Earth is EFI* in *LXF40* I was struck by the absence of one question, what happens when my HDD dies? Oops!**

**How many people remember the fate of IBM's micro architecture systems that had BIOS on their HDD? And who enjoys putting a new HDD in a Compaq, let alone an IBM Z50...**

**I may have missed the plot; *ie* will a HDD come with universal drivers or will I we be looking for an EFI for that obscure board? And as for virus writers – I'll leave that to your imagination.**

**How will the motherboard know what HDD it has, and where it is – will I have to resort to a sub BIOS in binary to fix this? They were not exactly user-friendly even with the Mboard manual.**

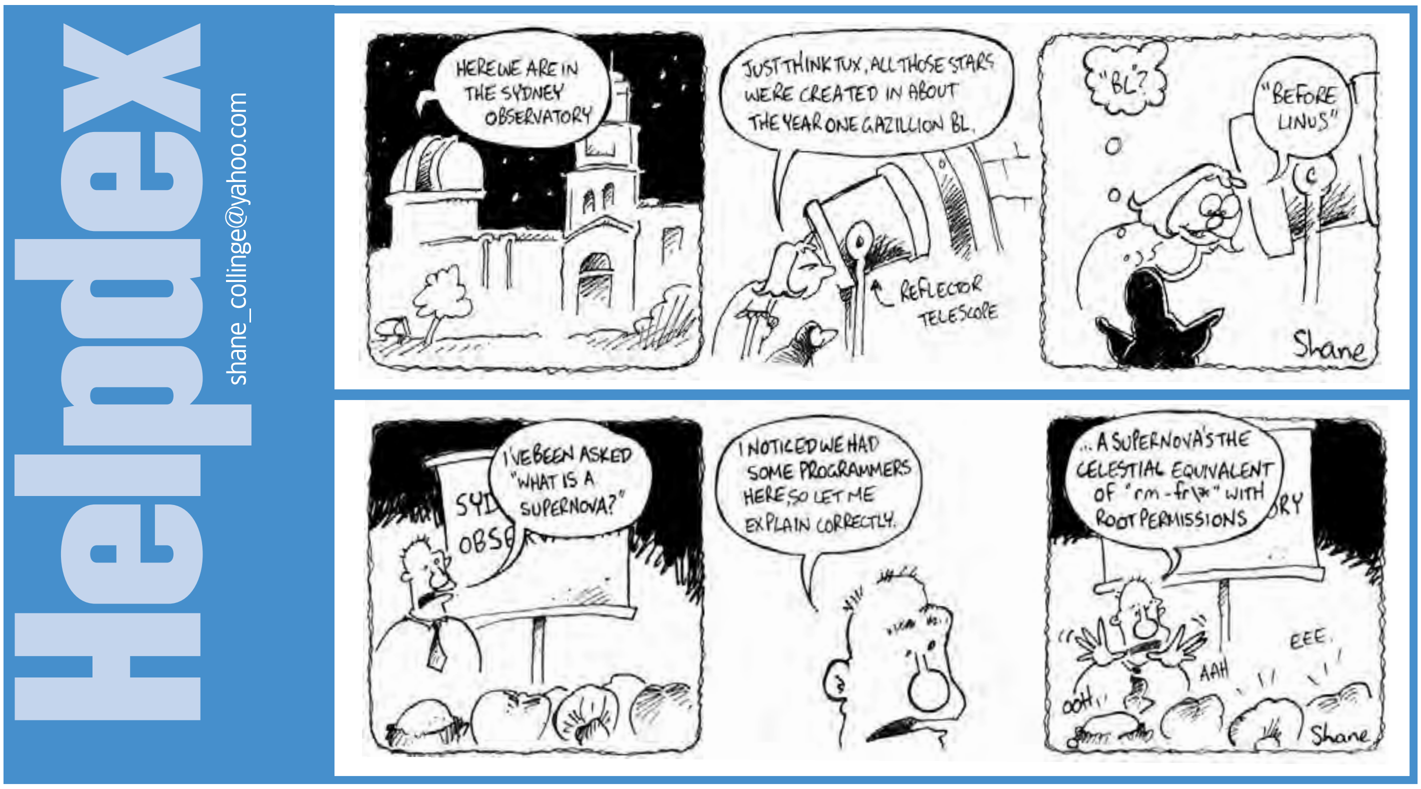
Helpdex
shane_collinge@yahoo.com
HERE WE ARE IN THE SYDNEY OBSERVATORY
JUST THINK TUX, ALL THOSE STARS WERE CREATED IN ABOUT THE YEAR ONE GAZILLION BL.
REFLECTOR TELESCOPE
"BL?"
"BEFORE LINUS"
Shane
I'VE BEEN ASKED "WHAT IS A SUPERNOVA?"
SYD OBS
I NOTICED WE HAD SOME PROGRAMMERS HERE SO LET ME EXPLAIN CORRECTLY.
...A SUPERNOVA'S THE CELESTIAL EQUIVALENT OF "rm -fr *" WITH ROOT PERMISSIONS
RY
OOH
AAH
EEE
Shane

**I would see this as an extension of the Microsoft licensing policy, Dead HDD = Dead motherboard = dead system = new system and NEW LICENCE required... or am I being cynical?**

**Simon Rogers,** *via email*

The EFI specifications are still being worked on (version 1.10 at the moment) which may address some of your concerns. The current specification includes generic drivers for all sorts of boot devices, including booting from ROM if required. I think the real question is which options will be implemented by manufacturers. There's no point in producing an EFI motherboard if you need to sell an HDD with it to get it working!

There's plenty more info on EFI at Intel's developer site – **http://developer.intel.com/technology/efi/index.htm**

## Upgrading

**Just a couple of observations about your Mandrake 9.1 review. You mention *ReiserFS* and *XFS* joining *ext3*. Surely they've been there since 8.1 at least?**

**You also say you didn't test the 'upgrade' option. You may have without knowing it... I can't tell the difference. I started by upgrading, but this carried forward a problem I developed with GNOME so I resigned myself to losing my bookmarks and stored mail and started again with install. When it was up and running I still had my bookmarks and mail... and the problem with GNOME.**

**I was intrigued to read about the new Mandrake galaxy desktop with all the commercials placed into one 'Welcome' folder. I got a page full of ads and no welcome icon. I did get the FirstTimeWizard though. But as usual it malfunctioned, so I never got the invite to join the Mandrake Club. Not that this is a problem since the Club, through no fault of its own, one of my first rules on Internetting – never to pay tax to a foreign country!**

**I've not had many problems attributable to the OS itself, but whereas under 8.1 through to 9.0 Galeon was the only browser to auto-destruct, now all of them have done it at one time or another, and Galeon no longer remembers everything I did right up to the second it crashed. And *Opera* has disappeared from the WWW menu, although it still resides on the GNOME taskbar.**

**I think you guys are just a little too expert. Mandrake's not quite ready for the guy in the street who expects everything to work like it should yet!**

**Deke Roberts,** *via email*

Mandrake's upgrade facility, which has existed in all versions for the last few years, at least works to some degree. The problem is, if you've substantially tinkered with the base install, the upgrade is going to be less effective, and sometimes a complete re-install is the saner policy (which is a very good reason for keeping /home on a separate partition). The problem isn't just limited to Mandrake you'll be pleased to know. Red Hat 9's upgrade feature is, in the experience of one reader who conatcted us, even worse. Failing to make sense of altered config files from 7.3, it merely left them alone resulting in an X server which didn't work, Application menus with no entries, a screwed up kicker panel and more.

From past experience, I have found it easier to 'upgrade' by pointing the package manager at the new distro tree and uploading packages. This usually results in about 400MB of downloads, but you can also use the CDs if you like. Then you get the most up to date packages (including Galaxy and whatever else) without causing so much damage to the installed (and working) version of Mandrake you already have.

When it comes to installing, a lot still depends on the hardware you are installing onto. Many people have a very straightforward experience, where everything is identified and configured correctly first time, whther they are using SuSE, Mandrake or Red Hat. LXF

### SUBMISSION ADVICE

**WHAT WE WANT:**

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

**WHAT WE DON'T WANT:**

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

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

LINUX DISTRIBUTION

# Red Hat Enterprise Linux (AS)

**Nick Veitch assesses the performance of Red Hat's high-end range of distributions.**

**BUYER INFO**

Designed for big servers on Itanium and x86 clusters, the only real alternative would be some version of United Linux.

- **DEVELOPER** Red Hat
- **WEB** www.redhat.com
- **PRICE** AS version $1499

Red Hat, one of the oldest and certainly popular Linux distros, recently adopted the strategy of effectively creating two tiers of distro to support the current market. The standard Red Hat distro, and a new, higher-end *Enterprise* version (formerly known as *Red Hat Advanced Server*). This has both been beneficial to users and Red Hat itself.

For a start, it more clearly defines the expected usage of the system. Back in the days of Red Hat 7.x, the end user could be a network admin installing a new server, or just a home user wanting to find out what this Linux lark was all about. Clearly, trying to be everything to everyone in one distribution would be problematic.

The more user oriented version of Red Hat has already appeared – Red Hat 8 and more recently Red Hat 9 have been transformed by the liberation of being able to add more cutting-edge software and features that would have been out of place in a server oriented distribution. There are other advantages to this strategy that we will touch on later.

*Red Hat Enterprise Linux* actually comes in three different editions – *AS*, *ES* and *WS*. These letters may actually stand for something, but Red Hat's marketing department denied any knowledge if they do (Advanced Server, Enterprise Server and WorkStation, perhaps?)

## Software

What constitutes a server environment then? In the server world, there is no such thing as 'old' software. Once software is no longer new (read 'untrusted') and people have stopped raving about it's excellent features (and presumably all of the security holes and more unpleasant bugs have been ironed out) it doesn't become 'old', it becomes 'reliable', 'trustworthy' and 'safe'. It is in this light that we should examine what software Red Hat have chosen to use as the basis for their *Enterprise* server versions.

Developers might be disappointed that the default version of *gcc* included is *2.96*, and the system is based on *glibc 2.2.4*. This is not the infamous *gcc-2.96* which caused concern in the release of Red Hat 7.3, so it will indeed work, though you will be without some of the nicer features of 3.x. Although 2.96 is the default, and the one used to build the system, Red Hat have included the RPM for *gcc 3*, as the documentation explains, for developers who want to take advantage of new features, Java support and so on. Bizarrely though, this is 3.0.4 – at least a year old. If they are going to include an optional piece of software with the caveat that it may not be reliable, the least they could do was include a more recent version where some of the new features worked more reliably.

The included kernel is based on the 2.4.9 release, though as with most kernels this has (presumably) been heavily tinkered with by Red Hat. Some of the enhancements include a lot of work on the I/O subsystem – more efficient memory usage, asynchronous support – and the scheduler, particularly important for SMP hardware

The inclusion of older versions of GNOME and KDE (2.2.2!) included aren't much of a problem – it was only in the spirit of giving the system a thorough test that we installed them at all. Servers don't even need X, unless you are determined to run some sort of VNC connection (which is included by the way).

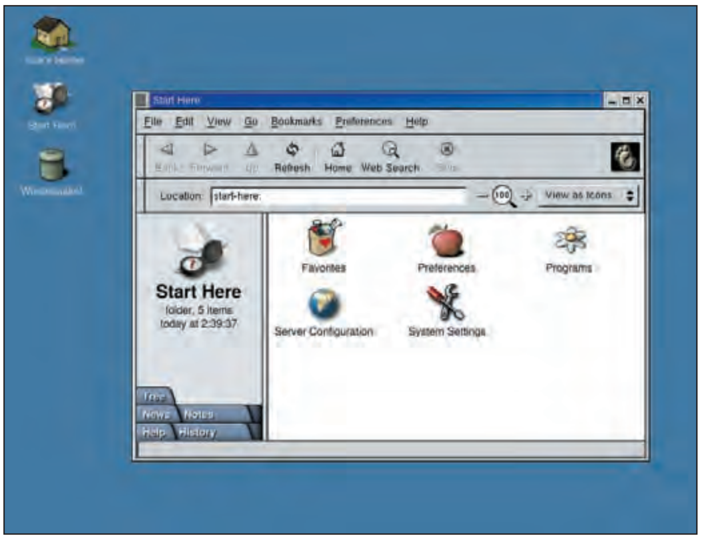

**This is the screen that greets you from a default install.**

## Documentation

### Are the instructions any good?

**As well as a CD containing loads of documentation (mostly copies of the Red Hat 7.2 guides), *Red Hat Enterprise Linux* also comes with two printed manuals. The first is presumably generic to the range and details install instructions and the usual getting started stuff.**

**The second is a well-written, extensive and helpful manual on installing the cluster management software, and the actual construction of the cluster itself. Various config files and settings need to be tweaked to make *AS* 'cluster-ready', but as the tasks that you need to undertake are clearly explained here, building your own shouldn't pose too much of a problem.**

If you are wondering how Red Hat came to decide on this particular collection of packages, it should come as not too big a surprise to discover that the *Enterprise Server* series is actually based on the last solid and trustworthy release of the standard Red Hat distro, 7.2. Admittedly, this has been patched and updated to overcome the bugs and vulnerabilities of that version. Impressively for example, *gcc-2.96* has reached patch level 116 according to the latest updates available. Key server software is similarly available on the update site – no doubt patched versions of *sendmail, PHP, MySQL, Apache* and so on address security problems that have been highlighted.

*Apache* may be worth a mention at this point. *Enterprise* server is using the old, reliable 1.3.x versions, but in this case there is no *Apache 2.x* option. This is a shame because multi-processor machines (which is what the *Enterprise Linux AS* version is aimed at) can take advantage of the better threading architecture of *Apache 2* resulting in a much more responsive web server.

Of course, you are perfectly free to install whatever extra software you like. Depending on what you may intend your server to be used for you'll probably have to. But in that case your software will fall out of the support and upgrade policy that is one of the reasons for buying a product such as this in the first place.

## Product Cycle

Support and a known product cycle are two of the key reasons for buying *Enterprise* server. Red Hat is basically guaranteeing that they will continue to support this particular version of Linux for at least 12 months, and more likely 18. This is important for people running systems who don't want to be forced to upgrade every 4 or 6 months just to keep support for systems. At *LXF* for example, we still run our webserver on (a much altered) Red Hat 7.2 distribution – why on earth would we want to 'upgrade'?

Red Hat also provides a whole range of support services, from the useful *up2date* system updater to full email and telephone support, including developer and detailed tech support if you subscribe to the right options. Red Hat's support is widely regarded as the best in the Linux space, so this should also figure in your calculations.

Another bonus of buying Red Hat is that it is the major platform that commercial software vendors certify on (another reason for longer lifecycles and the split from the 'normal' distribution), though United Linux has achieved a similar status now in the marketplace.

## Clustering

So far we have talked about features common to all the *Enterprise* versions, but there are some specifics of the top-end *AS* distribution that are worth mentioning (see table for other differences). The major component of this is Red Hat's cluster management software. Clustering, even in the Linux domain, has largely been a case of DIY constructs (see any European university) or customised hardware clustering (SGI etc.). The Red Hat cluster manager supports a wide panoply of commodity hardware and connection methods, and brings scalability and availability to a wide range of common applications including databases, file-servers, print servers and almost any other application you could run on Linux.

## Other version feature comparison

### The different varieties of *Red Hat Enterprise Linux*

| Features | WS | ES | AS |
|---|---|---|---|
| Systems with >2 CPUs | No | No | Yes |
| Supports for >4GB memory | Yes | No | Yes |
| Runs on Intel x86 architecture | Yes | Yes | Yes |
| Runs on Intel Itanium architecture | Yes | No | Yes |
| 12x5 services available | Yes | Yes | Yes |
| 24x7 services available | No | No | Yes |
| Includes network server applications | No | Yes | Yes |
| Includes High Availability Clustering | No | No | Yes |

With integrated monitoring and *failover*, and a sackful of other features to boot, the cluster manager software recently picked up a deserved *Best Cluster Solution* award at LinuxWorld Expo NY 2003.

Another useful technology, *Piranha*, is also included. Essentially this is a load-balancing solution designed to make a cluster appear as a single server to the outside world. This retains a high degree of availability (depending on how many nodes you have) and, with more than three nodes in the cluster (one routes requests, one acts as a backup and whatever remain act as the real servers) balances requests between the server power available.

*AS* is the only version of *Red Hat Enterprise Linux* which has SMP capability. Presumably this means the kernels for other versions have not been compiled with SMP, but if you want support on SMP, *AS* is the version to get.

This isn't an exciting release, but it is solid and reliable, and good value considering the included solutions. The most suspect area, and one we often have trouble with for any distribution, is the package choice and the default install (which for example will happily install various GNOME applications, but no database). A few more config tools would have been nice – especially for common setup tasks like *Samba/NFS*, printing etc. For the market this is aimed at – corporate clustering and high availability, it actually seems like a bargain. Next version though, can we please have *Apache 2*? LXF

## VERDICT

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

**Nothing too exciting, but predictable and reliable stuff.**

**LINUX FORMAT RATING** 8/10

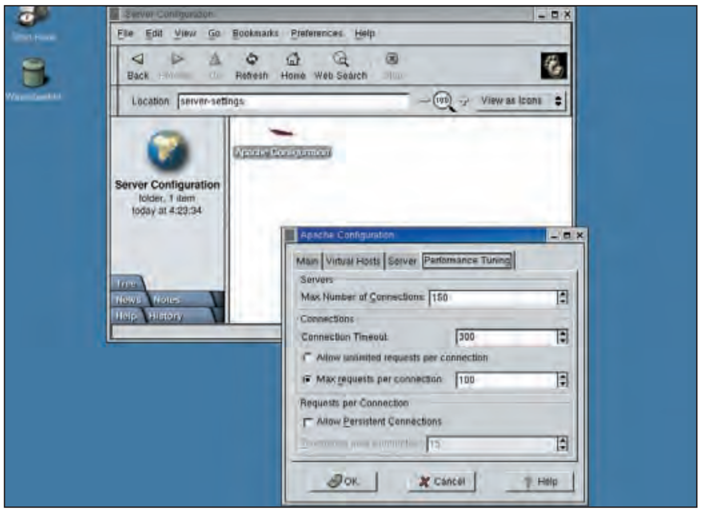

Servers
Max Number of Connections: 150
Connections
Connection Timeout: 300
○ Allow unlimited requests per connection
● Max requests per connection: 100
Requests per Connection
□ Allow Persistent Connections
15

**The server settings area only contains one graphical tool – but at least it works! This isn't a problem as this area's not likely to be used by novices.**

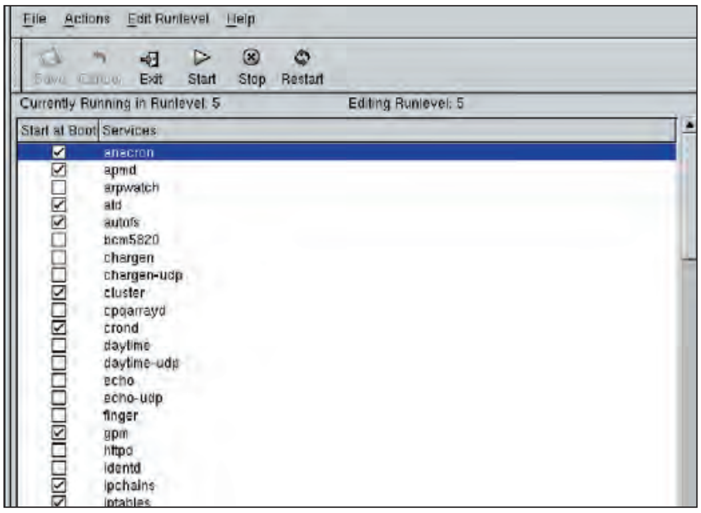
File Actions Edit Runlevel Help
Save Cancel Exit Start Stop Restart

Currently Running in Runlevel: 5
Editing Runlevel: 5

| Start at Boot | Services |
|---|---|
| ☑ | anacron |
| ☑ | apmd |
| □ | arpwatch |
| ☑ | atd |
| ☑ | autofs |
| □ | bcm5820 |
| □ | chargen |
| □ | chargen-udp |
| ☑ | cluster |
| □ | cpqarrayd |
| ☑ | crond |
| □ | daytime |
| □ | daytime-udp |
| □ | echo |
| □ | echo-udp |
| □ | finger |
| ☑ | gpm |
| □ | httpd |
| □ | identd |
| ☑ | ipchains |
| ☑ | iptables |

**You may want to check the service configuration as soon as you boot, as the installer doesn't tell you which are configured to run at boot.**

COMPILER

# GNU Compiler Collection 3.3

**GCC has always been one of the most crucial components of the Linux world, and in recent times it has gone through a great deal of change. Paul Hudson looks at what the latest release has to offer...**

On the CD and DVD

**BUYER INFO**

TheGNU Compiler Collection is the essential development tool

- **DEVELOPER** GNU
- **PRICE** Free
- **WEB** www.gnu.org

What does the C statement **i = i++;** do? It's actually not as obvious as you might first think – in fact, the return is undefined, which means that compiler writers are free to do as they please in this situation. If i started off as 4, it might end up as 4, or perhaps 5. There's even the chance of it becoming an embarrassed-looking elephant – as I said, the result is undefined, which means relying on it returning 4 isn't a smart move as it can be different in other compilers, or even newer versions of the same compiler. Up until the release of the *GCC 3.x* series, *GCC* had quite a few "features" in it that were relied up on to perform a certain task and to return a consistent value – with the release of *GCC 3.0, 3.1, 3.2*, and now *3.3*, the development team are slowly ridding themselves of these inconsistencies, adding heaps of new features, and providing much more standards-compliant C++ compilation.

If you have used *GCC 3.2*, you'll know it's a lot slower than *2.9x* – particularly when C++ templates were involved. So, what does *GCC 3.3* bring to the table?

## Fixes and features

This release has just under 200 bugs fixed over *GCC 3.2*, which is an amazing improvement by any standards. However, it's not all bug fixes – there are several new front-end features that you can make use of right now. Perhaps most importantly, a lot of work has been done to enable easier code profiling. Pavel Nejedly contributed a lot of work to enable much more accurate profiling of code, and this data is passed to the optimizer to help it identify hot spots in code on a global basis, which results in better code. Another excellent (and very complicated!) improvement is a new scheduler using deterministic finite automata, which should hopefully allow instruction re-ordering during compilation to provide an extra speed boost for run-time execution.

Several parts of *GCC* have been deprecated or removed, although none are surprising in light of *GCC*'s continuing quest towards standards. For example, the preprocessor no longer accepts multi-line string literals – a feature deprecated in *GCC 3.0, 3.1*, and *3.2*, and now totally removed in *3.3*. This makes the following C++ statement illegal:

```
MySTLMap["foo"] = "Hello,
 world!";
```

Removal of support for things like the above have been a long time coming, and it's good to see the *GCC* team aren't afraid to make such changes even now. While this might bite the few programmers who make use of such odd hacks, it's better in the long run.

## Performance tradeoff

I'll be blunt: *GCC 3.3* is no faster to compiler than *GCC 3.2*, which means it's *slow*. However, like *3.2*, you'll generally find that what you trade off in compilation time you'll get back with run-time performance, which is generally much more important. With *GCC 3.3*, you'll generally find some programs will really fly once recompiled, and this is thanks to the fact that *GCC* now supports both the *SSE2* and *3DNow!* instruction set extensions – perfect for programs that do a lot of multimedia work. There have also been some big steps forward on the x86-64 port, so Opteron early adopters will be particularly eager to get their hands on this release.

The general slow-down in speed isn't a bad thing, because it's easily justified by all the extra lengths GCC goes to make sure your code is correct. If you're only now switching from *2.9x* you'll find *GCC 3.3* issues many more warnings and errors than *2.9x*, which is very good in the long run as it makes programmers write better-structured code.

## More breakage

Pretty much as soon as *GCC 3.3* was out in the wild, programmers were complaining they couldn't compile the kernel in *GCC 3.3*, and the basic reason for this is that the Linux kernel has for a very long time relied on various bugs, holes, and misunderstandings in the *GCC* compiler, which means that as *GCC* is brought more and more into line with the C and C++ standards, more and more programs break. In this situation the problem is that *GCC* had a special **__inline__** extension that Linux developers took advantage of, which has now been removed.

While some say that *GCC* should just leave those kind of features alone, it's hard to draw the line – if *GCC* is careful not to break the kernel, should it also be careful not to break XFree86? Or *OpenOffice.org*? The answer is simply that GCC is much, much bigger than the Linux kernel. Yes, Linux is a wonderful OS, but *GCC* works on dozens of platforms, from Windows to Macintosh, and from Solaris to Amiga – it needs to stay true to the standards, and programmers who rely on undocumented features should come into line with everyone else.

*GCC* will hopefully soon get enough standards compliance that people need to worry about breaking programs when upgrading, but it's a two-part affair – programmers need to upgrade as soon as possible and check to make sure their code is compliant with the new release, and to make sure they take deprecation notices seriously.

## Three point four

*GCC 3.4* is already on the horizon, albeit still quite a way away, and is looking to follow up *3.3*'s feature bash with a solid focus on speed. Yes, you read that correctly – the *3.x* series is finally going to get some tuning. *3.3* is a great release – particularly to help developers get their code in line – but I wouldn't say it's necessary for everyone. While it does retain backwards compatibility with prior versions, and does add quite a bit of interesting new functionality, it's something primarily of interest for developers.

If you *are* a developer, then this is a definite upgrade. If you're thinking anything like "I'll stick *2.9x* for now and wait till they stop making changes", then you'll find that when you finally choose to upgrade (to version *3.x* or even *4.x*), you'll find you have thousands of errors and warnings. Chances are you'll fare better by upgrading now and fixing errors along the way.

So, to conclude, this is another solid step forward for *GCC*, keeping it firmly at the forefront of compiler technology. At this rate I don't think it will be that long before *GCC* comes close to rivalling the runtime performance of Intel's compiler – only time will tell. LXF

**VERDICT**

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

**The world's most popular compiler gets another worthy upgrade. Not essential, though – consider waiting for 3.4.**

**LINUX FORMAT RATING**

**9/10**

PROGRAMMING LANGUAGE

# Java 2 Standard Edition 1.4.2

**We haven't looked at Java for some time in *Linux Format*, despite the fact that it continues to march forward in terms of the technology it provides and acceptance among developers. Paul Hudson investigates...**

**BUYER INFO**

The definitive Java compiler and run-time environment.

- **SUPPLIER** Sun
- **PRICE** Free
- **WEB** http://java.sun.com

Java, if I could choose, would be the one language I wished I had more time to study. C, C++, PHP, Perl, and Delphi I use regularly and know like the back of my hand, and yet Java – arguably the most advanced language from the list – barely gets a look-see beyond me trying it out every so often to see what's changed. If your nostrils flared and you got tunnel vision when you read that I thought Java is the most advanced language out there, hold your horses – Java is quite different to both C and C++ because it includes built-in libraries to accomplish most basic tasks (XML parsing, 3D graphics, etc), and that's what I'm referring to.

Although 1.4.2 doesn't sound like a particularly huge jump up from 1.4.1, trust me – there are big changes going on. And if you haven't looked at Java for some time, you should definitely consider giving it a shot.

## Java is too slow!

After years of Java being around, "it's not fast enough" is still probably the biggest complaint. Since it was first released, Java has come a long way in terms of performance, but it's true that still doesn't operate at the speed of native code. Java compiles down to a cross-platform bytecode, that is then interpreted by the virtual machine to perform native instructions. A few versions again, Sun introduced a hotspot just-in-time JVM, which is to say that the JVM analyses the code it's going to execute, and converts it to native machine code for maximum performance. In this release, however, two things have combined to give Java a substantial speed boost. Firstly, a special effort was made to focus on performance in the virtual machine, which has resulted in a 30% speed increase in command-line application startup, and 15-20% in small Swing applications. Secondly, 1.4.2-compatible JVMs will hotspot emit SSE and SSE2 instructions if they are supported, which should provide a tremendous speed boost.

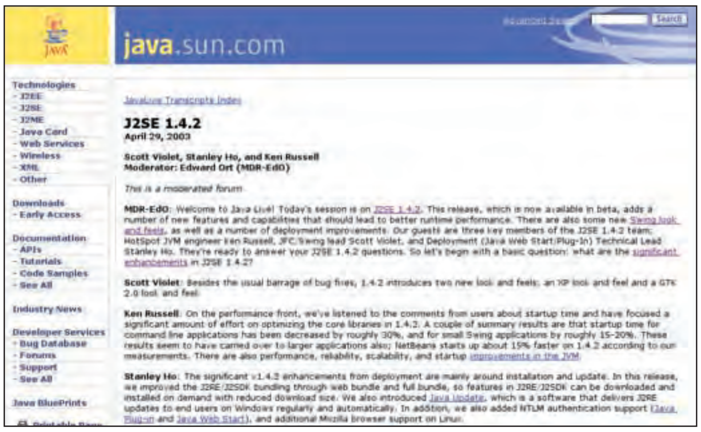

JavaLive Transcripts Index

**J2SE 1.4.2**
**April 29, 2003**

**Scott Violet, Stanley Ho, and Ken Russell**
**Moderator: Edward Ort (MDR-EdO)**

*This is a moderated forum.*

**MDR-EdO**: Welcome to Java Live! Today's session is on J2SE 1.4.2. This release, which is now available in beta, adds a number of new features and capabilities that should lead to better runtime performance. There are also some new Swing look and feels, as well as a number of deployment improvements. Our guests are three key members of the J2SE 1.4.2 team: HotSpot JVM engineer Ken Russell, JFC/Swing lead Scott Violet, and Deployment (Java Web Start/Plug-In) Technical Lead Stanley Ho. They're ready to answer your J2SE 1.4.2 questions. So let's begin with a basic question: what are the significant enhancements in J2SE 1.4.2?

**Scott Violet**: Besides the usual barrage of bug fixes, 1.4.2 introduces two new look and feels: an XP look and feel and a GTK 2.0 look and feel.

**Ken Russell**: On the performance front, we've listened to the comments from users about startup time and have focused a significant amount of effort on optimizing the core libraries in 1.4.2. A couple of summary results are that startup time for command line applications has been decreased by roughly 30%, and for small Swing applications by roughly 15-20%. These results seem to have carried over to larger applications also; NetBeans starts up about 15% faster on 1.4.2 according to our measurements. There are also performance, reliability, scalability, and startup improvements in the JVM.

**Stanley Ho**: The significant v1.4.2 enhancements from deployment are mainly around installation and update. In this release, we improved the J2RE/J2SDK bundling through web bundle and full bundle, so features in J2RE/J2SDK can be downloaded and installed on demand with reduced download size. We also introduced Java Update, which is a software that delivers J2RE updates to end users on Windows regularly and automatically. In addition, we also added NTLM authentication support (Java Plug-in and Java Web Start), and additional Mozilla browser support on Linux.

**http://java.sun.com is a great repository of resources for all things Java.**

## But Java looks awful!

Or perhaps more accurately: "Java looks awful on my platform!" This is probably the second most common complaint about Java – that end-users feel uncomfortable in Java's alien interface. Again, steps have been taken in 1.4.2 to address this, and JVMs now have two new pluggable look and feels (PLAFs) – GTK+ 2.0 and Windows XP. The GTK+ support is more than just adopting a certain look – 1.4.2-compliant JVMs support limited GTK skinning, which means it will attempt to look like the GTK skin you're using system-wide rather than the unskinned GTK default. As you may know, currently GTK skinning support is less than perfect – it only supports three or four skin types. However, better compatibility is one of the targets for 1.5.

The new Windows XP PLAF should come as a relief for developers deploying to Windows end-users. It works seamlessly, and is only available on XP systems – when you specify the Windows PLAF, the JVM will choose XP if it's running on Windows XP, and the default Windows PLAF on other versions.

In addition, there's a new Metal theme planned for 1.5, which should revive the ailing look and feel – after all, it hasn't changed for some time.

## Java updates too often!

This is the third and last most quoted complaint with Java - that Sun keep releasing new versions, which means people have to continuously download updates and fixes. Again, Sun have taken some good steps to correct this in 1.4.2, and it's a three-pronged assault. Firstly, 1.4.2 includes "Java Update" – a tool to automatically deliver J2RE updates to end users on a regular basis. It's currently only available for Windows, but as its usefulness becomes apparent it's likely to expand to other platforms. The second prong is that J2RE updates are now done using an install-on-demand system, which means that only the features of the J2RE that are required will be downloaded and installed. Finally, 1.4.2 supports incremental downloading – rather than downloading one large chunk of software, it will instead attempt to download a patch to bring the current version up to date. Three three features combined should result in users being able to keep up to date with the minimum possible effort and download time.

## Conclusion

So, there are quite a few excellent new features for those upgrading from 1.4.1 to 1.4.2. If you haven't looked at Java for some time, then you probably missed the entire 1.4 release - the super-fast new I/O APIs, the excellent XML support built in, the enhanced 2D support, or the Itanium support. Sun are still pushing things forward faster than anyone else out there, which is great to see – I think many developers will be delighted to see features like AES support in 1.4.2, even though the encryption standard was only finalised recently. Users and developers alike will also be happy to see hundreds more bug fixes, both in the compiler and in the virtual machine, to make the platform much more stable.

This is an excellent release, offering a slew of cool new functionality and automatic speed boosts through SSE, SSE2, and the other various performance improvements. It also offers a good taster of what can be expected in 1.5 – better PLAFs (including a new user-definable PLAF that reads its look from XML) and more performance. LXF

**VERDICT**

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

**Continues Java's tradition of being on the cutting edge of the cutting edge.**

**LINUX FORMAT RATING**

**9/10**

BUG TRACKER

# Seapine TestTrack Pro 5.1.1

**Knowing when your code is buggy isn't an issue – it just won't work! But tracking down where the bugs *are* is more difficult. Jono Bacon calls in the exterminator...**

**BUYER INFO**

Bug tracker that runs through your browser to document and manage corrections for coding teams.

- **DEVELOPER** Seapine
- **PRICE** From $295 (various rates)
- **WEB** www.seapine.com

These days, with ever-bigger software projects, quality control can be a difficult thing to keep track of. Not only must software be correctly developed, but bugs and problems must be identified and documented. Once these problems have been discovered, there is then the job of ensuring that the fix is correctly applied and documented. This entire procedure usually needs to happen in a consumer- and developer-friendly way.

*Seapine TestTrack Pro* is a software application to fulfil this very requirement. The software is designed to enable software defects, problems and bugs to be documented at every stage of the way.

## Installation

Installation is always an interesting issue when reviewing software due to the many different Linux distributions and dependencies that the software vendor needs to satisfy. Unfortunately many applications fall over at this stage and I was interested to see how this product would fare.

I received *TestTrack Pro* on a CD-ROM with a huge stack of manuals. The sheer amount of documentation that comes with the software is quite impressive and it shows *Seapine's* dedication to the quality of the application. I shoved the disc in the drive, mounted it and then proceeded to run the graphical Java installer. The installer popped up and installation was a smooth process, just asking for the locations of where the HTML and cgi-bin files should be stored.

After installation I then started the *TestTrack Pro* server which was available in /etc/init.d/ and then accessed the web interface. One important thing to point out is that I had to manually set up *Apache* to access the HTML interface. Although a pain to set up, this was understandable as you need to do this for any web based interface running on your server.

## First Impressions

When I fired up the Web interface, I checked with the included Installation Manual to see what the admin password was. I logged into the server admin page and I was presented with the usual array of options for configuring a server – users, security *etc.* After browsing through this page, I then proceeded to login to the main *TestTrack Pro* bug tracking area.

The first thing that really struck me was the layout quality. I tested *TestTrack Pro* in *Konqueror* and *Mozilla* and the interface looked perfect; no quirky HTML errors or layout problems. Seapine has obviously worked hard to ensure that their software is viewable in many different browsers.

The next thing I noticed was how the *TestTrack Pro* interface has been designed. The interface has a number of tabs along the top of the page, options down the left and the main bulk of the page showing the defect list of submitted bugs. I found this interface simple and effective right from the start. Another feature that I liked was the fact that the environment loads in its own window with no browser controls on it. Although a small feature, this gives the impression that *Test Track Pro* is not just a web page and is more of a sophisticated environment. Little features like this really add sheen to initial impression when loading up the interface.

## Features

*Test Track Pro* is essentially a professional bug tracking system. The aim of the system is to provide methods of submitting bugs, adding

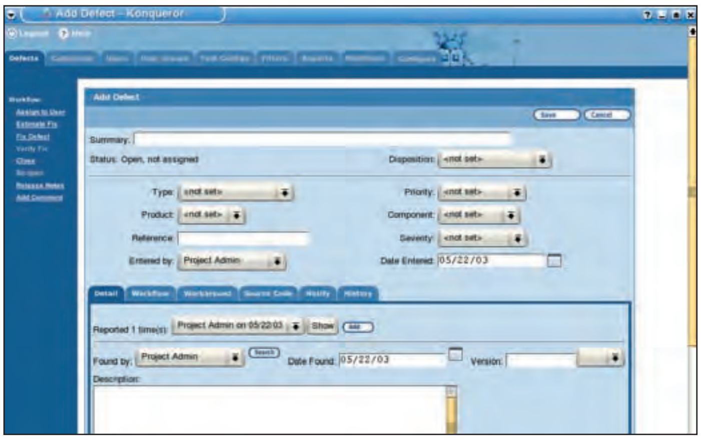

Summary:
Status: Open, not assigned
Disposition: <not set>
Type: <not set>
Priority: <not set>
Product: <not set>
Component: <not set>
Reference:
Severity: <not set>
Entered by: Project Admin
Date Entered: 05/22/03
Reported 1 time(s): Project Admin on 05/22/03
Found by: Project Admin
Date Found: 05/22/03
Version:
Description:

**Adding bug details to *Test Track Pro* lets it know what to look for.**

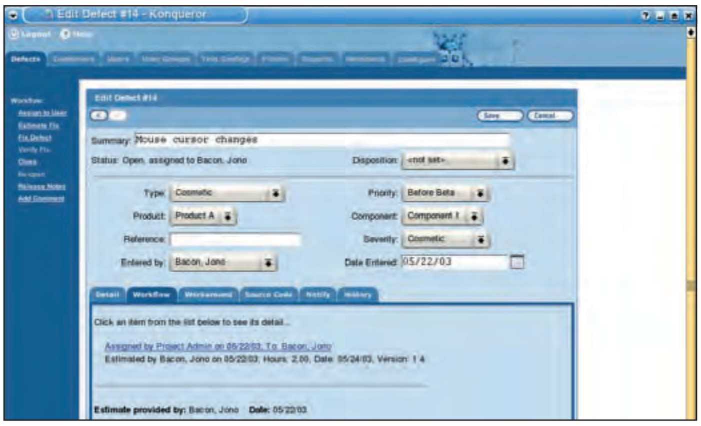

Summary: Mouse cursor changes
Status: Open, assigned to Bacon, Jono
Disposition: <not set>
Type: Cosmetic
Priority: Before Beta
Product: Product A
Component: Component 1
Reference:
Severity: Cosmetic
Entered by: Bacon, Jono
Date Entered: 05/22/03

Click an item from the list below to see its detail...

Assigned by Project Admin on 05/22/03, To: Bacon, Jono
Estimated by Bacon, Jono on 05/22/03, Hours: 2.00, Date: 05/24/03, Version: 1.4

Estimate provided by: Bacon, Jono Date: 05/22/03

**Bug estimation details can be used to assign team members to fixes.**

information, managing who the bugs are assigned to, and dealing with fixes for bugs. The system is designed to provide functionality at every step of the bug reporting and fixing process. A good method of testing this was by running through a short example of adding a bug.

When you click on the Add button you are taken to the bug reporting screen. Here you can add a bug summary, the type of bug, which component it is in, the severity, the date and the details of the bug. As well as this normal information you can add details on your system setup, reproduction information and any attachments that may be relevant to the bug report. Once the initial bug has been filed with the relevant information, you can then add additional information on how the bug is managed. The first thing that would happen in most projects is that the bug would be assigned to someone to fix. This can be done by clicking on the Assign To User option on the left-hand side of the interface. When assigned to the relevant user, the user will then be made aware of their bugs that they need to fix.

Although this is a common format of many bug tracking systems, the method it is implemented in with *Test Track Pro* is good. When the bug has been submitted, the developer can then estimate a date for a fix and enter some initial notes about the nature of the fix. After the estimation has been made, the final fix details can be entered. Once these further steps have been executed, then *Test Track Pro* will provide archive details of bugs that have been fixed and add them to any reports.

## Reporting

One of the features I really liked about *Test Track Pro* is the reporting functionality within it. There are a number of different reports available including Closed Reports, Feature Requests, Open Defects, Team Update/Assignment, Team Update/Priority and more. This gives the developer the ability to keep an eye on the general trends and patterns within the project development cycle.

As well as generic project reports, there are detailed reports for specific bugs. These detailed reports specify the reporting, estimation and fix stages and clearly summarise the history of the bugs life. This is very useful if your boss needs specific evidence of how a particular bug was fixed (if for example a client wishes to see how the bug they reported was handled).

## Configuration

The configuration of *Test Track Pro* is done right within the environment by clicking on the Configuration tab at the top. The configuration page allows you to edit the names of Types, Severity, Products, components, Dispositions, Reproductions and more. The configuration also allows users to configure resolutions and email templates in a straightforward way.

As well as configuring the specific terms and settings for the project, there are controls to set the users up for the system. Here you can add, edit and remove users as well as setting the different user groups that users fall into such as Administration, Customer, Engineers/Developers, Managers, Testers and more.

Another useful configuration setting is the Test Configurations. You can set particular system settings and record them to *Test Track Pro* so the developers are fully aware of the system/environment that the product is running on. This is a great feature and one that solves a major issue of bug tracking systems – getting users to clearly state their configurations. The configuration options allows you essentially control every portion of the *Test Track Pro* interface so it is tied in with your products and users. This configuration in combination with the Server Admin options makes the environment very flexible.

## Conclusion

*Seapine Test Track Pro* is an expansive and good quality product. Seapine has obviously worked hard to implement pretty much every feature you would ever need to make use of in a bug tracking system. Not only are these features to do with the reporting of bugs and how they are handled, but there's the capability to cover the reporting and administration of the bug tracking system.

It is clear that *Seapine Test Track Pro* was find particular use within an organisation developing software with a number of developers either on-site or distributed and working via the Internet. In this context I can foresee this software working well and efficiently maximising the productivity of the development team.

Aside from the confusion about the Windows program in the manuals (see below) this software is highly recommended if you need a solid bug tracking system.

## Databases

From my initial review and tests on *Test Track Pro*, it seems that Seapine has opted for its own database system to store bugs and defects that are accessed within the environment. This move was an interesting one and one that has benefits and disadvantages that you should be aware of.

The major benefit of including their own database is that Seapine has removed any dependency on a third party database such as MySQL. This means that *Test Track Pro* is easy to get up and running on your machine and get started straight away.

The only problem with having their own database though is that it limits the accessibility of the database. Although the *Test Track Pro* database is accessible, it is not as flexible as MySQL in terms of accessing from your own code (such as making your website interface with *Test Track Pro*).

## Windows components

When I received the review copy of *Test Track Pro*, there were two large manuals referring to a database configuration client that is included. The manuals appeared to be authoritative and well-written, but they showed countless Windows screenshots and there was no information whatsoever on accessing this client in Linux. I have therefore assumed that this client is not available for Linux and is just a Windows program. If this is the case then I feel that Seapine should release this client for Linux so the manuals are applicable for all Operating Systems. Unfortunately some companies treat their Linux versions differently to their Windows counterparts, and I hope Seapine is not doing this, and a Linux client is under consideration or soon to be available. LXF

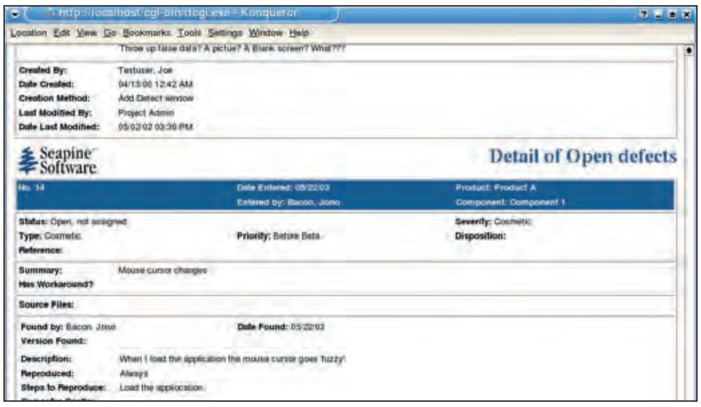
http://localhost/cgi-bin/ttcgi.exe - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help

Throw up false data? A picture? A Blank screen? What???

Created By: Testuser, Joe
Date Created: 04/13/00 12:42 AM
Creation Method: Add Defect window
Last Modified By: Project Admin
Date Last Modified: 05/02/02 03:36 PM

Detail of Open defects

No. 14 | Date Entered: 05/22/03 | Product: Product A
Entered by: Bacon, Jono | Component: Component 1

Status: Open, not assigned
Severity: Cosmetic
Type: Cosmetic
Priority: Before Beta
Disposition:
Reference:

Summary: Mouse cursor changes
Has Workaround?

Source Files:

Found by: Bacon, Jono
Date Found: 05/22/03
Version Found:

Description: When I load the application the mouse cursor goes 'fuzzy'
Reproduced: Always
Steps to Reproduce: Load the application

Generated reports cover both bugs themselves and the tracking process.

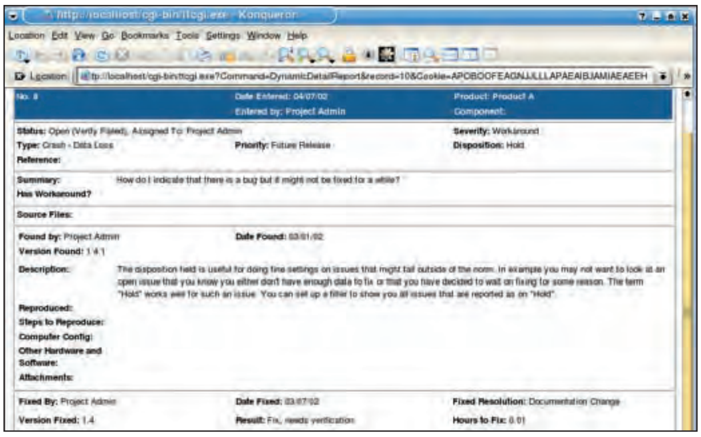
http://localhost/cgi-bin/ttcgi.exe - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help
Location

tp://localhost/cgi-bin/ttcgi.exe?Command=DynamicDetailReport&record=10&Cookie=APOBOOFEAGNJJLLLAPAEAIBJAMIAEAEEH

No. 8 | Date Entered: 04/07/02 | Product: Product A
Entered by: Project Admin | Component:

Status: Open (Verify Fixed), Assigned To: Project Admin
Severity: Workaround
Type: Crash - Data Loss
Priority: Future Release
Disposition: Hold
Reference:

Summary: How do I indicate that there is a bug but it might not be fixed for a while?
Has Workaround?

Source Files:

Found by: Project Admin
Date Found: 03/01/02
Version Found: 1.4.1

Description: The disposition field is useful for doing fine settings on issues that might fall outside of the norm. In example you may not want to look at an open issue that you know you either don't have enough data to fix or that you have decided to wait on fixing for some reason. The term "Hold" works well for such an issue. You can set up a filter to show you all issues that are reported as on "Hold".

Reproduced:
Steps to Reproduce:
Computer Config:
Other Hardware and Software:
Attachments:

Fixed By: Project Admin
Date Fixed: 03/07/02
Fixed Resolution: Documentation Change
Version Fixed: 1.4
Result: Fix, needs verification
Hours to Fix: 0.01

Full bug-reporting details are invaluable for large programming teams.

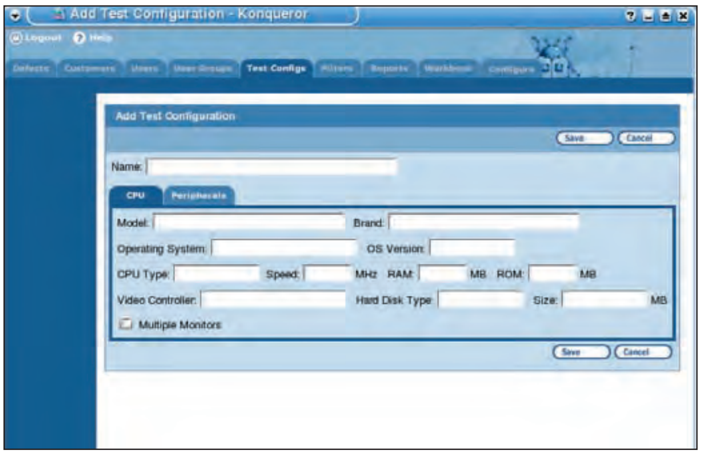

System settings recorded in a straightforward tabbed organisation.

### VERDICT

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

**A good quality, solid defect tracking system. Recommended for software projects with a large team of developers who need a good system for documenting bugs.**

**LINUX FORMAT RATING 8/10**

PROGRAMMER'S TEXT EDITOR

# VisualSlickEdit 8.0

LINUX FORMAT TOP STUFF AWARD

**Nick Veitch limbers up his typing finger with the latest release of this programmer's editor.**

### BUYER INFO

More than just a text editor, *VSlick* functions like an IDE for a variety of languages and environments.

- **DEVELOPER** SlickEdit
- **PRICE** $269(download edition)
- **WEB** www.slickedit.com

Coding is what *VisualSlickEdit* (*VSlick*) is all about, and there are an abundance of features to make your coding experience more productive. As with an IDE, *VSlick* can organise files into a 'project' space, so you know what files are related and how the relate to each other. There are menu entries (and keyboard shortcuts, naturally) for Compile, build, rebuild, Debug and Execute. If *VSlick* knows the nature of the projects (*eg* a Borland *JBuilder* or *GCC C++*) these menu items will be linked to the relevant commands. These are totally customisable though, so if you want you can use specific options. What is possibly unique to *Vslick*, in Linux environments anyway, is the ability to define different build configurations. You may want specific options for a release version build, or others for a debug/testing build (*eg* no optimisation for a quicker compile or whatever). Vslick can store these profiles to make it even easier to manage projects.

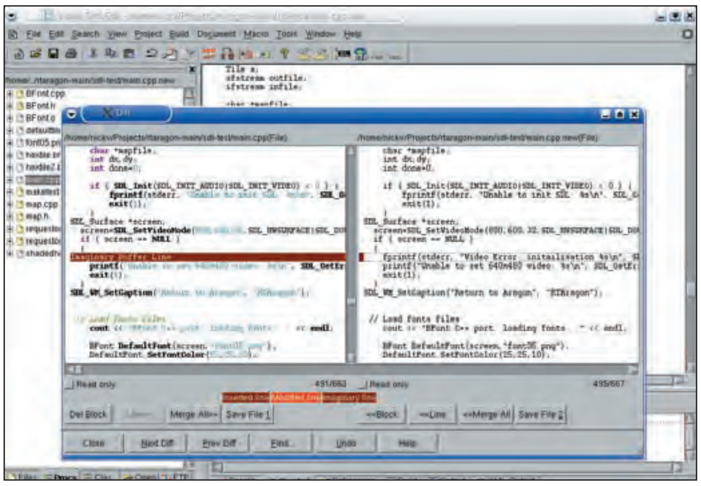

/home/nickv/Projects/rtaragon-main/sdl-test/main.cpp(File)

/home/nickv/Projects/rtaragon-main/sdl-test/main.cpp.new(File)

```
    char *mapfile;
    int dx, dy;
    int done=0;

    if ( SDL_Init(SDL_INIT_AUDIO|SDL_INIT_VIDEO) < 0 ) {
        fprintf(stderr, "Unable to init SDL: %s\n", SDL_G
        exit(1);
    }
SDL_Surface *screen;
  screen=SDL_SetVideoMode(800, 600, 32, SDL_HWSURFACE|SDL_DO
  if ( screen == NULL )
  {
    fprintf(stderr, "Video Error: initialisation %s\n", SD
    printf("Unable to set 640x480 video: %s\n", SDL_GetEr
    exit(1);
  }
SDL_WM_SetCaption("Return to Aragon", "RTAragon");

  // Load fonts files
    cout << "BFont C++ port. loading fonts ..." << endl;

    BFont DefaultFont(screen, "font36.png");
    DefaultFont.SetFontColor(25,25,10);
```

Imaginary Buffer Line

491/663 495/667

**The *DiffZilla* environment quickly highlights code changes and allows merging or manual alteration.**

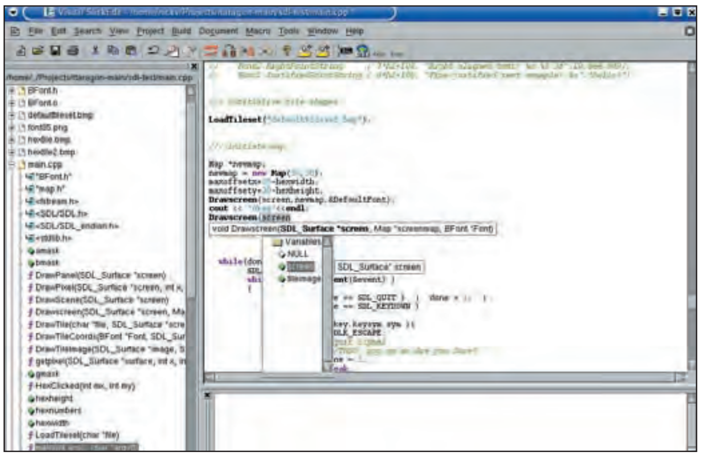
File Edit Search View Project Build Document Macro Tools Window Help
/home/../Projects/rtaragon-main/sdl-test/main.cpp
BFont.h
BFont.o
defaulttileset.bmp
font05.png
hextile.bmp
hextile2.bmp
main.cpp
"BFont.h"
"map.h"
<fstream.h>
<SDL/SDL.h>
<SDL/SDL_endian.h>
<stdlib.h>
amask
bmask
DrawPanel(SDL_Surface *screen)
DrawPixel(SDL_Surface *screen, int x,
DrawScene(SDL_Surface *screen)
Drawscreen(SDL_Surface *screen, Ma
DrawTile(char *file, SDL_Surface *scre
DrawTileCoords(BFont *Font, SDL_Sur
DrawTileimage(SDL_Surface *image, S
getpixel(SDL_Surface *surface, int x, in
gmask
HexClicked(int mx, int my)
hexheight
hexnumbers
hexwidth
LoadTileset(char *file)
Variables
NULL
Sbmage

```
LoadTileset("defaulttileset.bmp");

Map *newmap;
newmap = new Map(  ,  );
maxoffsetx=  -hexwidth;
maxoffsety=  -hexheight;
Drawscreen(screen,newmap,&DefaultFont);
cout << "  " <<endl;
Drawscreen(screen
void Drawscreen(SDL_Surface *screen, Map *screenmap, BFont *Font)
SDL_Surface* screen
```

**Code completion in the editor pops up a relevant variables list to choose from.**

For popular compilers, *VSlick* also understands the compiler error messages, so as in a dedicated IDE, it can jump to the relevant section of the code, so you aren't losing anything by building directly within *VSlick*.

As well as CVS, *VSlick* also supports other versioning systems such as *SourceSafe, TLIB, GNU RCS, ClearCase* and *SCCS*.

Describing *VisualSlickEdit* as a text editor is rather akin to describing the Taj Mahal as a building, or Bob Dylan as a musician. On a very superficial level it may be correct, but in no way communicates the various attributes that make it more than this simple definition. *VSlick*, as users are wont to call it, is a text editor that is finely tuned to one specific task – coding. It isn't strictly an IDE, as that tends to have associations with a particular language or environment, but it does provide a whole raft of tools of use to any programmer.

## Pricing

**The standard price for the download version on Linux (or Windows) is $269, but there is an academic discount version at $99. This is based on a single user licence. There are various prices available for concurrent or specific multiuser site licences, which are explained on the SlickEdit website.**

»

## Major new features in version 8

### VSlick's big guns

**JBUILDER SUPPORT** Easy setup of build and compile environments, support For *JBuilder* workspaces

**JAKARTA/ANT** Support for building Java projects with this popular environment

**IMPROVED CVS** Update, checkout, diff and history support using CVS servers

**SFTP** Configure SFTP tool for more secure transfer of files

**VISUAL BASIC** Enhancements to VB support

**C#** Develop C# on Linux. Supports Mono (see *Linux Pro* 41). Provides context tagging, syntax expansion, colour coding for C#

« On the downside, the GUI certainly isn't pretty, but it is functional. You will get a panelled and tabbed display familiar to any IDE you may have used (code on the right, files, classes etc on the left, output at the bottom). It does feel a little clunky, like some failed experiment in making *Motif* work nicely, but it gets the job done. At least you can change the font the various parts of the application use to save some strain on the eyes.

## Other Tools

**I don't know about you, but I'm a :**

```
 if (whatever) {
          j++;
 }
```

**kind of guy. However, there exist in the world people who insist that:**

```
 if (whatever)
          {
          j++;
          }
```

**is the way to do it. Crazy I know, but there you are – it takes all sorts to make a world! There are companies that insist on one style, and when you're collaborating on a big project, this makes sense. *VSlick's* code 'beautifier' can easily convert from one type to the other, or any other mess you have made. It will straighten up your indents, sort out your comments, pad out variables in brackets and a dozen other things to make your code consistent and easy to read. The beautifier works for C, C++, Java, JavaScript, C#, XML and HTML, so even your documentation can look nice. Now you can write the code how you like and have *VSlick* tidy up after you.**

**We mentioned versioning support briefly earlier on, but it's worth noting again that the versioning subsystem can cope with a range of different environments (no *Bitkeeper* though, sorry Linus). A setup dialog will extract all the information that you need to give *VSlick* to use local commands to effect checkouts and updates, and you can automatically create projects and workspaces from CVS retrieved modules. We tested the system with Sourceforge and it seemed to work fine.**

**The final tool we should mention is the *DiffZilla* file comparison utility. Actually, it does more than compare files – you can compare entire directory trees if you need to, or, through the use of wildcards, files of a particular type within that tree. Simple controls on the output allow you to step through the change points in the code, merge changes or even edit the files directly in the text view windows.**

## Gooey GUI

It's a shame because one of the great things about *VSlick* as a whole is that you can customise almost anything else you can think of. This extends from the kind of things you might expect (keybindings) to almost the tiniest detail (how to highlight the current line, where to jump to when matching parenthesis, how special characters are displayed). If you are used to a different environment, you can even set *VSlick* to emulate it, including *Vi*, *Emacs* and *Codewarrior*.

There is also a comprehensive macro environment that allows sensible recording of keys and clicks to automate certain functions. Actually, this short comment doesn't do the system justice. There is a whole language here, complete with dialog boxes and functions to access the deepest parts of the *VSlick* environment.

## Languages

*VSlick* covers everything from Ada to yacc. The extent of the support varies from just recognising and colour-coding keywords to understanding the build environment and the common language constructs, tags and library functions, class browsing and basically all the features you would expect of a customised IDE. Major languages like Java and C/C++ are completely supported, but even the likes of Python are recognised.

Function popups not only give you the prototype of the function, but also helps you fill it in by giving drop down lists of all the variables that fit the known parameters – a great boon if you can't remember variable names from one line to the next.

As well as standard tagging and completion features, *VSlick* also automatically generates a database of your own code tags. If you define a class with member functions, these too will appear in popups for completion and are colour coded correctly to avoid confusion.

## Cross-platform

Although some may disparage the appearance of the GUI, one of the major advantages of *VSlick* for large development environments is its cross platform uniformity. Not only does it look similar, it behaves identically on whatever platform you choose to use. For teams working on cross-platform projects, this is a great boon, and the fact that all the setup files and configurations, tags etc can be used on different platforms makes creating a uniform environment a lot easier.

## Support

### Workarounds save the day – from day one!

**The SlickEdit team obviously cares very deeply about their product. I had been testing the software for a while on Red Hat, then moved to a Mandrake 9.1 environment. *VSlick* didn't like it at all, I supposed because of the later version of *glibc*. Before going back to RH, I checked out the www.slickedit.com site, and there, in the support section was the workaround to the problem, and sample scripts to put in /usr/bin to make everything happy again. That this a) worked and b) was on the website about three days after Mandrake 9.1 was released is amazing.**

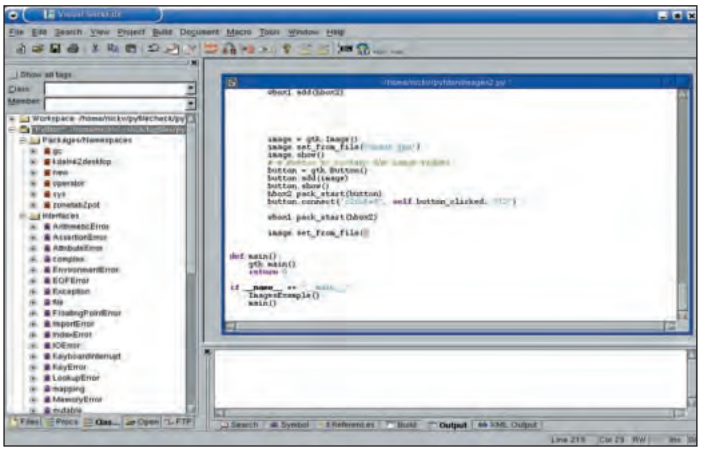

```
        vbox1.add(hbox2)

        image = gtk.Image()
        image.set_from_file(
        image.show()
        # a button to contain the image widget
        button = gtk.Button()
        button.add(image)
        button.show()
        hbox2.pack_start(button)
        button.connect("clicked", self.button_clicked, "1")

        vbox1.pack_start(hbox2)

        image.set_from_file(

def main():
    gtk.main()
    return 0

if __name__ == "__main__":
    ImagesExample()
    main()
```

File Edit Search View Project Build Document Macro Tools Window Help
Search
Symbol
References
Build
Output
XML Output

***VSlick* supports many languages, including python, C, SQL, tex, PHP, Perl, Java, Fortran, Cbol C# and more.**

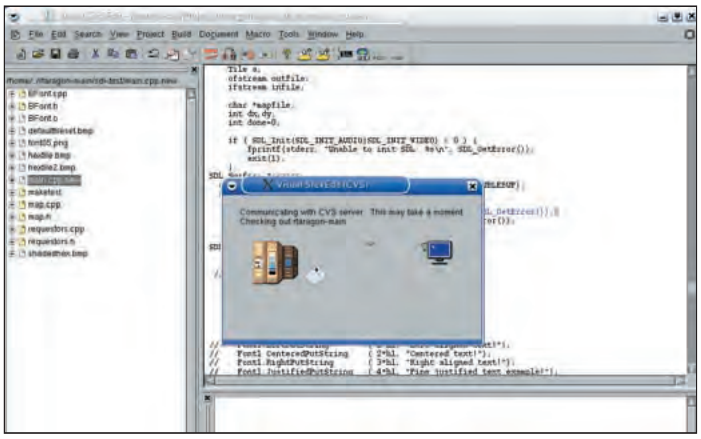

```
    File a;
    ofstream outfile;
    ifstream infile;

    char *mapfile;
    int dx, dy;
    int done=0;

    if ( SDL_Init(SDL_INIT_AUDIO|SDL_INIT_VIDEO) < 0 ) {
        fprintf(stderr, "Unable to init SDL: %s\n", SDL_GetError());
        exit(1);
    }

//      Font1.CenteredPutString    ( 2*hl, "Centered text!");
//      Font1.RightPutString       ( 3*hl, "Right aligned text!");
//      Font1.JustifiedPutString   ( 4*hl, "Fine justified text example!");
```

Communicating with CVS server. This may take a moment
Checking out naragon-main

Visual SlickEdit (CVS)
File Edit Search View Project Build Document Macro Tools Window Help

**Working with CVS repositories is quite straightforward once you have configured them properly.**

## Conclusion

This is a very mature product and has grown to encompass all sorts of features. If this were an end-user application, it would probably be guilty of offering too much in the way of options, but for the technically adept developer market it's great. The ability to work with many languages across many platforms makes it the most versatile programming tool you are likely to find. LXF

## VERDICT

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

**Open Source IDEs may be really good, but this is a professional all-round performer. Shame about the GUI...**

**LINUX FORMAT RATING**

**9/10**

OPEN-SOURCE DATABASE SYSTEM

# MySQL 4

LINUX FORMAT TOP STUFF AWARD

**With MySQL 3.23.x having been available for three and a half years, many thought that MySQL 4 would *never* get out the door. Paul Hudson eats his hat, and investigates why MySQL 4 was worth the wait...**

**BUYER INFO**

World's most popular Open Source database. Also consider PostgreSQL.

- **DEVELOPER** MySQL AB
- **PRICE** Free
- **WEB** www.mysql.com

Mention MySQL to other database administrators and you'll either receive a response along the lines of, "Pah – MySQL? It's a toy database", or you'll find the person you're speaking to is a long-time MySQL admirer. To put it simply, few people think MySQL is "OK" – you generally either love it or hate it.

For a long time MySQL has been the most popular Open Source database for any platforms, and it now averages over 27,000 downloads a day from the official web site alone. Its popularity is partly due to the strong co-operation between PHP and MySQL that has led most new PHP programmers to use MySQL for their database development work, but also largely on its own merit; when it's all said and done, MySQL is a remarkably speedy and powerful database management system that often outperforms its competitors – proprietary or otherwise – by quite a significant margin.

The one drawback often cited about MySQL, however, is that it achieves its incredible performance by only the supporting parts of the SQL language that are fast and easy. To a large extent, that's quite true – but how have the MySQL development team responded to this? With v4 of the product being in development for such a long time, did the developers take the opportunity to turn things around and push the technology envelope for the new release?

## MySQL 3.24?

Generally speaking, the MySQL developers only change version numbers when something major changes. For example, upgrading from v3.22 to v3.23 gave us the new MyISAM table type, the *unique* column restraint, and the *show table status* query to get back table info. Over the period of its lifespan, some fairly large fixes and additions were made to the 3.23 branch, and yet they only ever managed to get a 0.0.1 version increment. So, what has caused the MySQL team to jump versions from 3.x to 4.x? Well, not all that much, to be frank – at least not what many users were hoping for.

New features in v4.0 include a query cache, support for *InnoDB* in all versions, improved full-text searching, *union* support in *select* statement, dynamic server variables, and multi-table *delete* statements. Each of these are big improvements for the database as a whole, and will make life much easier for database admins who upgrade. However, as you can see there is still no sign of triggers, stored procedures, and sub-queries, or even foreign key support and transactions for MyISAM tables.

To be fair, the MySQL team have laid plans down for future versions to include much of the above functionality. Sub-queries, for example, are slated to be in v4.1, with stored procedures and triggers coming in v5. So, it's we should probably be happy with what we have currently, because the *really* good stuff is on the horizon.

## Performance improvements

For most users, the query cache will provide a gigantic performance immediately – with no tweaking whatsoever, you can see query performance double or triple. Query caching works by simply storing the results returned by a *select* query in a special cache, alongside with the query given to generate that result. When the same query comes in (it's compared byte-for-byte with the original query), and if the data returned hasn't changed since the first query, MySQL simply returns the data held in its cache – no need to get table locks, do any calculations, check indexes, or even stray from RAM. Whenever a write is made to a query-cached table, any cached results that rely on that table are flushed so that stale information is never returned. If you're thinking "Well, caching all those results must take up a lot of CPU time – what if I do far more *updates* and *deletes* than *selects*?", then you needn't worry – caching takes no extra effort on the server's part, because it has already done the hard work in the actual query calculation. Furthermore, through the use of special modifiers to the *select* statement, you're able to force MySQL to not cache certain queries, which means if you're selecting some data that you know changes regularly you can stop it filling up useless space in the query cache.

Of course, MySQL exceeds at doing *selects* - the MyISAM table is lightning fast if your *selects* outweigh your *updates* and *deletes* by a large margin, and the query cache only extends this performance.

## Querying improvements

The improved full-text searching is wondrous – MySQL now has built-in support for boolean operations in the full-text search criteria, which means if users type **+** before a word, MySQL only returns rows that include that work. There's also support for **–** (exclude words), brackets (sub-expressions), ~ (downgrade word), and even double quotes to allow users to enter exact phrases. This is some very powerful technology made available for everyone to use, and this should really prompt large-scale adoption.

## Enterprise MySQL

### Pushing the corporate envelope

**MySQL development continues to drive forward, and recently two big announcements were made – firstly, MySQL has been ported to the AMD64 architecture, which should make for a substantial speed improvement given the amount of 64-bit mathematics in databases, and secondly, MySQL AB is in the process of developing a MySQL certification program, that would certify successful learners as being able to make the best use of MySQL features. *"The MySQL database with Linux and the AMD Opteron processor deliver high performance computing power and the scalability that customers require to run the enterprise and stay competitive."*** – Rich Heye, AMD

**The certification will come at two levels: *MySQL Core Certification* and *MySQL Professional Certification*. A user who holds the *MySQL Core Certification* must have demonstrated solid proficiency in SQL and MySQL maintenance, as well as knowing the best way to extract data from the system. Holders of the more advanced *MySQL Professional Certification* are further qualified to install, secure, and manage a MySQL system, as well as being able to optimise the system for maximum performance.**

Also added is support for the *union* keyword in *select* statements, which is a rarely used command to allow queries to be combined together into one result before being returned. This shouldn't be confused with sub-queries – *selects* tied with *unions* are run separately and have no impact on each other, and then returned as if it were one group of data.

## And that's not all...

Two changes that seem minor on the surface are the addition of hot variable changing and the inclusion of *InnoDB* support in the standard release of MySQL (not just MySQL-max). Hot variable changing allows administrators to change server variables *while the server is running*. Previously this required a restart of *mysqld* to make the changes kick in, however that's old news now, and system adminstrators the world over can rejoice.

*InnoDB* support has always been something of a peculiar thing in MySQL, and hopefully inclusion "as standard" (that is, in all releases) will clear up a lot of the confusion surrounding it. The problem is that when MySQL uses the *InnoDB* table type it supports transactions, foreign keys, row-level locking, and other such goodies - all the key data technologies that people claim MySQL lacks. Previously people often just didn't realise MySQL had all this functionality available to them, partially because *InnoDB* support came only in the special "Max" version of the server, and also because MySQL defaults to its own, ACID-incompliant MyISAM table type. With *InnoDB* support being standardised, everyone who uses MySQL 4 will have the option of using InnoDB tables, and thereby take advantage of its extra functionality – hopefully this should drive the point home that MySQL isn't as undeveloped as many think.

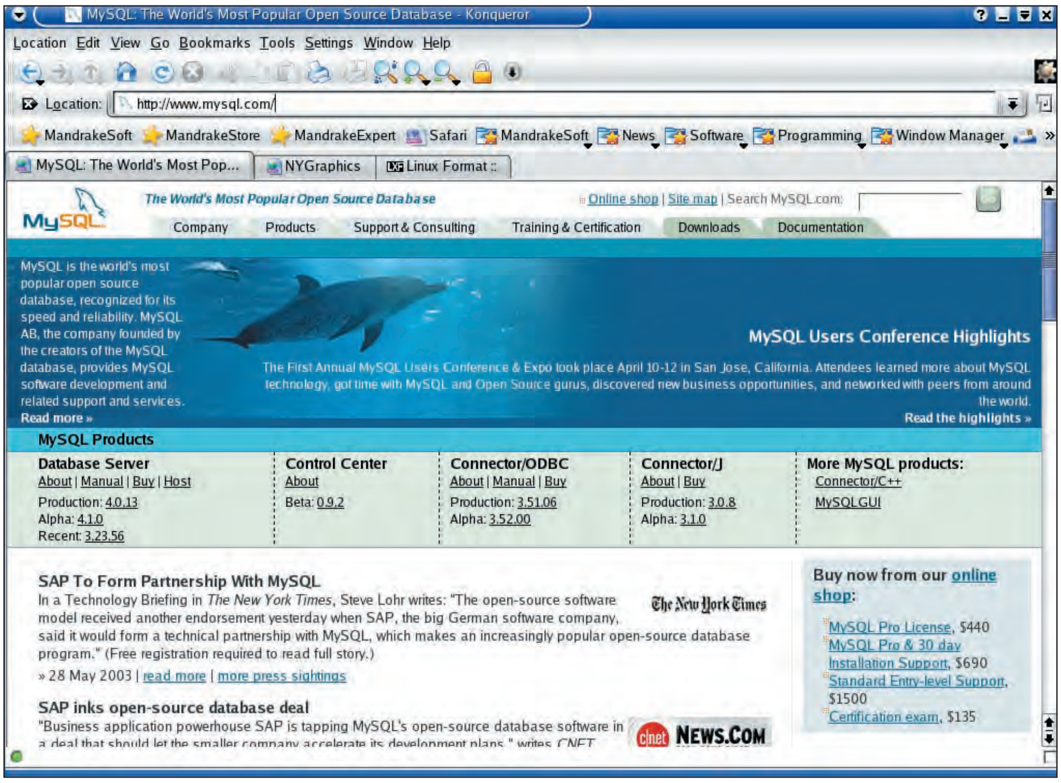
MySQL: The World's Most Popular Open Source Database - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help
Location: http://www.mysql.com/
MandrakeSoft MandrakeStore MandrakeExpert Safari MandrakeSoft News Software Programming Window Manager
MySQL: The World's Most Pop... NYGraphics Linux Format ::
MySQL
Online shop | Site map | Search MySQL.com:
Company Products Support & Consulting Training & Certification Downloads Documentation

*The World's Most Popular Open Source Database*

MySQL is the world's most popular open source database, recognized for its speed and reliability. MySQL AB, the company founded by the creators of the MySQL database, provides MySQL software development and related support and services.
Read more »

**MySQL Users Conference Highlights**

The First Annual MySQL Users Conference & Expo took place April 10-12 in San Jose, California. Attendees learned more about MySQL technology, got time with MySQL and Open Source gurus, discovered new business opportunities, and networked with peers from around the world.
Read the highlights »

**MySQL Products**

| Database Server | Control Center | Connector/ODBC | Connector/J | More MySQL products: |
|---|---|---|---|---|
| About \| Manual \| Buy \| Host | About | About \| Manual \| Buy | About \| Buy | Connector/C++ |
| Production: 4.0.13<br>Alpha: 4.1.0<br>Recent: 3.23.56 | Beta: 0.9.2 | Production: 3.51.06<br>Alpha: 3.52.00 | Production: 3.0.8<br>Alpha: 3.1.0 | MySQLGUI |

**SAP To Form Partnership With MySQL**
In a Technology Briefing in *The New York Times*, Steve Lohr writes: "The open-source software model received another endorsement yesterday when SAP, the big German software company, said it would form a technical partnership with MySQL, which makes an increasingly popular open-source database program." (Free registration required to read full story.)
» 28 May 2003 | read more | more press sightings

The New York Times

**SAP inks open-source database deal**
"Business application powerhouse SAP is tapping MySQL's open-source database software in a deal that should let the smaller company accelerate its development plans," writes *CNET*

cnet NEWS.COM

**Buy now from our online shop:**
- MySQL Pro License, $440
- MySQL Pro & 30 day Installation Support, $690
- Standard Entry-level Support, $1500
- Certification exam, $135

**The MySQL website continues to go from strength to strength, and you're well advised to make good use of it.**

## Building for the future

A lot of the key changes in MySQL 4 aren't user-visible, but are important nonetheless. Indeed, MySQL 4 has been touted as the stepping stone towards oft-requested functionality such as stored procedures – that is, MySQL 4 has restructured how the engine works behind the scenes in order to allow big changes to be made in subsequent releases. As such, there's already an alpha version of MySQL 4.1 available, supporting sub-queries no less, hot on the heels of the 4.0 release. The MySQL developers have never been known to rush releases out, and that's reflected in the fact that they produce an incredibly stable and bug-free product. While sub-query support and stored procedures are still sorely lacking, most of us should be able to live with the promise that they are coming "real soon now".

## Conclusion

MySQL 4 doesn't add a single one of the most requested features for the product, but it still manages to be a great improvement over its predecessors, which says a lot, really.

If you're using *MyISAM* tables under heavy load, the query cache will be your best friend: give it 100MB of space in your RAM, and it will give you performance you would never have believed possible. All being well, InnoDB table support by default should make non-*MyISAM* tables more popular, which should in turn help legitimise MySQL in the eyes of its critics. Although if you're a regular user, consider the table types carefully – *MyISAM* is often the fastest table type by a long shot.

An important consideration is that there are no backward compatibility issues to worry about when upgrading from 3.23 to 4.0 – you can make the change immediately, keep your old data, and see immediate performance improvements.

So, overall, MySQL 4 is a big improvement over earlier versions, and you'd be crazy not to make the switch. It's this kind of attention to detail and thoughtful tweaks that has put MySQL where it is, and means it's likely to stay the world's most popular open-source database for quite some time to come. No, it doesn't give you a few things you were probably hoping for, but they're *en route* – honest! LXF

## An excellent reference guide for an excellent database

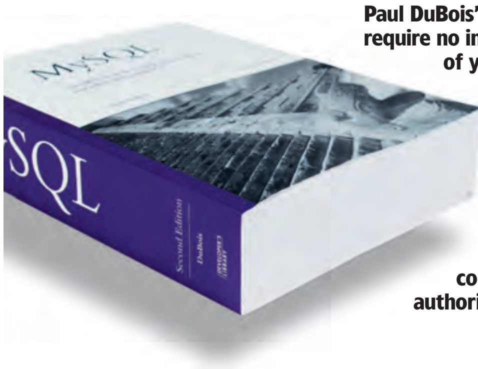

**Paul DuBois' *MySQL* book will require no introduction for those of you who were smart enough to purchase the first edition, and now, in its second edition, the book has been entirely revised and updated for MySQL 4.0. The book continues to be the authoritative reference to all things about the MySQL server – from basic newbie information like how to run queries, through how the MySQL optimiser works, and even onto how to use the MySQL API in your own programs. For more information, read the full review from last issue, however the short version is this: if you're a regular MySQL user at any level, this book is for you – it's no wonder that it won our coveted *Top Stuff* award in issue 41!**

## VERDICT

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

**It's MySQL 3.23 with go-faster stripes painted on, plus enough assorted extras in there to make us love it.**

**LINUX FORMAT RATING**

**9/10**

2U SERVER

# Xinit Sharq SPS234

LINUX FORMAT TOP STUFF AWARD

**What do you get when you cross a dual 2.8GHz Xeon, 2GB of RAM, over 200GB of hard disk space, and Linux? When 1U systems just aren't up to the job, Paul Hudson calls in a heavyweight...**

**BUYER INFO**

Powerful, high-end 2U server designed to rock your world wherever you choose to put it. For smaller uses, consider the less expensive SR113 (reviewed *LXF 40*)

- **SUPPLIER** Xinit
- **WEB** www.xinitsystems.com
- **PRICE** £4,640 (inc. VAT)

Altogether, 1U servers are powerful little things, and generally come in at such a great price as to make them irresistible for admins wanting a quick-fix "set and forget" box. We've reviewed two excellent 1U machines in the last two issues, and both have been quite impressive in their own right, showing that your money can be stretched a surprisingly long way nowadays. However, 1U can only go so far – after all, you're only able to fit three hard drives in there, which, for some, can be a big limitation. So what do you do if you need more power – if you want to push the technology envelope even further? Well, the answer's simple: make a computer that's twice as big!

A 1U server is a pretty chunky piece of kit, however a 2U server, being twice the size of a 1U server, is about as big as you can get before you start needing to put things on wheels. The Sharq SPS234 unit is big and black, with an array of flickering green lights at the front that do as much to calm in an "everything's alright" way as they do to hypnotise you. Being a 2U box, the Sharq has a total of six 36GB disks inside chained together with an Ultra-320 SCSI controller with 64MB of on-board cache – the kind of hardware that manufacturers should really paint Racing Green to reflect the kind of speed at which these things move. At the front there's a serial port for terminal piping, two USB ports so that plugging in a mouse and keyboard needn't be a chore, and also the mandatory CD-ROM/floppy drive.

Naturally, the focus is the hard drives – all of them 10000RPM. They work smoothly in unison and together should ensure an excellent amount of hard disk throughput – we'll have to give them a good going over when it comes to the benchmarks!

2U servers are generally considered the domain of high-end power machines that can take quite a heavy load without complaining, because once you get into the domain of large hard drive stripe sets, performance can get very high very quickly. Generally, smaller companies don't have the space for these things, particularly as "mini-racks" are becoming more popular, and also because many more companies have now implemented 1U firewall units to protect their Net connections.

## Compatibility

Xinit are one of the few companies who happily say "we support Linux without question." And not just *one* distro – they will happily pre-configure servers for Red Hat, SuSE, or even Debian! Why is it that more companies don't offer this? As Xinit guarantee compatibility with three quite different distributions, it shows they have consistently quality, well-supported devices inside the box, which makes the server seem all the more worth its price tag. As far as we're concerned, if a machine can run Red Hat, SuSE, *and* Debian flawlessly, it will almost certainly run any Linux distribution that's fairly recent

The machine we had bundled came with Red Hat 8 (Psyche) pre-configured to take advantage of the multiple CPUs, hyperthreading, and the RAID array. It also came installed with pre-configured copies of *Apache* and MySQL, although they weren't switched on by default so as not to breach security.

To further sweeten the machine, Xinit provide every server they ship with a 3-year next business day on-site service warranty at no extra cost, which can further be upgraded to same-day engineer callout if you want absolute peace of mind.

## Noisy but speedy

This level of Linux support shows that Xinit has clearly got its finger on the pulse of the marketplace and are trying hard to make servers that fill practical requirements.

I had this thing under my desk, and I think most of the office knew when I powered it on – and most of the city knew when I started running benchmarks on it. However as is always the case with rack-mounted servers, they aren't designed to be either ergonomic or silent. So, quiet it isn't, but who cares - this thing performs like a cheetah on steroids.

To kick off, hard drive performance: oh my goodness. This thing returned a score of 8.55, which makes it over eight times faster than our standard machine when pushing the hard drive system to the maximum. What kind of system could possibly require such speed? Well, given that hard drives are generally the bottleneck behind any system, *every* system could really do with this kind of I/O performance. Of course, it comes at quite a hefty price, but, *wow* do you get some real bang for your buck!

The excellent hard disk performance can't help but have a knock-on effect on most of the other tests (excluding oggenc), so it's not surprising to see the *Apache* test return 3.66 and MySQL return an even higher 3.92. Both of these take extensive advantage of the dual-Xeons inside the machine, which further helps the score up. Remember that

this thing uses dual-Xeons, so hyperthreading kicks in automatically making the kernel report four CPUs instead of just two. Naturally, adding hyperthreading to the mix does provide quite a performance boost, but certainly not by 100% – much of this machine's processing performance comes from the fact that the CPUs are simply *fast*: 5,600,000,000 operations a second should be enough for anybody in today's market!

The MySQL score equates to just under 5500 *update* and *select* queries a second, which is a mind-bogglingly high amount. Not content with just having mind-bogglingly fast results, we wanted to push the server to its limits – bring on MySQL 4! If you read the review on page 28, you'll have seen MySQL implements a new query cache to make the system *really* fly. So, we installed it, and did a straight *select* query without the cache - 18,000 queries a second (gulp!). Fast enough for you? Next we installed the *query cache*, and the Sharq server handled almost 35,000 queries a second. That's such a massive figure it needs saying again: the Sharq SPS234 can handle thirty-five thousand queries a second – that's so fast it makes me feel guilty getting up to make coffee every so often!

Compilation performance isn't as high as it could be, but remember that GCC isn't very optimised for the P4/Xeon CPU, and also it doesn't take advantage of the multiple CPUs. CPU-heavy performance on the Sharq *is* good, as can be seen in the exceptional oggenc result.

With an overall score of 4.34, this is definitely a monster of a machine. While I can't think of all that many businesses that need 35,000 queries a second from just one box, it's a safe bet to say that this machine will keep system administrators happy for quite a few years to come.

Not everyone needs this level of performance, however, and as such Xinit allow you to customise your server when you purchase it – you can even go for a uniprocessor option if you want to get the price down as much as possible.

## Extras and weight

Around the back we've got two Intel gigabit network cards built-in, which switch down to 100 and 10-BaseT when they need to. There's also a SCSI-III VHDCI connector to allow fast interconnecting to the external world – not a feature you see all that often these days, sadly. Still, this is a bit of a luxury server, and it's good to see that Xinit haven't taken the opportunity to hike the price up unnecessarily.

With regards to weight, this server clocks in at a massive 22kg – enough to make anyone's back sore. It is designed to be mounted on a rack, and not a flimsy rack at that; it's not worth risking this machine by just leaving it under a desk (like we did!), or, worse, putting it on an old rack that's had its day.

## Conclusion

With such excellent disk speeds and raw CPU performance, this is a real beast of a machine. Coupled with no less than 2GB of RAM, this is starting to look like a bargain for the price. This is the kind of machine you need if you want to be 100% certain that your servers will be able to cope with your load – there's no grey area with this machine where you might think "Hmm, this isn't performing so well"; it's all positive, and all best-of-class.

At the price it's not the kind of machine anyone would buy at the drop of a hat, and with competition at such a high right now, it's probably worth checking what the competition are offering at the same price. As Dell are competing so strongly for every sale available, it's possible that Xinit may consider a small cut if you ask nicely. However, even if they aren't eager to make any deals, the server is still good value at the price. A benefit that shouldn't be overlooked is Dell's consistent refusal to support anything other than Windows, which is where Xinit have a large advantage – no one knows their systems better than the people who built them, so it's a big advantage to have those same people pre-configure and pre-tune your server for maximum performance.

So, on the whole, this is a heck of a speedy machine which represents the pinnacle of quality 2U server technology. If you find yourself requiring three or four cheaper 1U systems, the SPS234 will almost certainly handle the equivalent workload in half the rack space. I wonder if Xinit have any spares they might like to donate to me... LXF

## BENCHMARKS

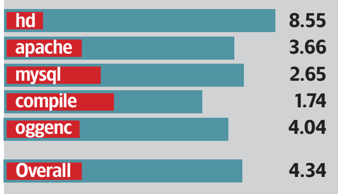

## VERDICT

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

**Performance through the roof with a price tag that won't require you to re-mortgage.**

**LINUX FORMAT RATING**

**10/10**

NETWORK ACCESS HARDWARE

# Neoteris Instant Virtual Extranet Access 3000

**As a sysadmin, giving users secure external access to your internal network is a nightmare, so why hasn't anyone made it easy? Paul Hudson reckons this new Linux appliance changes the extranet paradigm...**

**BUYER INFO**

1U Linux appliance providing 'set and forget' extranet capabilities.

- **SUPPLIER** Neoteris
- **PRICE** £25,000 (inc VAT)
- **WEB** www.neoteris.com

So, the untrusted Internet - full of bad guys – comes into your router, your router goes through your firewall, and then your firewall splits off into a trusted domain full of your business-critical systems and also into your De-Militarized Zone made up of web servers and other insecure functionality. Sound familiar? This has been the networking paradigm for so many years now that people often just take for granted that this is the best layout. Under the watchful eye of its chairman, Jim Clark (of SGI and Netscape fame), Neoteris is a fledgling company with a whole new idea for handling the situation. Their aim is to provide a Linux appliance, designed and developed by them with large chunks implemented using Open Source software, that provide secure extranet functionality out of the box. Their Instant Virtual Extranet (IVE) range is split into three distinct models: the 1000, the 3000, and the 5000, in order of power.

## What's in the box?

Each machine is kitted out with some fairly standard hardware inside, and one of the primary differences between the ranges is the amount of horsepower inside each box. Also inside the box is a folded A3 page of instructions that gets your server up and running in four steps. I'm a stickler for documentation as you may well know, so I was feeling a bit cheated by this flimsy amount of documentation – a £25,000 machine should come backed with some *serious* treeware.

Then, I turned the machine on, and, if I were less honest, I would delete the last sentence and replace it with, "The documentation is far too long, and could almost have been summed in one step, 'Turn it on. Job done.'" More on that soon – first, the rest of the box contents. For the first time in a long time – and, really, more companies should do this – I was happy to find a complete print-out of the GPL in there, which shows Neoteris isn't afraid to pay homage to its foundations. Also in there is the usual warranty and support information, plus a short Neoteris licence.

Enough of that – I think you all want to know how it works. Shame, really, because I could happily write hundreds of words about the beautiful design of the unit – "Neoteris" is clearly a strategy as well as a name!

## Configuring & using

You can configure this thing all you need, if you want to, however it's not really necessary. Configuration is done with a direct terminal connection from the serial port on the front of the unit and is simply a matter of answering questions as prompted.

Once the text-based configuration is done and the server is plugged into your network, you can get to it through your web browser, which is a far easier way to change options. One particularly welcome feature is a Cobalt-like patch installer that uses your browser to snag pre-packaged service patches that are automatically applied to the box to keep it up to date. This always worked well for Cobalt users, so it's good to see Neoteris are making use of good ideas from the past.

Administering the box is made very simple, and it lives up to its aim of being an appliance very well. Like with your toaster or washing machine you don't need to know or care about how this thing works inside – you just plug it in and it will handle the rest. There *is* a lot of value-added stuff here, which is good as it helps justify what many would think is quite an exorbitant price tag. To be blunt, it *doesn't matter* what hardware is inside the machine, or what version of software is installed, because the box is designed to perform one task and perform it to a specification. In the same way that people don't buy a fridge because it's got a branded "Ultra 3000+" condenser, you don't buy an IVE based upon its specs – all that matters is that it smoothly handles a set number of users.

This is some *very* high-end hardware, no doubt about it. Plugging one of these things in lets your outside users connect to through the IVE to Citrix or Oracle, or even their Exchange server, without the need to worry about encryption or security. The Access 3000 is a solid example of what Linux was designed to do, and it is to work in the background and get things done. While the price of this machine might make some baulk, it is a great deal less than most companies are already paying for the equivalent of this technology, with the difference being that the Neoteris solution has no further hidden costs.

If you've already committed a lot of cash to your extranet solution, this is certainly something to consider – particularly if you find money is still leaking away due to ongoing costs. It's a powerful box that's "fire and forget" in every sense of the word. While the price tag might make you think twice, you should definitely consider an evaluation – the chances are this thing beats your current extranet's total cost of ownership hands down. LXF

**VERDICT**

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

**Powerful, predictable, reliable, and lots of other praiseworthy "ables" too.**

**LINUX FORMAT RATING 9/10**

# Maximum Security

**So hard-hitting and hacker-esque that the author is "Anonymous"! Paul Hudson is impressed...**

LINUX FORMAT TOP STUFF AWARD

**BUYER INFO**

- **PUBLISHER** Sams
- **AUTHOR** Anonymous, *et al*
- **ISBN** 0-672-32459-8
- **PRICE** £36.50

MAXIMUM SECURITY
A HACKER'S GUIDE TO PROTECTING YOUR COMPUTER SYSTEMS AND NETWORK
FOURTH EDITION
SAMS

Now in its fourth edition, *Maximum Security* continues to best a bestseller in the field of "hacker's guide" security books. Quite rightly, too – despite this book being marked for user level "Intermediate to Advanced", it's wholly accessible for readers at all levels. If you're a beginner and want to become a seasoned security veteran in an evening, this is the book to read!

It starts by explaining various security concepts: how to assess risk, how to design an incident response policy, what tools hackers and crackers use, all the way up to how to judge your own internal security level. The first 100 pages makes for powerful reading, and gives readers a solid diving board from which they can dive into other chapters as they wish.

Of course, readers would be foolish to miss chapter two, which discusses technical details about TCP/IP, how spoofing works, and it also, crucially, makes a great effort to dispel security myths – things on which system administrators often waste their time (and money of their employers).

To be honest, pretty much the whole book is the same – "oh, you *must read*...", "readers would be *crazy* if they missed..." – this book is one you can read from start to finish without feeling you've wasted a moment. *Two hundred pages* are devoted to operating system-specific problems and fixes, and there really is nothing missing there. Other chapters discuss wardriving (including an excellent section titled "Practical WEP Cracking"), firewalls, mailbombing, virii, and DOS attacks.

Overall, I have only one complaint about *Maximum Security*, and it is very serious indeed: this book is so damn good at explaining security to any reader with time to spare that it runs the risk of maybe putting some security professionals out of work!

Sams has really cracked the market with this book, through and through – if you've ever had a worry about your system security, this will certainly be the best £40 you'll ever spend on a way to find a solution.

**VERDICT**

**Very, very, very, very, very, very, *very* good. I wish there was a superlative that covered books as good as this!**

**LINUX FORMAT RATING** 10/10

# Red Hat Linux 8 Unleashed

**1000 pages of text can't be wrong, can it? Paul Hudson gets stuck in.**

**BUYER INFO**

- **PUBLISHER** Sams
- **AUTHOR** Bill Ball, Hoyt Duff
- **ISBN** 0-672-32458-X
- **PRICE** £36.50

Hoyt has been writing for *Linux Format* for twenty issues now, and still manages to write a most cutting (and often controversial!) column for our news pages – as such, he's more than qualified to co-write such a known-authoritative text as Sams' *Red Hat Linux Unleashed* series. The other author, Bill Ball, has written several books, including previous versions of *Red Hat Unleashed*, and so the two make a very strong writing pair that rarely if ever say "might", "could", or "should" – this a book full of certainties.

At over a thousand pages, this book is very clear on most topics; as I read it through I occasionally thought "oh, they could have said more about that", only to turn over the page to find a special note about precisely what I was thinking. Generally speaking you only get that kind of quality out of a second or third edition of a book, which reflects very well on the technical editor, Dave Taylor.

By the time you reach Chapter 5, *First Steps with Linux*, you've already gone through about a hundred pages of introduction and install help, which even manages to cram in info about business considerations for system administrators.

The book has quite a strong focus on system administrators, but seems to mix in general home user information somewhat clumsily. For example, there are two pages on IRC stuck between configuring a local news server and programming in C! Furthermore, information on choosing an email reader is wedged at the end of the chapter *Handling Electronic Mail*, the first half of which is dedicated to configuring *sendmail* and *fetchmail*. I think it would be perhaps better in subsequent releases of this book series if they cleaned things up so that home users didn't have to wade through information they weren't interested in. For instance: how to read your email should probably come *before* how to configure your PC as a DHCP server, before *Editing httpd.conf*, and, yes, even before *ACID Compliance in Transaction Processing to Protect Data Integrity*. Furthermore, if I was a new reader, I would not look up *email, MUAs* to find what programs are available – "email, reading", perhaps, or, if I was a particularly switched on Windows migrant, "email, clients".

Don't get me wrong – *Red Hat Linux 8 Unleashed* is a top book if you are a power user, but it probably isn't the best bet if you're looking for a book to get started with Red Hat 8.

The fact that Red Hat themselves have superceded this work with the release of Red Hat 9 shouldn't be too much of a worry as the changes are fairly superficial.

**VERDICT**

**"Red Hat Linux 8 for Experts" would have been a better name – beginners should stay well clear!**

**LINUX FORMAT RATING** 7/10

# Practical Unix & Internet Security

**Paul Hudson reviews an in-depth expert guide to Unix security**

**BUYER INFO**

- **PUBLISHER** O'Reilly
- **AUTHOR** Simon Garfinkel, Gene Spafford, and Alan Schwartz
- **ISBN** 0-596-00323-4
- **PRICE** £38.95

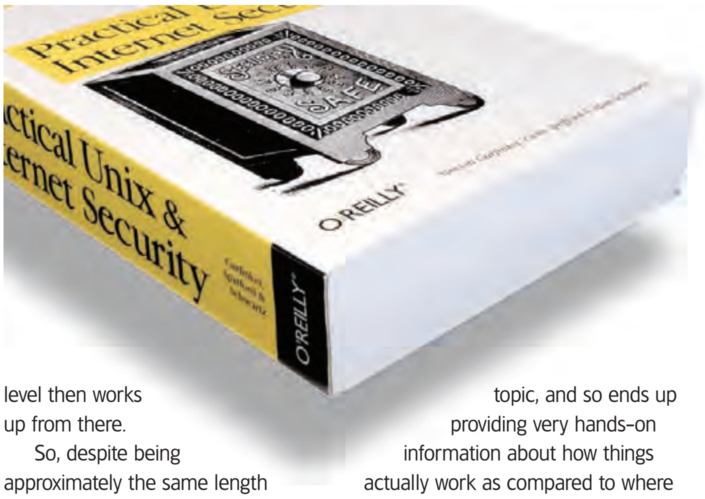

A quarter of a million people have found earlier editions of this book useful, or at least that's what the blurb of this book would have us believe. Quite frankly, I wouldn't be surprised if that were a conservative estimate!

If you read the *Maximum Security* review on the opposite page, you'll have seen a product that was an excellent all-round introduction into computer security - indeed, as a general book, I have yet to see one better. This book takes a different angle on things. Firstly, it's Unix only, so don't expect coverage of anything but Unix and Unix-alike (Linux) operating systems and tools. Secondly, there's little if any pandering to beginners - the book starts at an intermediate level then works up from there.

So, despite being approximately the same length as *Maximum Security*, this book is an entirely different kind of animal – and that's *not* just a cheap pun on O'Reilly's tradition of using critters on their covers!

With the subject matter being much more tightly controlled (more advanced readers only, and Unix-only as well), the book has much more room to go into great depth about its topic, and so ends up providing very hands-on information about how things actually work as compared to where potential weaknesses are. Yes, weaknesses are covered in-depth too, but a lot of the message of this book is "be a good system administrator", which is no bad thing – there aren't anywhere near enough books that try to curtail your sloppy coding or admin practices as they go along.

If you want a really hands-on guide about your computer, and aren't afraid of having to really get into a great deal of depth, then *Practical Unix & Internet Security* is a good read, and a great reference. Unlike *Maximum Security,* which has quite a chatty overall approach to the subject and is structured in quite a linear fashion, you really can keep sliding back into this book again and again whenever you want to find out something new.

If you own any O'Reilly books already, you'll be familiar with the style employed and the way in which information is presented in an easy-to-find manner. This book certainly won't disappoint, which for many will sum it up entirely.

**VERDICT**

**A great reference guide to *nix security, but generally expects quite a competent reader**

**LINUX FORMAT RATING**

▮▮▮▮▮▮▮▮▮▯ **9/10**

# Webmaster In a Nutshell (2nd Edition)

**Amias Channer dives into O'Reilly's updated quick reference for all things web.**

**BUYER INFO**

- **PUBLISHER** O'Reilly
- **AUTHOR** S Spainhour and R Ekstein
- **ISBN** 0-5960-0357-9
- **PRICE** £24.95

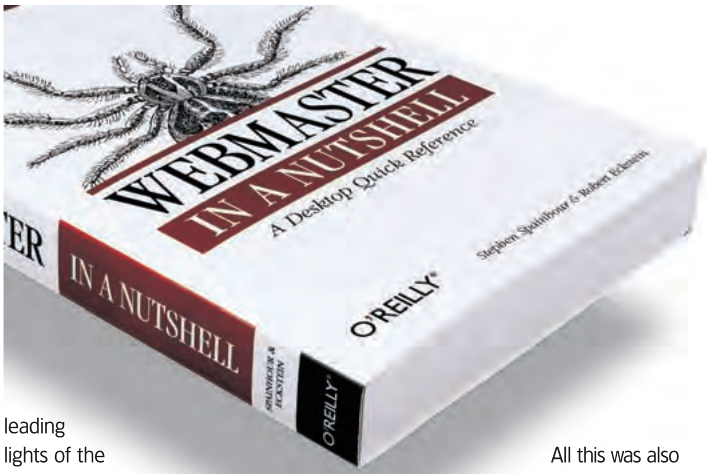

Most computing books are hefty reference tomes that don't yield their knowledge without a serious level of commitment – this usually means they don't get much desktop action. Good solid reference work has a place; when learning something new, explanations of the concepts of programming are important, but once read, they're just not convenient for quick lookup . What is needed is an easily portable quick-reference guide for Open Source web development. O'Reilly has achieved this admirably with the ideal combination for web programmers: HTML, PHP, JavaScript, XML, Apache, CSS , HTTP, Perl CGI and mod_perl .

Each having its own section giving brief but concise descriptions from the leading lights of the Open Source development community of how to do things the right way. When combined with API listings and expansive breakdowns of important functions and constants, it makes this book very useful indeed; and the O'Reilly brand is evident throughout. There's lots of extra touches like tab marks for each chapter, clearly spaced contents and a useful website.

All this was also the case two years ago for the first edition, and the O'Reilly website errata page for the second edition has grown quite busy. If you have the first ed, it's worth updating – the technologies have also matured since the last version and with all of the *Apache2* changes, so the original 1st ed ultimate quick reference is showing its age. The updated *Apache* configuration section now contains information on both versions , sadly the mod_perl section doesn't document mod_perl2, but given that its interface is still quite sparse, you could argue that this is fair. The XML section has been expanded to cover XSLT and Xpath which is a essential given their complex syntaxes. The PHP section covers version 4 but lists functions for both versions. The other sections are mostly unchanged .

A deluxe version of this book was available with CD-ROM versions of five related O'Reilly titles, but it seems to be in short supply having been snapped up on sight by anyone who knows! Almost all of the O'Reilly books are available to browse online using its 'Safari' pay-per-view system, so you can get any other information you need quite easily.

**VERDICT**

**The internet programmers bible reincarnated!**

**LINUX FORMAT RATING** ▮▮▮▮▮▮▮▮▮▮ **10/10**

# Web Server Programming

**Amias Channer hears the strange and unearthly voices from the server-side...**

**BUYER INFO**

- **PUBLISHER** Wiley
- **AUTHOR** Neil Gray
- **ISBN** 0-470-85097-3
- **PRICE** 29.95

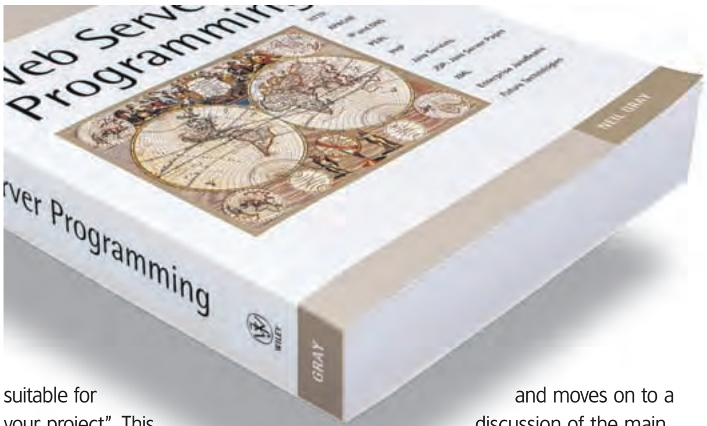

Server-side programming paradigms have, over recent years, grown into the kind of monster a Greek tragedian could only have dreamed of – hack off one protocol and another sprouts back in its place! Choosing the right technologies is always important, and for Internet businesses this can be a life or death issue. This blur of competing technologies has created regular tribal battles on the Net and made objective comparisons of their relative merits very difficult.

Unfortunately, this book doesn't seem to manage to crystallise this situation or even offer any conciliatory humour to soften the inevitable maxim: “only you will know whether this is suitable for your project”. This fatal flaw can be put down to the book being written by an academic who, while very knowledgeable, just doesn't temper the sheer weight of all the information with enough real-world experience or examples to appear truly authoritative in such a rapidly changing subject.

The book starts on the basic infrastructure for any web application and moves on to a discussion of the main languages in use on the server-side: Perl , PHP, Java (servlets, JSP and EJB) , XML (SOAP, XSLT, DOM). Some ASP and .NET information is included as appendices. Seeing as most of the languages described are Open Source or open-standards, the complete lack of information on what it means in specific cases is a glaring ommission.

The back cover suggests its target market is “existing practitioners and students” but the poor quality indexing, dense paragraphs and large chapters are not suited to people needing information in a hurry – students or recent graduates however will probably find the lecture note style and questions more useful than professionals seeking a desk-side reference. Read chapter by chapter, this book provides quite a useful introduction to the general concepts but rarely points to appropriate man pages, and going any further would require dedicated books for each language.

**VERDICT**

**Great idea in theory, but not specific enough to any of the technologies nor target markets to be worth the money.**

**LINUX FORMAT RATING** ▮▮▮▮▯▯▯▯▯▯ **4/10**

# HotPicks

The best new open source software on the planet!

On the CD and DVD

**Mike Saunders**

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

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

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

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

## HotPicks at a glance

| | |
|---|---|
| MPlayer | 38 |
| Twin | 39 |
| KMyFirewall | 39 |
| Quick Spam Filter | 40 |
| Pan | 40 |
| Nessus | 41 |
| Dillo | 41 |
| Freelords | 42 |
| Barrage | 42 |

## HOTPICKS AWARD

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

MEDIA PLAYER

# MPlayer

LINUX FORMAT HOTTEST PICK

■ **VERSION** 0.90 ■ **WEB** www.mplayerhq.hu

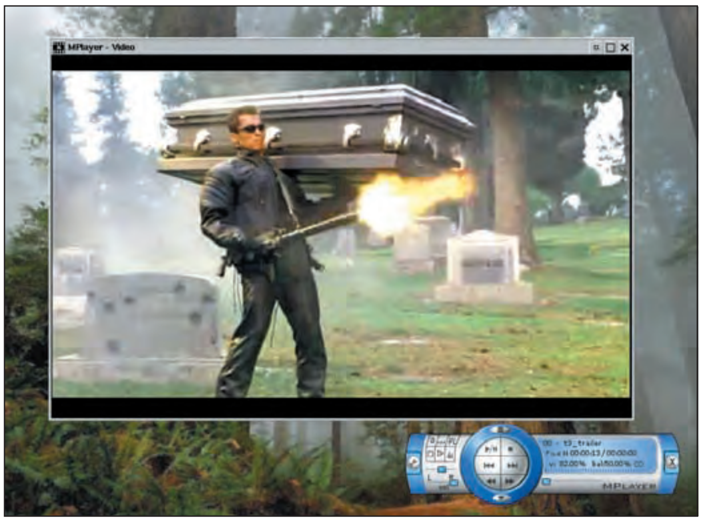

**Media playback is no longer a problem for Linux – here's Arnie shrugging off the heart surgery in the latest *Terminator 3* trailer, with the GUI beneath.**

As Linux users, we're forced into a compromise when it comes to the latest flashy gadgets and technology – hardware support generally isn't as immediate as with Windows, and using certain file formats can be troublesome. Of course, we make this trade-off for openness, stability and security, but Linux's growing popularity on the desktop has resulted in some fantastic new software to co-exist with other OSes and file formats. *Mozilla* and *OOo* are two major projects spearheading the desktop assault; equally, *MPlayer* is now a huge part in making Linux attractive to end users.

*MPlayer* is an Open Source media player tool which supports outrageous numbers of file formats, either through its in-built libraries or external Windows codecs. It can currently play (deep breath): MPEG, AVI, QuickTime, RealMedia, VOB and Microsoft's ASF/WMA/WMV formats, along with some more esoteric types like VIVO, NuppelVideo, RoQ and PVA. Also, VCD, DVD and DivX (3, 4 and 5) movies are catered for, along with Ogg and MP3 sound files. It's an enormously impressive range, and covers just about all popular media formats floating around on the Internet today.

## MPlay it again, Sam

Although ready-built RPMs are available, compiling from source provides the most flexibility; support for various formats and features can be toggled on or off, and GCC '2.96' users will have to pass a special flag to the configure script (following a row between the *MPlayer* developers and Red Hat). Before building, you'll need to grab the appropriate codec packages from the project's website – extract these files (mostly DLLs) into /usr/lib/win32, and the configure script will spot them and tailor the player to match. Finally, if you want a GUI instead of the vanilla command-line interface, make sure to pass the **--enable-gui** flag as it's disabled by default.

Once up and running, simply entering **mplayer <file>** at a shell prompt will start the app and begin playing the specified media file – no frills, no snazzy splash screens. If you've built the GTK-based GUI version, **gmplayer** is the command you'll need (a skin needs to be extracted in /usr/local/share/mplayer/Skin first). By default the output device will be a plain X window (supporting fullscreen mode), but it can also use SDL, OpenGL, DGA, XV, VESA, DirectFB and even AAlib (yes, ASCII-art rendering!). For sheer speed, some specialised graphics card drivers have been thrown in too.

*MPlayer*'s default GUI theme is Blue, which presents a borderless window littered with the usual control buttons. It's reasonably attractive and gives access to a decent prefs box where many of the command-line options can be enabled visually. Grumblingly, the native GTK file selector hasn't been used and there are a few other glitches to be found in the interface, but on the whole it's a pleasant show.

Among the program's vast array of options are OSD (on screen display, *ie* messages and subtitles) support, audio and video filters (sound resampling, image scaling/flipping/rotating/cropping etc.) and keyboard bindings, along with stacks of others. All are explained very thoroughly in the heavy-going but competently written manual page, and the supplied HTML documentation goes into even greater depth. Aside from the player itself, an extra *mencoder* utility is included for taking a movie file and encoding it into a different format.

Ultimately, *MPlayer*'s importance as a Linux desktop application can't be overstated; after a lot of hard graft and code snapshots, the developers are approaching version 1.0 and have provided the community with a fast, free and incredibly flexible media player. It's a doddle to use and we found no serious problems with stability; a bit more spit-shine on the interface would be welcome, though. Most importantly, *MPlayer* means that Linux users don't have to reboot into Windows (or mess around with WINE) just to watch certain video files. Superb.

## TEXT-MODE WINDOW SYSTEM

# Twin

■ **VERSION** 0.4.6 ■ **WEB** http://linuz.sns.it/~max/twin/

XFree86 (see our cover story this month) and the commercial X Window System implementations are used on just about every non-server installation of Linux – it's the *de facto* standard GUI system. Still, it's not always appropriate for every situation, and the wealth of excellent console apps means that some still prefer to work in a few virtual terminals with *screen(1)* and friends. *Twin* is a windowing system for the console, supporting a client-server architecture (*a la* X) and themeing.

Who would want such a tool? Old 486/P1 boxes with stuffy graphics chips often lack the oomph to run X adequately, yet can perform very well with text-based apps. Similarly, the X protocol can be too sluggish over dialup or slow broadband links, making SSH (or Telnet at a push) the better choice.

With no unusual dependencies, *Twin* is a breeze to build and will pick up on the X and GTK (1.2.x) libraries if installed to make a somewhat snazzier X front end. After running **./configure** you may find that **make** dies with an error; simply enter **make** again to resolve this. Similarly, if the man page isn't installed correctly, you'll find it in the docs/ directory.

When started at the text console, *Twin* throws up an empty screen with a menu bar along the top – hitting the Pause key or holding the right mouse button brings this up *Amiga Workbench*-style. You'll need GPM running at the console to make use of the mouse there, but under X it'll work regardless and the titlebars are enhanced with a fully graphical style. Windows can be moved, resized and rolled-up, and UI elements can be tweaked.

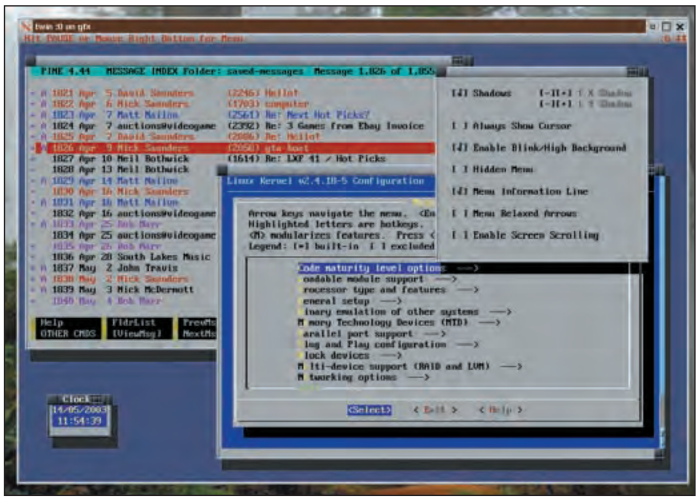

```
twin :0 on gfx
PINE 4.44   MESSAGE INDEX Folder: saved-messages  Message 1,826 of 1,855
  1824 Apr  7 auctions@videogame  (2392) Re: 3 Games from Ebay Invoice
  1827 Apr 10 Neil Bothwick       (1614) Re: LXF 41 / Hot Picks
  1828 Apr 13 Neil Bothwick
  1832 Apr 16 auctions@videogame
  1834 Apr 25 auctions@videogame
  1836 Apr 28 South Lakes Music
  1837 May  2 John Travis
  1839 May  3 Nick McDermott

Linux Kernel v2.4.18-5 Configuration
Arrow keys navigate the menu.
Highlighted letters are hotkeys.
<M> modularizes features.  Press
Legend: [*] built-in  [ ] excluded
  Code maturity level options  --->
  Loadable module support  --->
  Processor type and features  --->
  General setup  --->
  Binary emulation of other systems  --->
M Memory Technology Devices (MTD)  --->
  Parallel port support  --->
  Plug and Play configuration  --->
  Block devices  --->
M Multi-device support (RAID and LVM)  --->
M Networking options  --->

[✓] Shadows
[ ] Always Show Cursor
[✓] Enable Blink/High Background
[ ] Hidden Menu
[✓] Menu Information Line
[ ] Menu Relaxed Arrows
[ ] Enable Screen Scrolling

Clock
14/05/2003
11:54:39
```

***Twin* running under X, with a few terminals and the Options box.**

Twin sports some X-like networking code for running certain utils remotely, and a full session runs well over SSH/Telnet from another Linux text console (keyboard only). Those trying it on an old box will want **vga=ask** or similar on their kernel command line (or a framebuffer) to get something better than the 80x25 default, though. It's an intriguing little project, executed well, and having overlapping windows and a sweet interface makes a nice change from full-screen VCs.

## FIREWALL CONFIGURATOR

# KMyFireWall

■ **VERSION** 0.9.6 ■ **WEB** http://kmyfirewall.sourceforge.net

Last month in *Hot Picks* we looked at *Jay's Iptables Firewall*, a text-based dialog-driven firewall config tool to help in securing a system. Since then we've had a few requests to look at KDE's graphical alternative, *KMyFirewall*, which is attracting much attention and is on its way to being an official KDE program. At the moment it has worked its way into the 'Extra Gear' add-on bundle, which hopefully means we'll see it fully supported in the upcoming KDE 3.2 release.

Providing you have an up to date KDE and *Qt* installation (at least 3.x for each), kernel 2.4 and a recent iptables binary, compilation should be straightforward enough. Some functionality will only be available if your kernel has certain features enabled, but most stock distro-supplied kernels should be OK. The tool assists in editing rules and chains for the filter, NAT and mangle tables, based on interface, protocol, IP and MAC addresses, state and more.

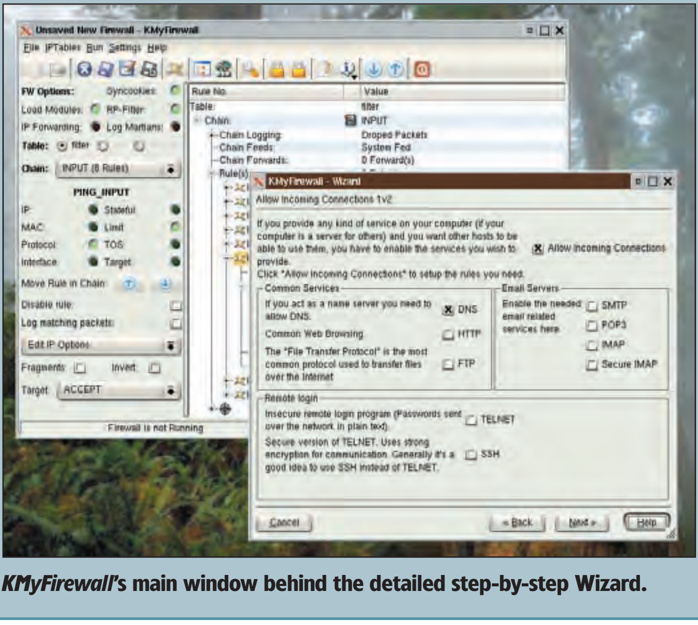

Unsaved New Firewall - KMyFirewall

| Rule No | Value |
|---|---|
| Table: | filter |
| Chain: | INPUT |
| Chain Logging: | Droped Packets |
| Chain Feeds: | System Fed |
| Chain Forwards: | 0 Forward(s) |

FW Options: Syncookies ● Load Modules: ● RP-Filter: ● IP Forwarding: ● Log Martians: ●

Table: ⦿ filter ○ ○

Chain: INPUT (8 Rules)

PING_INPUT

Target: ACCEPT

Firewall is not Running

KMyFirewall - Wizard

Allow Incoming Connections 1v2

If you provide any kind of service on your computer (if your computer is a server for others) and you want other hosts to be able to use them, you have to enable the services you wish to provide.
Click "Allow Incoming Connections" to setup the rules you need.

☒ Allow Incoming Connections

Common Services

If you act as a name server you need to allow DNS. ☒ DNS

Common Web Browsing ☐ HTTP

The "File Transfer Protocol" is the most common protocol used to transfer files over the internet ☐ FTP

Email Servers

Enable the needed email related services here. ☐ SMTP ☐ POP3 ☐ IMAP ☐ Secure IMAP

Remote login

Insecure remote login program (Passwords sent over the network in plain text) ☐ TELNET

Secure version of TELNET. Uses strong encryption for communication. Generally it's a good idea to use SSH instead of TELNET. ☐ SSH

***KMyFirewall*'s main window behind the detailed step-by-step Wizard.**

Following the splash screen and disclaimer (the authors rightly state that the program shouldn't give a false sense of security!), *KMyFirewall* probes the system for its startup scripts (spotting Gentoo's non-standard setup), network interfaces and required binaries, and pops up a box to alter the findings if necessary. Then appears the slightly overcrowded main window – it could prove to be somewhat intimidating to the less experienced, but the supplied documentation does well in explaining raw *iptables* commands and good security practices.

*KMyFirewall*'s most notable feature, though, is the Wizard which steps through phases of configuration; each option is admirably well explained and they're grouped together understandably. Logging, incoming and outgoing connections, trusted and totally blocked hosts are all dealt with, and handily the generated script can be viewed before finishing. Additionally, the script can be exported for use on other boxes as a time-saver.

In all, *KMyFirewall* is satisfyingly comprehensive with its range of support and the tree layout keeps things manageable. It's a smart choice for visually tuning all but the most advanced firewall setups, and the Wizard is great for newcomers who need a quick and simple solution. Recommended.

The inclusion of this program was a result of reader requests. If there's some software that you'd like us to include on the coverdiscs, post your requests in the appropriate forum at **www.linuxformat.co.uk** or email us at the usual address.

SPAM FILTERING TOOL

# Quick Spam Filter

■ **VERSION** 0.5.0 ■ **WEB** www.ivarch.com/programs/qsf.shtml

Hormel Foods Corp, makers of the popular 'spiced pork and ham' tinned meat, don't have a problem with their product's name being used to describe junk email – all publicity is good publicity, as they say. However, most of us wouldn't mind receiving thirty tins of tasty food in our mailboxes every day; instead though, we have to suffer Nigerian bank scams, Viagra adverts and other unwanted and often offensive nonsense.

And the problem's getting worse – according to spam research/protection outfit MessageLabs, one in every 199 UK emails could be classified as spam in January 2002. By June 2002 the figure was one in 36 and now (May 2003) it stands *below* one in five.

Free email accounts and endless open SMTP relays on the Net combine to make the situation quite dire, so spam filtering software like *QSF* is always good to see.

Many previous attempts at spam filters have been based on searching for keywords, giving an alarming number of false positives (proper emails incorrectly marked as spam). *QSF* puts into place a system described in a recent paper by Paul Graham called *A Plan for Spam*. This is based on a Bayesian filtering system: you initially gather together a large bundle of spam messages, and another of bona fide emails, run them through the filter, and it builds a database. After some time, it should become quite accurate at identifying spam and marking it as such.

The small program (written in C for speed) takes a single email, tries to identify if it's spam or genuine, and then marks it as such with an 'X-Spam' header. Consequently, it's best used by procmail (mail processor); in a typical setup, you might retrieve your messages from a POP server with fetchmail, pass it through procmail and make it dump the spam in a separate folder. Before that, you'll need to train it with **qsf --train spam non-spam** for two mbox (not maildir) folders.

We tested it with an initial 75-message spam folder, and even after the first training run it proved successful – it identified some (but not all) spam messages when we later grabbed our new emails. Over time, with more training, it should become even more accurate, and best of all it's fast, unintrusive and mostly transparent to work with. If you're drowning in a sea of 'anti-ageing pill' adverts, give it a look. If you're stuck for gift ideas for your computer-using friends and family, why not visit **www.spamgift.com/store/commerce.cgi**?

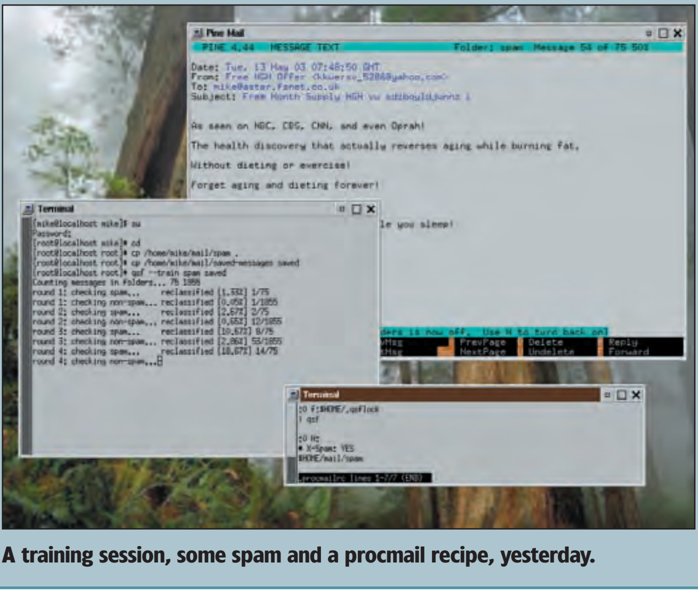

```
PINE 4.44   MESSAGE TEXT                    Folder: spam  Message 54 of 75 50%

Date: Tue, 13 May 03 07:48:50 GMT
From: Free HGH Offer <kkuersv_5286@yahoo.com>
To: mike@aster.fanet.co.uk
Subject: Free Month Supply HGH vu adibouldjunnz i

As seen on NBC, CBS, CNN, and even Oprah!

The health discovery that actually reverses aging while burning fat,

Without dieting or exercise!

Forget aging and dieting forever!

le you sleep!

ders is now off.  Use H to turn back on]
```

```
[mike@localhost mike]$ su
Password:
[root@localhost mike]# cd
[root@localhost root]# cp /home/mike/mail/spam .
[root@localhost root]# cp /home/mike/mail/saved-messages saved
[root@localhost root]# qsf --train spam saved
Counting messages in folders... 75 1855
round 1: checking spam...      reclassified [1.33%] 1/75
round 1: checking non-spam...  reclassified [0.05%] 1/1855
round 2: checking spam...      reclassified [2.67%] 2/75
round 2: checking non-spam...  reclassified [0.65%] 12/1855
round 3: checking spam...      reclassified [10.67%] 8/75
round 3: checking non-spam...  reclassified [2.86%] 53/1855
round 4: checking spam...      reclassified [18.67%] 14/75
round 4: checking non-spam...
```

```
:0 f:$HOME/.qsflock
| qsf

:0 H:
* X-Spam: YES
$HOME/mail/spam

.procmailrc lines 1-7/7 (END)
```

**A training session, some spam and a procmail recipe, yesterday.**

NEWS CLIENT

# Pan

■ **VERSION** 0.14.0 ■ **WEB** http://pan.rebelbase.com

Despite the ever-growing popularity of Web-based forums, IRC and email, newsgroups are still some of the best places to engage in friendly banter or violent flamewars. Covering such diverse topics as food and pet care, USENET is still frequented by many thousands of Netizens - but to use it, of course, you need a newsreader. These clients exist in all shapes and sizes within the open source community, and perhaps the most popular graphical one is *Pan*; we've looked at this some time ago as veteran *LXF*ers will recall, and with the progress it's making, it's certainly worth an update.

Earlier *Pan* releases were dependent on the GNOME libraries, but now the coders have moved to pure GTK 2. You'll also need the *gnet* library built and installed to compile Pan, but that aside there are no other unusual dependencies, and the optional *gtkspell* bits will enable a spelling checker if the script finds it.

Thanks to its long development history, Pan has acquired a very polished front end – the Setup Wizard that appears on first start is a sweet touch, and it steps through basic configuration (identity, NNTP server etc). The main window follows the familiar three-pane convention – each pane is resizable and the icons are suitably GNOME-esque, while colours are used liberally for message headers, quotes and signatures. It's fairly straightforward to get started with, even though the huge number of options and menu items maybe appear a bit threatening to the newsgroup first-timer.

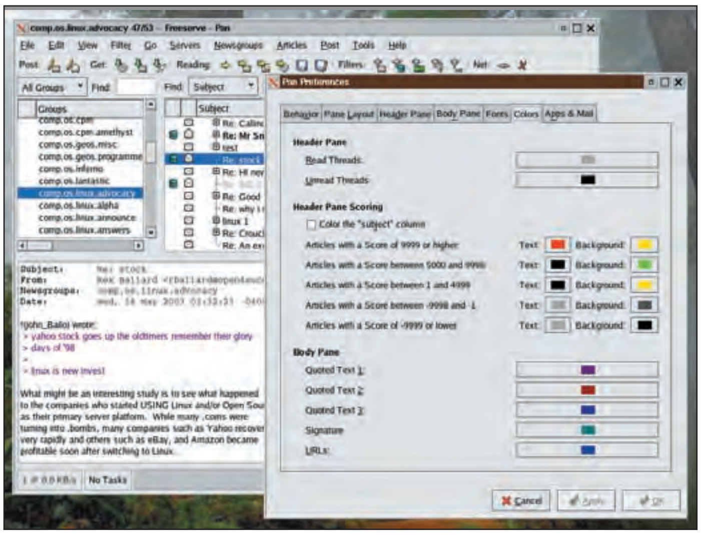
comp.os.linux.advocacy 47/53 – Freeserve – Pan
File Edit View Filter Go Servers Newsgroups Articles Post Tools Help
Post Get Reading Filters Net
All Groups Find Find Subject
Groups
Subject
1 @ 0.0 KB/s
No Tasks
Pan Preferences
Behavior Pane Layout Header Pane Body Pane Fonts Colors Apps & Mail
Cancel Apply OK

comp.os.cpm
comp.os.cpm.amethyst
comp.os.geos.misc
comp.os.geos.programme
comp.os.inferno
comp.os.lantastic
comp.os.linux.advocacy
comp.os.linux.alpha
comp.os.linux.announce
comp.os.linux.answers

Re: Calin
Re: Mr Sn
test
Re: stock
Re: HI nev
Re: Good
Re: why i
linux 1
Re: Crouc
Re: An ex

```
Subject:     Re: stock
From:        Rex Ballard <rballard@opensource
Newsgroups:  comp.os.linux.advocacy
Date:        Wed, 14 May 2003 01:33:31 -0400
```

John_Baloi wrote:
> yahoo stock goes up the oldtimers remember their glory
> days of '98
>
> linux is new invest

What might be an interesting study is to see what happened to the companies who started USING Linux and/or Open Sou as their primary server platform. While many .coms were turning into .bombs, many companies such as Yahoo recover very rapidly and others such as eBay, and Amazon became profitable soon after switching to Linux.

**Header Pane**

Read Threads: ▨
Unread Threads: ■

**Header Pane Scoring**

☐ Color the "subject" column

| | | | | |
|---|---|---|---|---|
| Articles with a Score of 9999 or higher | Text: | 🟥 | Background: | 🟨 |
| Articles with a Score between 5000 and 9998 | Text: | ■ | Background: | 🟩 |
| Articles with a Score between 1 and 4999 | Text: | ■ | Background: | 🟨 |
| Articles with a Score between -9998 and -1 | Text: | ▨ | Background: | ▨ |
| Articles with a Score of -9999 or lower | Text: | ▨ | Background: | ■ |

**Body Pane**

Quoted Text 1: 🟪
Quoted Text 2: 🟥
Quoted Text 3: 🟦
Signature: ■
URLs: 🟦

**comp.os.linux.advocacy is home to some very disturbing flamewars.**

Boasting full threading, sorting and filtering (using *regexps*) features, alongside message caching, multiple servers and external Web and mail clients, *Pan* is jam-packed with enough goodies to make a power-user's day. It'll happily download MIME and unencoded binary attachments, with images being shown inline in the message, and the task manager keeps track of large download operations.

All in all, *Pan* is a wonderful newsreader with more bells and whistles than almost anyone could need, yet it remains friendly and simple through its good design. If you enjoy perusing newsgroups but want a slick upgrade from *Slrn* and its friends, try it out.

## SECURITY SCANNER
# Nessus

■ **VERSION** 2.0.5 ■ **WEB** www.nessus.org

Making sure a server is fully locked-down after installation can be a messy and complex job. It gets harder if a box is offering multiple services, but thankfully there are tools around to hunt down any possible weak points and gauge a machine's overall vulnerability. *Nessus* is a mature contender with some interesting features – most notably a plugin system and in-built scripting language, and uses the popular Nmap for the port scanning side.

Although vanilla tarballs are available for building from source, the project's site offers a script containing the compressed code - just **chmod +x** it and run as root. After a few questions it'll build and install the program, plugins and additional tools, and if you have X and GTK on the machine it'll build a fancy front-end for the client. Once installed, you'll need to create a user account and SSL certificate as described, before finally running **nessusd -D**'

*Nessus* is based around a client/server architecture; essentially, one box does all the scanning work, and you connect to it from a client anywhere. This is particularly useful when working over a dialup/slow link - the network intensive stuff can be left to a more endowed box somewhere else. Once logged in, you're offered a list of plugins to enable during the scan. These cover various services and vulnerabilities, with more detailed features tweakable, and they're all explained thoroughly with a 'Risk factor' level and possible solution.

The particularly aggressive tests (which may cause problems on the server being probed) can be disabled, and the sheer range is overwhelming - enough to fill half of this mag. After the portscan (Nessus makes no assumptions about services running on their expected ports) and vulnerability identification, the Report box offers an overview which can then be exported in text, LaTeX, HTML and other formats.

Overall, *Nessus* is an amazingly comprehensive security scanner with enough tests to satisfy the ultra-paranoid. Apart from some slowdowns we found when managing the plugins, it performed solidly and the documentation for each option is excellent. A great choice for tightening a server.

You should refer to the *Nessus* website if you're serious about using this program, and includes instructions on how to make Nessus work with Mac OS X. The project welcomes any and all constructive input from users.

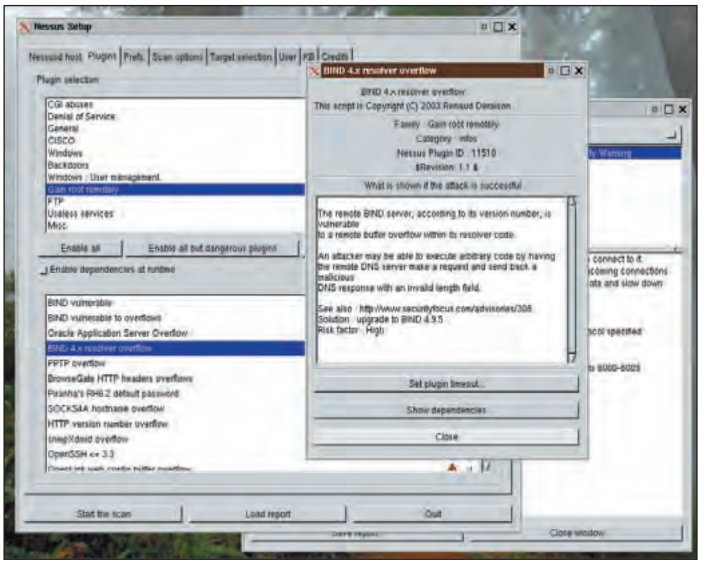

BIND 4.x resolver overflow
This script is Copyright (C) 2003 Renaud Deraison
Family : Gain root remotely
Category : infos
Nessus Plugin ID : 11510
$Revision: 1.1 $
What is shown if the attack is successful :

The remote BIND server, according to its version number, is vulnerable
to a remote buffer overflow within its resolver code.

An attacker may be able to execute arbitrary code by having the remote DNS server make a request and send back a malicious
DNS response with an invalid length field.

See also : http://www.securityfocus.com/advisories/308
Solution : upgrade to BIND 4.9.5
Risk factor : High

**The GTK-based client in action, showing the plugins list and details of a potential problem.**

## WEB BROWSER
# Dillo

■ **VERSION** 0.7.2 ■ **WEB** http://dillo.auriga.wearlab.de

Long gone are the days when *Netscape Navigator* was the only real choice for a graphical browser on Linux. Now we have three first-rate clients in the form of *Mozilla, Konqueror* and *Opera*, with many smaller projects also showing promise. Of course, the main problem with that bunch is their heaviness on system resources; *Opera* is the zippiest, but there's still room for a lightweight fully graphical browser (*Links* and *Lynx* are top-notch text-moders).

*Dillo* has been around for some time now, but releases are becoming more regular and it's starting to win many fans. Being based on the *gzilla* rendering engine – in the same way *Mozilla* is built on *Gecko* - it's small, speedy and a tiny download. The only major dependency you'll need to build from source is the GTK 1.2.x toolkit libraries, which are installed by just about every modern distro. *Dillo* has even been compiled on handhelds like the Psion 5 and iPAQ – living proof of its diminutive overhead.

Starting up amazingly quickly, Dillo's browser window holds the usual furniture expected of a WWW client; a location bar, button bar and status line are all present. Since earlier releases the bookmarks have been moved from a menu into a separate page, and they can be edited on-the-fly too. To the right of the toolbar is a progress box which shows how many images have been loaded and the size of the page.

*Dillo* can render JPG, GIF and PNG images, and copes well with tables and small forms. Cookies are supported too, although there's no JavaScript implementation yet and frames (now rarely used, admittedly) won't work either. Altogether, this makes *Dillo* fine for light browsing of Slashdot, eBay and other sites of moderate complexity, but for FTP and HTTPS connections you'll need to look elsewhere.

Naturally, *Dillo* is no replacement for *Mozilla* or *Konqueror*, but works marvellously well for simple browsing tasks and the rendering engine rivals Links for pure speed. We've had it running for some time on an old Pentium 75 laptop, and it has been perfectly usable on such a low-spec box – just the ticket for your second or third machine, maybe.

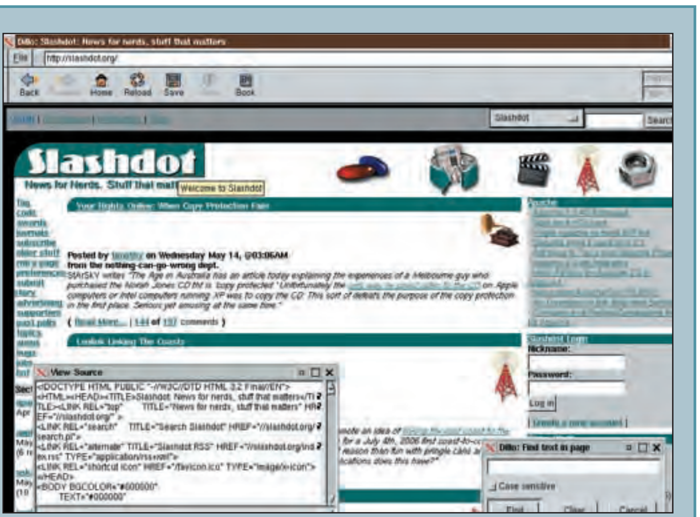

Posted by timothy on Wednesday May 14, @03:06AM
from the nothing-can-go-wrong dept.
StArSKY writes "The Age in Australia has an article today explaining the experiences of a Melbourne guy who purchased the Norah Jones CD that is 'copy protected.' Unfortunately the only way he could listen to the CD on Apple computers or Intel computers running XP was to copy the CD. This sort of defeats the purpose of the copy protection in the first place. Serious yet amusing at the same time."
( Read More... | 144 of 197 comments )

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 3.2 Final//EN">
<HTML><HEAD><TITLE>Slashdot: News for nerds, stuff that matters</TITLE><LINK REL="top"      TITLE="News for nerds, stuff that matters" HREF="//slashdot.org/" >
<LINK REL="search"    TITLE="Search Slashdot" HREF="//slashdot.org/search.pl">
<LINK REL="alternate" TITLE="Slashdot RSS" HREF="//slashdot.org/index.rss" TYPE="application/rss+xml">
<LINK REL="shortcut icon" HREF="/favicon.ico" TYPE="image/x-icon">
</HEAD>
<BODY BGCOLOR="#000000"
      TEXT="#000000"
```

**Astoundingly, there are no dupes on slashdot.org's front page today...**

STRATEGY GAME

# Freelords

■ **VERSION** 0.2.3 ■ **WEB** www.freelords.org

Hundreds of years ago, huge battles with opposing armies and a bit of pillage and plunder was a great way to pass the time. These days it's not so romantic though; conflict is messy and time-consuming, and it's harder than ever to believe that God is totally on your side. Still, turn-based strategy games have remained steady in popularity since olde tymes (well, the 8-bit era anyway), throwing out notable gems like the superb *Castles II* and *Warlords*. And as with most famous titles, it wasn't long before some coders knocked together an imitation in their spare time – *Freelords*, as the name suggests, is a GPLed clone of *Warlords*, and can be played across a network or multiplay on the same PC.

Even though it's still heavily in development, *Freelords* is already looking good as a multiplayer online strategy romp. It's also multi-Platform – BeOS, Linux, Mac OS X and Windows. Compiling from source should be mostly problem-free, providing you have the right dependencies: *ParaGUI* is the most significant, being a GUI toolkit based on SDL, but we've provided that on the coverdisc too (and it may already be supplied with your distro).

At the title screen, you can begin a new quest or get involved in an online game, setting yourself up as the game server or just joining an ongoing game as a client. Starting a new game offers a choice of computer or human players (using names from prominent kernel hackers) along with terrain settings. From here you're thrown straight into the fray, with a scrollable map, command buttons, and menu items up the top (most of which aren't implemented yet). Just take land, build up an army and keep control of your money.

Battles are represented by static character icons with depleting health bars, and although there's little in the way of animation or other superfluous effects, the sprites and scenery are clear and well-drawn. Much of the gameplay mechanics is present and the overall design is explained well in the docs, but whether or not *Freelords* floats your boat purely depends on what you enjoy: no cathartic gun shenanigans here, just planning, cunning and patience.

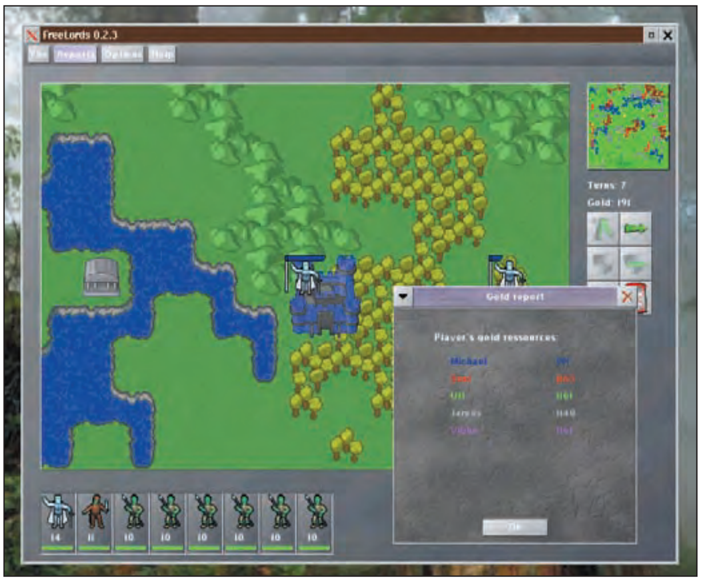

**Gadzooks! Low funds situation = more justification for mindless attacks.**

SHOOTING GAME

# Barrage

■ **VERS** 1.0.1 ■ **WEB** http://lgames.sourceforge.net/index.php?project=Barrage

War, then. "When you have to kill a man it costs nothing to be polite," said Churchill. "So long as there are men there will be wars," quipped Einstein. Even George W Bush seems to concur with this view: "The reason we start a war is to fight a war, win a war, thereby causing no more war!" Luckily, we don't have to be politically powerful to drift into our own fantasy world of maiming soldiers and exploding tanks thanks to games like Barrage.

By his own admission, the author Michael Speck doesn't regard the use of just one gun as being a proper barrage, but he just liked the word. Fair enough – as he says, the game is "rather violent", who cares? Barrage is a mouse-driven action game in which you have to destroy as much of the opposition as possible within a time limit of 3 minutes. It's built around SDL, the cross-platform toolkit favoured by games coders, and will also make use of the *SDL_mixer* library for sound effects too if installed. Other than that, it should compile without hassle and we've provided a Red Hat 8 RPM on the disc (should work on 9 and some other distros).

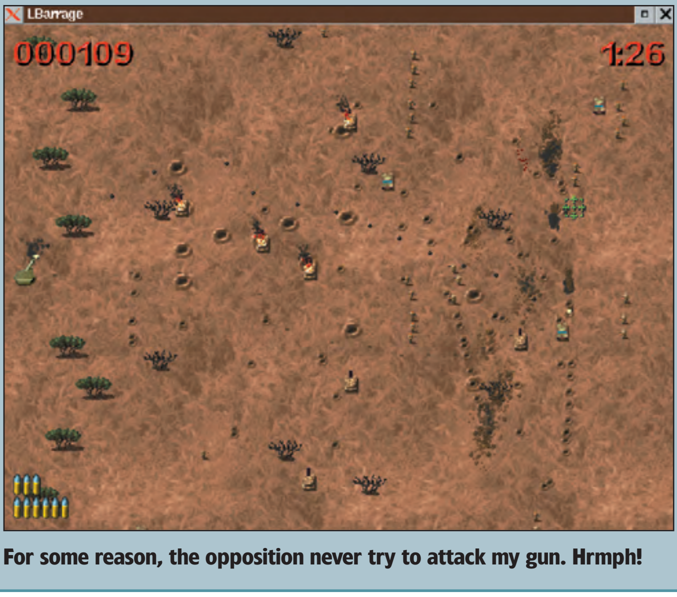

**For some reason, the opposition never try to attack my gun. Hrmph!**

*Barrage*'s gameplay takes place over a rugged desert-ish terrain; peppered with a few trees and craters from your projectiles, it's left wide open for some frantic blasting antics. You control a large grounded gun on the left-hand side of the playing area, and use the mouse to point the crosshair at anything that moves. Jeeps, tanks and infantry wander up and down the screen, so the distance and timing of shots has to be ultra accurate. Imagine *Cannon Fodder* crossed with *Worms*. Well, sort of...

With the left and middle buttons firing large and small grenades respectively, and the right button used for reloading, it's essential to use the correct weapons on the appropriate victims and 'clean house' as efficiently as possible. You'll lose points for hanging around too. Even though it's more of a diversion than a fully-fledged game, *Barrage* is well crafted, polished and a good test of reflexes and hand-eye coordination. Best of all, it's crammed with entertainingly nasty impact noises.

The author's website is well worth a visit for Linux versions of some classic game clones. There's *Ltris*, *LBreakout*, *LGeneral* and links to tons of other time-stealers! LXF

# EVERYTHING YOU EVER WANTED TO KNOW ABOUT X

Work started on the X Window System two decades ago. Can it still cut it at the heart of a modern Linux desktop? **Richard Drummond** gives X some scrutiny...

cover feature

When people discuss the suitability of Linux as a desktop platform, they generally focus their attentions on the features and capabilities of one of the major Open Source desktop projects such as KDE or GNOME. That's not surprising; it is these projects and the applications that are built upon them that provide the direct user experience for desktop Linux. Many desktop users don't know or don't care that layered beneath the KDE and GNOME application frameworks on every Linux installation is a vital piece of infrastructure that creates and manages the graphical, windowed environment in which these desktops run: the X Window System (or X for short). It's actually a testament to X's reliability and ubiquity that it can be safely ignored by the end user. X is middleware, plumbing. And who doesn't take their plumbing for granted, as long as it works? Events of earlier this year, however, have thrust X – and especially its most popular implementation, produced by the open-source XFree86 project – unwillingly into the limelight.

In late February, the release of XFree86 4.3.0 was announced, and all seemed well. This was the first XFree86 release in nearly six months and brought some long-awaited new features (see the box *What's new in XFree86 4.3*). Barely a month after, though, the XFree86 project's Core Team announced that one of their most prolific developers, Keith Packard, the man behind many of the advancements in the XFree86 4.x series, had been dismissed due to alleged misconduct. He had, it was claimed, been seeking support from the X community to fork the XFree86 project; that is, set up an independent project, led by himself, that would develop a separate X release based on XFree86. Shortly afterward Packard responded with a message to a public XFree86 forum stating that he hadn't intended to undermine the XFree86 project but rather he was seeking support for a change in the way that XFree86 was governed. This post, *A Call for Open Governance of X Development*, cited XFree86 failings such as slow release schedules and lack of co-operation with other open-source projects. Packard said, 'Persistent problems in XFree86 development have become widely recognised within the X community... The key issue is that XFree86 is not a community-governed project... Decisions appear to be arbitrary and are not seen to reflect the will of the community.'

## The controversy

The debate that has ensued on the mailing lists and forums that constitute the online Open Source community has often created more heat than light. The raised emotions are understandable: X is such a crucial piece of technology for Linux and Unix, and the XFree86 developers have expended considerable amounts of time and effort on creating and maintaining their implementation. But, amidst all the flaming and trolling, some serious questions have been raised concerning the way X is developed, the role the XFree86 project should play, and the technical problems which need to be solved for X to meet the demands of the contemporary desktop. A few have even asked whether X should be ditched entirely and replaced with something that's more modern.

In the light of all this controversy, this article will examine the current state of the X Windows System and, in particular, its most popular implementation: XFree86. Just what is X and what advantages does X offer to Linux's assault on the desktops of the world? Is X even necessary? Like any piece of engineering, X is not without its problems, so we'll look at

»

> "A few have even asked whether X should be ditched entirely and replaced with something that's more modern."

« the technical challenges facing X and the solutions that have been proposed. We'll also discuss some of the management problems surrounding the XFree86 project. To understand all the issues at stake, however, we need to start at the beginning.

## What is X?

Simply put, the X Window System is the *de facto* standard for providing a graphical environment on Unix and Unix-like platforms, including Linux. (How it became established as such is a story we'll look at shortly and is due largely to a couple of fortuitous design decisions and the availability of its source code, rather than any grand master plan. In this, the popularity of X mirrors the rise of Unix itself.)

Although X is superficially similar to rival graphical platforms such as Windows and Mac OS, the environment that X creates differs in two important ways.

## What's new in XFree86 4.3.0?

### Significant changes and updates

**XFree86 4.3.0 is arguably is the most significant XFree86 release since 4.0.3, which first introduced the *Render* extension. Besides the usual round of bug-fixes and driver improvements – in particular 4.3.0 includes updated Radeon drivers supporting the newer 9000 and 9100 models – it also includes some important technology updates.**

**Firstly, XFree86 4.3.0 includes the new *Xft2* and *fontconfig* libraries to improve font rendering and vastly ease font configuration, both for developers and users. *Xft2* is the X-based wrapper for the Truetype font rasterizer and is part of the new client-side font-handling system for X which replaces the old and inefficient font handling of the core X protocol. It was *Xft* that introduced anti-aliased fonts, based on the alpha-blending capabilities of the Render extension. New in *Xft2* is support for X servers without the *Render* extension, so that it's no longer necessary to fall back on core font handling when *Render* isn't available.**

**Also new is a partial implementation of the new X *RandR* extension, which allows an X display to be dynamically re-sized, rotated and reflected without the need to restart the X server.**

**Finally, support for full-colour, alpha-blended, animated cursors (mouse pointers) is included for the first time (previously, only monochrome pointers could be used), and this is backed up by a new client-side library, *Xcursor*, for handling cursor configuration.**

**Much to the delight of all those that love to theme, XFree86 4.3.0 now supports full-colour alpha-blended, animated mouse pointers.**

Firstly, X was designed from the ground up to be network transparent. That is, X allows you to run a program on one computer and interact with it via its graphical interface displayed on another. This magic is due to the client/server nature of X (see box *X architecture*, right). In X terminology, the display and input devices, such as the mouse and keyboard, are managed by a piece of software called an X server. The programs that you interact with via this X server – known as X clients – may be running on the same computer or on any other host on the network. An X client communicates with the X server via a mechanism called the X protocol, and it is this that lets the client remotely create its interface on the user's display and receive user input.

The other great difference with X is that X really does just provide the core infrastructure for a windowing system. X doesn't specify how applications should look and feel: it's the job of an X client to create the user interface elements such as menus and buttons, either directly via the primitives that the X protocol provides or via an X toolkit. X doesn't even provide a window management policy; that's up to a special X client called a window manager. All this flexibility can be a double-edged sword, however. One the one hand, this approach of 'mechanism not policy' – as the mantra goes – has been key to the endurance of X; on the other, the lack of a consistent look and feel is just one of the plethora of reasons why Linux is having such a hard time on the desktop now.

Before we can understand these differences and the pros and cons they bring to Linux on the desktop, we need to look at a little history.

## The genesis of X

The X Window System began life in 1984 at Project Athena, a research project in distributed systems at MIT whose aim was to make computing resources easily accessible to students. The name X reflects that it drew inspiration from an earlier project called 'W' (presumably short for 'window' – sorry, all you recursive acronym fans!) A platform-independent windowing system was required because MIT couldn't afford to buy all the workstations needed for the project themselves, nor was any one vendor willing to supply sufficient workstations. Thus from day one, X was designed to provide network-transparent windowing across a variety of heterogeneous hardware.

Initial X development was a joint effort by MIT and Digital Equipment Corporation (DEC at the time supplied Unix for their VAX and VAXstatation hardware). Progress was rapid at first, with Version 8 being the first to support colour display hardware and Version 10 the first to be widely deployed. Development culminated with release of X version 11 (or X11) in 1987 which offered vastly better performance, an improved rendering model and support for 32-bit deep colour displays. Note that the X version number mentioned here refers to the version of the X protocol implemented, with different version numbers being incompatible with each other. For instance, an X10 client won't work with an X11 server and *vice versa*. However, the core X protocol has remained largely unchanged since 1987. Only occasional minor revisions have been made in the intervening period, the last significant one being X Version 11 Release 6 or X11R6 in 1994. X11 clients written sixteen years ago will still work on modern X servers, an impressive achievement in backwards compatibility.

DEC began shipping X with its flavour of Unix, Ultrix, and thanks to the portability of X11 and the availability of the source code, other Unix vendors followed suit. The rest, as they say, is history. The X Windows Systems was an open-source project long before that term was coined. Like the UC Berkeley version of Unix – BSD Unix – the X source code was made available under a very liberal license. And, as with the BSD license, the MIT X License allows the unrestricted use, modification and distribution of the X source code and the even the sale of commercial products based on it. However, to ensure that X development wouldn't splinter into incompatible factions – as had happened with Unix – MIT set

## X architecture

### Client/Server, the X protocol and X toolkits

The X Window System is a network-transparent windowing system based on a client-server architecture. The X server is the program which manages the display hardware, input devices and so on, while programs which a user interacts with via the X server are known as X clients. This terminology may seem confusing at first, but it makes sense if you look at things from the point of view of the X client. The client uses services provided by the X server to display a graphical interface and get input from the user.

An X client communicates with the server via a stream-based channel known as an X wire using a protocol called the X wire protocol or just X protocol. This is a device-independent protocol and enables the client to send commands to the server to open and close windows, draw geometric shapes, render text, receive inputs and so on without caring about the specific hardware upon which these requests take effect.

The X protocol can be transported over any suitable medium, and this is what gives X its network-transparency. Usually, any X clients you are using will be running on the same machine as the X server – on the computer or X terminal that you are sitting in front of – and some local method of IPC (Inter-Process Communication) will be employed to construct the X wire. For example, on Unix hosts this will commonly be the Unix socket. But, since the X protocol will work equally well (albeit more slowly) over a network link such as a TCP socket, you can run an X client on any machine on your local network (or even on the Internet, security issues not withstanding, of course!) and still operate it via your X server. Moreover, since all communication takes place via the device-independent X protocol, the X client doesn't need to be running on the same type of host as the X server. It can have a completely different CPU architecture, and the remote host upon which the client is running doesn't even need its own any display hardware. This network-transparency is a unique and compelling feature of X and a major reason for its popularity.

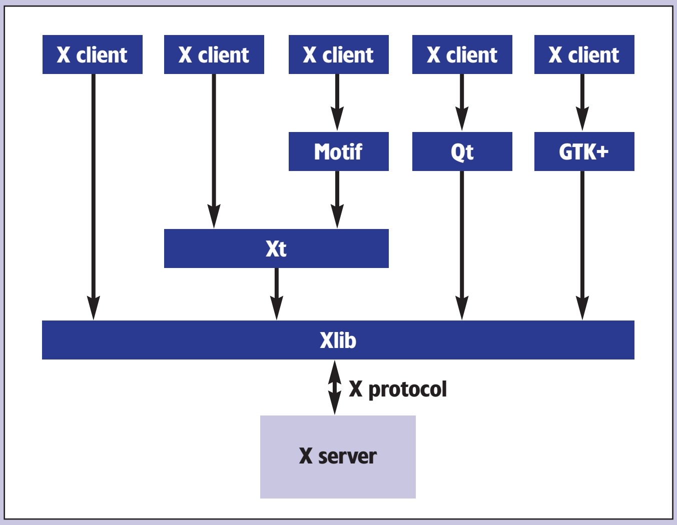

Toolkits – *Xt* is better than *Xlib* for building a GUI as it's higher-level.

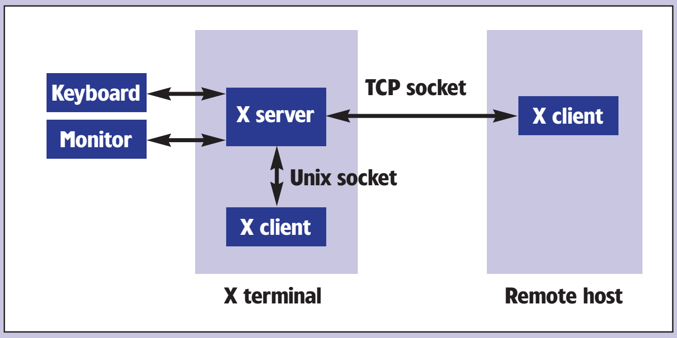

X architecture – device-independent protocol.

### X toolkits

The X Windows System provides a shared library of functions called *Xlib* which can be used by X clients to communicate with an X server without actually having to understand the X protocol. The problem with *Xlib* is that it is very low-level. The functions it offers directly correspond to the requests supported by the X protocol, and thus it is an awful lot of hard work to create a user interface just using *Xlib* on its own.

The other problem is that there is no consistency of look and feel between applications that employ *Xlib* directly, since they have to implement the entire interface themselves. To overcome these difficulties, various GUI toolkits have been created over the years to ensure a measure of continuity in the user experience.

To make it easier to develop GUI toolkits, X includes an additional library called the *X Toolkit Intrinsics* or *Xt*. This provides an object-based infrastructure for implementing user interface elements (known as widgets) in C and C++ and contains many utility functions to ease development. Once popular toolkits based on *Xt* include *Xaw* (the *Athena Widget* library) and *Motif*.

Motif was once the industry standard for GUI programming on X, but as a proprietary system it has fallen out of favour in recent years to Open Source and more modern-looking toolkits such as *Qt* and *GTK+*. (Incidentally, as you can see in the diagram above left, *Qt* and *GTK+* themselves don't go through *Xt* but use *Xlib* directly to construct their widgets.)

up the X Consortium as a standards body to maintain the X11 protocol and to provide a reference implementation. Eventually, as interest in Project Athena at MIT waned, in 1988 the X Consortium was spun off as a separate not-for-profit company, which was finally merged under the wing of the Open Group around 1996 and renamed X.org.

## Getting X for free

Unsurprisingly, as far as Linux is concerned, our story really begins with the rise in popularity of the PC. In the early 1990s, the increasing performance-to-price ratio of the average PC began to make it a platform worth reckoning with. In terms of performance, it was fast catching up with the commodity hardware peddled by the proprietary Unix vendors. This, coupled with the internecine squabbling of the Unix vendors, enabled Windows and especially Windows NT to gain a foothold in the workstation and server marketplace – much to the detriment of Unix and the X Window System. (This was a trend that would continue throughout the 90s until the free Unices and Linux began to reclaim the ground lost to Microsoft.)

The first popular port of X Windows for PC versions of Unix was done in 1990 by Thomas Roell and was released under the name X386. These changes were fed back to the X Consortium and became part of the reference distribution with X11R5. Unfortunately, or perhaps fortunately, this release was rather buggy. It prompted several developers – independently at first – to begin fixing the various problems with X386.

> "MIT couldn't afford all the workstations, thus from day one X was designed to have network-transparent windowing across a variety of heterogeneous hardware."

»

Initially, this group of four – David Wexelblatt, David Dawes, Glenn Lai and Jim Tsillas – released their changes under the name X386 1.2e, where the 'e' stood for 'enhanced', but in April 1992 they set up their own project, under the name XFree86 (a pun on X386 as Roell had taken X386 commercial). Their aim was to produce the best X implementation – commercial or otherwise – for the x86 architecture. Which they did, and then some.

XFree86 was ported to Linux in 1992, and the growth of the XFree86 project owes much to the growth of Linux. But XFree86 soon spread to other platforms, not just Unix, including OS/2 and Windows (the latter thanks to the *Cygwin* library, which implements a Unix-like API for windows) and to other architectures, not just the x86. Many consider today that XFree86 really is the best implementation of X for any platform; it's certainly the most popular.

When commercial Unix lost its direction in the second half of the 1990s, so interest in maintaining the X standard declined. As the X Consortium and then X.org became more and more ineffectual, it fell upon the reluctant shoulders of the XFree86 project to drive forward X development. XFree86 became in the eyes of many the real reference implementation of X and the testing ground for advancements.

"XFree86 was ported to Linux in 1992, and the growth of the XFree86 project owes much to the growth of Linux."

## The X advantage

As we have seen, X is the *de facto* graphical standard for Unix, but is it still relevant to the Linux desktop? Well, X does have its share of detractors. X is, they say, over-engineered, over complex and uses too many system resources. Those whose concern is only in Linux on the desktop naively demand that X's network-transparency be dropped; others say that X has had its day and should be replaced entirely. Should such claims be heeded?

Throughout its life, X's network transparency has been its killer feature. While this may not appeal to the home user running Linux on a single machine, that is no reason to drop such an important advantage. The network-transparency of X is crucial in the commercial sector, where it is put to good use in thin-client computing and for remote administration. Consider this: X's main competitors, Windows and Mac OS, have both recently had a remote access capability tacked on. X does it much more fluidly and naturally than either proprietary competitor.

Moreover, claims that any poor performance of X is due to its networked nature are generally unfounded. When you run client and server on the same Unix machine, the X protocol is transmitted via a Unix socket, which, on modern implementations such as Linux, is a very efficient means of Inter-Process Communication. In addition, when running X clients local to the server, there are techniques for by-passing the X-wire entirely for certain tasks and such techniques often find application in games programming. Complaints that X uses too much memory are often mistaken, too. Novices fail to recognise that the memory usage figures they quote for X also includes the on-board memory of their graphics card – which is routinely 32MB or more for modern cards.

Criticisms that X is old-fashioned are perhaps more substantive. But the fact that the core X technology has remained constant for sixteen years can also be seen as an advantage. X is a well-proven, reliable and widely-deployed system. The longevity of X and the fact that it runs on such a variety of hardware and operating systems means that there is a wealth of software out there written for the X platform. This would be foolish to carelessly throw away. As evidence, bear in mind that when any new graphics architecture is proposed for Linux – such as DirectFB or GGI, for instance – an X server built on that architecture is often quick to follow. Apple now even distributes an X server (based on XFree86) for Mac OS X. They would be unlikely do that if there weren't demand!

The secret to X's longevity is its flexibility. As we have seen, the approach of 'mechanism not policy' means that X doesn't present a barrier to creation of modern graphical toolkits, for instance. The other significant choice the original designers of the X11 protocol made was to allow for the protocol to be easily extended. So-called X extensions can be bolted on to an X server to implement optional new features and introduce new behaviour – without changing the core protocol. This mechanism is crucial to X's ability to adapt to changing requirements (see box *X Extensions* on the right).

## The need for change

Back in 1987 when X11 was designed, true-colour display hardware was rare and scalable font technology – although used for printing (for example, in the PostScript page description language) – was not appropriate for on-screen fonts. These facts are reflected in the design of the core X protocol. While it's suitable for monochrome and low-colour displays, it doesn't match the abilities of today's accelerated, true-colour graphics hardware well or the demands of modern applications. The X protocol has poor support for handling images, and the way it handles fonts is inefficient especially for fonts with large numbers of glyphs (such as non-Latin scripts and Unicode fonts).

Part of the problem is X's old-fashioned rendering model, that is, the functions it provides for drawing (known as geometric primitives) and the way that it combines what you draw with what's already on screen. X's rendering model is inspired by PostScript's model – state of the art for the day – and provides a well-packed toolbox of primitives (but alas not PostScript's support for drawing splines or arbitrary curved paths). The result of drawing with each primitive is combined on-screen with a series of bit-wise logical operations, known as a raster operation or ROP. While this works well for monochrome and palette-mapped displays, it's cumbersome with true-colour displays and doesn't lend itself to techniques such as anti-aliasing (anti-aliasing is the process of blurring the edges between two shapes drawn on screen by averaging out the colours along their boundary and thus making them look less jagged).

An *ad hoc* solution to such problems is for an application to perform any advanced rendering required at the client side and transport the result over the X-wire to the server as an image. Such an approach is inefficient, not only because it fails to take advantage of hardware acceleration but also, due to a lack of standardisation, it leads to the endless re-invention of the wheel.

## The Render Extension

It was clear that what was needed was an extension to the X protocol to modernise X's rendering model, but the X community was forced to wait a long time for this. The demand was eventually met by Keith Packard's *X Rendering Extension* (*Render* extension for short), which he first proposed in 2000. The *Render* extension provides a compositing operator which allows any rendering to be blended with the screen using an optional level of translucency (known as an alpha channel); support for alpha-blending leads naturally to the implementation of anti-aliasing. In addition, the *Render* extension provides a much more basic set of geometric primitives, which map more closely to what can be accelerated by hardware in modern graphics chipsets. The *Render* protocol defines primitives for drawing only trapezoids and triangles, but this is not anywhere near as limiting as it might sound, since any other primitive can be decomposed into a sequence of trapezoids and triangles.

Not content with simply creating a new rendering model for X, Packard also revolutionised font-handling in X. Along with *Render*, he created a library for the client-side rasterization of scalable fonts, employing the popular TrueType font engine. This library, *Xft*, provides an X client much more power and flexibility in the handling of fonts, and uses the Render protocol's ability to upload rasterized font glyphs to the X server for rendering with optional anti-aliasing. At last a solution to how to support anti-aliased fonts in X!

The *Render* extension first appeared in XFree86 4.0.3, but it wasn't until the recent release of 4.3.0 that the new font system came of age. With *Xft*, when X clients were running on servers without the *Render* extension they had to fall back on the old core X protocol font-handling. *Xft2*, as included in XFree86 4.3.0, supports the new font model even on *Render*-less servers. Moreover, 4.3.0 also includes Packard's fontconfig library, which is a system-wide solution for the configuration and location of fonts – not just for X.

With the release of XFree86 4.3.0 most of the infrastructure is in place for a modern rendering system capable of meeting the demands of current application developers. The missing pieces of the puzzles are hardware acceleration for the *Render* extension and a fuller and more complete Render-based, client-side API for rendering. Little progress has been visible on the former. Currently only the Matrox XFree86 driver supports hardware-acceleration for the *Render* protocol and that's largely experimental. A full implementation will probably have to wait until the XFree86's acceleration architecture (XAA) is overhauled which is not expected until XFree86 5.0.

The other requirement, an advanced 2D-drawing library to take advantage of the *Render* extension, is much closer. Keith Packard and Carl Worth are developing the *Xr* library for just this purpose. Actually, *Xr* is device-independent rendering toolkit,which provides a rich set of graphics primitives similar to the PDF (Portable Document Format) and SVG (Structured Vector Graphics) models. Its supports the compositing abilities of the *Render* extension directly, but will also work on servers without the *Render* extension via a separate compositing toolkit, *Xc*, which Packard and Worth are also developing.

## Development issues

Despite some of the debate that has flared up following the Packard incident, the technological challenges facing X have never really been an issue: any such problems are being solved or are solvable. More difficult questions to answer are how the XFree86 project should be managed and what role the project should play.

Many of the points that Keith Packard raised in his *Call for Open Governance* post are valid. But claims that XFree86 is not open are perhaps misleading. Maybe the XFree86 team could do more to encourage new developers, yes, but they have limited resources. Besides, important progress has been made, such as the opening up of the XFree86 CVS repository and mailing lists, and availability of regular source code snapshots. A critical problem for the XFree86 is a lack of developers, especially when it comes to driver development. Few new developers really have the expertise to tackle such a difficult task as writing and maintaining the code that drives the hardware of modern graphics chipsets.

Driver development is definitely an area that needs to be addressed. The XFree86 release schedules mean that end users often have to wait at least six months to be able to get updated drivers to support new hardware or fix bugs; this can be especially frustrating when the fixes appear much more quickly in CVS. A solution to speed up the distribution of drivers updates to the end user is to take advantage of XFree86's module-loading ability and finalise the hardware driver interface, so that binary drivers can be built and released independently of XFree86

»

## X Extensions

### Some common extensions and what they do

**DGA** The XFree86 Direct Graphics Architecture extension gives an X client direct access to the X display's framebuffer. Naturally, this works only when the client runs locally to the server. The DGA extension is frequently used for games programming, where the performance gained by direct access is useful, but inherent security problems make it unpopular.

**GLX** The GLX protocol was created by SGI as means for sending OpenGL graphics primitives over an X wire and thus allowing X to support 3D graphics. XFree86 4.0 introduced hardware-accelerated 3D graphics with its Direct-Rendering Infrastructure, which boosts 3D performance by by-passing the X wire when the X client is local to the server.

**MIT-SHM** The MIT Shared Memory extension allows X images to be stored in memory shared between the X client and server and so avoids the need to send the image over the X wire offering a performance boost with large images. Client and server obviously must be on the same host.

**RENDER** The *Render* extension provides a new compositing render model for X. This is discussed in more detail on this page under the heading *The Render Extension*.

**SHAPE** The Non-rectangular Window Shape extension, as its unabbreviated name suggests without any ambiguity, permits X windows and borders to be non-rectangular.

**VidModeExtension** – An XFree86 extension which allows a client to modify monitor settings or change the display mode (unlike the *RandR* extension it doesn't change the size of the root window). It's used by the *xvidtune* tool to tweak monitor settings and is often employed in conjunction with the DGA extension by games software to give the appearance of running 'full screen'.

**XVIDEO** The XVideo extension allows X to make use of video-oriented hardware features supplied by a graphics chipset, such as hardware overlays, colour-space conversion, image scaling, and video capture.

releases. This would also be good news for hardware vendors, such as NVidia and ATI, who ship binary-only drivers for their cards. Although encouraging closed-source drivers will not please Free Software advocates, the truth is the bulk of desktop computer users doesn't care about such issues; they just want their hardware to work. (This is why SciTech's SNAP driver architecture may do well, especially with distro vendors who have no problems shipping proprietary software.)

> **"XFree86 was never intended to be a standards body, nor is its goal the global domination of Linux as a desktop system."**

Some suggest that XFree86 such be made more modular throughout, to enable the separate release of its various components; others say the splitting off the driver development as a separate sub-project and promoting its use (without the full X distribution) as a hardware-layer for embedded systems would revitalise driver development. A simpler solution would be for the distro vendors to devote as many of their developers to the XFree86 project as they do to the Linux kernel at present. After all, it's just as critical.

Complaints that the XFree86 project is not doing enough to develop the technology and standards that the various Linux desktop projects need for interoperability are unfair. XFree86 was never intended to be a standards body, nor is its goal the global domination of Linux as a desktop system. It simply aims to create the best X implementation on the platforms it supports. The XFree86 developers can be decidedly bullish in this regard: if you don't like the way we do things, do it yourself. As an Open Source undertaking, the XFree86 project can be forked. But a fork may not be a bad thing if it is done properly and with a definable aim – such as, for example, as a testing ground for experimental X extensions. The danger is a duplication of effort and splitting an already very shallow pool of developers. Notably, Keith Packard, appears not to be in favour of an XFree86 fork, despite this being the reason he was ousted from the XFree86 core team. If official standards are what are needed for desktop Linux to flourish – along the lines of the work that is being done unofficially by the freedesktop.org project – then this should really be done under the umbrella of a standards bodies like X.org or the Free Standards Group (who maintain the Linux Standard Base) and not put on the shoulder of XFree86.

## X Alternatives

**Some other graphical environments to try**

**DIRECTFB**
www.directfb.org/

The DirectFB project provides an open-source, hardware-accelerated graphics system based on the kernel framebuffer layer and was originally designed for embedded use. It provides an advanced compositing model, much like the X *Render* extension. XFree86 has been ported, and XDirectFB now supports most X extensions, and, uniquely, translucent windows. Newly added to DirectFB is 3D hardware-acceleration, based on an embedded version of *Mesa* and a modified version of the kernel direct-rendering manager, and this can be employed by XDirectFB's implementation of the GLX extension. The range of hardware supported by DirectFB is currently limited.

**SCITECH SNAP**
www.scitechsoft.com/

Scitech is in the process of porting its well-respected 2D graphics architecture, called *SNAP*, to Linux. Beta versions are currently available for the x86 architecture. This also include an XFree86 module to enable XFree86 to use *SNAP* drivers. Notable features are ease-of-use (*SNAP* drivers are completely plug-and-play), support for a wide range of hardware, and good 2D performance.

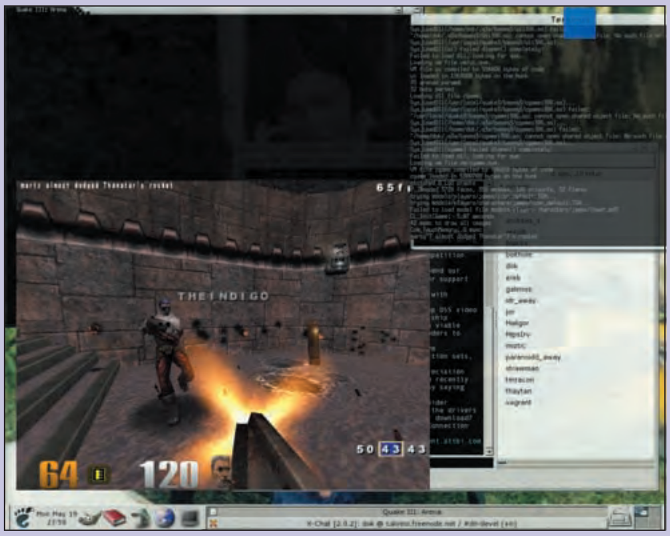

The DirectFB project's XFree86 port is becoming a viable replacement for the vanilla XFree86, providing your graphics chipset is supported.

Hardware-accelerated 3D and the XVideo extension aren't currently implemented. *SNAP* is a proprietary system, but this means support for new chipsets can often be more complete than in the Open Source alternatives.

**GGI**
www.ggi-project.org/

The General Graphics Interface is a device-independent graphics system. It is cross-platform and can output to a variety of hardware abstraction layers including *X11, SVGAlib, fbdev* (the kernel framebuffer) and *libSDL*. Ultimately, on Linux, the GGI project aims to target the Kernel Graphic Interface (KGI), a kernel-based accelerated graphics framework, but little progress has been made on this project. GGI's core API provides a basic rendering model, but this is extended by various auxiliary libraries. The port of XFree86 to GGI is known as XGGI.

**FRESCO**
www.fresco.org/

Fresco is a long-established project to make a modern, network-transparent windowing system with a powerful structured graphics toolkit to replace X. It has changed names several times (it was previously called Berlin) and seems to have had problems with direction at times. It still very much at the experimental stage, but its interesting features include CORBA for an IPC system, an advanced rendering model, and the fact that it uses GGI as a hardware layer rather than supplying its own drivers.

## The future

Whether you love or loathe the X Windows System, it's role as a fundamental component of the Linux desktop is assured. As X itself dramatically proves, standards and APIs are key, not any one particular implementation. The X API is simply too important to disregard. It may be that one day XFree86 will prove to be not the right implementation for Linux on the desktop. Fine. Other solutions are available (see box, *X Alternatives*). But note that all these other graphics systems provide or intend to provide a means for X compatibility.

The future for XFree86 itself looks bright, so there's no need to replace it just yet, however. XFree86 5.0 will contain may of the developments we have talked about here, such as the powerful new *Xr/Xc* rendering toolkit, a new acceleration architecture, and perhaps even – at last – real support for translucent windows. Hopefully, some of the solutions proposed for improving driver development will be pursued, as this is the one area of XFree86 that can cause the desktop user real problems. LXF

# The subtle flavours of WINE

A trio of tasters, messrs **Veitch, Channelle and Hudson**, explore the many varieties of WINE and how it can help you run Windows applications on Linux.

CodeWeavers' *Crossover Plugin* allows the use of common Windows web plugins like *Trillian* (left) and even Microsoft's own *Windows Media Player* (far left).

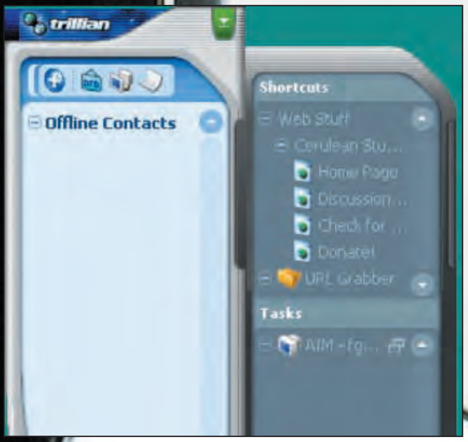

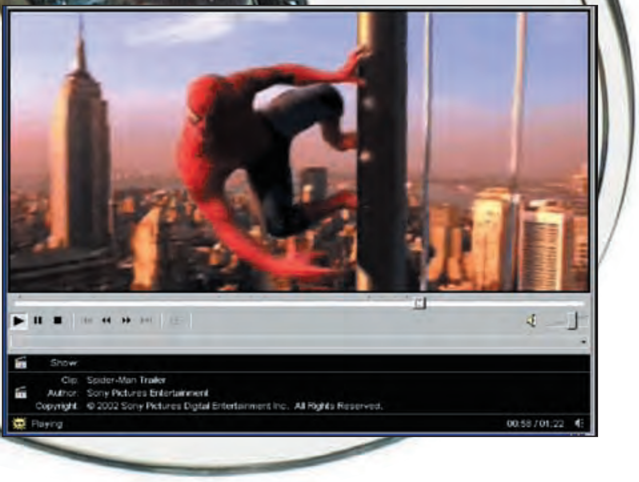
Show
Clip: Spider-Man Trailer
Author: Sony Pictures Entertainment
Copyright: © 2002 Sony Pictures Digital Entertainment Inc. All Rights Reserved.
Playing
00:58 / 01:22

By and large (and it's a fair bet someone will write in to criticise this generalisation), software is not a commodity. What is meant here is that all software is not created equal. Adobe's *Photoshop*, *The GIMP*, *Corel Photopaint* and *Kpaint* are all graphics packages, but they are by no means interchangeable. *Photoshop*, for example, has a rather dodgy PNG loader, some *LXF* staff hate the user-interface of *The GIMP*, and feature-wise, *Kpaint* is only suitable for the most simple tasks.

As the Linux software base has grown and matured, there are often several examples of virtually any common applications available, but some of the more niche markets aren't as comprehensively explored. Where would you turn to for a Linux equivalent of Curious Labs' *Poser* for example? Added to that, with often years of experience of non-Linux software behind them, many converts to the way of the penguin are reticent to forget everything they knew and learn a different application instead. These are two of the main reasons a lot of Linux users keep a version of Windows knocking around on a spare partition. It isn't an ideal solution – what if you want to create something in *Poser* to use in *The GIMP* for example – reboots galore.

Wouldn't it be great if there was some way to continue using Windows applications while running Linux? Machine emulators exist that enable Linux users to install Windows on a piece of virtual hardware, but while that removes the need for rebooting, it's hardly akin to running the software on the same OS. The answer to this conundrum is WINE. We have covered WINE several times in *Linux Format*, but this time we'll also be looking at what is actually going on, and examine some of the commercial variants that guarantee you will be able to run specific applications.

## Introducing Wine

First things first. As is customary at this juncture, we should explain that WINE is not an Emulator. Indeed, in a typical GNU recursive fashion, that's what the name stands for. So if it isn't an emulator, what exactly is it?

There are two parts to WINE and most of its relations. The first part is a loader, which enables the Linux system to decipher and understand a binary executable that was compiled to run on a native Windows system. Without the loader, Linux just wouldn't be able to make sense of a Windows application binary.

The second, and by far the hardest part to create, is a library which reproduces the standard API calls used in Windows. Think of it like this: if you were to compile a KDE application and try to run it on Windows, even assuming it understood the file, it wouldn't work. The binary will make all sorts of references to libraries (such as *Qt* and *glibc*) and systems (other KDE apps it expects to be running) and possibly filespaces (user directories) that simply don't exist on Windows. The reverse is also true. Windows provides its own standard libraries and systems (like the registry) to create an environment for applications to run on. They expect these services to be there. The task of the WINE library is by and large to reproduce this environment, contained in a Linux library, so that the binaries don't notice the difference. The important distinction here is that because WINE is not an emulator, you don't need a version of Windows to actually run Windows applications.

## Get WINE

Most distros will have a version of WINE supplied, and for most cases, it will be best to use this version. Red Hat, SuSE, Debian, Mandrake and others produce their own WINE packages for use on these distributions. If you can't find one applicable to you, you could try **www.rpmfind.net**. Source is available from the main WINE website and its mirrors, **www.winehq.org**. You can download the latest stable release or one of the nightly snapshots. Bear in mind that WINE is very big these days. The source is around 8MB and it can easily take half an hour to compile.

## Other vintages

While WINE is and will remain a project licensed under the LGPL, that will certainly continue in development for the indefinite future, there are also different projects built around WINE. Two of these are detailed in the coming pages – both are specifically targeted releases that tackle two of the common areas of interest for those wishing to run Windows applications – office software and games...

Adobe Acrobat - [Skies_IF.pdf]

Skies & Cloudscapes
100 PHOTOGRAPHIC IMAGES VOL. 8

To order call
1-800-438-2276
Page 1 of 5

**Sometimes WINE and its cousins are the only way to view proprietary document formats under Linux.**

## A brief history of WINE

### A decade of development

**The WINE project originally began back in 1993. In those days the aim was merely(!) to support the API for Windows 3.1 and allow the emerging graphical software to run on Linux.**

**Bob Amstadt was the original coordinator, but Alexandre Julliard took over fairly early on, and has continued to oversee the project ever since. As the years progressed, WINE expanded both in terms of the platforms it ran on (with ports to many popular flavours of Unix and BSD) and the scope of it's mission (most notably to support other versions of Windows too). Although WINE did run happily on many flavours of Unix, later versions use kernel-level threading, not currently implemented in the code for non-Linux platforms, but it is likely that this will be addressed at some stage.**

**Corel took a big interest in WINE in 2000, when they used it to help release Corel software, such as *Corel Draw* and *Photopaint*, for Linux. Corel also provided a fair amount of resources to the WINE team.**

**WINE is still very much in development, and ten years on from its inception, is still not regarded by the development team as being ready for general use. However, that notwithstanding many 'ordinary' people do use it every day, running all manner of software not originally intended to work under Linux.**

> "Many 'ordinary' people use WINE every day to run all manner of software not originally intended to work under Linux."

## WINE resources

**WineHQ**
**www.winehq.org**
**The main WINE website with FAQs, forums and tons more info**

**Frank's Corner**
**frankscorner.org**
**Amazing site with lists of apps that run on various shades of WINE, plus tips on how to get them working.**

**Transgaming**
**www.transgaming.com**
**Home of WineX, a commercial version aimed at gamers (see overleaf).**

**CodeWeavers**
**www.codeweavers.com**
**As well as producing *Crossover Office* and the *Crossover Plugin*, CodeWeavers also produces an easy to install and configure version of Wine, available for free download.**

»

# WINEX3: GET YOUR GAME ON

**As we're all serious coders, there's no way that fripperies like computer games hold any interest whatsoever for Linux users, right?**

Native Linux games are few and far between, particularly now that Loki is no longer with us. While the odd title does come out, such as *Unreal Tournament 2003*, many companies still do not wish to invest the time and effort in producing a port for Linux users. However, all is not lost thanks to the efforts of Transgaming, which is developing its own port of WINE that allow many games developed for Windows to work almost flawlessly with Linux. The third release of their software, WineX, has added support for some top of Windows' top games, such as *Medal of Honor, Battlefield 1942, SimCity 4,* and *Everquest*. Add to that its already extensive list of supported titles and most people should have enough games to keep them going for some time!

"WineX has superior support for Windows technologies – programs failing to work in WINE will often work flawlessly in WineX."

WineX is made available in two forms: a wodge of source code direct from the developers, or as pre-packaged files available for popular Linux distributions. The former is made available for free, but naturally many want the latter as it makes life much easier. Transgaming charges a fee of $5 a month for access to their prepackaged files, however there are other benefits to being a subscriber.

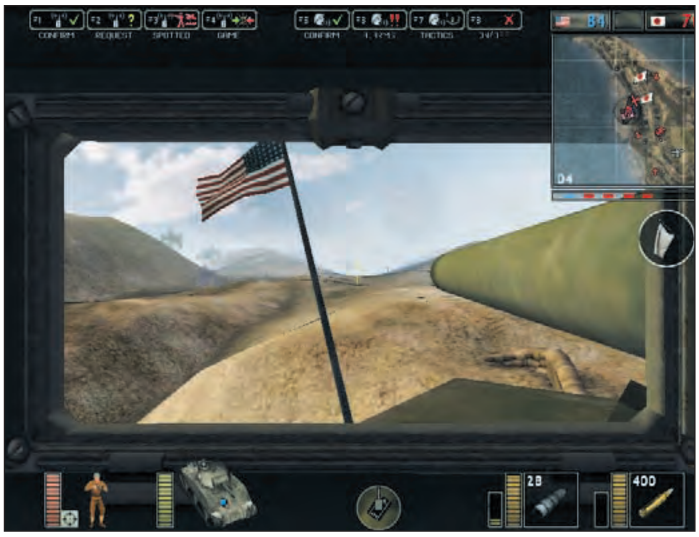

*Battlefield 1942* **is just one of the many popular games that can be made to run under WineX3. Now, where did the enemy get to?**

For example, all subscribers are able to vote for what they'd like the Transgaming developers to work on next, whether that be faster 3D action, a smoother installation process, or support for a particular game. Also available are special support options for subscribers – if you're experiencing problems getting WineX to work on your machine, you can get in touch with Transgaming support staff who will do their best to help you out. One advantage that should not be underestimated is that by subscribing you're supporting a company who are driving forward WINE's development greatly - Transgaming gives large amounts of its source code directly back to the WINE project, which benefits everyone who wants to use MS Windows on Linux.

## How good is it?

We downloaded the prepacked files for Debian, and they installed smoothly with dpkg. WineX is run simply by typing **winex3** before the name of the WIndows executable you want to run, so we installed *Medal Of Honor: Allied Assault* – see the box, *How to get your games working* on the right to see how "hard" it was.

The speed of WineX is incredible - we tested *Medal Of Honor* on an 800MHz PIII and it ran very smoothly indeed. While there were a few hitches with textures on walls failing to appear on screenshots, the game ran consistently quickly – even when there were complicated events on screen. Installation of the game looks and feels just like Microsoft Windows, which should make many migrants feel at ease – WineX even smoothly maps drives to Windows-like letters, and places a friendly WineX shortcut on the desktop to launch each game as it's installed.

## Superior support

Because making games work perfectly has required Transgaming to work on more than just DirectX/ OpenGL support, WineX has superior support for various other Windows technologies. As such, very often programs that fail to work in WINE will work flawlessly in WineX. Many of these technologies are in the process of being incorporated into the main WINE tree, thanks to donations from Transgaming, and so WINE should catch up sooner or later.

WineX 3 brings more to the table than just a new list of supported games. For example, there's a new ALSA sound card driver available, early support for force-feedback

## Games you can play

**At the time of writing, WineX supports just shy of three hundred games and the list is growing daily. These include (and this list is by no means exhaustive – just some of *LXF*'s favourites): *Battlefield1942, Civilization 3, Command & Conquer: Red Alert 2, Curse of Monkey Island, Deus Ex, Diablo II, Everquest, Grand Theft Auto 3, Imperium Galactica 2, Jedi Knight 2, Master of Orion 3, Max Payne, Medal of Honor, Nascar Racing 3, Nox, Return to Castle Wolfenstein, Soldier of Fortune1 and 2, SimCity 3000 and SimCity 4, Sudden Strike 2, SWAT 3, The Sims,* and *Warcraft 2* and *3*.**

**Once you subscribe, you can vote for the new games that the development team should work on – if you don't see it in the list on the Transgaming website, your vote will likely count!**

joysticks in games that take advantage of the hardware, and a new graphical game installer called *Point2Play*. *Point2Play* also handles automatic updating to new versions of WineX as they become available.

## The bad bits

It's not all perfect in WineX, sadly. Firstly, it does require some "pollution" of your Linux box – WineX isn't Free Software with a capital F, and neither are the Nvidia drivers. In fact, while installing the Nvidia kernel module, you'll even be given a warning that it will taint the kernel with non-free code. If you're not a GPL purist, then naturally this isn't a problem.

Another problem is that Linux distributions are so wildly different today that WineX isn't guaranteed to work on them all. For example, Red Hat 9 is quite different to Red Hat 8 when it comes to threading libraries, so WineX has a lot of extra code in there to try to sort these problems out. While we didn't experience any problems ourselves, there are quite a few messages on the Transgaming support messageboards from subscribers experiencing problems here and there.

One final problem is in the availability of drivers for your system. WineX is very much Nvidia-specific, although it can work to a degree with other graphics cards. As Nvidia's card drivers are proprietary, you're pretty much at their whim if you find a game not working quite correctly on your machine. This is negated somewhat by the fact that Transgaming do a great deal of work to make sure games are as compatible as possible – very often you'll find recently released games "just work" in WineX, particularly if they re-use code from prior games such as the *Quake 3* engine.

## Summing up WineX

While it would obviously be ideal if popular games all had native Linux ports, it's sadly not likely to happen in the near future. In the meantime, though, we have WineX, and it does a remarkably good job of making Windows games run unaltered on Linux. Yes, there are flaws here and there – its pickiness regarding graphics cards might bite quite a few users, for example - but that's not enough to drag it down too much. The reality is that Transgaming have successfully made Windows games available to Linux users, and it's as hassle-free as anyone could want.

»

## How to get your games working

### Dropping dual boot becomes a reality!

**Here's a step-by-step guide to getting *Medal Of Honor* working on your machine. The process is pretty much the same for other WineX-supported games, and really isn't tricky at all as you'll see...**

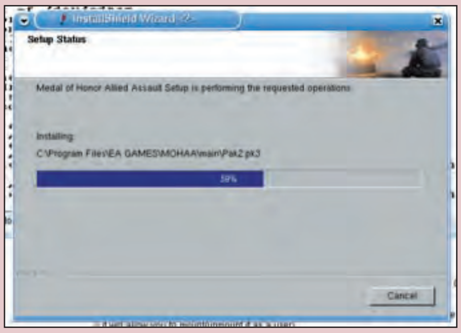
InstallShield Wizard <2>
Setup Status
Medal of Honor Allied Assault Setup is performing the requested operations
Installing
C:\Program Files\EA GAMES\MOHAA\main\Pak2.pk3
59%
Cancel

**STEP 1: INSTALLATION**
**To kick off the installation process, mount your CD ROM, then run winex3 /path/to/cdrom/Setup.exe. Note that you should *not* do this from inside the /cdrom directory, as this will keep the drive mounted even when WineX prompts you to insert the second install CD. You may find it helpful to tick 'Create shortcut on my desktop' at the end of installation, as WineX re-writes this so that it calls the winex3 executable. Transgaming recommends you always patch your games to the latest release, so we installed the *MoH* 1.11 patch.**

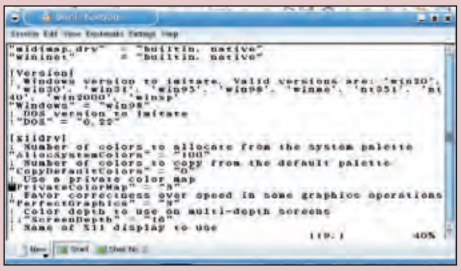

```
"midimap.drv"   = "builtin, native"
"wininet"       = "builtin, native"

[Version]
; Windows version to imitate. Valid versions are: 'win20',
'win30', 'win31', 'win95', 'win98', 'winme', 'nt351', 'nt
40', 'win2000', 'winxp'
"Windows" = "win98"
; DOS version to imitate
;"DOS" = "6.22"

[x11drv]
; Number of colors to allocate from the system palette
"AllocSystemColors" = "100"
; Number of colors to copy from the default palette
"CopyDefaultColors" = "0"
; Use a private color map
"PrivateColorMap" = "N"
; Favor correctness over speed in some graphics operations
"PerfectGraphics" = "N"
; Color depth to use on multi-depth screens
;;"ScreenDepth" = "16"
; Name of X11 display to use
```

**STEP 2: CONFIGURATION**
**WineX *Medal of Honor* works just fine with default settings, but WineX's configuration file is very easy to read and toy with if you want to enable specific settings. For our screenshots, we edited the file to force WineX to run games inside a window, rather than in the default full screen fashion.**

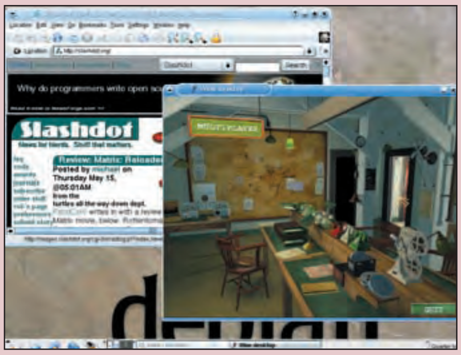

**STEP 3: START THE GAME UP**
**cd to the directory of the game executable, then type winex3 <game.exe>. Once you're into the game, you should make sure it has detected your hardware correctly and that you don't try and push your computer too hard – remember, WineX does its best but isn't perfect! We tweaked our *Medal Of Honor* settings to get maximum performance.**

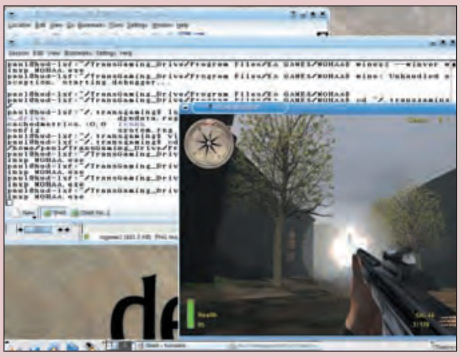

**STEP 4: GET PLAYING!**
**Dakka dakka dakka dakka squawk! Our Reviews Editor takes time out to enjoy his favourite sport of Pigeon Machinegunning.**

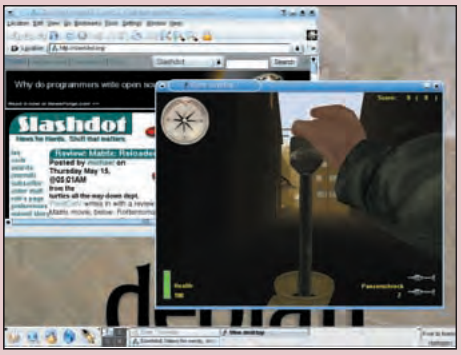

**STEP 5: CALL IN THE HEAVY METAL**
**The rocket launcher in the *Medal Of Honor* is the bee's knees, however as you can see in the picture, it takes some time to reload!**

# CROSSOVER OFFICE 2.0

**It is often said that one of the 'barriers to entry' for Linux in the corporate world is the lack of support for the *de facto* standard productivity applications. Andy Channelle tests one solution to the Linux users' perennial 'just one app missing' syndrome.**

For many users in offices and homes across the world, Linux is simply not worth considering, often due to the fact that there is one or two applications deemed irreplaceable. These may be obvious candidates such as *Microsoft Office* or *Adobe Photoshop* for professional buyers and *Quicken* for home users, or more peripheral applications. And while Linux has its fair share of competitors in the form of *OpenOffice.org*, *The GIMP*, *Gnucash* etc. the lack of 'industry standard' applications is a major stumbling block.

To combat this situation, the Open Source WINE Project was initiated which would add a new application layer to Linux, translating Windows API calls instantaneously and allowing the installation of many Win32 applications. As the project matured more installation successes were reported, but whether or not an application would work on a given system was still a little hit and miss, often requiring some fiddly work.

Enter CodeWeavers and its groundbreaking suite of products that allowed Linux users to access web content previously limited to Windows and Mac PCs (*CrossOver Plugin*) and install a range of Windows productivity applications through a simple to use interface (*CrossOver Office*). *CrossOver Office* (*COO*) was well received in its first incarnation but there were limitations and the range of supported applications was quite narrow. Also, as most people have come to realise, Microsoft stands still for no man and the latest release of *COO* not only improves support for *MS Office 97* and *2000* it also brings, with a few caveats, *Office XP* to your Linux desktop. Furthermore there is also full support for *Adobe Photoshop 7*, which should gladden the hearts of professional users – especially in the SFX/film industry – currently struggling with *Photoshop* on *VMWare* or other virtualisation software.

## The menu structure

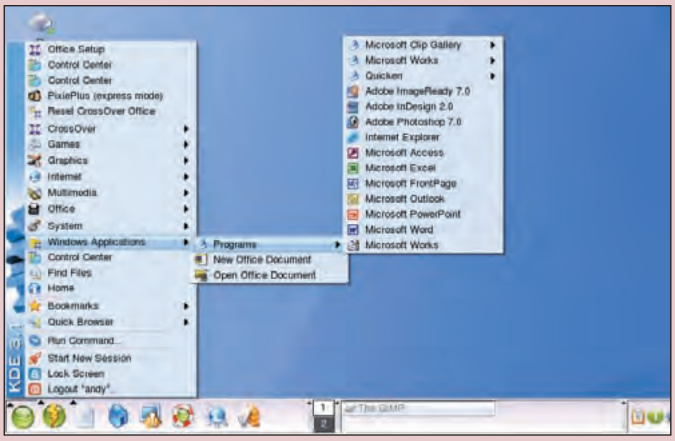

***Crossover Office* is a software tool that brings *Photoshop, MS Office, Quicken* and other Windows applications to Linux. The menu structure mirrors Windows completely, making it seem closer to more traditional virtualisation software such as *VMWare* or *Win4Lin*.**

## Installation No.1

After downloading and burning COO to a CD (the latter is not essential) I ran the single script that installed and configured the system. As well as installing the main set-up application and a couple of subordinate programmes, the script creates a 'fake Windows' directory containing the full Windows file structure including the likes of 'Program Files' and 'System' and which is seen by Windows apps as the C drive. You can also select a directory from any partition to act as 'My Documents' and this defaults, handily, to /home.

The two accessory applications allow you to either reset *COO* in case of a crash, or simulate a Windows reboot, which may be needed during the installation of some applications. Under the *COO* menu heading – the script adds entries to both GNOME and KDE – you can also access the online documentation or uninstall everything. Finally, you will also notice a new 'Windows Applications' entry on the K/GNOME panel, which is where you can launch applications if you don't want to go through the set up user interface.

The installer worked flawlessly on both Red Hat and SuSE systems, and didn't interfere with the already present WINE installation on the latter. Office Setup is where most of the work takes place; here you can install, remove and repair applications, configure menu entries, set up file associations and update the supported applications list. It's all so simply put together, thanks to the limited remit of the application, that the newest of new users should be able to find their way around with no difficulty.

## Installation No.2

To install a new Windows application you simply start up Office Setup and, from the opening screen, click on the 'install...' button. This will take you to the installation dialog, which lists the supported applications. Choose one of these, insert the relevant CD-ROM or select a location on your hard disk and you're away; from here on in it is the same as installing on a basic Windows machine, without the need to reboot halfway through. On the installation screen there is also an option for using 'unsupported' applications but here, as with a vanilla WINE setup, your mileage will vary.

The list of supported applications is quite good, comprising *Microsoft Office (97, 2000, XP), Lotus Notes, Quicken, Microsoft Visio, Adobe Photoshop 7* and *Internet Explorer* (which, when selected, is installed directly from the Microsoft site unless you specify otherwise).

It is likely that most users will be evaluating or buying *COO* specifically to run *Office*, so this is where *we* started, bypassing the *97* version and going straight for *Office 2000* instead.

Purists will no doubt cry out in shame at the sight of the Microsoft install dialogs cluttering up a clean KDE desktop, but those who need this software but have no desire to run Windows will find satisfaction in its familiarity. Ironically, thanks to the simulated reboot, putting *Office* onto

Linux takes about half the time of a native install. Any online components (such as DCOM95, which was needed for *Internet Explorer*) are automatically downloaded and installed, file associations are configured and, within a few minutes, the full suite is available through the Windows Applications menu.

The most significant addition for many people in *COO 2.0* will be support for *Access*; indeed, CodeWeavers regards this as one of the major selling points. While Linux has some incredibly powerful database systems, they say, it hasn't yet been blessed with a fully featured, intuitive front end. I tested first by building a simple address book to house a small amount of personal data then installed and played with the sample database included with the suite. There is an inevitable performance hit when using an emulator (yes, we know WINE Is Not an Emulator), but on a fairly powerful machine it makes little difference. Occasionally an application would stop and ponder for a few seconds, but whether this is a problem of *COO* or the actual application is difficult to pin down. One thing I did note – and this appears to be common across the core Office 2000 applications – is *COO*'s displeasure with the Office Assistant. More often than not, trying to query 'Clippy' caused the application to hang, requiring the use of **xkill** and **xrefresh** to sort it out. While Clippy is the cause of much teeth-grinding among *Office* users, he is the only conduit to the help system, so pretty vital if you don't have the manuals to hand (and who in an office does?).

While *Office 97* and *2000* are both fully implemented, XP only has partial support, you won't be able to use either *Outlook* or *Access,* but everything else works as expected.

The next big thing is *Photoshop*. CodeWeavers claim the application can manage everything up to the latest release, version 7. To test this, I successfully installed *Photoshop 6* and then upgraded to *7*. All went fine, but to double check I took it all off and did a straight installation of *Photoshop 7*. Again, I encountered no problems. Every tool, filter and option worked as expected, though with a tiny degradation in raw performance. Running the 'Sepia' action on a 1024x768, RGB image, for instance, the difference between the Linux and Windows 'version' was negligible. The situation was similar with *Quicken* and *Internet Explorer*, both installed and ran exactly as expected.

Finally I tried a few unsupported applications with varying success. Adobe's DTP package *InDesign* works great, but *Dreamweaver* is a non-starter; *Kazaa* works, *Bryce* doesn't; and I couldn't get any of my daughter's CD-ROMs installed, which is a shame as she's the only one left using Windows!

## Conclusion

A complex task like emulating every aspect of a closed system is obviously going to be tricky, but *COO* does an admirable job of it. It's not a flawless application: WINE is notorious for its screen redraw problems and *COO*, though improved, still suffers from the occasional glitch. For instance, the menu items in *Adobe ImageReady* (part of the *Photoshop* suite) appeared inconsistent and sometimes failed to materialise at all.

The applications 'supported' by CodeWeavers work well, slight Help hiccup in *Office* notwithstanding, and it was gratifying to get a few apps that have consistently failed to work with WINE up and running here. Complete support for *Office XP* is promised for the next edition, so if you have upgraded beyond *Office 2000*, you'll be waiting around to use *Outlook* and *Access*. The ability to run *Photoshop 7*, and run it so well, is the real kicker for me. *The GIMP* is great, but it lacks CMYK support, essential for the print industry, and the esteem that Adobe's application is accorded by employers. Similarly *Scribus* comes nowhere near the features of *InDesign*, so it's nice to finally have those resources available without having to reboot. Similarly, banks don't appear, on the whole, to acknowledge the existence of anything outside of the Windows sphere, so being able to access you details through *IE* on Linux – though a compromise – is sometimes your only option. But it *is* an option.

And that's what it comes down to. If there is one or two applications that leave you tied to Windows, and they're supported here, you can forego the need to dual-boot or, more significantly, forgo buying a bundled OS with your next hardware purchase. LXF

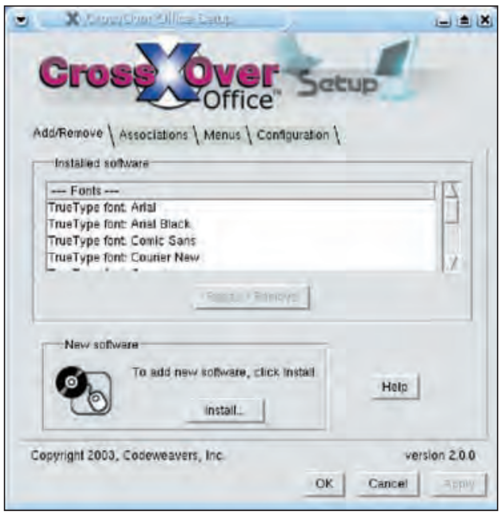

--- Fonts ---
TrueType font: Arial
TrueType font: Arial Black
TrueType font: Comic Sans
TrueType font: Courier New

To add new software, click Install

Copyright 2003, Codeweavers, Inc.
version 2.0.0

Installation is a very simple procedure.

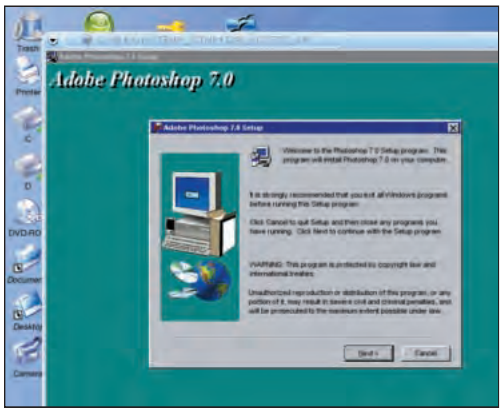

Welcome to the Photoshop 7.0 Setup program. This program will install Photoshop 7.0 on your computer.

It is strongly recommended that you exit all Windows programs before running this Setup program.

Click Cancel to quit Setup and then close any programs you have running. Click Next to continue with the Setup program.

WARNING: This program is protected by copyright law and international treaties.

Unauthorized reproduction or distribution of this program, or any portion of it, may result in severe civil and criminal penalties, and will be prosecuted to the maximum extent possible under law.

*Photoshop* support is a major selling point.

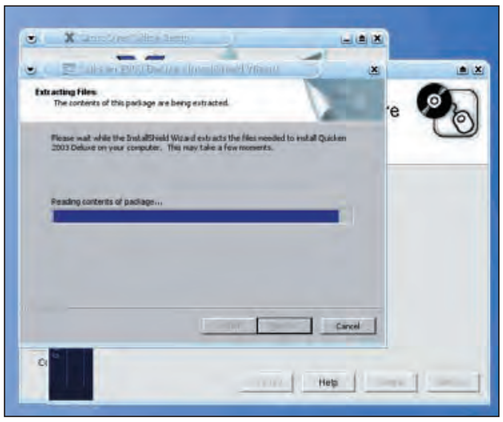

Extracting Files
The contents of this package are being extracted.

Please wait while the InstallShield Wizard extracts the files needed to install Quicken 2003 Deluxe on your computer. This may take a few moments.

Reading contents of package...

The installation routine puts *Quicken 2003* into KDE.

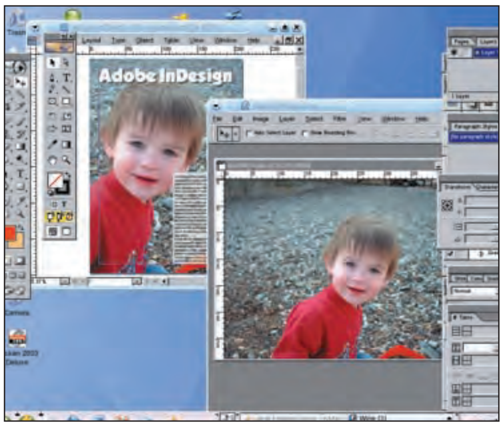

*Photoshop* and *InDesign* pose no problems for *CrossOver Office 2.0.*

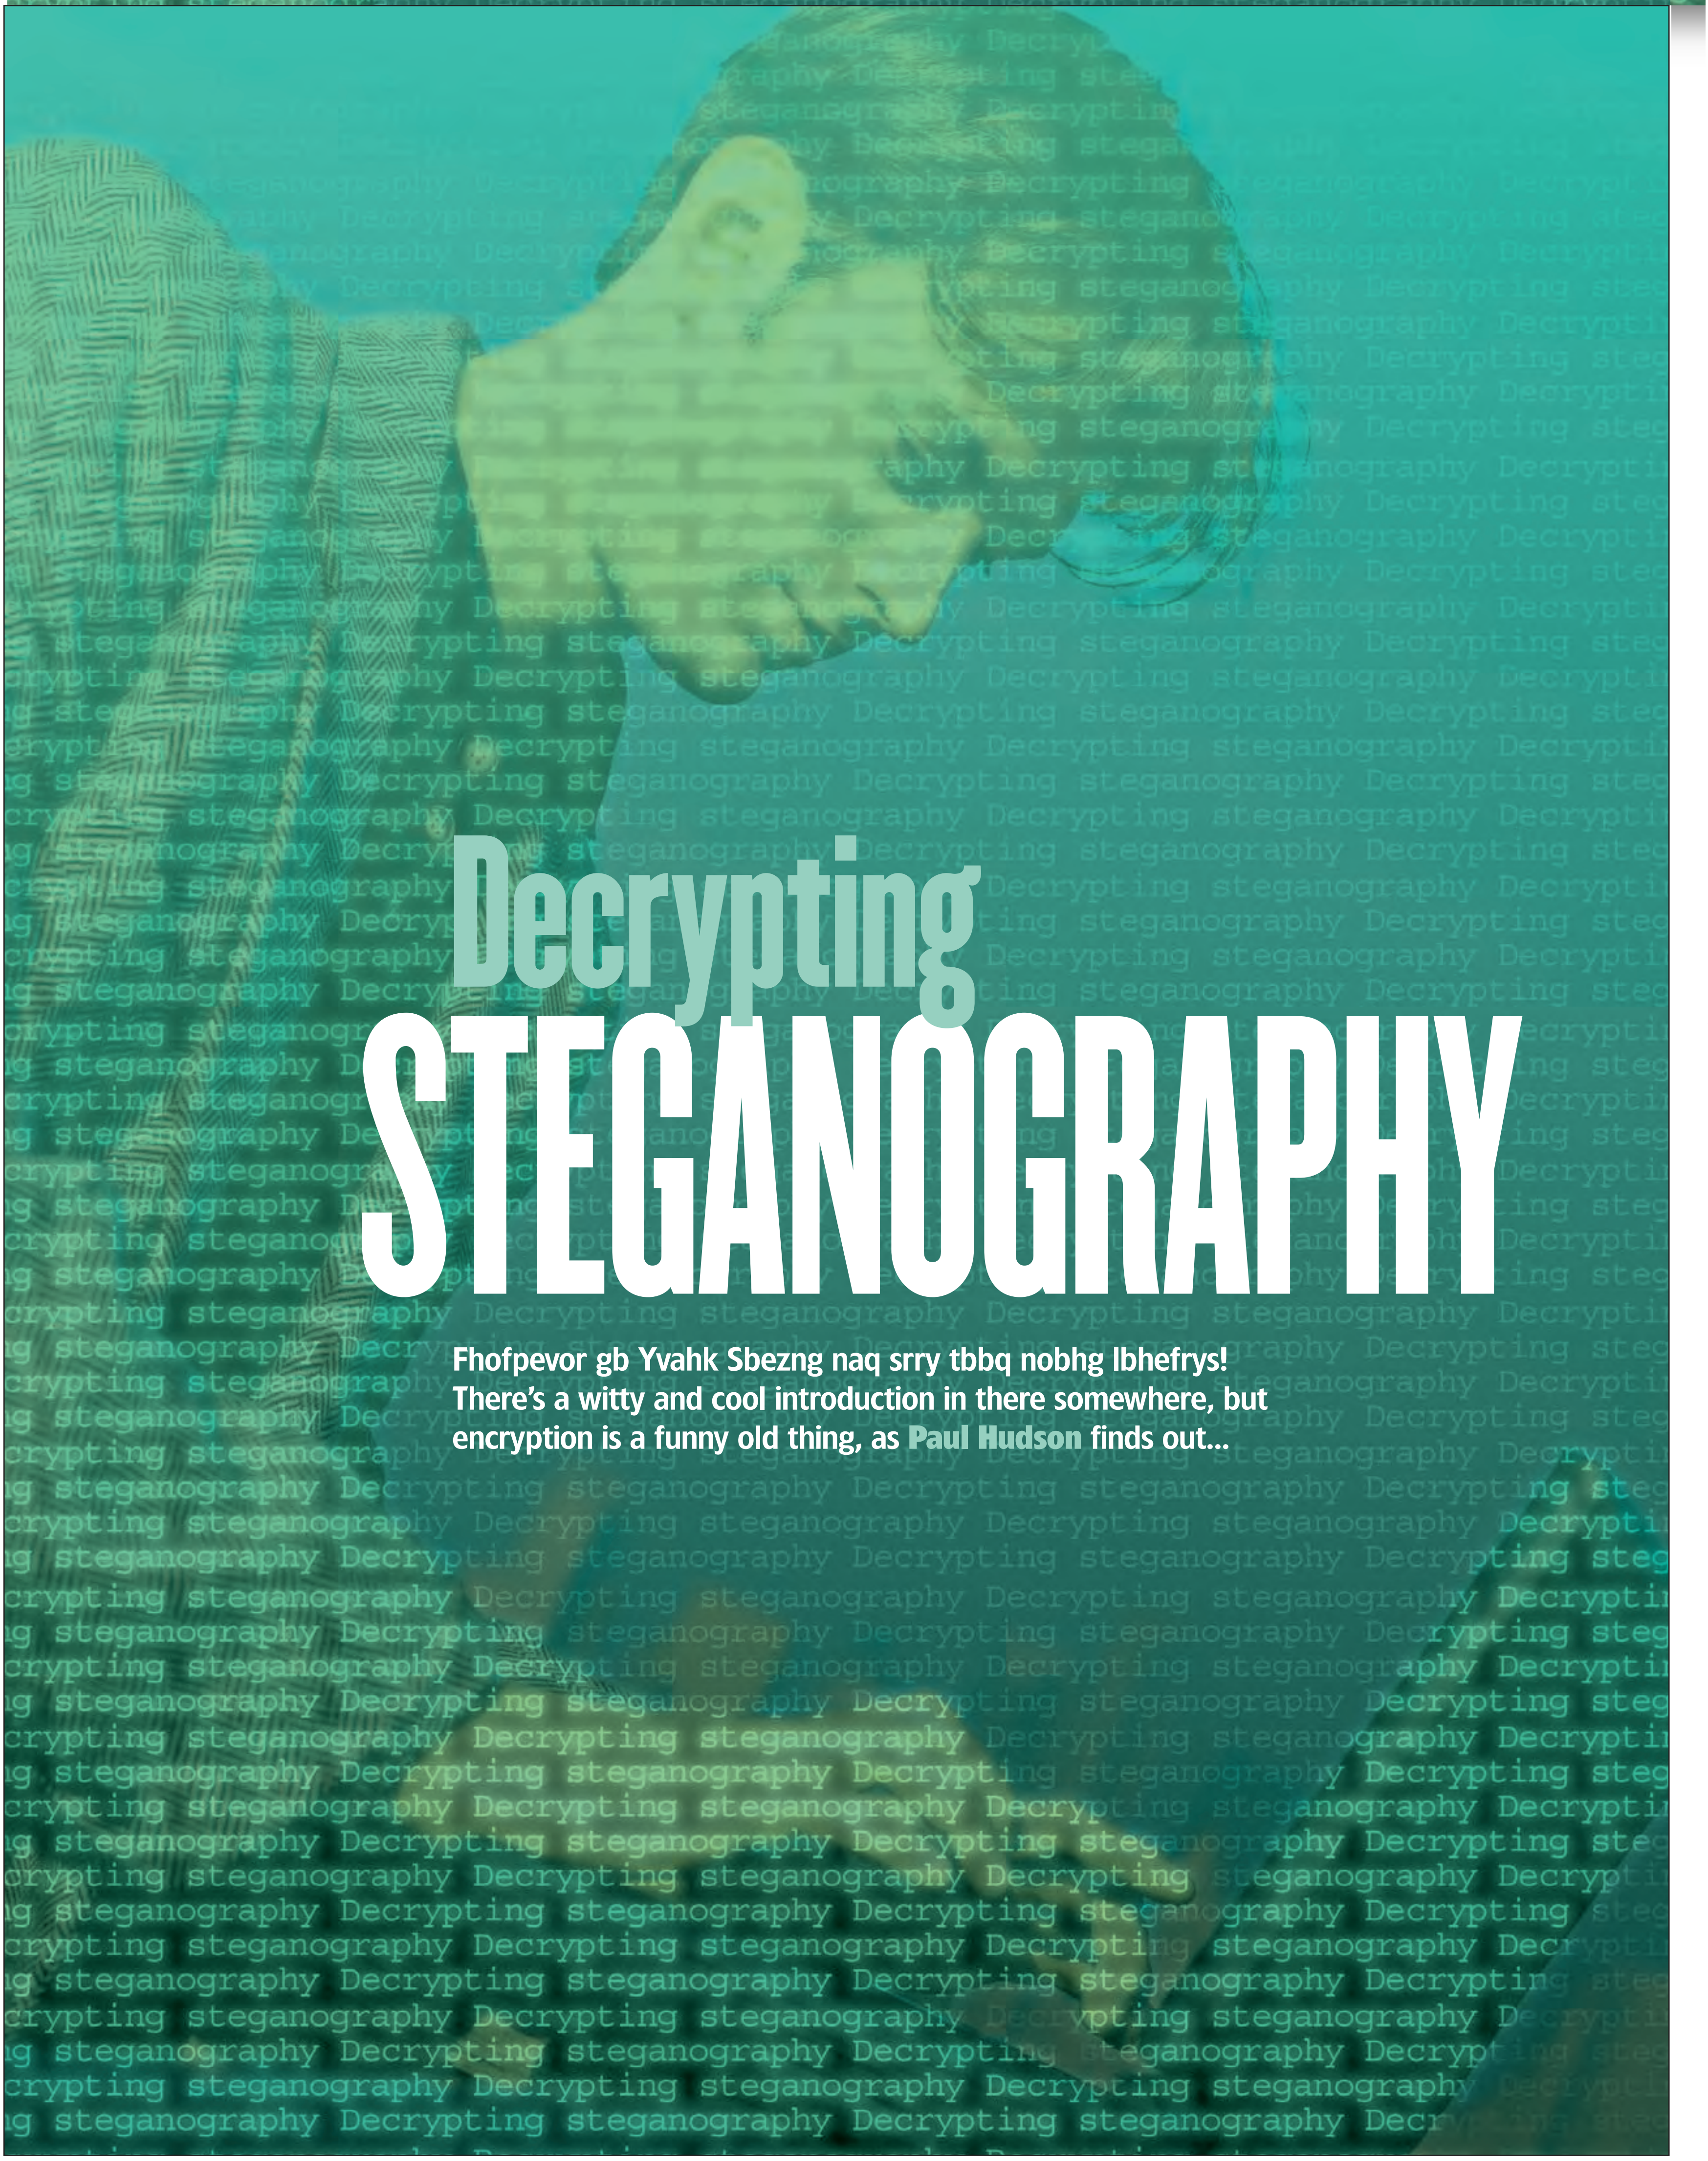

# Decrypting STEGANOGRAPHY

**Fhofpevor gb Yvahk Sbezng naq srry tbbq nobhg lbhefrys! There's a witty and cool introduction in there somewhere, but encryption is a funny old thing, as Paul Hudson finds out...**

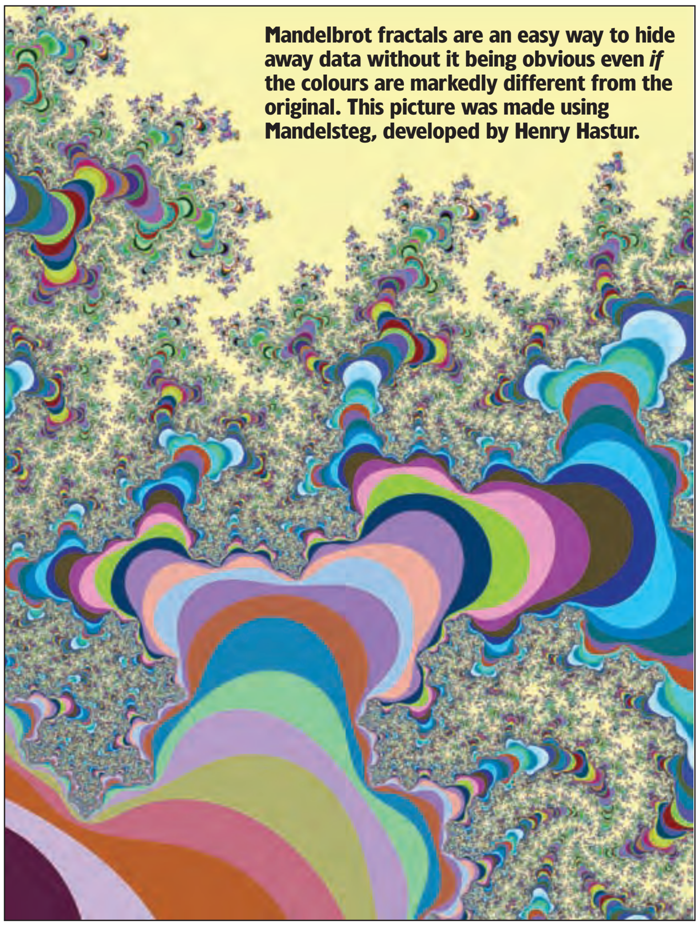

**Mandelbrot fractals are an easy way to hide away data without it being obvious even *if* the colours are markedly different from the original. This picture was made using Mandelsteg, developed by Henry Hastur.**

During the wars over which he had command, Julius Caesar was purported to use a very early cipher, known today as *rot 13,* which is short for "rotate characters by 13 places". So, "a" becomes "n", "n" becomes "a", etc. It's a very simple cipher, and one generally only used today so that the meaning of a message isn't immediately apparent. A common example of this is on Usenet, where messages containing movie spoilers and offensive material are often posted in rot13 so that visitors can't read them by accident.

Encryption has come a long, long way since the days of the Roman Empire; particularly so in recent years. However, one problem with encryption is that it is generally obvious – if you receive a sheet of paper with seemingly random writing on, you're bound to be somewhat suspicious of it.

A new technique has arisen in recent years that attempts to hide data, encrypted or otherwise, inside perfectly innocuous-looking documents. This is the technique of *steganography*, which comes from the Greek "hidden writing".

## How it works

Steganography is the process of hiding data inside data, and so can be quite a tricky process. Some types of file lend themselves well to the process, however – particularly pictures. If you take an original picture and steganographise some data into it, it *will* be different from the original - the key is *how* different it will be. Generally speaking, you need to have a picture of at least eight times the data size of the data you wish to put inside if you want the finished product to be indistinguishable from the original to the human eye.

Think of steganography like watermarking. If you look at a sheet of quality paper, you'll see only the text written upon it; but when held up to the light (in steganography's case, a special scanner) the "paper" shows hidden text.

To move from the general to the specific, here is a technical example of how a steganographic encrypter might work. As you may know, your screen probably runs in 24-bit colour mode (32-bit just adds 8 bits of alpha to the 24-bits of colour), where red, green, and blue all make up eight bits of the 24. This means that a photographic 24-bit JPEG on your desktop will be composed of thousands and thousands of RGB triplets - that is, "pixel 1, 1: red:255, green: 0, blue: 128; pixel 2, ......". There's one RGB triplet for each pixel on the screen, and sometimes one for each pixel in a file. Bitmap files, for example, work like this – they don't do any sort of run-length encoding.

Now, given that each pixel has a range from 00000000 (0) to 11111111 (255) for red, green, and blue, there's a lot of potential slack space available for hiding things in. For example, the difference between blue 11111110 and blue 11111111 is minimal, which means that we could use the least-significant bit of each blue pixel for our encrypting purposes – it's simply a matter of turning letters into 8-bit ASCII, and distributing each letter across eight blue pixels. A picture at 1600x1200 resolution has 1,920,000 blue pixels available for encryption, which, if we use ASCII, gives us 240,000 letters, or approximately 234Kb of text - plenty of room for text messages, with almost no loss from the image. Now imagine we use all red, green, and blue, and used the *two* least-significant bits in each channel – we now have enough space to store an entire floppy disk of information inside a picture, with only slightly more image quality loss.

Of course, you don't need to work with a photographic image. For instance, the above picture of a mandelbrot fractal – believe it or not – contains the entirety of the text in this whole magazine article in an encrypted state.

»

"512-bit private keys are hard to hack, but the hardest encryption to hack is the one you never even notice"

## Stega-spam

**You can encrypt your text in more ways than you can imagine...**

***Spam Mimic* is an excellent online steganography encrypter that takes sentences you provide and converts them into messages that appear to be spam. If you want to try it yourself (it encrypts and decrypts online), visit the homepage at www.spammimic.com. Here's what it turned out when we encrypted the text "secret info" – naturally that could have been "home safe" or "attack now", depending on the intentions of the person sending the message.**

**"Dear Friend;**

**"We know you are interested in receiving cutting-edge news. If you are not interested in our publications and wish to be removed from our lists, simply do NOT respond and ignore this mail. This mail is being sent in compliance with Senate bill 1626 , Title 3; Section 304! This is NOT unsolicited bulk mail. Why work for somebody else when you can become rich inside 96 weeks! Have you ever noticed nearly every commercial on television has a .com on in it, plus people will do almost anything to avoid mailing their bills. Well, now is your chance to capitalise on this. We will help you decrease perceived waiting time by 170% and increase customer response by 160%! You can begin at absolutely no cost to you . But don't believe us. Mr Anderson who resides in Florida tried us and says "My only problem now is where to park all my cars"! We are licensed to operate in all states! We implore you – act now! Sign up a friend and you get half off. God Bless !"**

## What are its uses?

Many have argued that encryption on the whole has no place in today's society, particularly due to the apparent threat of terrorism. After all, while privacy is important for trusted people, in the wrong hands it could cause chaos – or at least so we're told. Despite that, encryption is widely available, and quite a few people take advantage of it – how does steganography help them?

> "US officials believed that Osama bin Laden and his followers were using steganographic techniques to hide terrorist plans inside pictures."

To be frank, it doesn't. Steganography is used to pass data around without unintended recipients seeing the data realising it has an encoded message inside, and most people who use encryption have little need for that. After all, if I want to send an encrypted message to someone, I don't have any requirement to hide the fact that it is encrypted!

One field where steganography has a valid use is that of espionage and undercover operations, and this is where a great deal of controversy has been caused. Back in May 2001, the USA Today newspaper ran a story on how US officials believed that Osama bin Laden and his followers were using steganographic techniques to hide terrorist plans inside pictures, then uploading them to various websites with seemingly innocuous intentions. Whether or not this is true, it certainly provides food for thought – strong encryption techniques render cracking terrorist messages hard enough, but if they are further made *almost invisible* with steganography, then tracking such messages down becomes near-impossible.

While all this sounds very dire, a study conducted by researchers at the University of Michigan studied over two million pictures on Ebay and failed to find any with hidden messages. While you might think "Sure, but that's the idea – steganography wouldn't be hidden if it was easily found!", and that's true – to some extent, as you'll see…

## Strong encryption?

Steganography is the art of *hiding information*, and is a very weak form of encryption. If I were to show you a room in which I had hidden a secret message written in plain English, once you had found the message you would be able to read it immediately. Consider further how easy it would be if I left the message in a fairly obvious place. Steganography is much the same – it places whatever data you give it directly into the destination file, and all steganography programs currently available make little if any attempt to hide their tracks.

What this means is that if you know a picture (for example) has encoded data in there, it's very easy to find and extract it. Sometimes it's quite straightforward – if the example of every least-significant bit was used, unencoding would be a snap. With predictable random numbers, with the seed being used as a password of sorts, data can be distributed pseudo-randomly across a file, which makes decryption much trickier Despite that, it's still quite straightforward to scan a group of images to check whether they have steganographic information inside.

The key is in understanding that steganography *isn't designed* to have strong cryptography; after all, it can encode any data inside any other data, which means you can encrypt your needle using PGP, blowfish, AES, or any encryption you choose, then place it inside a haystack using steganography.

The one piece of control steganographers *do* have over the security of their data is deciding how much they "dilute" their haystack with their needle. That is, if you have a 100KB image and you want to hide 50KB of text in there it will be quite obvious that there is something more to the picture than just a photograph. Naturally there is quite a clear-cut trade-off between how obvious the presence of encrypted data is and the size of the haystack used – if you put the same 50KB of text inside a 50MB file, it will be almost completely invisible, particularly if the data isn't placed sequentially.

## Conclusion

Think of steganography as a means to an end as opposed to an end itself – while it is technically a way to encrypt data, it isn't strong enough to stand by itself beyond a visual scan. Pretty much anyone who actively searches for cryptographic data inside a steganographised file will be able to extract data fairly easily.

On the other hand, plain encrypted data is glaringly obvious in the wild. Mix the two together, however, and suddenly secret things become much more hard to find. Yes, 512-bit private keys are hard to hack, but the hardest encryption to hack is the one you never even notice. Even *if* someone notices a steganography-encoded private key-encrypted file and manages to extract the private key-encrypted data, they still have a lot of work ahead of them. The two technologies work together very well indeed – especially if security is your game. A further, somewhat ironic use for mixed encrypted and steganographical data is to send a picture with some text. The text, encrypted, can be anything whatsoever – perhaps a copy of a thread from Slashdot. The picture would contain the *real* data, and the chances are in your favour that those who intercepted the image would spend their time trying to break the strong (and irrelevant) encryption on the text, ignoring the simple image alongside.

Steganography is definitely the way forward for encryption, and its use is certain to increase as time goes by. So, next time you downloaded an MP3 and hear a mysterious hiss part way through, just think for a moment – that noise might not be as random as it sounds!

## Learn more...

### Get ahead of the game

Wiley has recently published a book called *Hiding in Plain Sight: Steganography and the Art of Covert Communications* by Eric Cole (ISBN: 0-4714-4449-9). The author, having worked for the CIA for years, discusses various ways how steganography can be, will be, and *is being* used, drawing from his extensive experience and research in the area. If you're interested in learning more about steganography, how to detect it, and, perhaps most importantly, how to counteract it, then there are 360 pages of solid material here waiting for you.

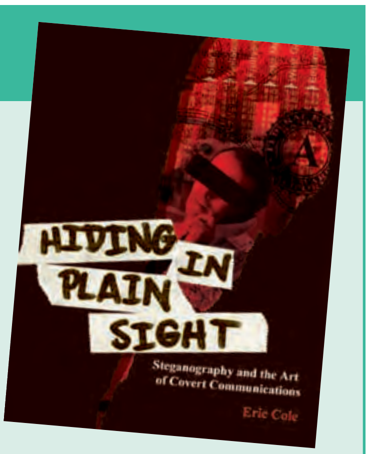

Authoritative and interesting, this book should keep you busy for some time.

# What on Earth is... DISTRIBUTED COMPUTING?

Without access to a huge computer, **Jono Bacon** shows that a lot of littler ones may be the answer...

**Don't just sit there... Tell me what Distributed Computing is.**

Well, distributed computing is the concept that you hook up a stack of computers together to process big and complex tasks.

**What kind of tasks do you mean?**

Well, we will get into this a bit later so kindly hold your horses. The types of tasks however are jobs that require huge levels of processing that would require extremely beefcake computer systems to process. These systems are expensive, loud and the size of fridges.

**OK, so how do lots of computers beat the beefcake beasts?**

Well it is quote simple. Let's assume you have 5 computers; A through to E. You would set the network a task, and a client on each machine would download a small portion of the task. Each machine would then crunch away at their portion, an when it is finished they would send the result back to the main server. When the result is sent back, another chunk is sent out to the computer. This keeps going until all of the chunks have been completed and the results are assembled together to form the final result. Nifty eh?

**Not exactly up there with sliced bread, is it? I'm sure that you can convince me some more.**

OK, well the idea with distributed computing is that you make use of your spare computing cycles to process the chunks of data.

**Whoa...hang on...my computer does not have a bike!**

No, not bicycles you muppet, your computer's processor. A processor cycle is a single measurement of how a processor works. Your computer does incredible amounts of these every second, and when you are not using your computer for something they simply get wasted; in other words your computer sits there doing nothing. The idea with the distributed processing concept is that a program will make use of these spare cycles to crunch while you're at lunch.

**Funny man. OK, so how does it know when I am not using my computer to use those spare cycles?**

Well, this is the clever bit... many distributed processing projects use a... wait for it...screensaver. It is the natural choice for running something useful when your PC is idle.

**Hang on a sec, you said "many distributed processing projects" - does that mean I need a different screensaver for each project?**

Well, for most projects you need a different client for each project. It isn't too bad as most people choose a single project and just contribute towards that.

**OK, this is sounding interesting, and I might spend tomorrow morning at work wasting a couple of hours googling about it, but I don't understand one thing – what if I wanted to run the client all the time, even when I am on the machine?**

You can certainly do this, and I run it at home on my webserver as it happens. Most projects have a command line program you can run that processes all the process is running.

**Hang on, how do you keep the client running when you log out of the machine, but keeping the machine on?**

You need to make use of the nohup command and use:

```
nohup clientprogram > /dev/null
```

**Why the /dev/null bit?**

This is because most clients display information to the screen and the **nohup** command puts information in a log also. This just sends it to the bit bucket so it doesn't exist.

**OK, so I can run this program on my machine full time. Does the machine need to be on the Internet all of the time?**

Not necessarily. Most clients just get new blocks of information off the net when you have finished your blocks, hence requiring a connection. Some clients do have workarounds to this though, so you would need to check with your client.

**Hey, that reminds me, earlier you said you would tell me what these things process.
Go on...tell me.**

Well this is where it is interesting. There are so many projects out there that you can choose one that suits. One of the biggest is SETI@Home.

**And what the heck would that be? A furniture developer?**

SETI is the Search for Extraterrestrial Intelligence. They have a bunch of massive telescopes that suck in radio information from outer space and they use distributed processing to sift through the information to see if any alien signals are in there.

**You must be kidding me?**

No. I never joke about aliens. SETI@Home is a hugely successful project with thousands of people and teams involved in the project.

**Teams? You never mentioned that. Please explain...**

Well, the whole concept of distributed computing has taken quite a competitive stance and people have formed teams to process the most amount of

chunks of data. It is quite fun and anyone who gets involved with a project can usually join one of the many teams available. When you join the team your processed chunks then add up to form a total for the team as well as the total for yourself. The more people you have in the team, the more chunks are processed and the further rise up the league table of teams for the project.

### So what kind of credentials would I need to be in the winning team?

Well theoretically you just need to join the right team. What you need to be aware of through is there are some incredibly powerful machines on some of the highest ranking teams that are dedicated to the project 24/7.

### Ahhh! Do you mean some guy who has dedicated his flashy new 2.4GHz Athlon to it?

No, no, no. I mean seriously powerful, trouser-shaking machines that are huge, unwieldy and generally scare cats. These machines are enormous boxes with multiple processors, hundreds upon hundreds of megabytes of RAM and a permanent Internet connection.

### Why on earth would someone want to dedicate a machine like that to this distributed computing thing? Surely it would be better to use the machine for *Unreal* or something?

Well there are two things here. Firstly, the machine is not a normal desktop machine with fast graphics card...this is a beefy server that rarely sees graphics but has the ability to tread on bytes like no ones business. The other thing is that the people who run these machines are dedicated to the end with the projects and hence have the enthusiasm to dedicate machine such as this to it. There are also a number of companies who dedicate machines to the projects. With some projects there is cash involved also.

### Cash? You said the magic word. What is all that about?

Well, some projects (such as the RC5 project) offer cash prizes to the people who find particular solutions to the problem. As an example the RC5 project winner will get $1,000, $1,000 goes to the winners team (or the winner if he/she isn't on a team), $6000 goes to a non-profit organisation of the winners choice and $2000 goes to the distributed.net project for hosting the networking and supplying the code.

### Hmmm, so there is money involved, do you think I could win it?

Well that depends. You could theoretically log on for the first time and your first chunk could contain the solution, but it shares the metaphor of a haystack with the mandatory needle well and truly lost in it. Give it a go though because you may be lucky, and if you contribute anyway you are doing a good thing.

### So what is this RC5 thing anyway?

It is best if I let the website describe it. They say: "The 'Bovine' RC5 effort was formed to take the responsibilities of coordinating and maintaining the »

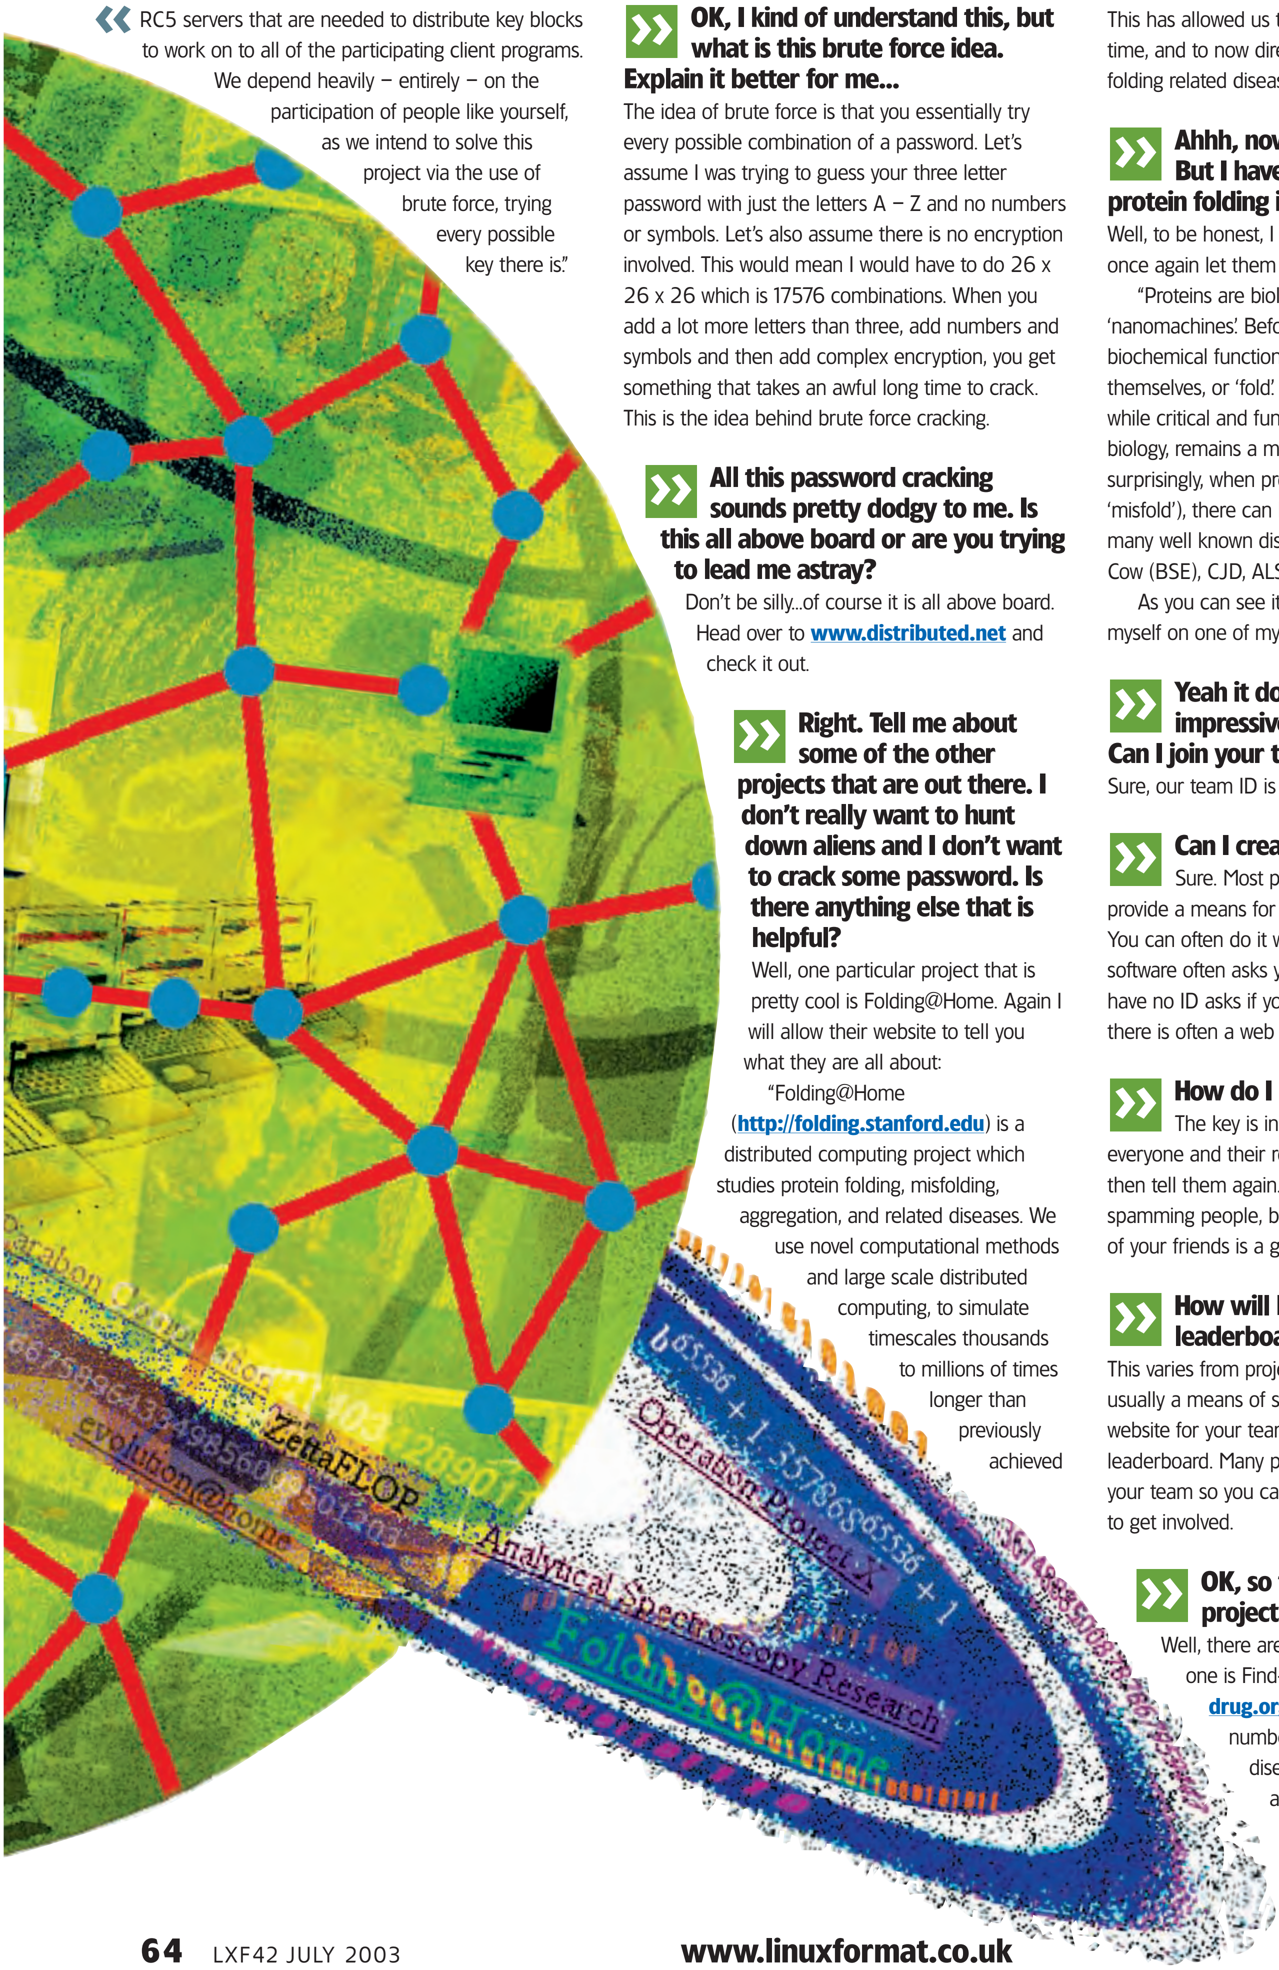

RC5 servers that are needed to distribute key blocks to work on to all of the participating client programs. We depend heavily – entirely – on the participation of people like yourself, as we intend to solve this project via the use of brute force, trying every possible key there is."

### OK, I kind of understand this, but what is this brute force idea. Explain it better for me...

The idea of brute force is that you essentially try every possible combination of a password. Let's assume I was trying to guess your three letter password with just the letters A – Z and no numbers or symbols. Let's also assume there is no encryption involved. This would mean I would have to do 26 x 26 x 26 which is 17576 combinations. When you add a lot more letters than three, add numbers and symbols and then add complex encryption, you get something that takes an awful long time to crack. This is the idea behind brute force cracking.

### All this password cracking sounds pretty dodgy to me. Is this all above board or are you trying to lead me astray?

Don't be silly...of course it is all above board. Head over to **www.distributed.net** and check it out.

### Right. Tell me about some of the other projects that are out there. I don't really want to hunt down aliens and I don't want to crack some password. Is there anything else that is helpful?

Well, one particular project that is pretty cool is Folding@Home. Again I will allow their website to tell you what they are all about:

"Folding@Home (**http://folding.stanford.edu**) is a distributed computing project which studies protein folding, misfolding, aggregation, and related diseases. We use novel computational methods and large scale distributed computing, to simulate timescales thousands to millions of times longer than previously achieved This has allowed us to simulate folding for the first time, and to now direct our approach to examine folding related disease."

### Ahhh, now this sounds interesting. But I have no idea what this protein folding is all about. Explain.

Well, to be honest, I am no biology student so I will once again let them tell you.

"Proteins are biology's workhorses – its 'nanomachines'. Before proteins can carry out their biochemical function, they remarkably assemble themselves, or 'fold'. The process of protein folding, while critical and fundamental to virtually all of biology, remains a mystery. Moreover, perhaps not surprisingly, when proteins do not fold correctly (*ie* 'misfold'), there can be serious effects, including many well known diseases, such as Alzheimer's, Mad Cow (BSE), CJD, ALS, and Parkinson's disease."

As you can see it is a very worthy project. I use it myself on one of my home PCs.

### Yeah it does seem pretty impressive and helps real people. Can I join your team?

Sure, our team ID is 31805 and we are Wolves LUG.

### Can I create my own team?

Sure. Most projects that support teams provide a means for you to create your own team. You can often do it within the client software. The software often asks you for a team ID, and if you have no ID asks if you wish to set a team up. If not, there is often a web interface to create one.

### How do I get people on my team?

The key is in publicising it. Make sure everyone and their robot dog knows about it and then tell them again. I certainly don't mean spamming people, but putting up ads and telling all of your friends is a good start.

### How will I know how far up the leaderboard of teams I am?

This varies from project to project, but there is usually a means of searching on the projects' website for your team and a list of teams on the leaderboard. Many projects have a constant URL for your team so you can tell people to go to that URL to get involved.

### OK, so treat me. What other projects are there?

Well, there are stacks of them. Another good one is Find-a-drug (**www.find-a-drug.org**). This project hunts for a number of solutions to various diseases including SARS, Cancer and Bio-terrorism diseases. Another is the Rothberg Institute for Childhood

Diseases (**www.childhooddiseases.org/community_tsc.html**). Their project looks for Tuberous Sclerosis Complex (TSC). Keeping in tune with the health theme there is also Parabon Computation's (**www.parabon.com**) project which helps to fight Cancer.

### What other types of projects are there?

How about the Climate Prediction project (**www.climateprediction.net/index.php**) that will predict our climate in 50 years? **www.grub.org** project uses distributed computing to create an up to date search engine by using the machines to trawl the Net. ChessBrain (**www.chessbrain.net/cbindex.html**) is another project that uses all the machines to create an enormous chess computer. There are plenty of different projects about – you may even make some money!!

### Hang on, you have done it again Bacon. What do you mean about making money?

Sorry about that. Some projects have been set up in which you can make some money from your contributed processing. You are usually paid a certain amount of money for a certain amount of time or data processed.

### Cool, fill me in with the projects where I can get paid!

Well here are a couple of projects:

**Gomez PEERCommunity** **https://www.porivo.com/peernetwork/jsp/programDetails.jsp**

This project will evaluate the performance of Internet sites and find bottlenecks.

**CapacityCal** **www.capcal.com/**

Project measures in real-time to loads of sites.

### So have any of these huge distributed computing projects been solved?

There are a number of projects that have reported successes. One of the most prolific has been the RC5 project. The project has so far had two successes; the RC5-56 key was cracked in 256 days, and RC5-64 key was cracked in 1,757 days. These are considerable processing challenges and really demonstrate the sheer power of the technology.

### Interesting stuff. It strikes me that the distributed computing concept has real potential, particularly for health-related projects.

That's right. Many drugs are developed from extensive simulations, and these simulations are usually run on expensive machines that require lots of processing power and memory and are very costly. Distributed computing has opened up the possibility of spreading this load out considerably and reducing the costs to just a server that pumps out chunks to be processed and sucks in the results from the clients that have processed them.

### So what other uses does this distributed computing have?

Well there a number of other projects and uses. One of the major uses is for rendering 3D animations and images. In many cases 3D animations take hours to render and multiple machines can be used for this.

### Is this how they rendered *Toy Story* then?

That's right. In fact, each frame (and there are between 24 and 30 of them a second in cartoons) took anything from 6 hours to 90 hours to render in the Pixar rendering farm.

### Jeez, yet more jargon.

No, jargon is something a Norfolk guy does to warm up before a marathon. A rendering farm is just a batch of machines using distributed computing to render the images or animations.

### Cool, can I do this in Linux?

Well, there is work afoot to create a network rendering system for the *Blender* 3D Modelling tool, and if you have been clever, you will have been reading my series on *Blender* in *Linux Format*. As the *Blender* files get bigger and the animations get more complex, the distributed rendering daemon will be useful. As a side note, there is also a distributed processing project called the Internet Movie Project (**www.imp.org**) to create a networked render farm.

### OK, so what about the future of distributed processing?

Well, distributed computing seems to have a bright and thriving future. There are a number of projects that are making use of the technology, and the projects that are using it to find cures and solutions to medical disorders are proving that the technology has a very real impact on people.

The great thing is that as technology improves and people get faster and faster machines, the distributed networks across the world will have more and more powerful resources available.

### Surely this distributed processing could be used by Open Source projects for compiling code for other platforms?

That's right. Theoretically, there is great potential for distributed computing being used for a variety of purposes such as sharing the load compiling daily snapshots for example. There is also great potential for distributed computing being used for artificial intelligence and other huge computing tasks – in fact any area of computing that needs large amounts of data-crunching can be performed using the distributed computing model.

### I heard that there are other types of distributed computing also. Is this true?

Well, in many ways any kind of shared processing or resource is a distributed resource. Theoretically file sharing networks are also distributed networks, but not distributed processing networks as they merely distribute files and don't compile files.

Another form of distributed processing that is under consideration by a number of companies is the idea of sharing game processing data across multiple machines. This is a concept Sony is rumoured to have developed for its up and coming Playstation 3 games console. The idea is that the console will spread the load to other consoles with spare processing cycles. Although not officially announced, this technique could be used to maximise the resources of people consoles. As an example, a console playing *Tetris* is not going to be as busy as a console playing *Tekken15*.

### Surely though if Sony did this the quality of the processing would vary depending on what people are playing at the time?

Well, I was thinking about that as well. I suppose the true test is in the detail and we will have to wait and see if anything happens.

### I'm feeling brave. How would I set my own distributed network up?

Well first of all you need to think of a mammoth processing project. If you have it, you will need to hack one of the existing clients to process your project. You will also need to set up a website to handle teams and users.

Not only would you need to set up the client and website, but you would also need considerable bandwidth to send out chunks to be processed and receive results from people. I suspect most ISPs would not be happy with this level of bandwidth required.

It is a big challenge and one that I have never done. I refer you to my good friend Google...

### Well thanks for the information. I feel clearer on what this whole distributed computing thing is all about. Any finishing points?

Not really, just that you should really check all this out. Distributed processing is not only incredibly cool to be a part of in the sense of the bigger picture, but is also a really good cause for some of the projects. You could think of adding a machine as your contribution to a charity; instead of donating money, donate clock cycles to help cure a disease.

If you've got the space, I would recommend that you go to a computer fair, get a cheap machine, install Linux on it and use it as a dedicated machine. I have done it and it is always pushing me higher and higher up the leaderboard in the teams that I am on. Good luck! LXF

# Tutorials

Our experts offer help and opinions on a whole host of Linux applications

## YOUR GUIDE TO GETTING THINGS DONE!

Whether you are just starting out in Linux, or an experienced veteran, there's always more to learn. Every issue of *Linux Format* is packed full of practical advice, and nowhere is it more concentrated than in our tutorials pages.

Here you'll find expert guides to all sorts of things, from Basic Linux usage to understanding and deploying network solutions, from simple script coding to the complexities of Perl regular expressions, Java server apps and more. We aim to bring a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact **linuxformat@futurenet.co.uk**, or log on to our website at **www.linuxformat.co.uk** to post your suggestions in our forums? Or use snail mail... Hope to hear from you soon!

**Nick Veitch** EDITOR

## HOW CODE IS REPRESENTED

**Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:**

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
```

**otherwise, there is usually a gap between lines:**

```
begin
  mniWordWrap.Checked := false
end;
```

**Usually, you'll find the code on our CD/DVD too.**

## THIS MONTH TEACH YOURSELF...

### Beginners' Guide »

NEW SERIES! Taking the plunge into Linux? Start here! This month: Choosing Your Distro **p68**

### DHCP

Strore all your network settings on a single server with Dynamic Host Configuration Protocol **p70**

### The GIMP

Bad photos? Sort out indecent exposure problems without a trip to the police station **p72**

### PHP

Writing extensions and file manipulation needn't be a hassle with **zval** and **ext_skel** **p76**

### Python

Using packaging mechanisms such as modules, packages, importing, the module search path and module reloading **p80**

### Blender

Let there be light! It's all very well building great 3D models, but you can't see them in the dark... **p84**

### Runtime Revolution

Bringing a sense of scale and proportion to our self-built image-viewing application **p88**

## TIP OF THE MONTH! WGET

Ever found yourself wanting to snag a file from somewhere on the Internet without having to go through the fuss of opening a browser/FTP client and entering the URL? Or, have you ever found yourself wanting to be able to fetch files as part of a script? If so, **wget** is for you.

For basic purposes, just type **wget http://server.com/foo.tar.gz**, and **wget** will copy it locally. **wget** can also be configured to analyse HTML pages it downloads and follow links/download associated pictures by adding the **-r** option before the URL to grab, allowing you to snag a local copy of entire websites this way – you can even add the **-k** switch to have **wget** convert links in the HTML to work locally!

### So many uses

**wget** has so many uses that we don't have space to list them all here. However, he's a quick taster of the most important options you can use: try the **--delete-after** switch to have **wget** download all the files it finds then delete them immediately afterwards – perfect if you want to pre-fill your proxy server with content. You can use the **-H** switch to allow **wget** to download files from other hosts when it's in recursive mode (otherwise it won't go beyond the original domain). The **-b** switch is available to make **wget** work in the background, which is good if you use it as a backup tool. Use **-c** to make **wget** try to continue a download from where it left off. Use **--limit-rate** to specify how much bandwidth **wget** is able to make use of. For example, **--limit-rate 5k** would stop **wget** using more than 5KB/s to download files. Finally, use **-w** to specify how much time **wget** should wait in between file downloads – this is crucial if you don't want to overload small servers!

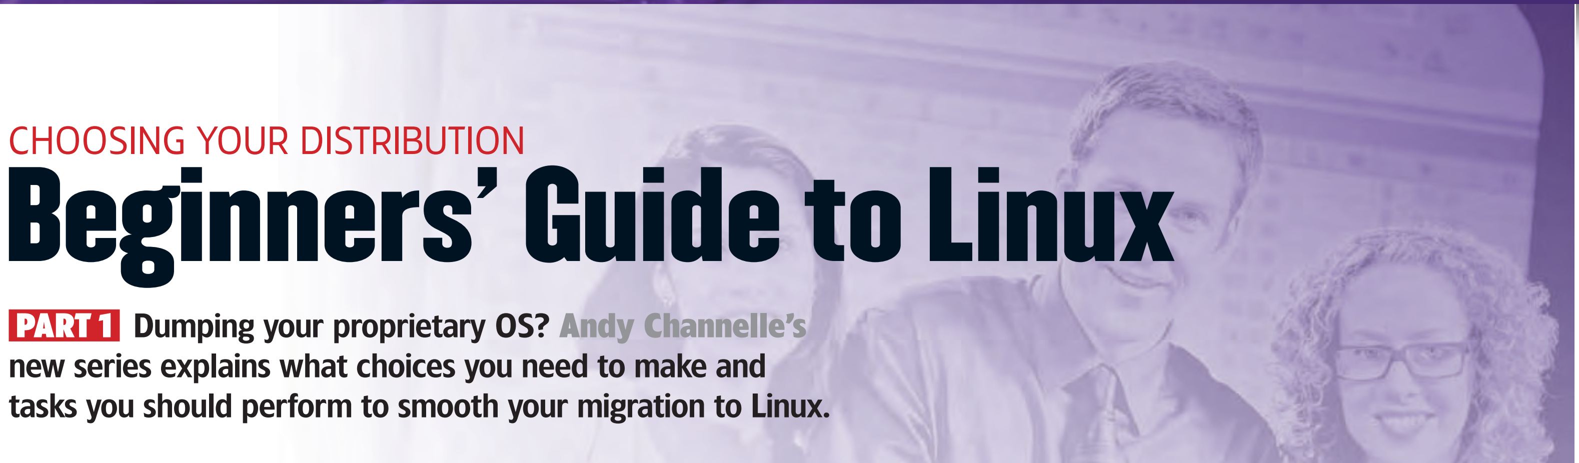

CHOOSING YOUR DISTRIBUTION

# Beginners' Guide to Linux

**PART 1 Dumping your proprietary OS? Andy Channelle's new series explains what choices you need to make and tasks you should perform to smooth your migration to Linux.**

Linux is all about choice, and one of the earliest choices you will be called on to make concerns the most fundamental aspect of the system: the distribution. The website www.distrowatch.com, which is dedicated to posting the most up-to-date news on new releases, currently lists some 130 LInux distributions in active development. While this is an awful lot, the number can be broken down into various sections. At the bottom are the projects such as *tomsrtbt*, a distro designed to rescue broken systems. At the other end of the scale there is *Red Hat Linux Advanced Server* which is aimed at enterprise users running large file, print and web servers. In between we have everything from firewall appliances such as *SmoothWall* to packages aimed at home and business desktop users, which is what we're looking at here.

Right, so we've narrowed the field down to the home/office sector, now it is time to decide whether we want a fully featured distribution that has a great deal of software included (a lot of which we'll probably end up ignoring) or a more stripped package intended to cover the basics, which is usually defined as email, web browsing, word processing, MP3 ripping/playback, digital photo management and games.

## Want it all?

The market leaders in the 'full service' sector are Red Hat, which dominates North America; SuSE, which holds a similar position in mainland Europe; and Mandrake, which is challenging everyone all over the world. As with Microsoft's Windows OS, it is not unusual for these companies to target different markets with a range of packages; all offer at least a Personal and Professional edition of their distros. Often there is also a free version available which can be downloaded and burned to a CD or installed across the Internet which, if you have a broadband connection, makes choosing your distro a lot easier: try before you buy! The free versions of distros typically come with no support, so if you're concerned about installation hassles (and want to support your favourite developer) you'll find a range of boxed products at online stores such as www.dabs.com and www.simply.co.uk or even in PC World. As well as support, these retail products usually come with an invaluable collection of manuals, a full set of CDs and, in most cases, a selection of stickers to decorate your PC!

Prices range from £29.99 for SuSE 8.2 Personal to £120.00 for Red Hat 9 Professional, and the latest releases have all received glowing reviews over recent months in *Linux Format* so you know that whatever you choose, you're getting quality kit. Like the prices, the facilities on offer will vary, though at the very least you can expect to get a fully featured office suite capable of opening and saving the most common formats, a full selection of graphical Internet applications, photo manipulation/management tools, file managers, media players, CD burning utilities, home/business finance packages and a collection of games. More often than not there will be more than one of these elements to choose from. The Professional versions are usually distinguished by the inclusion of fully licensed commercial software or more fully featured administration or server tools. SuSE Linux Professional, for instance, includes a copy of *MainActor*, a high-quality video editing package. All the main distro sellers include a complete list of applications included in all their packages, so visit the websites, hunt around and choose the one that suits your needs.

Some of the names to look out for include:

**The GIMP** excellent Photoshop-like image creation and editing package

**MOZILLA** Open Source offspring of the old Netscape web browser

**OPENOFFICE.ORG** extensive productivity suite featuring Word Processor, Spreadsheet, Illustration and Presentation tools.

**EVOLUTION** email/groupware client in the Outlook vein.

**XMMS** capable – and skinnable – MP3 player.

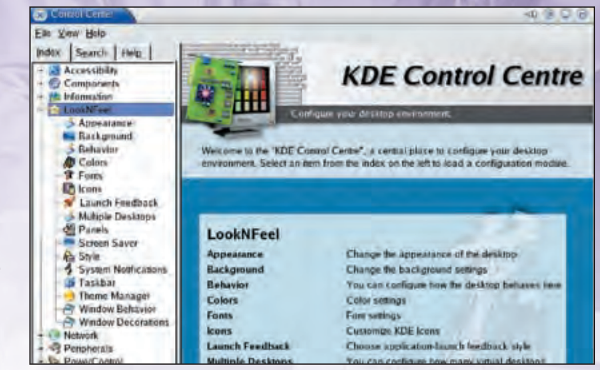

Welcome to the "KDE Control Centre", a central place to configure your desktop environment. Select an item from the index on the left to load a configuration module.

LookNFeel

| | |
|---|---|
| Appearance | Change the appearance of the desktop |
| Background | Change the background settings |
| Behavior | You can configure how the desktop behaves here |
| Colors | Color settings |
| Fonts | Font settings |
| Icons | Customize KDE Icons |
| Launch Feedback | Choose application-launch feedback style |
| Multiple Desktops | You can configure how many virtual desktops |

***KDE Control Center* can tailor the behaviour of related applications to your exact needs.`**

## Linux Basics

**What On Earth is a distro?**

**WHAT IS A DISTRO?**

A Linux distribution, or distro, is a collection of software put together by a company or organisation. The core of this is the operating system (OS), itself a compendium of parts, that acts as a platform for the applications which we use to do things like send an email or remove 'red-eye' from a digital photo.

**HOW IS THAT DIFFERENT FROM JUST AN OPERATING SYSTEM?**

As well as containing the thing that makes a computer work – the OS – a distro includes a selection of software. In many cases this could be upwards of 2000 applications covering everything from basic everyday uses to highly specialised jobs.

**SO, ARE MICROSOFT WINDOWS AND APPLE OS-X DISTROS?**

Well, sort of. While these OSes may feature integrated web browsers, movie editing software etc, the breadth of software is usually limited in both scope and functionality. The big Linux distros, however, will include applications to do almost any PC-based task you can think of, and some you probably couldn't! Some distros let you install only the components that you want.

**KDE** integrated desktop environment including office, connectivity and educational applications.
**GNOME** alternative to KDE
**TUXPAINT** happy, noisy paint package aimed at children.
**PIXIEPLUS** fantastic photo management application that can format collections of images into web pages.
**WINE** package designed to install and run SOME Windows applications.
**APACHE** industry standard web server equally at home running personal pages as it is serving up enormous corporate websites.

## Want it easy?

If all this seems a little complicated, a new breed of Linux developers are endeavouring to make the process of installing and managing a Linux system foolproof, as well as forging links with hardware companies to get their products pre-installed. The three names making an impact here are Lindows, owned and run by Michael Roberts (formerly of MP3.com); Xandros, which rose from the ashes of Corel's abortive entry into the Linux market; and Lycoris. Ironically these three systems, though less 'complete' than those mentioned above, are more expensive – putting a premium on easy installation and friendliness to users of other operating systems. Lindows, for instance, can go from a bare system to completely installed in just eight minutes. But this simplicity must be balanced, in most cases, by the drastically reduced number of included applications. While Lycoris sells its Desktop/LX system for US$29.95, this price doesn't include the companies integrated version of *OpenOffice*, for which you'll need to shell out an additional $39.00, though you can always get the 'official' *OpenOffice* suite which costs nothing. Desktop/LX does, however, include useable selection of applications and there is an unsupported 'Community Edition' available. Lycoris also has an edition tailored toward games players featuring *WineX*, a special version of *Wine* capable of playing many Windows games.

With Lindows the situation is more serious: it comes with the bare minimum. No office suite, no calculator (!), no graphics application, to get these you'll have to join Click-'N'-Run which adds an annual charge (currently $99) to the initial charge of US$129 for the actual OS. This was fine because the price included unlimited access to the 'Warehouse', which included gems such as *StarOffice* and *TuxRacer*, but now Lindows has begun adding premium content, which obviously, attracts a further cost. You can, with a little bit of effort, configure Lindows to use the installation system from Debian, which is widely regarded as the best package management system around.

The situation with Xandros is a little different again, as the US$99 Deluxe version includes *CrossOver Office*, an application which can install and run the *Microsoft Office* productivity suite, *Quicken* and other Windows software.

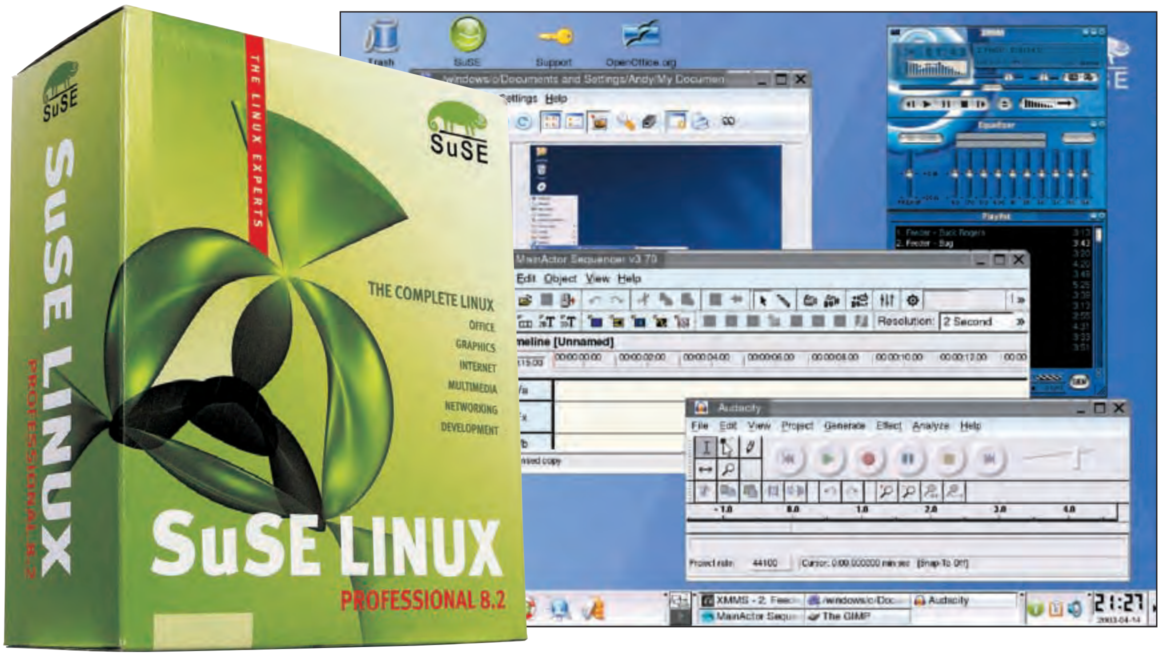

**Apart from the icons looking a bit different, SuSE is no different from Windows in terms of the desktop functionality that is available.**

These are all good choices, and as you're using Linux – regardless of the company to which you give your custom – there will always be an alternative available.

## Best of both worlds

While the hand-holding and simplification of Xandros, Lindows and Lycoris can be tempting, the packages available from the 'traditional' Linux sellers provide far better value for money. While each of these products will grow with your experience, the products from SuSE, Red Hat and Mandrake are updated more frequently, contain a wider selection of applications and give these others a run for their money in the usability stakes. The use of KDE or GNOME across the board also means that moving from one system to another won't trigger the (potential) culture shock of your first experience with Linux.

Read the reviews, browse the forums of distro sellers – but don't labour over your decision too much, picking up a dud is actually quite difficult– and then take the plunge! LXF

## NEXT MONTH

**Is Linux really difficult to install? We'll show you how a little preparation can make what is regarded as a notoriously difficult job into a walk through Elysian fields.**

## Glossary

### Linux terminology

**OPEN SOURCE** Software which is distributed in both binary and source form. The binaries are pre-packed and ready to install while source must be compiled. The value of including the source code is that users are not limited to the capabilities of a given package, they can add or take elements away - if they have the right skills. One of the most common licenses (the rules governing what can and can't be done with a package) is the GNU General Public License (GPL).

**KERNEL** The most current series is 2.4. This is the very heart of the system.

**XFREE86** This is what drives the Graphical User Interface (GUI), handling the resolution and colour depth of your screen. The most recent versions include the ability to change these elements 'on the fly'.

**KDE/GNOME** The two main desktop environments. Current versions are 3.1 and 2.2 respectively. The desktop sits between you and the computer, so if you don't already have a preference, take your time and play around before you opt for one over the other. Both Red Hat and Mandrake have made efforts to make the two main desktops work in a similar fashion, while Xandros, Lindows and Lycoris are pure KDE-based distros.

**RPM** The Red Hat Package Management system which is used to install, update and remove software from your PC. Mandrake, SuSE, Red Hat (obviously) and many other rely on RPM to maintain their systems. The alternative is the .deb format used by Debian. RPMs have a .rpm file extension.

**DEBIAN** A free (as in speech) software distribution that encapsulates the Open Source ethos. While not the easiest thing to install, Debian is one of the most complete Linux distributions you can get and once installed, management is very simple.

NETWORK ADMINISTRATION

# Configuring a DHCP server

**Richard Drummond** and **Nick Veitch** ease us through a bread-and-butter sysadmin task.

The Dynamic Host Configuration Protocol (DHCP) is a system, based on the earlier BOOTP or Bootstrap Protocol, that lets a network host query a server for the information it needs to set up a network interface. DHCP extends and enhances this system by implementing a scheme whereby network addresses can be allocated dynamically to hosts as and when they are needed, rather than relying entirely on static Internet Protocol (IP) numbers.

As you can imagine, DHCP can make network management much simpler. All the network settings can be stored on a single server rather than requiring each host's settings to be stored locally. This is much easier to administer – and the ability to allocate IP addresses dynamically makes adding and removing hosts to a network much simpler.

DHCP is defined in RFCs 2131 and 2132 and basically works as follows. The DHCP client sends out a DHCPDISCOVER message as a network broadcast to look for a DHCP server. Any local DHCP servers will reply with a DHCPOFFER message, offering an IP address to the client. The client may have retained an IP address from a previous DHCP session (in its cache, depending on the client being used) and can suggest this IP number in the DHCPDISCOVER message in an attempt to get this number reassigned.

The client chooses which offer it wants, and broadcasts a DHCPREQUEST message selecting the server that made the offer; any other DHCP servers will interpret this as declining their offers. The selected server will bind the offered IP address to the client and return a DHCPACK message to the client containing configuration parameters.

There are two methods of allocating IP addresses in DHCP: automatic and manual. For manual allocation, the administrator must assign an IP address to a specific host. With automatic allocation, the server assigns - or leases - an IP address from its pool of addresses for a finite period. When that period is up, the client has to extend its lease with a further request. A client can ask for lease period when it makes its DHCPREQUEST; however, this may not be granted if its longer than the server's maximum lease period. If a client doesn't specifically request a lease period, then the lease will then be set for the server's default. When a client is finished with an address, it can send a DHCPRELEASE message to relinquish it and return it to the pool.

## Using DHCP with Linux

Many implementations of DHCP – both server and client – exist for Linux. The most commonly used is the Internet Software Consortium's (ISC) reference implementation (see the details at **www.isc.org/products/DHCP**). Version 3.0 patch version 2 is the current stable release, while version 3.0.1 is currently being beta-tested. Some distros such as older versions of Red Hat and Debian also make use of the minimal DHCP client, *pump*.

At the client side of things, DHCP requires little if any configuration, so we won't cover that here. Usually it is simply a matter of running the client, and specifying a network interface to be configured. In virtually all cases, ISC's DHCP client will work with an empty config file – and all modern distros offer the option of automatically using DHCP to get network settings during the installation procedure.

Instead, we will focus on configuring a DHCP server, in particular the ISC's implementation. All the major Linux distributions offer this as a pre-built package – and it's far easier to use this instead of compiling it yourself.

Using DHCP on older Linux kernels can cause some subtle routing problems with some DHCP clients, but with recent 2.2 and 2.4 series kernels you shouldn't encounter any difficulties (see the mini-HOWTO or the docs). The only significant requirement is that you have multicasting enabled with the CONFIG_IP_MULTICAST option – which all precompiled kernels should do. Check by calling **ifconfig** and look for the word MULTICAST in diagnostics for your network interface.

## Configuring

The ISC's DHCP server is called, unsurprisingly, *dhcpd*. It is typically launched as a daemon during system startup. Its configuration is stored in /etc/dhcpd.conf and any leases granted are recorded in /var/dhcp/dhcpd.leases (although the latter may vary according to your distribution). See the dhcpd.conf and dhcp.leases man pages for a comprehensive examination of the syntax of these files. There is also a sample config file supplied from which most of the options can be determined.

The DHCP daemon doesn't read this config file dynamically, so you will need to remember to restart the service if you make any changes to it. If the thought of manually configuring *dhcpd* is too much for you, then you should also note there are many tools to ease the job. For example, the remote admin tool webmin includes a DHCP module which provides form-based configuration of a DHCP server from any client with a web browser (see **www.webmin.com/webmin**).

The file /etc/dhcp.conf contains any default settings and any DHCP options (options are the configuration parameters you wish to supply to a client) to apply globally, followed by a series of statements describing the topology of your network. You have to tell dhcpd the network address and netmask of any subnets that you want to offer a DHCP service on, and the pool of addresses to allocate IP numbers from on each subnet. This done with the subnet statement. For example:

```
subnet 192.168.1.0 netmask 255.255.255.0 {
 range 192.168.1.2 192.168.1.254;
}
```

A subnet statement can additionally include any settings and options you wish to apply to clients within that subnet. These will override any global options. Have a look at this excerpt:

```
option domain-name "mylocalnet.com"
subnet 192.168.2.0 netmask 255.255.255.0 {
 range 192.168.2.10 192.168.2.240;
 option domain-name "workstations.mylocalnet.com";
}
```

Here the domain name for clients is specified globally, but hosts on the subnet 192.168.2.0 will be given the domain name 'workstations.mylocalnet.com'.

Hosts that require special configuration, such as those that are assigned fixed IP numbers manually, or hosts that require BOOTP services must be identified with the host statement. This can include options to apply to that host. The hardware declaration is used to identify a particular host by its hardware address (MAC):

```
host anaximander {
 hardware ethernet 00:01:02:B7:35:B6;
 option fixed-address 192.168.1.1;
}
```

Here the fixed-address options assigns this fixed IP number to the machine with specified ethernet address. If you miss out this option, the DHCP will try and assign an IP address dynamically as per usual. If you do need a fixed IP though, it is often better to configure that host manually without DHCP – the Ip is assigned to the card (via the MAC address) so if you switch hardware, it won't work anymore.

Note that you can group hosts (and subnets) together with a group statement and apply options to the group.

## More depth

Time for a fuller, although still fairly simple, example for setting up a DHCP client. For a complete list of DHCP options, consult RFC 2132, but for now, try this:

```
# global domain name for this network
option domain-name     "mynetwork.org";
# default gateway
option routers         172.16.0.1;

# default DNS servers
option domain-name-servers 172.16.0.1 172.16.0.2;

# default network broadcast address
option broadcast-address 172.16.255.255;

# default lease period = 12 hours
default-lease-time 43200;

# max lease period = 24 hours
max-lease-time 86400;

# Our subnet. Note the pool has two ranges
# with a break in the middle.
subnet 172.16.0.0 netmask 255.255.0.0 {
 range 172.16.0.10 172.16.0.49;
 range 172.16.0.60 172.16.0.240;
}
# Static IP address for mail server.
# On subnet above, but uses IP address not in pool
host smtp {
 hardware ethernet 00:01:02:B7:35:B6;
 option fixed-address 172.16.0.50;
}
```

The options used here in this tutorial are the ones you will encounter most often and are sufficient for a client to get a network interface up and running; they are all commented and should be fairly self-explanatory. LXF

## Which IP numbers can I use?

### The Class system

Obviously, if you start just assigning any old numbers to the IP addresses of computers on your network, without a doubt, chaos will ensue. Fortunately, several sets of IP numbers have been specially reserved for use in private networks to take the guesswork out of network administration. The first of these is the range of Class C addresses which should be used by small networks. The full details of these classes are available in the RFC1918 document.

Typically, you would use sub-sets of these to form the pool of available addresses for your DHCP server. Remember you should always leave a number of these addresses free for static IPs you may want to use (*eg* for gateway machines, mailservers etc.).

**Class A:** 10.0.0.0 to 10.255.255.255
**Class B:** 172.16.0.0 to 172.31.255.255
**Class C:** 192.168.0.0 to 192.168.255.255

We start with two images, both of which are very colorful. We've placed the puppy's image
The opposite of Darken mode is Lighten mode. In this case, if the pixel
puppy layer
lighter than
responding
pixel in the
balloon
then the
Value
in a
man
The GIMP
File Xtns Help
Fill this layer with a dark, reddish-brown color. Try a hex triplet value of #a46711 in the Color Selection window. Your tastes will vary here. Next, add noise using the Noisify
Start with clouds. The clouds should be relatively soft, little to no harsh edges if possible. Less cloud fill is probably better for a grunge look
The scanned image of this complex shape is not completely black and white. We first run Image->Colors->Levels Auto

IMPROVE YOUR PHOTOS

# Subtle Changes

**PART 5 Michael J Hammel introduces the photo-editing capabilities of this great art application.**

Underexposed images can occur because of poor lighting or by improper film processing. Fixing the mistakes isn't as hard as you might think using a few basic tools of *The GIMP*.

There are many problems in photography that are difficult to solve with the mechanics of the shoot but easy to fix with digital imaging software. Problems like underexposed portions of an image, washed out colour from poor processing or poor lighting, and redeye can all be addressed with minimal effort using tools like *The GIMP*.

This month's tutorial looks at all three of these common problems and how you can fix them with only a few minutes of work.

## Scanner Support

### Where to get your drivers

| | |
|---|---|
| **SANE** | www.mostang.com/sane/ |
| **XSANE** | www.xsane.org/ |
| **Vuescan** | www.hamrick.com |
| **Image Scan!** | www.epkowa.co.jp/english/index_e.html |
| **QuiteInsane** | http://quiteinsane.sourceforge.net/index.shtml |
| **Epson KOWA** | www.epkowa.co.jp/english/linux_e/lsd_e.html |
| **Linux USB Project** | www.linux-usb.org/ |

## CONTRAST, LEVELS AND CURVES: washing the colour back in

Washed out images are often the result of using the wrong speed film, overexposing the shot (which means using too much light), or because of poor film processing. In all three cases, much of the information available to you is still in the image – it just needs to be scanned at a high resolution and then post-processed using *The GIMP*. This tutorial goes step by step through fixing one image shot with the wrong speed film.

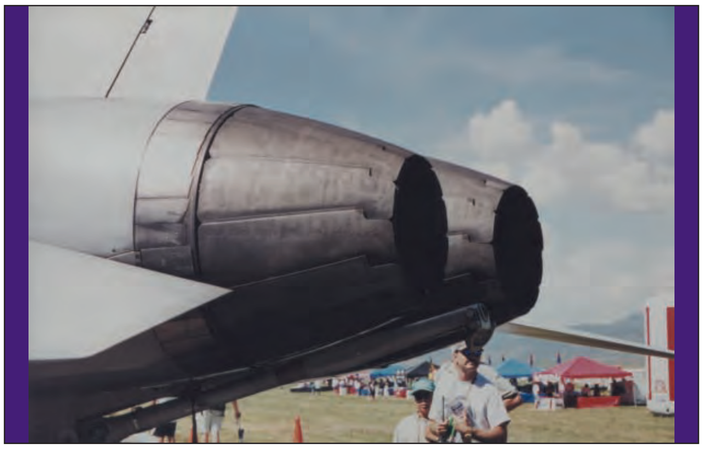

**1 We start with an image with a fair amount of detail hidden by our poor choice of film. The image content may be a bit on the boring side, but lets see if we can pull some colour – and a bit of detail – back out.**

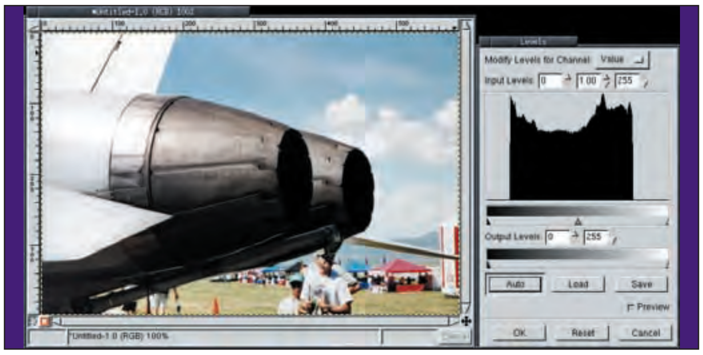

Untitled-1.0 (RGB) 100%

Modify Levels for Channel: Value

Input Levels: 0 1.00 255

Output Levels: 0 255

**2 The Levels tool is our first stop in correcting this image. This dialog shows us, through the histogram at the top, that the darkest pixel is not truly black and the lightest is not truly white. To fix this, we'll spread the range of colours so the darkest becomes black and the lightest becomes white. To do this, just press on the Auto button. The picture shows the Levels dialog and the histogram before the adjustment is made and the image after it has been Auto corrected. The result is good, but we can improve on this.**

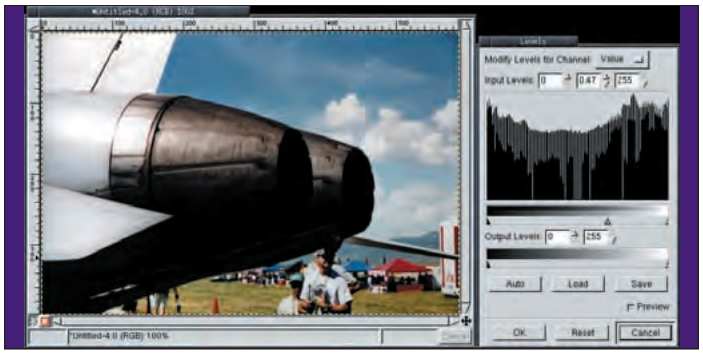

Untitled-4.0 (RGB) 100%

Modify Levels for Channel: Value

Input Levels: 0 0.47 255

Output Levels: 0 255

**3 One option here would be to manually adjust the Levels once again, this time moving the center point slider (beneath the gradient bar under the histogram) off to the right. The picture shows the image adjusted in this way. This reduces the white range and increases the black range, effectively bringing out more colour in the image by making it darker. Unfortunately though, this process is less than perfect. A better tool to use here is the Color Correction plug-in.**

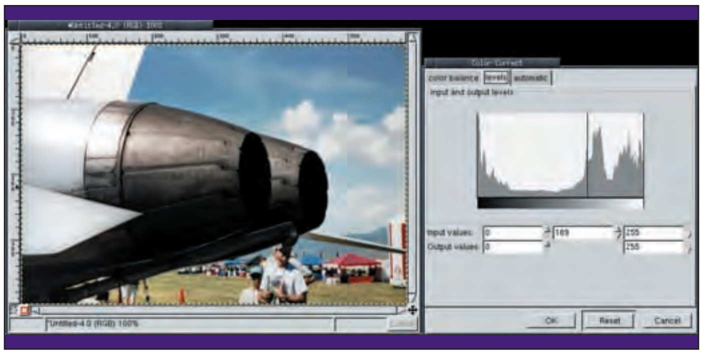

Untitled-4.0 (RGB) 100%

Input values: 0 169 255

Output values: 0 255

**4 The Color Correction Plug-In, available from the *GIMP* Registry provides a more detailed histogram in its Levels page. Here we see that there is a slight skew of the histogram towards the right – more pixels with heavier white content than black. We can adjust this by choosing one of the peaks to the right (we chose the first peak) to add more black to the image. The result is not as dark as the manually adjust Levels version, but more vivid than the Auto-adjusted image. It's better, but we still want more colour!**

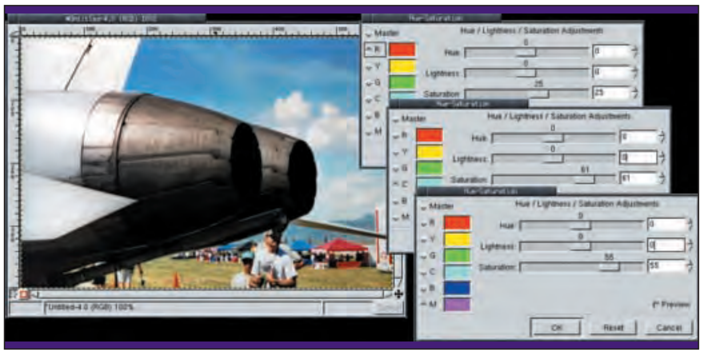

Untitled-4.0 (RGB) 100%

Hue / Lightness / Saturation Adjustments

Hue: 0 Lightness: 0 Saturation: 25

Hue: 0 Lightness: 0 Saturation: 61

Hue: 0 Lightness: 0 Saturation: 55

**5 Now come the most subtle changes of all – we want to improve the colour of the sky and tents using the Hue-Saturation tool. We're going to increase the saturation levels of three colour channels – the Red, Cyan and Magenta. These were chosen through a little trial and error, but adjusting these will bring out the blue in the sky and tents (from Cyan adjustment), the Red in the tent and cone (from the Magenta and Red adjustments) and the green and yellow in the grass (from the Red adjustment).**

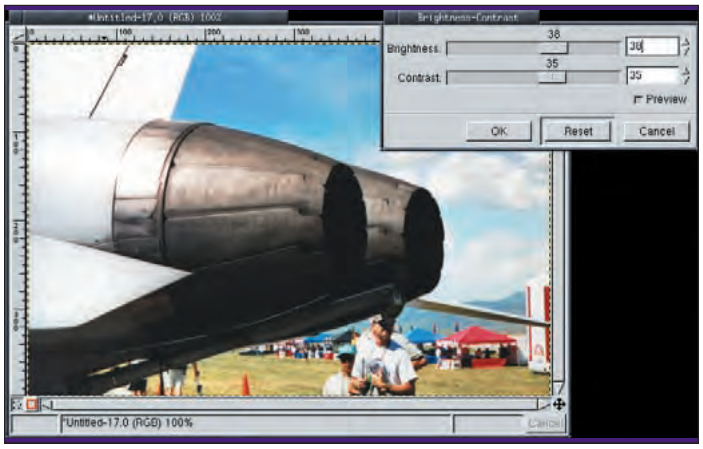

Untitled-17.0 (RGB) 100%

Brightness: 38

Contrast: 35

**6 One final adjustment to the Brightness-Contrast of the image to bring out a little more detail in the jet engines and we've quickly and easily turned this faded and unremarkable photo into something respectable!**

»

## BRINGING OUT LOST DETAILS: tweaking poorly exposed images

The last tutorial added a snap of colour to a poorly exposed image, but a closer look shows some detail is still not visible. How can we pull real detail from a poorly exposed image? By using the image data itself to brighten the areas that are too dark, that's how. Experiment with photos of your own to get a feel for it.

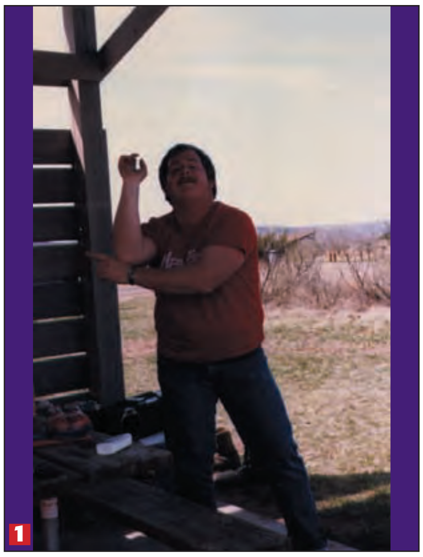

**1 Our original image shows a dancing fool on the plains of Texas (cows and oil wells not included). Standing in the shade of an outdoor shelter on a bright day left him and the shelter underexposed. But there is detail in those shadows! We start by duplicating the background layer.**

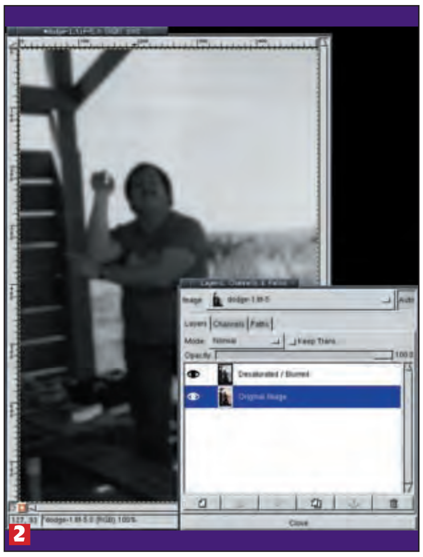

**2 The duplicate layer will be desaturated (Image> Colors>Desaturate) first. This layer will be added (actually, Screened) to the original to brighten the image. Desaturating it assures that we add even amounts to each colour channel. Next, blur this layer just a bit. We do this so the Screening process doesn't create harsh edges in the final image. The amount to blur depends on the image content, but in this case we only blurred about 8 pixels.**

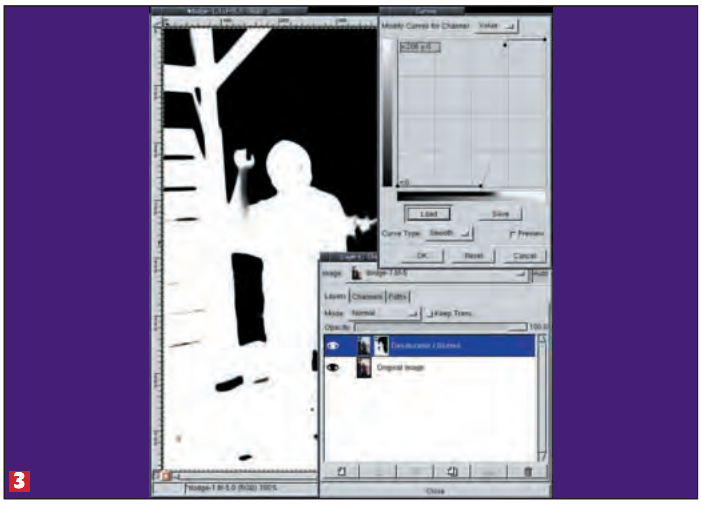

**3 With the duplicate layer active (its layer name is white on blue in the Layers and Channels dialog), use Edit>Copy to save a copy of the layer. Add a layer mask (right click on the layer name in the Layers and Channels dialog to get the Layers menu) to the duplicate layer. The layer mask will start out white. Click on the new layer mask to make it active and then paste a copy of the layer by using Edit>Paste (or Ctrl-V). Anchor the floating layer into the mask. With the mask still active, select Image>Colors>Invert to produce a negative mask. Adjust the Curves for the mask to isolate the area under the shelter and our dancing fool. You might have to manually touch up the mask using the airbrush or paintbrush. Once you're satisfied with the mask, blur it slightly – 10 pixels were used here but we came back later and tightened up that blur by adjusting the Curves yet again.**

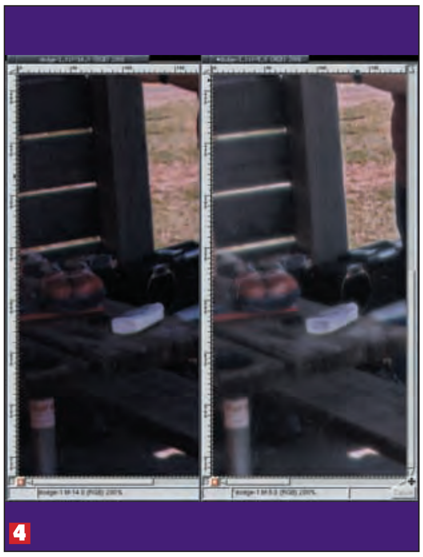

**4 After cleaning up the mask just a bit more around the raised arm by air brushing it white (so it would be lightened the same amount as the rest the body), we changed the layer blend mode to Screen. And that's it. A close up here shows how the detail in the final version compares to the same are in the original. The blur in the final image can be adjusted even more by reducing the blur used in the layer mask.**

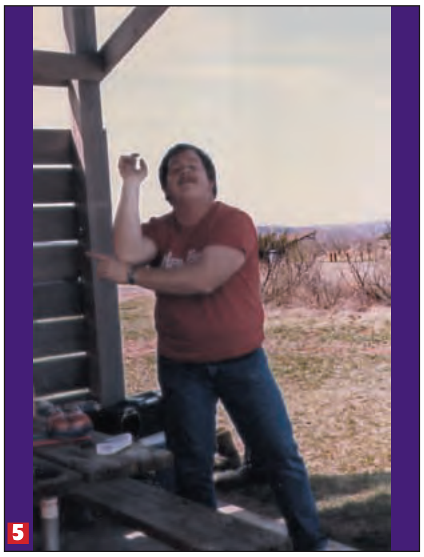

**5 The resulting image, while suffering from poor content, is at least clearer and filled with more detail. We could apply the steps from this month's first tutorial to bring out the colours if we really wanted, though here it might be safer to keep from bringing out too much detail in our dancer.**

## REMOVING REDEYE: make a proper selection

Probably the most common touchup for amateur photographers is the removal of redeye, the effect where pupils glow red when reflecting flash off the back of the eye. Removing this effect is fairly simple, though its not the colour removal where you'll get tripped up – its in making a proper selection.

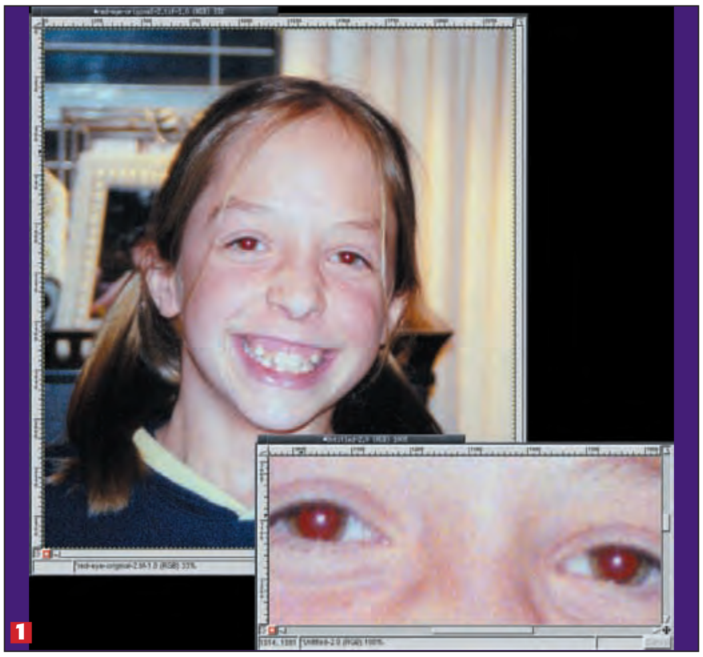

**1 Our original image isn't terrible, but those red eyes a definitely distracting. A close up shows that the red area actually bleeds a little in this photo, possibly due to poor scanning or even poor film development.**

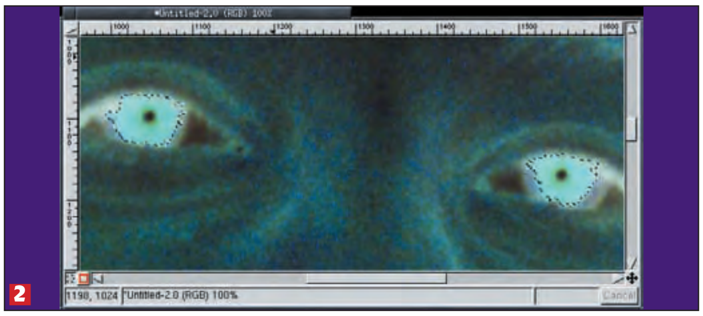

1198, 1024 Untitled-2.0 (RGB) 100%

**2 If this were a truly good photograph we might be able to get a quick selection of the pupils using just the Fuzzy select tool. But this photo is not that great, and we need a little trick to isolate those pupils. First, duplicate the layer so we don't corrupt the original layer. Next, invert the colours (Image> Colors>Invert) in the duplicate layer. This lets us see the regions of interest a little better. In the Layers and Channels dialog, choose the Channels page. All three channels should be active. Click on the Red and Blue channels to deselect them, leaving just the green channel. Now choose the Fuzzy Select tool from the GIMP Toolbox. Click inside the pupils. The selection you get may not be that good – you might have to adjust the Threshold level in the Fuzzy Select's Tool Options dialog (double click on that Fuzzy Select icon to get this dialog). We had to manually add the center points of the eye after our Fuzzy Select by holding the Shift key down while doing an Elliptical selection. Once you have the selection, feather it lightly, perhaps 3 pixels or so.**

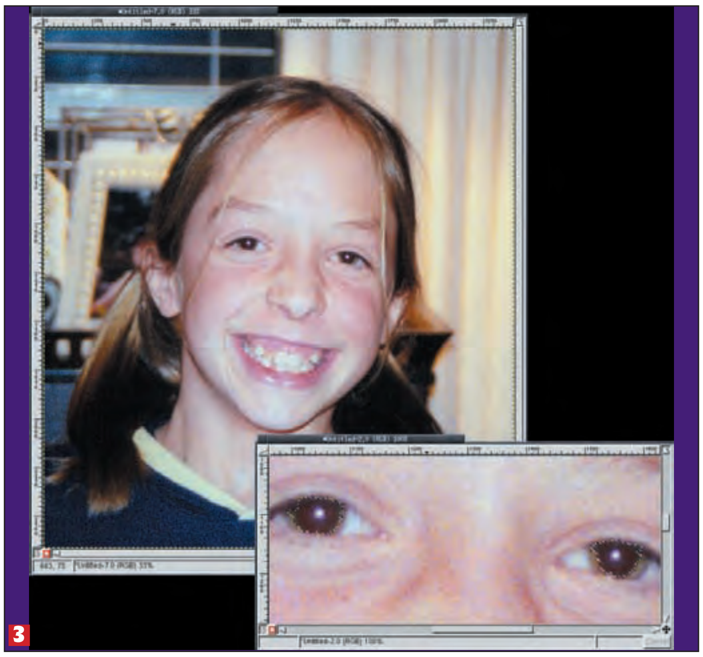

**3 Delete the duplicate layer – we don't need it anymore. Make the Red channel active in the Channels page of the Layers and Channels dialog but turn off the Green and Red Channels. Then choose Image>Colors> Desaturate. That will remove the red from the eyes.** LXF

## Treeware

### Books on *The GIMP*

*Grokking The GIMP*, Carey Bunks, New Riders ISBN: 0-7357-0924-6
*GUM – GIMP User Manual*, Kylander and Kylander, Coriolis; ISBN: 1-5761-0520-2 (also online at www.gimp.org)
*GIMP for Web Professionals*, Hammel, Prentice Hall ISBN: 0-1301-9114-0

## Printer resources

### For your *GIMP* masterpieces

www.cups.org/
CUPS provides networked printing support and a spooling interface.

http://gimp-print.sourceforge.net/
*GIMP-Print* goes way beyond just a *GIMP* print plug-in. It's a full-featured set of drivers for many printers. Its drivers are so well done they are considered better than some of the original manufacturers drivers and are being used by many Mac OS X users now as well.

www.linuxprinting.org/foomatic2.9/
Foomatic is the glue between print spoolers and printer drivers.

www.linuxprinting.org
This is the place to start to find out whether there is a driver available for your particular printer.

## URLS

### Online *GIMP* Resources

*The GIMP*
www.gimp.org
*GIMP* Registry (hundreds of plugins)
http://registry.gimp.org/index.jsp
Graphics Muse Tools CD
www.graphics-muse.com/gfxmuse/gfxmuse.html
Tutorials and other useful links:
www.graphics-muse.com/cgi/gmcat.pl?id=11

## NEXT MONTH

Cleaning up your old photos isn't all the difficult with *The GIMP*, as long as you know how to isolate the damaged portions. Next month we'll look at creating selections using various *GIMP* selection tools, including the often overlooked Quick Mask.

FILE MANIPULATION

# Practical PHP programming

**Paul Hudson** continues production of a PHP extension to handle tar files

When you followed the PHP tutorial in the last issue of *LXF*, you'll recall that our extension can now be compiled into PHP correctly. From here on, the process is simply a matter of editing the tar.c and tar.h file, and recompiling.

So, as this extension is designed to handle tar files, the first step is to create the functions available to PHP users. Although there are lots of things you can do with tar files, space is limited here, so we'll stick to three functions: **tar_list()**, **tar_add()**, and **tar_extract()**. These names are in line with the PHP coding standards (see CODING_STANDARDS included with the PHP source code) which dictate that functions should be lowercase, with words separated by an underscore, and should start with the "family" name of the function, in this case "tar". Here are the PHP function prototypes we'll work to:

**array tar_list(string tarfile)**
**bool tar_add(string tarfile, string file_to_add)**
**bool tar_extract(string tarfile, string location)**

It would be technically better to have an object-oriented solution, as it wouldn't require using **string tarfile** in each function. However, it would also require much more explanation than there is space for here.

## Listing tar contents

This is the easiest part of manipulating tar files, which is why it's first. As you can see in the PHP prototype, the function should accept a string of the tar file to read from, and should return an array of the files in that tar file.

Open up tar.h into your favourite text editor (one that includes C syntax highlighting is preferable). Look for the line **PHP_FUNCTION(confirm_tar_compiled);** and change it to **PHP_FUNCTION(tar_list);**. Now that we know our extension compiles into PHP properly, we can replace the debug function confirm_tar_compiled with our own fnuction, tar_list.

Now open up tar.c, and scan to line 42 or thereabouts. Change the line **PHP_FE(confirm_tar_compiled, NULL)** to **PHP_FE(tar_list, NULL)**. This merely mimics the change made in tar.h, and both are used to inform PHP of what functions this module offers.

Further down the file, you'll see the line containing this: **PHP_FUNCTION(confirm_tar_compiled)**. This is where the actual **confirm_tar_compiled** function lies. Change that line, and the two lines above it, to read:

```
/* {{{ proto array tar_list(string arg)
   Return an array of files in a given tar file */
PHP_FUNCTION(tar_list)
```

As you can see, PHP source code is self-documenting – a habit that's good to keep. Those changes above merely rename the **confirm_tar_compiled** function to **tar_list()**. As you can guess, each function needs an appropriate **PHP_FUNCTION()** line in tar.h, a **PHP_FE()** line near the top of tar.c, and also a **PHP_FUNCTION()** line somewhere in tar.c.

Save your changes, cd to the PHP source directory, and run **make** and **make install** again. **make** should take substantially less time now, because only part of the project has changed.

Once the new version of PHP is installed, next you need to execute this command:

```
php -r “echo tar_list(‘foo’), \”\n\”;”
```

That will call the **tar_list** function, passing in the parameter **foo**. As you can see by the output, the actual contents of the function remain the same as **confirm_tar_compiled** – let’s change that.

## The C implementation

Here’s the new code for tar_list: please remember that the emboldened number prefixes reproduced here are purely for reference purposes and *should not* be typed in.

```
1   PHP_FUNCTION(tar_list)
2   {
3       TAR *t;
4       char *tarfile = NULL;
5       int tarfile_len;
6       int i;
7       if (zend_parse_parameters(ZEND_NUM_ARGS()
TSRMLS_CC, “s”, &tarfile, &tarfile_len) == FAILURE) {
8           return;
9       }
10      if (array_init(return_value)==FAILURE) RETURN_FALSE;
11      if (tar_open(&t, tarfile, NULL, O_RDONLY, 0, TAR_GNU)
== -1)
12      {
13          php_error(E_WARNING, “%s”, strerror(errno));
14          RETURN_FALSE;
15      }
16      while ((i = th_read(t)) == 0)
17      {
18          add_next_index_string(return_value, th_get_
pathname(t), 1);
19          if (TH_ISREG(t) && tar_skip_regfile(t) != 0)
20          {
21              php_error(E_WARNING, “%s”, strerror(errno));
22              tar_close(t);
23              RETURN_FALSE;
24          }
25      }
26      if (tar_close(t) != 0)
27      {
28          php_error(E_WARNING, “%s”, strerror(errno));
29   RETURN_FALSE;
30      }
31   }
```

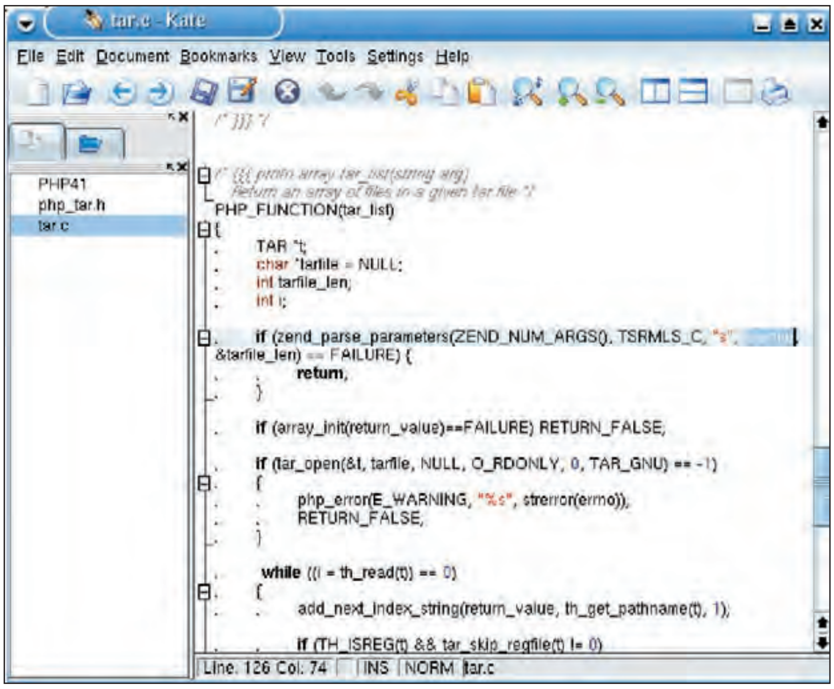

tar.c - Kate
File Edit Document Bookmarks View Tools Settings Help
PHP41
php_tar.h
tar.c
Line: 126 Col: 74 INS NORM tar.c

```
/* }}} */

/* {{{ proto array tar_list(string arg)
   Return an array of files in a given tar file */
PHP_FUNCTION(tar_list)
{
        TAR *t;
        char *tarfile = NULL;
        int tarfile_len;
        int i;

        if (zend_parse_parameters(ZEND_NUM_ARGS() TSRMLS_C, "s",
&tarfile_len) == FAILURE) {
                return;
        }

        if (array_init(return_value)==FAILURE) RETURN_FALSE;

        if (tar_open(&t, tarfile, NULL, O_RDONLY, 0, TAR_GNU) == -1)
        {
                php_error(E_WARNING, "%s", strerror(errno));
                RETURN_FALSE;
        }

         while ((i = th_read(t)) == 0)
        {
                add_next_index_string(return_value, th_get_pathname(t), 1);

                if (TH_ISREG(t) && tar_skip_regfile(t) != 0)
```

**Using a text editor that includes C syntax highlighting will make your life much easier.**

If the code seems confusing at first, relax – half the new code is PHP-related and half the code is tar-related; you’re learning two new things at once here!

## How it works

Lines 3 to 6 define our variables for this function. **t** is a pointer to a tar file, which is defined as type TAR (the built-in libtar definition of a tar file). **tarfile** is a pointer to the filename we’re going to load – this is passed in as our first parameter. Finally, we have two integers; **tarfile_len**, to store the length of the filename to load, and **i**, which is used later.

Lines 7 to 9 deal with parameters being passed into the function, and was inspired by Python (see? Competition is never bad! See page 80 for more on Python).

The **zend_parse_parameters()** function takes the parameters passed into our function, then places them into C variables we define, automatically performing type conversion where possible.

**zend_parse_parameters()** takes a variable number of parameters. The first parameter is nearly always **ZEND_NUM_ARGS() TSRMLS_CC**, which is a little confusing in itself. The first part, **ZEND_NUM_ARGS()** is quite straightforward – it returns the number of parameters actually passed in by the end-user to our function. **TSRMLS_CC**, however, is a little bit of magic – it’s a macro standing for “Thread Safe Resource Manager Local Storage Call with Comma”, and ensures thread safety for your extension. The reason it’s “with Comma” is because it’s actually passed in as parameter *two* to zend_parse_parameters, because it’s preceded by a comma. However, it’s easier just to treat **ZEND NUM_ARGS() TSRMLS_CC** as one parameter, and let *Zend* perform the magic.

The third and subsequent parameters are the important part of the function. Parameter three is a list of the types of »

## Modules made easy

### How to use ext_skel for maximum efficiency

**Extension writing for PHP needn’t be a hassle, and it’s important not to make your life harder than it should be. Using ext_skel is the best way to start an extension, as it creates lots of code for you to handle all sorts of eventualities – using global variables, reading php.ini values, adding entries to phpinfo() and more.**

**ext_skel can be passed parameters to make it produce a more customised extension skeleton, and these options are generally used by veteran extension writers.**

**Particularly of interest will be:**

**--xml** Generate XML documentation that can be added to the PHP documentation.

**--no-help** Don’t add comments throughout the files, and also don’t create the initial test function.

**--stubs=file** Leave out all module-specific inforation and just write function stubs.

parameters you expect to receive in mnemonic format – **l** is a long, **d** is a double, **b** is a boolean, **a** is an array, **r** is a resource, *etc.* **s** is special, because it stands for "string" and receives *two* parameters – a char* string of characters and also the length of the string as an integer.

So, our third, fourth, and fifth parameters are **s** (we want to receive a string), **&tarfile** (the location of our char* ready to receive the contents of the string parameter), and **&tarfile_len** (the location of our int ready to receive the length of the string parameter). If we were to receive a string and a boolean, we would use **sb**, then add a variable to store the boolean. If we request a string and the user passes an integer, **zend_parse_parameters()** will automatically convert the string to an integer before giving it to us, and it will perform similar conversions for most other types. Important exceptions to this are arrays, objects, and references, which cannot be converted because of their arbitrary nature.

Note that **zend_parse_parameters()** returns FAILURE on failure, and SUCCESS on success. In our example, we bail out with **return** if the function call fails. It's important to note that **zend_parse_parameters()** will automatically flag up warnings if the incorrect number of parameters are received or if a type conversion isn't possible.

Continuing on with the code, line 10 sets up the value we wish to return from the function call, which is an array. Returning values in PHP is remarkably simple on the whole, however returning arrays and objects is trickier due to their naturally complex nature. Returning values from functions is handled by a special zval called return_value. *zvals* are how the Zend Engine stores variables – see the box *Anatomy of a zval* below for detailed information. The simple definition of a zval is "a multi-type variable that handles refcounts".

When we use PHP_FUNCTION(somefunc), we're actually using the PHP_FUNCTION macro, which internally sets each function up with basic information. Particularly, it ensures that each function accepts several parameters, one of which is a zval called return_value. Also passed, just so you get a more complete picture, is an integer called return_value_used – if this is 1, the calling script needs a return value, and it's 0 if no return value is needed (*ie* function return values aren't assigned to a variable or used in another function).

So, line 10 sets up **return_value** to be an array through the **array_init()** macro, which sets the zval to be an array and allocates it some space. Lines 11 to 30 deal with the libtar part of the extension, with a little *Zend* code in there. As the focus of this tutorial is PHP, I'll have to keep libtar-specific information to a minimum – sorry! These lines of code can be split into three distinct sections: opening a tar file, reading from the tar, and closing the tar file. This functionality is split over lines 11 to 15, 16 to 25, and 19 to 30 respectively.

## Anatomy of a zval

### Declaration and definition

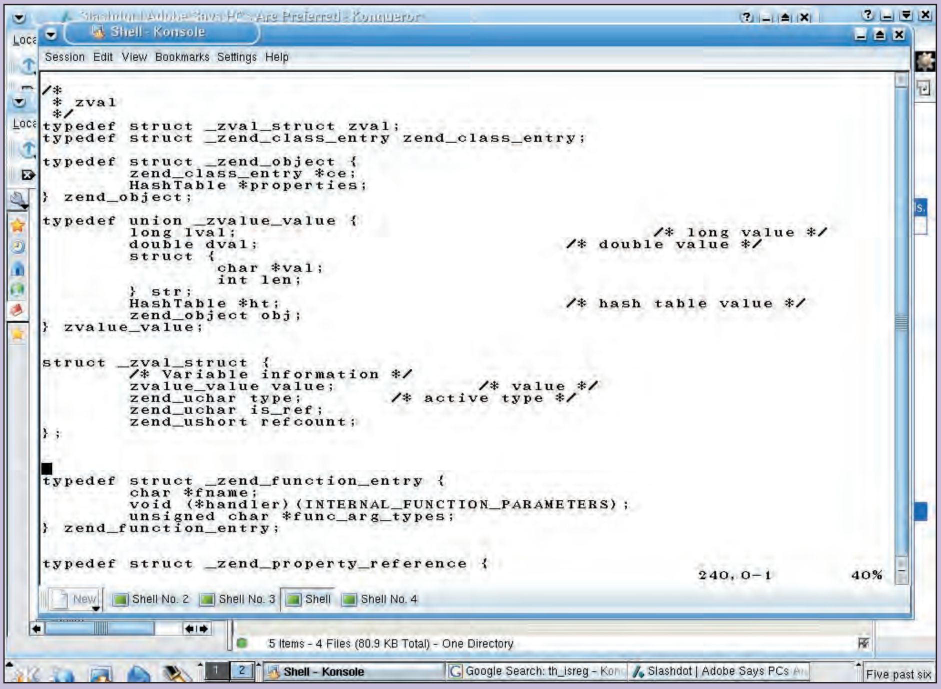

Slashdot | Adobe Says PCs Are Preferred - Konqueror
Shell - Konsole
Session Edit View Bookmarks Settings Help
New
Shell No. 2
Shell No. 3
Shell
Shell No. 4
5 Items - 4 Files (80.9 KB Total) - One Directory
1
2
Shell - Konsole
Google Search: th_isreg - Kon
Slashdot | Adobe Says PCs Ar
Five past six

```
/*
 * zval
 */
typedef struct _zval_struct zval;
typedef struct _zend_class_entry zend_class_entry;

typedef struct _zend_object {
        zend_class_entry *ce;
        HashTable *properties;
} zend_object;

typedef union _zvalue_value {
        long lval;                                      /* long value */
        double dval;                            /* double value */
        struct {
                char *val;
                int len;
        } str;
        HashTable *ht;                          /* hash table value */
        zend_object obj;
} zvalue_value;


struct _zval_struct {
        /* Variable information */
        zvalue_value value;             /* value */
        zend_uchar type;        /* active type */
        zend_uchar is_ref;
        zend_ushort refcount;
};


typedef struct _zend_function_entry {
        char *fname;
        void (*handler)(INTERNAL_FUNCTION_PARAMETERS);
        unsigned char *func_arg_types;
} zend_function_entry;

typedef struct _zend_property_reference {
                                                        240,0-1         40%
```

The PHP source code is, of course, the easiest way to see how things work. Well worth a read.

Each zval contains a value, a type, a reference count, and whether or not it's a reference to another variable or not. The reference count is used for garbage collection – the count is incremented for each reference that exists to this variable, and decremented when a reference is lost. When the count hits zero, the zval is automatically freed.

The value is a union structure in itself, so it's a little more complicated. A value can either be a long (lval), a double (dval), a string and its length (str), a hash table value (ht, for arrays), or an object (obj). How the value is read from a variable depends on what kind of variable it is, naturally. However, it's important to remember that reading one type of a union that had its data set as another type will result in garbage being returned.

To read a value from a zval that holds a string, use *Zend*'s built-in macros, Z_STRVAL(zval) and Z_STRLEN(zval). While you can access the value of a zval directly, it's much easier just to use the macro. In the same way, Z_LVAL(zval) reads the long value of the zval, Z_BVAL(zval) reads the boolean value of the zval, *etc.*

Writing to a zval can also be done by hand, but, because of the need to set the *type* field of zvals, it's easier again to use the Zend macros for setting variables: ZVAL_LONG, ZVAL_DOUBLE, ZVAL_STRING, etc. In the case of setting a string variable, the *Zend* macro also calculates the length of the string for you.

ZVAL_LONG, ZVAL_DOUBLE, etc, all use the same format for setting their value – parameter one is the zval to set, and parameter two is the value to assign to it, *eg*:

```
ZVAL_LONG(new_long, 10);
```

Owing to the more memory-reliant nature of strings, the ZVAL_STRING macro takes a third parameter, which should be a 1 if the string should be duplicated before being stored, or a 0 if not. Generally this should be a 1, otherwise data may go out of scope and cause problems. The primary occasion where a 0 needs to be passed is when you want to create a new variable referring to a string that's already allocated in *Zend*'s internal memory.

There's one final thing you need to know about zvals before getting started using them, and that's how they are initialised. MAKE_STD_ZVAL is a special macro that should be called after each zval has been declared, and does basic house-keeping tasks such as allocating memory for the zval, setting the reference count to one, *etc.*

So, to conclude, a complete zval declaration and definition would be:

```
zval *myint;
MAKE_STD_ZVAL(myint);
ZVAL_LONG(myint, 10);
```

**php_error()** is a macro pointing the **zend_error()** function, which automatically outputs various errors to users. It takes two parameters – the kind of error to issue, and the text of the error. Internally, **zend_error()** sets up extra information such as the line number where the error occurred, the file in which it occurred, etc.

Generally speaking, only three error types should be issued: E_ERROR, E_WARNING, and E_NOTICE. The difference between the three is that E_ERROR halts execution of the script, and E_NOTICE is generally disabled as it's relatively minor. E_WARNING should be the most commonly used error type.

RETURN_FALSE is one of the many ways to return values to users, and simply returns the value "false". Other basic return types include RETURN_BOOL, RETURN_NULL, and RETURN_STRING. As you can see, sending non-array/object values back to users is remarkably easy!

Line 11 calls the libtar function **tar_open()**, which takes six parameters: the address to save the tar file handle to, the pathname to the tar file to open, NULL, how the file should be opened (O_RDONLY for readonly or O_WRONLY for write only), 0, then special flags that can be ORed together.

Parameter one needs to be of type TAR* (see line 3), and the special flags can consist of one or more of TAR_GNU (enables GNU extensions to the tar format), TAR_VERBOSE (send status messages to stdout), TAR_IGNORE_CRC (skip validation of the tar CRC), and others. **tar_open** will return 0 on success, or -1 on failure.

So, lines 11 to 15 translate as "if opening the tar file as read-only fails, flag up a PHP warning, then exit the function". If the function succeeds, our tar file will be ready to work with in the variable **t**.

Moving on, lines 16 to 25 contain the code to handle reading files from the tar file, and adding the name of each file to the array we intend to return. **th_read()** is a function that reads one file header block from the TAR variable passed as parameter one. A file header block describes each file inside a tar file, and there is one header block for each file. **th_read()** returns -1 on error, 0 on success, and 1 when it reaches the end of the tar file, so we use a loop to continue iterating through file header blocks until a value other than 0 is returned.

**add_next_index_string()** is a function that handles adding strings to an array in an ordered fashion. The **next_index** part of the function name means that the string is added in the next available numerical slot in the array, so, as we're using this function exclusively to add array elements, our array will contain elements numbered from 0 to n-1, where n is the number of files in the tar file.

The function takes three parameters – the array to add to, the char* to add, and a 1 or a 0 depending on whether the string needs to be copied into *Zend*'s memory (1) or whether the *Zend Engine* can use the existing pointer (0). 1 is generally used, owing to 0's tendency towards instability. **th_get_pathname()** extracts the pathname of the current tar file header, and returns it in char* format – perfect for parameter two.

Line 19 calls **tar_skip_regfile()** to move onto the next file in the tar archive – this "just works" and is best left alone. However, there is a little bit of error handling in there – if libtar is unable to move onto the next file in the archive, it will return -1, and hence run our standard error reporting code. Note, though, that this time we close the tar file with **tar_close()** (see later) so that we clean up properly.

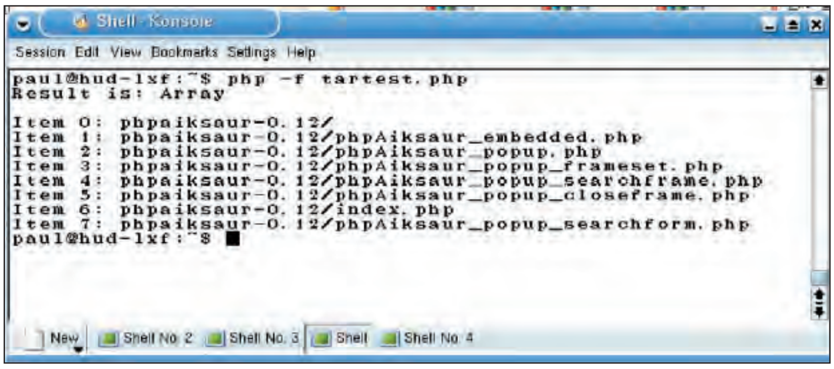

```
paul@hud-lxf:~$ php -f tartest.php
Result is: Array

Item 0:  phpaiksaur-0.12/
Item 1:  phpaiksaur-0.12/phpAiksaur_embedded.php
Item 2:  phpaiksaur-0.12/phpAiksaur_popup.php
Item 3:  phpaiksaur-0.12/phpAiksaur_popup_frameset.php
Item 4:  phpaiksaur-0.12/phpAiksaur_popup_searchframe.php
Item 5:  phpaiksaur-0.12/phpAiksaur_popup_closeframe.php
Item 6:  phpaiksaur-0.12/index.php
Item 7:  phpaiksaur-0.12/phpAiksaur_popup_searchform.php
paul@hud-lxf:~$ █
```

**It works! All that sweat and RSI was worth it!**

We're onto the last chunk of code now, at last. **tar_close()** takes a TAR file parameter, and closes it down and frees any associated memory. If libtar fails to close the tar file for some reason, it will return -1, which will thereby run our standard error handling code.

## Phew!

If you've made it this far, you're very brave indeed! Hopefully I've gone into enough depth to give you a full understanding of just *how* things work, but also *why* they work that way. Anyway, the hard work is now done – let's make use of the new extension!

The first step is to rebuild PHP by executing make from the PHP source directory – this will only rebuild the new extension, so it shouldn't take long.

Once it's built, run make install as root, then we're ready to test. Here's an example PHP script that makes use of the new function. You'll need to find or create a tar file for testing purposes – make sure it's just a tar file (.tar) and not a gzipped tar file (.tar.gz or .tgz).

```
<?php
  $result = tar_list('/path/to/your.tar');

  echo "Result is: $result", "\n\n";

  while(list($var, $val) = each($result)) {
    echo "Item $var: $val\n";
  }
?>
```

Save that as tartest.php, then run:

```
php -f tartest.php
```

If you've followed my steps precisely, you should see something similar the screenshot at the top of this column, although, of course, the files inside the tar will be different.

## Conclusion

Scanning over this article and part 1 last month, you should be able to see that there's actually not all that much code involved – for instance, editing the m4 file is something that you just get used to and is mostly the same for every extension you write. Furthermore, learning how to make use buildconf, configure, and make for PHP extensions is really just a one-off – it'll become second nature given practice.

This tutorial has hopefully taught you how PHP extensions work, how their code is laid out, how *Zend* works with variables, and also perhaps introduced just a little of how PHP and *Zend* work internally. We're not finished yet though – there's more to come in the next issue! LXF

## NEXT MONTH

**We'll be putting the finishing touches to the libtar topics covered here and polishing up your extension-writing skills in general. Comments, ideas or suggestions for the PHP series to paul.hudson@futurenet.co.uk please!**

MODULES AND PACKAGES

# Organising Python source code

**Patrick O'Brien** shows how big things can come in little Python packages.

Last month we let our love flow, focusing on the action in Python programs. In particular, we saw how to control the flow of our Python programs: how to make decisions, how to repeat an action a number of times, and how to iterate over the contents of a container, making use of each item in that container. And to see how these concepts related to the real world, I showed you how all these techniques have been used in some recent developments involving the *wxPython* GUI toolkit.

Until now, we've had a bit of fun with our "love" theme. That's about to change. Don't get me wrong – I still want you to love Python as much as I do. But a healthy relationship has to have its serious side as well, and it's about time we got serious. So this month we're going to look at ways to organise our affairs, using Python's packaging mechanisms, such as modules, packages, importing, the module search path, and module reloading.

Even though Python is an agile language, it's still important for you to be organised in your approach to programming with Python. Modules and packages are the fundamental organisational units in Python, and this tutorial will show you how to organise your Python code to make it modular, reusable, and easier to maintain. And, as promised last month, I'll bring you up to date on some changes to my *PyCrust* project, including a new addition to the "Py" family – those wonderful programs with whimsical names.

## Python source code

Python source code files are ordinary text files. They just happen to contain code that Python understands. That means you can create and edit them using the text editor you prefer. Of course, many programmers like to use an editor that recognises Python syntax and can do things like colourise the code to make it more readable. Most of the popular editors have support for Python code, so use the editor you prefer (I happen to like *Emacs*). If you have no strong feelings for any particular editor, you should try the IDLE development environment that's included with Python, which includes a source code editor, interactive shell, debugger, and other tools.

## Modules

Python source code files are called modules. A typical Python program is made up of several modules, each containing code for a particular purpose. Some modules are designed to be used by one or more programs (by being imported), others are designed to be executed from the command line. The same source code file can serve both purposes by checking its __**name**__ attribute and responding accordingly. We've seen this in previous tutorials, but let's review it again.

Modules are objects, just like all the other Python data types. That means modules have attributes. One of those module attributes is a __**name**__ attribute, that is automatically assigned to a module by the Python interpreter.

When a module is imported, its __**name**__ attribute is the same as its filename (without the extension). You can see this for yourself in this example from an interactive shell session where we import the **sys** module (which is actually a special, built-in module for which there is no **sys.py** file on disk):

```
>>> import sys
>>> sys.__name_ _
'sys'
>>>
```

When a module is being executed, rather than imported, Python overrides the default naming convention and sets the module's __**name**__ attribute to the string value of **'__main__'**, giving us the chance to behave differently in this situation. A typical approach is to call a function defined earlier in the module. While not a requirement, it is quite common to name that function `main()`, as I have done here:

```
def main():
    # Do whatever we want when executed...
    print "I've been executed, not imported."

if __name__ == '__main__':
    main()
```

## Making a Python program executable

If the first line in a Python source code file is a "shebang" line, the file is capable of being executed directly from the command line. What's a "shebang?" It looks like a Python comment, which it is, but on some operating systems, such as Linux and Unix, the "shebang" tells the operating system how to find the interpreter that will execute the program file. For Python, a typical "shebang" looks like this:

```
#!/usr/bin/env python
```

If a program file containing a "shebang" is then given executable privileges (using the **chmod** command, for example), you can run the program from the operating system command line by simply supplying the program name:

```
% ./example.py
```

In fact, we could even drop the .py filename extension, making our Python program practically indistinguishable from programs written in any other language. In short, the "shebang" line is a convenience for Linux folks that's simply ignored on other platforms.

## Module documentation

We mentioned earlier that modules are objects, and that they have attributes. One of those attributes is a documentation string (docstring). When the first statement in a module is a string, that string becomes the docstring for the module and is stored in the module's __**doc**__ attribute. It's good practice to include a docstring in all of your modules.

Here's an example module (named 'example.py') that illustrates the use of a "shebang", a module docstring, and several of the other concepts we've just discussed:

```
#!/usr/bin/env python

```

```
“””This is a module docstring.”””

import sys

def main():
    “””Main function for example.py.”””
    # Do whatever we want when executed...
    print “I’ve been executed, not imported.”
    print “This is the sys.path:”
    print sys.path

if __name__ == ‘__main__’:
    main()
```

You can access module docstrings in your code, in some development environments, and even in the Python interpreter running in interactive mode:

```
>>> import example
>>> print example.__doc__
This is a module docstring.
>>>
```

Docstrings are but one example of Python’s powerful introspection capabilities, and I encourage you to provide them for modules, classes, methods, functions, and any other place where Python supports a docstring. When you use tools like *PyCrust*, you’ll see just how handy your docstrings become.

## Python standard library

Python provides an extensive library of modules that you can use in your own programs. These modules have support for everything from networking to threads to regular expressions to random numbers and XML. There are literally over 150 source modules in the standard library, and more get added with each release of Python.

In fact, an entire book has been published about it: the *Python Standard Library* written by Fredrik Lundh, is a handy reference that includes example code illustrating how to use all the modules that come with Python. Beyond saving time and effort, the modules in the standard library are good examples of Python programming techniques. So in addition to importing and using them, you should also review and learn from their source code. (When you import a module in *PyCrust* and select the module in the filling tree, the module’s source code is displayed in the filling text window.)

## Importing modules

The way to make use of one module from another is to import it, using the **import** keyword. There are a variety of ways to import a module, or import something contained in a module, but the simplest way is to import the module, then access the attributes of the module by using the module name as a prefix. Note that when we import a module, we refer to the module by its filename, without the .py suffix.

Later we’ll look at some of the nuances of importing, but for now let’s look at the basics, importing some modules out of the standard library into a Python shell session. Earlier we saw an example using the **sys** module, which provides details about the environment that have to do with the working of the Python interpreter. We’ll use the **sys** module again in the following examples:

```
>>> import sys
>>> sys.executable
‘/usr//bin/python’
>>> sys.maxint
2147483647
>>> sys.version
‘2.2.2 (#2, Feb  5 2003, 10:40:08) \n[GCC 3.2.1 (Mandrake
Linux 9.1 3.2.1-5mdk)]’
>>>
```

In this session, we imported the **sys** module and referenced three of its attributes: **sys.executable** is a string identifying the path to the Python interpreter, **sys.maxint** is an integer identifying the largest integer value supported on the current platform, and **sys.version** is a string containing information about the version of Python that is currently running.

Sometimes we only need one or two items from a module, and would rather not have to use a module prefix each time we referenced those items. For that, we can use the following import variation, using the **from** keyword:

```
>>> from sys import maxint
>>> maxint
2147483647
>>>
```

This second form of import may look like it would lead to more concise code, and it can. However, importing a module and then using the module name as a prefix has the benefit of making it clear where the objects in your program are coming from. So keep that in mind as you choose names for your modules. The rule of thumb is to use short, descriptive, lowercase names for modules.

## Module search path

Now that we’ve seen how to import modules, you may be wondering how Python knows where to find the modules that we are telling it to import. And, you may also be wondering how to get Python to import modules that you create. The answer to both involves the module search path.

The module search path is a list of directories that Python searches for modules. It is stored in the **sys.path** attribute:

```
>>> sys.path
[‘’,
‘/home/pobrien/Code’,
‘/usr/lib/python2.2’,
```

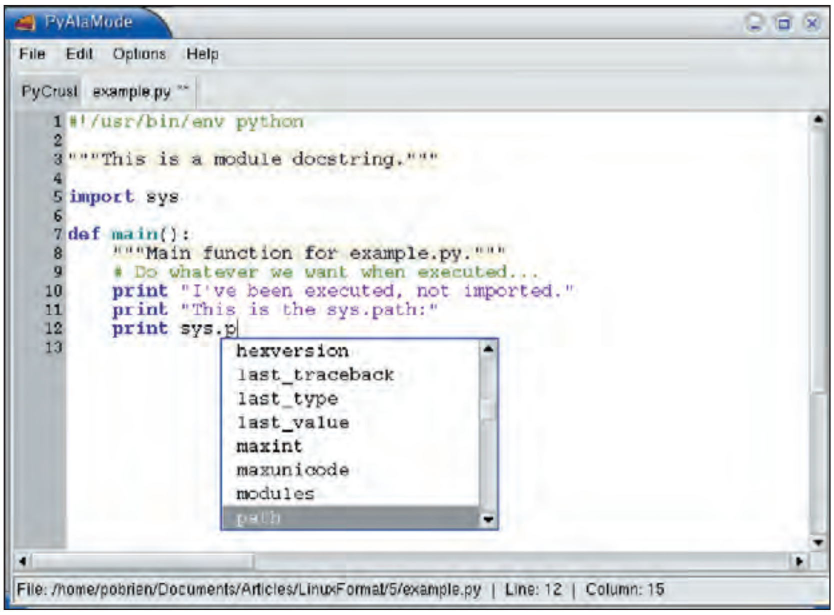

```
#!/usr/bin/env python

"""This is a module docstring."""

import sys

def main():
    """Main function for example.py."""
    # Do whatever we want when executed...
    print "I've been executed, not imported."
    print "This is the sys.path:"
    print sys.p
```

File: /home/pobrien/Documents/Articles/LinuxFormat/5/example.py | Line: 12 | Column: 15

**Editing example.py using *PyAlaMode* with an autocompletion list active.** »

```
'/usr/lib/python2.2/plat-linux-i386',
'/usr/lib/python2.2/lib-tk',
'/usr/lib/python2.2/lib-dynload',
'/usr/lib/python2.2/site-packages',
'/usr/lib/python2.2/site-packages/Numeric',
'/usr/lib/python2.2/site-packages/PIL',
'/usr/lib/python2.2/site-packages/gtk-1.2']
>>>
```

The first directory in the list is the directory of the program that was executed, or, in the case of an interactive session (or a program like *PyCrust* that mimics the regular Python shell), the first item will be a blank string, which tells Python to search in the current directory first.

The next few items in the list are determined by the **PYTHONPATH** environment variable. In this case, I have one directory assigned to the **PYTHONPATH** in my .bashrc file:

```
PYTHONPATH=~/Code
export PYTHONPATH
```

This path, which expands to **/home/pobrien/Code**, appears as the second item in **sys.path**. The subsequent items are automatically added by Python upon startup. First Python adds the path to its standard library, and a few critical directories, which varies by platform. Then it adds any directories identified in **.pth** files that it finds in its home directory (typically **site-packages**). These **.pth** files are usually installed by Python add-ons, such as the GTK, *Numeric*, and *Python Imaging Library* that I have installed on my machine.

One way to dynamically control Python's search for modules is to modify the sys.path attribute. It's an ordinary Python list, so any of the list operations (appending, inserting, deleting) you learned about in one of our previous tutorials can be applied to the sys.path.

For example, I like to put unit tests into a "tests" directory directly below the modules to be tested. My unit test modules then include code that adds their parent directory to the front of sys.path, guaranteeing a successful import of the module to be tested. Here is an example from a module that tests my **introspect** module:

```
import unittest

# Import from this module's parent directory.
import os
import sys
sys.path.insert(0, os.pardir)
import introspect
del sys.path[0]
del sys
del os
```

The os module used above is also from the standard library, and provides a platform-independent way of specifying things, such as the current directory, parent directory, path separator, line separator, etc. On Linux, these values appear as follows:

```
>>> os.curdir
'.'
>>> os.pardir
'..'
>>> os.sep
'/'
>>> os.linesep
'\n'
>>>
```

When Python is running on Windows or the Mac, those same attribute of the **os** module would have different values, appropriate to those operating systems.

## Packages

Python modules are often organised into packages, which are little more than specially designated directories. You turn a directory into a Python "package" by placing a specially named module (**__init__.py**) into that directory. Packages can exist within other packages, forming a hierarchical structure.

In order to import a module contained in a package, you simply use an import statement with a dot notation, specifying the package hierarchy to traverse. For example, imagine we had the following directory structure:

```
home/pobrien/Code/example/
    __init__.py
    one/
        __init__.py
        foo.py
    two/
        __init__.py
        bar.py
```

In this situation, the following import statements are valid:

```
import example.one.foo
import example.two.bar
```

However, those forms are rather cumbersome, as attributes of **foo** and **bar** would have to be referenced as follows:

```
example.one.foo.something
example.two.bar.whatever
```

So the more common approach is to import a module from its location in a package hierarchy, as in this example:

```
from example.one import foo
from example.two import bar

foo.something
bar.whatever
```

One nice benefit of packages is that it reduces the need to have lots of directories in your module search path (sys.path). As you can see in the previous examples, I only have one directory (home/pobrien/Code) of my own specified in my PYTHONPATH environment variable. Into that directory I place all my Python projects. If the directory contains modules that I need to be able to import, I simply make sure the directory is also a Python package.

## Importing revisited

Now that we've discussed the syntax for importing and referencing modules and packages, we need to talk about what actually happens when a module or package is imported by Python. If this is the first time a module is being imported by any module in the current process, Python will execute all the code in the body of the module. If there were no errors in the code that kept it from importing, the module is bound to the name specified in the import statement, as we saw in previous examples.

Python also keeps track of all imported modules internally. If the same module is imported elsewhere in the program, Python does not create a new module object, nor does it execute the code in the module again. Instead, it simply returns the existing module object that it has stored internally. What that means is that no matter how many times a module is imported, all imports but the first execute almost instantaneously. Like much of Python, the structures used internally by the Python interpreter are not hidden from view, or from manipulation by you. In the case of modules, all currently imported modules are kept track of in **sys.modules**, which is an ordinary Python dictionary. You can see exactly what modules have been imported by examining the contents of **sys.modules**:

```
>>> import sys
>>> sys.modules

{'stat': <module 'stat' from '/usr/lib/python2.2/stat.pyc'>,
'__future__': <module '__future__' from
'/usr/lib/python2.2/__future__.pyc'>, 'copy_reg': <module
'copy_reg' from '/usr/lib/python2.2/copy_reg.pyc'>, 'posixpath':
<module 'posixpath' from '/usr/lib/python2.2/posixpath.pyc'>,
'UserDict': <module 'UserDict' from
'/usr/lib/python2.2/UserDict.pyc'>, 'signal': <module 'signal'
(built-in)>, 'site': <module 'site' from
'/usr/lib/python2.2/site.pyc'>, '__builtin__': <module
'__builtin__' (built-in)>, 'sys': <module 'sys' (built-in)>, 'posix':
<module 'posix' (built-in)>, 'types': <module 'types' from
'/usr/lib/python2.2/types.pyc'>, '__main__': <module
'__main__' (built-in)>, 'exceptions': <module 'exceptions'
(built-in)>, 'strop': <module 'strop' from '/usr/lib/python2.2/lib-
dynload/strop.so'>, 'os': <module 'os' from
'/usr/lib/python2.2/os.pyc'>, 'readline': <module 'readline' from
'/usr/lib/python2.2/lib-dynload/readline.so'>, 'os.path': <module
'posixpath' from '/usr/lib/python....

>>>
```

## Reloading a module

There are times when you'd like to get the latest version of a module that you've already imported, for example, a module whose code you are revising and then interacting with from a Python shell. You can't simply import it a second time, because we've seen that Python's default behaviour is to use the module as it was the first time it was imported. So we need a way to force Python to import the module anew. The way to do that is using the built-in **reload()** function:

```
>>> reload(example)
<module 'example' from 'example.pyc'>
>>> example.main.__doc__
'Main function for example.py.'
>>>
```

## Import this!

The last point about importing is rather whimsical. It began as a slogan for one of the Python conferences, was added to the Python library as a prank, and remains as a reminder of what it means to be Pythonic. I'm talking about the famous "import this" statement. Try it from the Python shell and you'll see a display of a poem by Tim Peters, a well-known Python Zen master:

```
>>> import this
```

**THE ZEN OF PYTHON, BY TIM PETERS**

Beautiful is better than ugly.
Explicit is better than implicit.
Simple is better than complex.
Complex is better than complicated.
Flat is better than nested.
Sparse is better than dense.
Readability counts.
Special cases aren't special enough to break the rules.
Although practicality beats purity.
Errors should never pass silently.
Unless explicitly silenced.
In the face of ambiguity, refuse the temptation to guess.
There should be one – and preferably only one – obvious way to do it.
Although that way may not be obvious at first unless you're Dutch.
Now is better than never.
Although never is often better than **right** now.
If the implementation is hard to explain, it's a bad idea.
If the implementation is easy to explain, it may be a good idea.
Namespaces are one honking great idea – let's do more of those!

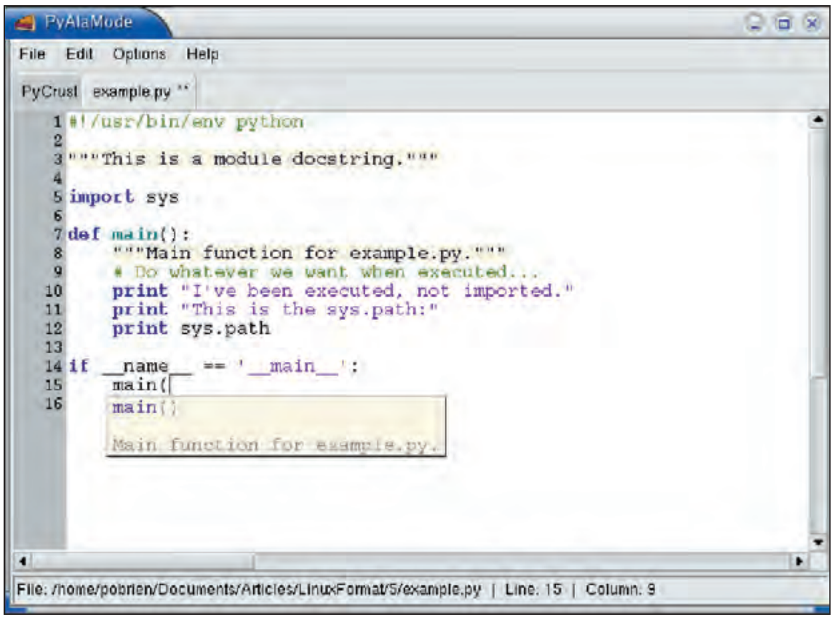

```
#!/usr/bin/env python

"""This is a module docstring."""

import sys

def main():
    """Main function for example.py."""
    # Do whatever we want when executed...
    print "I've been executed, not imported."
    print "This is the sys.path:"
    print sys.path

if __name__ == '__main__':
    main(
    main()

    Main function for example.py.
```

File: /home/pobrien/Documents/Articles/LinuxFormat/5/example.py | Line: 15 | Column: 9

**Editing example.py using *PyAlaMode* with a calltip active.**

## PyCrust is dead! Long live Py!

Yes, the *PyCrust* package is dead. *PyCrust* began as a graphical, interactive shell. And for a long time I kept that single focus. But every so often I long to have some of *PyCrust*'s features available as I'm editing Python source code. So I decided to build on the *PyCrust* foundation and create a program with source code editing capabilities. The result of those efforts was called the *PyAlaMode* editor.

Suddenly, the *PyCrust* package name felt confining and too limited for this new program. Around the same time that I began working on *PyAlaMode*, the *wxPython* folks asked me if I'd be willing to move all my *wxPython* projects into the *wxPython* CVS repository, to make it easier to distribute them with *wxPython*. So I decided to make the move and change package names at the same time.

Long live *Py*! Due to the great tradition of pie puns that I began with *PyCrust*, I was unable to think of a better package name than "py". So now all the wonderful *Py* programs (*PyAlaCarte, PyAlaMode, PyCrust, PyShell,* and *PyWrap*) are available in the **py** package, a subdirectory under the new **wx** *wxPython* package. Most users of *Py* programs will simply run one of the executable scripts provided by the *wxPython* installer, such as **pycrust** or **pyalamode** But developers who embed any of the *Py* (formerly *PyCrust*) modules into their own programs will now be able to use the following import statement:

```
from wx import py
```

Short and sweet, just the way I like it. LXF

ILLUSTRATION BY SCOTT WILKINSON – WWW.BLENDER3D.ORG

3D MODELLING

# Blender lighting

**Being able to model a simple scene is all very well, says Jono Bacon, but to give your 3D art the impact that it deserves, you need to know how to light it properly.**

Last month we introduced the basic primitive shapes in *Blender* and the various methods of placing them in our scene. This month we begin taking a look at some basic lighting. To start with, let's take a look at some of the key concepts in modelling shapes, the first being the Extrude key. The concept of an extrude is that you essentially stretch part of a shape out further.

First, start with an empty *Blender* project, and then add a cube using the toolbox (**Spacebar**). When you have added the cube, split the windows so you can see the Camera, Top View and Side View. Your interface should look something like that of **Fig 1**. Next, in the side view, use the zoom button to zoom into the cube a little further to see things better. Now press **Tab** to enter edit mode, and you will see the different vertices that form the shape; they will be pink. You can now use the Bounding Box key (the **B** key) to draw a box around the top vertices in the side view; they will be yellow now.

Fig 1: ***Blender* with a variety of views.**

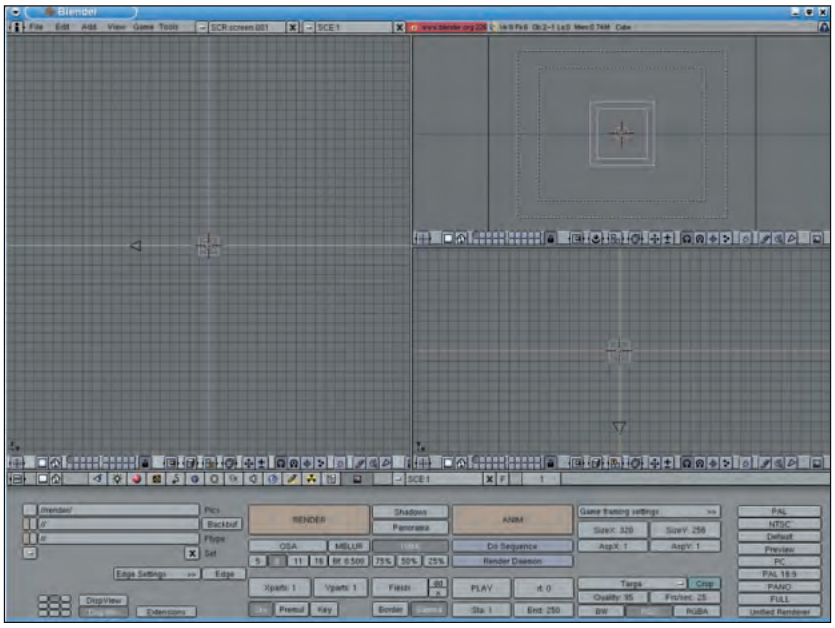

To extrude these vertices, hit the **E** key and click on Extrude in the small popup window that appears. When you move the mouse now, the vertices will stretch, maintaining the correct aspect ratio. Move the vertices up 6 squares and then click the left mouse button to place them – make sure that the lines are straight to ensure you have a square extrude. You will now see how it looks like we have a longer block sitting on our cube. With our vertices still selected, we can now deform this shape with the scale key.

Now, move your mouse over the top view and press the **S** key; this takes you into Scale mode and the vertices will be resize; notice how the shape is resized when looking at each view. Make the size of vertices a little smaller than the base and left-click.

We can now make use of the Rotate key to continue deforming this shape. Press the **R** key with the vertices selected (which they should already be, and move the mouse to rotate the vertices. Rotate them slightly and then finish by clicking the left mouse button. To finish this odd looking ornament we will extrude one more time, rotate again and then scale to a point. Your final shape should look rather like that shown in **Fig 2**.

## Setting the scene

We will now take our unusual structure and set it in a scene. First, add a plane in the top view to act as our ground that the statue will stand on. When you have added it, resize it to make it bigger (**S** key) and then move it in the side view to the bottom of the statue.

Next, move the camera by first leaving edit mode, right-clicking it and moving it with the **G** key. Move it over and use the

Rotate key to move the camera so it can see the entire statue (this will mean moving the camera both upwards and left). Your screen should look something similar to that of **Fig 3**.

Now use the toolbox to add a number of other shapes to the scene. I have added a Cone, a Cylinder and an ICOSphere.

## Adding Light

Currently our scene is in total darkness, and although we discussed some rudimentary lighting in the first part of this series, we will recap and look into lighting further.

Lighting is an essential part of *Blender* and general 3D graphics, and in many ways the lighting is one of the key elements to creating atmosphere in a scene. Adding a light is a simple procedure by simply selecting Add>Lamp from the toolbox. You will then see the light added to the scene and dotted lines showing the direction of the light. You can then hit **F12** to see the rendered scene with the lighting.

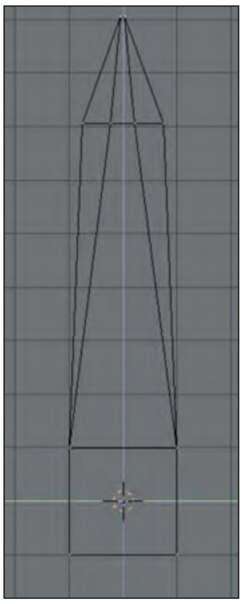

FIg 2: **Bizarre shape using a number of techniques.**

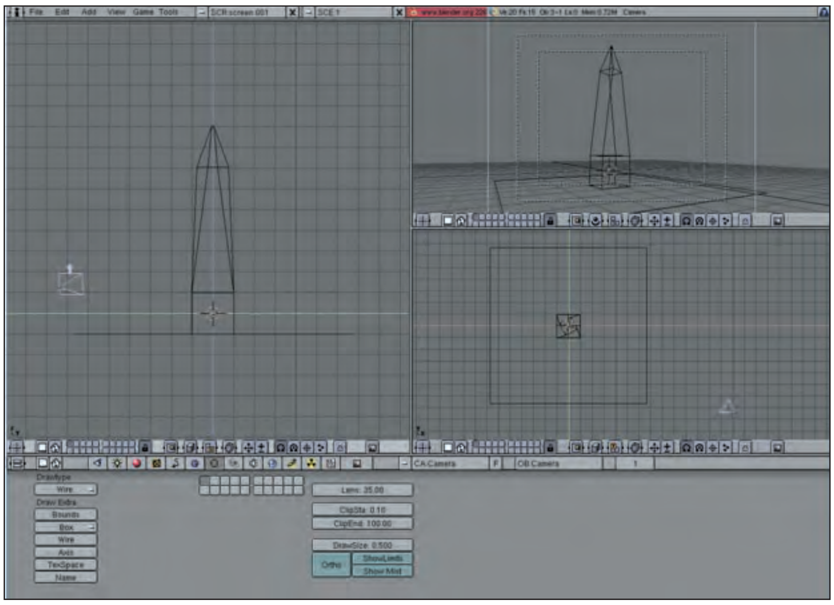

FIg 3: **Our scene develops.**

Though our simple light illuminates the scene, it would be wise for us to delve into lighting in more detail. *Blender* has four basic lights. These are listed below with examples.

»

## Blender tips

### Set Smooth and Lighting types

**Remember that to make your models look smooth you need to use the Set Smooth within the edit buttons. Right-click on an object when not in edit mode and click the Set Smooth button. The objects will then look much smoother. We will study smoothing more in-depth in a later issue.**

**Within *Blender* there are four types of lighting. Here is a typical scene lit by each type of light as an example.**

**1 Lamp**
**A central source of light that emminates in all directions.**

**2 Sun**
**A general light source... similar to the real sun.**

**3 Hemi**
**General light source cast from a half sphere source.**

**4 Spot**
**A spotlight that acts like a normal spotlight pushing light outwards from a central source. This light also casts shadows, which the others don't.**

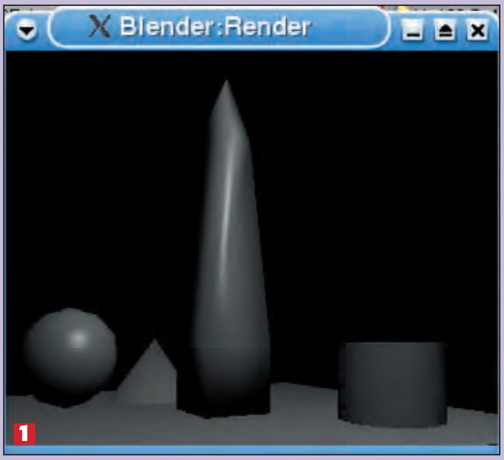

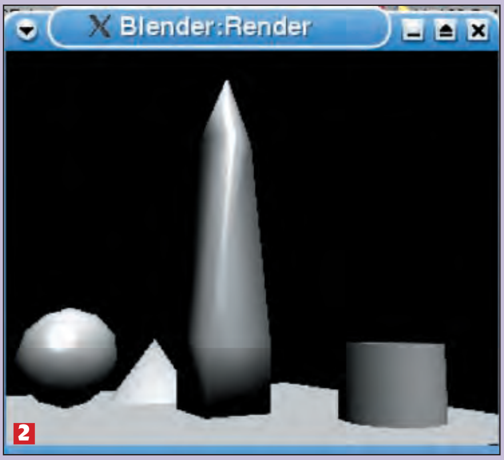

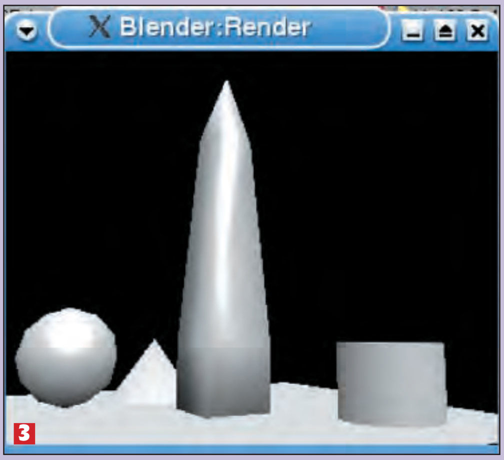

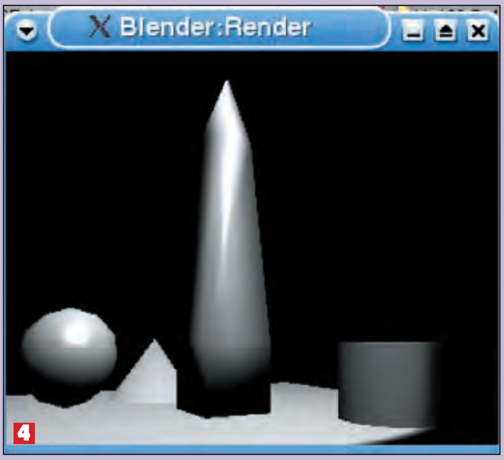

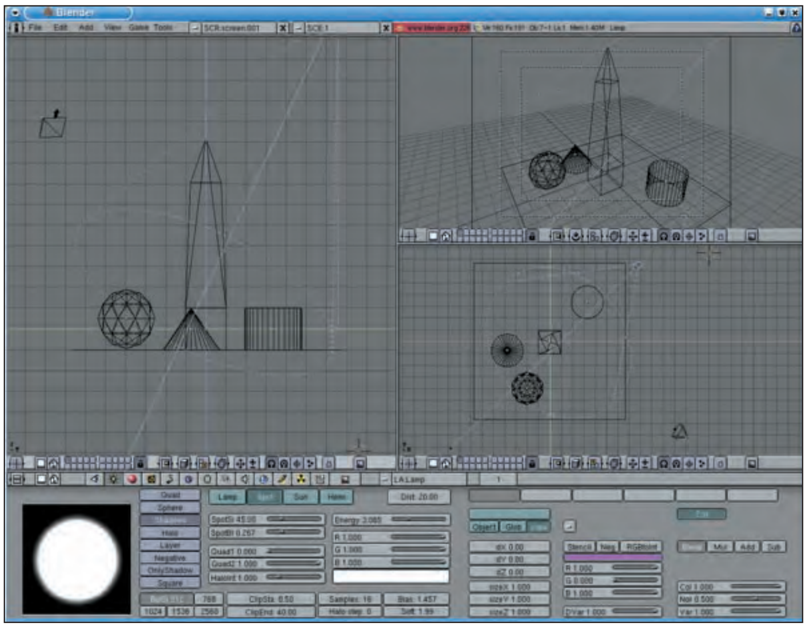

Fig 4: **Lighting a scene with a spotlight.**

One of the most common types of light in *Blender* work is the Spot; mostly because this light can cast shadows, and shadows are an important part of making a scene look authentic.

First, right-click the light to select it and click on the Spot button in the light buttons section (hit **F4** to show the buttons). You will then see the spotlight source and the light projection in each of the views. You can now make use of the normal Blender controls for scaling, moving and rotating objects to position the light so it shows at an angle on the scene (it is always to remember that to cast shadows, light must be from an angle.

When you have positioned the light we can begin to configure how our light shows on our scene. The first item to check is in the Display Buttons page (**F10**). Ensure that the Shadows button is pushed in; this will allow shadows in your render. Now hit **F12** and do a test render. You should notice some shadows on the scene protruding. If you press **F4** and go back to the Lighting Buttons page we can look at some of the options we can set for the spotlight. If you are looking for a basic rundown of the main buttons in Blender, please refer to the first tutorial in this series in issue 40 of *Linux Format*.

The first button to look at is the Energy button. This adjusts the amount of light that the spotlight emits. The slider allows you to adjust the light, and it is recommended to play around with the setting and do some test renders. The next setting to look at are the R, G and B sliders. These sliders adjust the amount of Red, Green and/or Blue in the spotlight. These settings can cause interesting effects with a number of different coloured lights.

At the bottom of the Lighting Buttons are some buttons to control the quality of the shadows that the objects are causing. The first button is the BufSi group of buttons. These buttons allow you to set the size of the shadow buffer; different buttons cause different render quality, and it is recommended to try each button with your scene. The Bias button is another interesting button and caused a different level of shadow on the objects. The Soft button can also be used to control the size of the shadow area. The final button of interest in configuring shadows is the Samples button; this controls the amount of shadow map samples being used and affects the quality of the shadow.

Two important buttons that should be made clear are the SpotSi and Spotbi sliders. These sliders are used to adjust the size of the spot projection (SpotSi) and the softness of the light (SpotBi). These sliders can be used in conjunction with multiple spotlights for creating better-lit scenes.

An important thing to point out at this part of the series is how important experimenting with *Blender* is. I have just discussed a number of buttons and what they do. Many of these buttons have terse technical descriptions, and most people want to be able to see the practical effect of each button. The only way to achieve this is to create a test seen (such as the one we have now) and fiddle with each button and play with different combination of buttons. It is this kind of experience that really pays dividends in not only understanding *Blender*, but also in understanding the fundamental concepts of light that *Blender* is simulating in its 3D space.

## Multiple Lights

Although we have put one light on our scene, we will now add more lights to give our scene more definition. A good method of trying to visualise light is when you are on a road driving at night, or being driven. Look at the streetlights and how the light cascades down on the objects. If you can visualise the light you can better represent it in your *Blender* scenes.

To get started, first set the size of your current spot to quite small, so the projection on the scene is fairly small. You may wish to change the Energy setting for the spot also. Next, add another Spot and set a similar size of projection. You can now manoeuvre the new light into another position, preferably so the projection overlaps each other. Now do a test render and you will see the two light projections and how they overlap.

Once you have added a number of lights, use the R, G, B buttons for each light to change the colour of each light. With these settings you can make quite subtle or blatant effects of light; an example being a live situation where you typically have a lot of Red, Green and Blue lights and their combination colours. Again, experimentation is the key.

## Conclusion

We have looked at a number of simple tools such as Extrudes, Scaling and Rotation. These tools when used together can be used in combination with our lighting technique to create already good looking scenes, even before we have got into the 'meat' of *Blender*. These simple techniques will be built on the coming moths to prepare you with a comprehensive set of skills.

If you have any comments. or requests about this *Blender* series, please email them to **linuxformat@futurenet.co.uk**.

## Blender Resources

**Get assistance from the community**

***Blender* 3D www.blender3d.org/**
**A showcase site for *Blender*. This site also acts as a centralised and official point for *Blender* downloads and documentation.**

***Blender* Community www.blender.org/**
**A focal point for main *Blender* development. Houses a number of third-party projects such as the network renderer, different plug-ins. If you want to know the latest about *Blender*, this is the site.**

**Elysiun www.elysiun.com/**
**A computer art site with a heavy community of *Blender* users. The main benefit of this site is its popular forum and its wide range of tutorials and documentation on using *Blender*.**

## How to make:

### A mountain scene using Proportional Vertex Editing

Making a mountain scene is a simple process and involves some new skills that are very useful for modelling complex shapes.

First, start with a new *Blender* project and then resize the default plane so it is bigger. When the default plane is 18 squares wide, split the view to see the Side View, Top View and Camera View. When you have split the views, exit edit mode by pressing Tab.

Next we will make use of the Subdivide function. First, select all Vertices by pressing the A key or by using the Bounding Box (the B key) to select all of the vertices. When the vertices are yellow, press the **W** key and a menu will pop up. Click on the Subdivide option and you will see the plane subdivide. Repeat this process a number of times and you should have something such as that in **1** on the right.

Next, press the tab key to go out of edit mode and then back in to deselect all of the vertices. Proportional Vertex Editing basically allows you to move a selection of vertices, which in turn move other vertices in a connected manner. You can think of this tool in the same way as the smudge function in an image program... it simply pushes vertices around. The button that does this is shown at the top of this box since its use is so integral to this exercise.

When you click on the button, two more icons show up; one with a high triangular point and one with a lower triangular point. These two buttons adjust how much the related vertices are pushed. Again, to understand these buttons you need to play with them, so lets give them a test.

Click on one of the vertices in the top view and click on the higher point of the two icons. This is the Sharp falloff button. Now use the **G** key to move the vertice and you will see how the vertices are deformed. Your top view should now look something like **2** on the right.

To create the mountain peaks, select a vertice in the top view and then move the mouse to the side view where you can move it with the **G** key. Using a combination of the two views and the two falloff peak buttons in this mode, you can create quite spectacular views. Finally, add some spot lights, move the camera view and render the scene and you will have something such as that in **3** on the right.

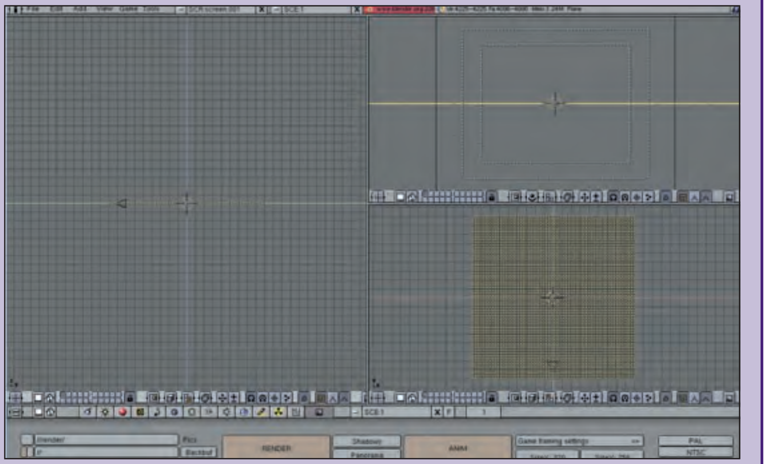

**1** A subdivided plane.

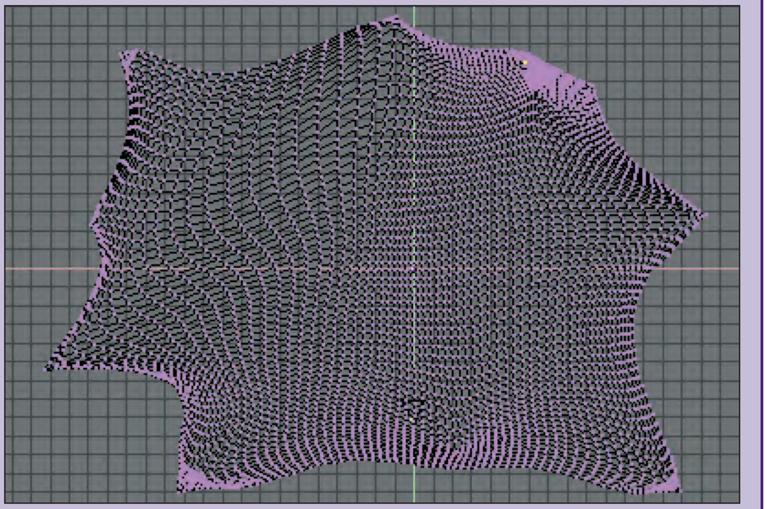

**2** Top view of the mountain.

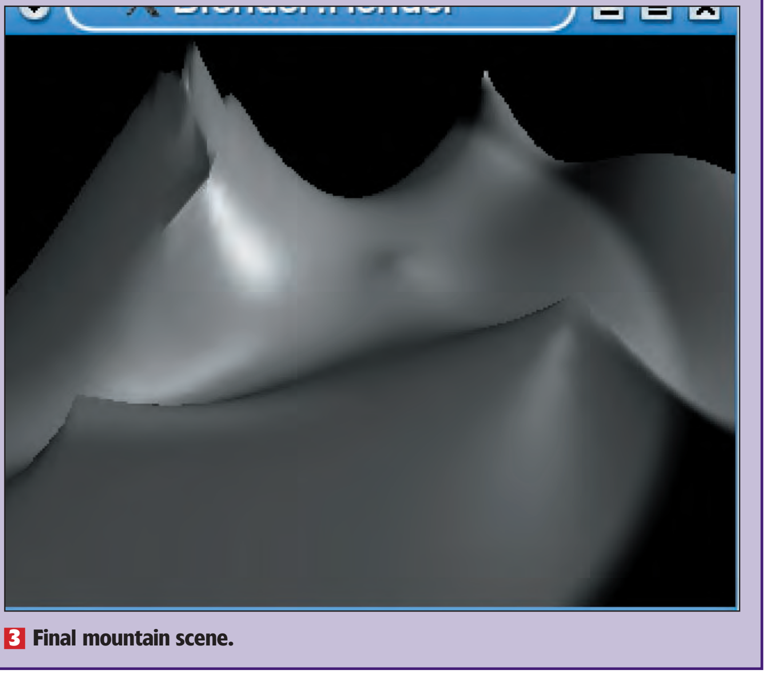

**3** Final mountain scene.

## NEXT MONTH

We will take our first look at Materials and Textures while still broadening our basic modelling skills. Try practicing, playing with the tools, experimenting and keep Blending!

RAPID DEVELOPMENT ENVIRONMENT

# Runtime Revolution

**PART 3** **Richard Drummond helps you maintain a sense of scale and proportion.**

*Runtime Revolution*, the rapid development suite that we've been learning to use in the last two issues, really scores when it comes to multimedia features. This fact should have started to become obvious last time, when we saw how easy it was to construct a program in *RR* to view images. The application that we built was very simple, however, and at the end of last month's instalment, I pointed out some of its shortcomings. But problems are there to be solved, after all. So, here we'll address and fix these issues and hopefully become more familiar with some of *RR*'s abilities.

If you remember, the first problem with our image viewer was that, if you cancelled the open image dialog, the image currently loaded would disappear. This is child's play to correct. Load up the image viewer program we've been writing, open the script editor on the 'File' menu button, and edit the script to read:

```
on menuPick pWhich
  switch pWhich
  case "Open"
    answer file "Select an image"
    if it is not empty then
      call "loadImage it" of card "ImageCard"
    end if
    break
  case "Quit"
  close defaultStack
    break
  end switch
end menuPick
```

All we've done here is to introduce a test before we send the **loadImage** message to the **ImageCard** object to tell it to load the selected image. The constant **empty** in the test is simply a more readable representation of the empty string **""**. Remember that the implicit variable **it** holds the result the from the file dialog.

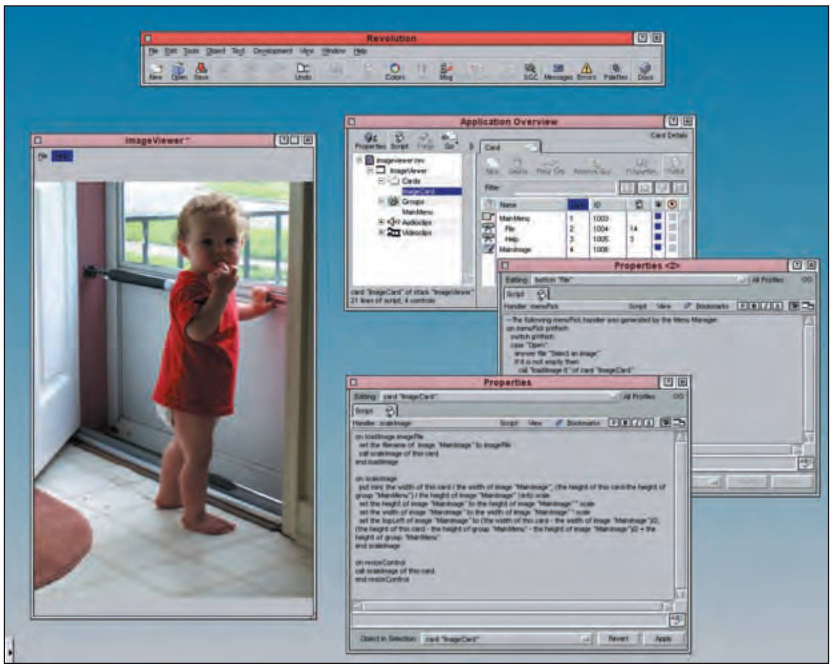

```
--The following menuPick handler was generated by the Menu Manager
on menuPick pWhich
  switch pWhich
  case "Open"
    answer file "Select an image"
    if it is not empty then
      call "loadImage it" of card "ImageCard"
```

```
on loadImage imageFile
  set the filename of image "MainImage" to imageFile
  call scaleImage of this card
end loadImage

on scaleImage
  put min( the width of this card / the width of image "MainImage", (the height of this card-the height of group "MainMenu") / the height of image "MainImage" ) into scale
  set the height of image "MainImage" to the height of image "MainImage" * scale
  set the width of image "MainImage" to the width of image "MainImage" * scale
  set the topLeft of image "MainImage" to (the width of this card - the width of image "MainImage")/2, (the height of this card - the height of group "MainMenu" - the height of image "MainImage")/2 + the height of group "MainMenu"
end scaleImage

on resizeControl
call scaleImage of this card
end resizeControl
```

**We've created a usable image viewer in *Runtime Revolution* that correctly scales pictures, in not much more than a dozen lines of code.**

## A question of geometry

The next problem is more difficult to solve – or at least more difficult to solve well! When we open a new image in our image viewer, *RR* centers that image based on the coordinates of the image that was previously being displayed. Alas, the results are generally visually unpleasant.

First, some background. The position of the image on screen - within its containing card – is governed by two properties, **left** and **top**, which specify the x and y coordinates, relative to the top left-hand corner of the card, of the top-left hand corner of the image. Its width and height on screen are similarly governed by the properties **width** and **height** respectively. The geometry of every control – whether it's a button, a field, or a scrollbar – has its own set of these properties which states where and how big it is drawn on screen.

The initial width and height of the **MainImage** object in our image viewer is derived from the width and height of the image file that we load from disk. The **left** and **top** properties of **MainImage** were initially set in *RR*'s form editor when we constructed the application; after that, it is reset every time we load a new image file, as *RR* centres the new image on the centre of the old. If you have tried loading many images into the application, you'll see the results are very messy, especially if you try re-sizing the application's window often. Luckily, we can take complete control over an object's geometry.

A partial solution to our problem is to reset the **MainImage** object's position after a new image is loaded. We can do this by editing the **ImageCard** object's script as follows:

```
on loadImage imageFile
  set the filename of image "MainImage" to imageFile
  set the left of image "MainImage" to 0
  set the top of image "MainImage" to the height of group
"MainMenu"
  set the width of image "MainImage" to the width of this card
  set the height of image "MainImage" to the height of this
card - the height of group "MainMenu"

end loadImage
```

This forces the image that we load into the **MainImage** object to be scaled and rendered to the size of the containing card (minus the area taken up by the menu bar).

While this solution works, it's still far from perfect: most obviously, the aspect ratio of the picture isn't preserved. (An image's aspect ratio is the ratio of its width to height. Changing the aspect ratio makes the picture look squashed – in one direction or the other).

## Maintaining your aspect

Let's tackle the aspect ratio problem by modifying the script again:

```
on loadImage imageFile
  set the filename of image "MainImage" to imageFile
  put the width of this card / the width of image "MainImage" into xscale
  put (the height of this card - the height of group "MainMenu") / the height of image "MainImage" into yscale
  put min( xscale, yscale ) into scale
  set the width of image "MainImage" to the height of image "MainImage" * scale
  set the height of image "MainImage" to the height of image "MainImage" * scale
  set the topLeft of image "MainImage" to 0, the height of group "MainMenu"
end loadImage
```

This may look complex, but it's really not. Transcript, *RR*'s scripting language, can be verbose, but that's what makes it so readable.

What we're doing here is to find the best way to scale the image to match the shape of the containing card. So we calculate two separate scale values: one from the width of the image to the width of the card, and one from the height of the image to the height of the card. Then we choose the minimum of these two scale factors (with Transcript's **min()** function) and use that to scale both the width and height of the image. Thus the aspect ratio of the image is maintained, and, by using the minimum of the x and y scale factors, we ensure that width and height of the scaled image will always fit within the dimension of the card in which it is drawn.

A couple of additional points are worth noting in the above. Firstly, we've used local variables to temporarily cache the results of the calculations that we perform. The statement **put a into b** stores the value of **a** into the property **b**. Here, we've used properties **xscale**, **yscale** and **scale** which – because they have not been declared anywhere – are temporary variables local to this handler. Secondly, instead of setting the image's **left** and **top** properties separately, we've used the 'compound' property **topLeft** to set both simultaneously. (Similarly, we could access the property **rectangle** to get or set an object's **top**, **left**, **width** and **height** properties as a set of four numbers.)

## Size matters

The next problem to deal with is how to adapt the position and dimension of the **MainImage** object when the size of the **ImageCard** it lives in changes – for example, due to the window (the stack) being resized by the user. *RR*'s Geometry Manager could be used for this. The Geometry Manager lets any object be configured so that either it's position can be translated or its dimensions scaled – automatically – in proportional response to changes in the size of its containing card. This provides a powerful and easy-to-use mechanism, but it doesn't suit our purposes because it doesn't let us preserve the aspect ratio of the image. Our solution will be to adapt our image scaling code above so that it can be called in response to resize events.

It's time to alter the **ImageCard** script again.

```
on loadImage imageFile
  set the filename of image "MainImage" to imageFile
  call scaleImage of this card
end loadImage

on scaleImage
  put min(the width of this card / the width of the image "MainImage", (the height of this card - the height of the group "MainMenu") / the height of the image "MainImage") into scale
  set the width of image "MainImage" to the height of image "MainImage" * scale
  set the height of image "MainImage" to the height of image "MainImage" * scale
  set the topLeft of image "MainImage" to (the width of this card - the width of  image  "MainImage")/2, (the height of this card - the height of group "MainMenu" - the height of image "MainImage")/2 + the height of group "MainMenu"
end scaleImage

on resizeControl
  call scaleImage of this card
end resizeControl
```

Here we've separated out the image scaling code as a separate event handler – which gets invoked in response to custom **scaleImage** events. This is now called by the new **loadImage** handler above, so that – just as before – the newly loaded image is scaled to fit the window. What's new here is the **resizeControl** handler. This handler is invoked in response to **resizeControl** events, which are events sent to every control when it is resized by the user. More on this in a moment.

The other changes we've made are to compress things a little; in the evaluation version of *RR*, an object's script may be only 10 lines long, so we've combined the scale calculation into a single line to save space. Also, we have tweaked the code which repositions the image, so that the image is centred in the card, rather than stuck in the top left corner.

Back to re-sizing. You should first realise that an object's **resizeControl** handler is not called by default when the stack it lives in is resized – for instance, when the user drags the stack window's re-size handle. By default, an object's **resizeControl** handler is only ever invoked when the user directly manipulates that object in *RR*'s form editor. Thus to make our new **resizeControl** method useful, we need to implement a **resizeStack** handler for the application's stack – which does get called when the user re-sizes the window of the running application – and add some glue. So, finally, open up the **ImageViewer** stack and enter:

```
on resizeStack
  call resizeControl of this card
end resizeStack
```

This in effect passes along resize events from the application window (the stack) to its currently displayed card, which, in our case, is the **ImageCard** object displaying the user's selected image, nicely scaled to find the size of the window. LXF

## Up the Revolution!

### Where to go for more information

**This is the end of our short series on *Runtime Revolution*. For more help on using the system, first consult the extensive set of tutorials and documentation shipped with the system. The built-in Transcript dictionary is also a great boon to the *RR* developer, and provides a handy cross-reference of *Runtime Revolution* keywords, messages and properties. Finally, the Revolution website at www.runrev.com is the canonical source of all information *Runtime Revolution*-related and contains some useful additional tutorials ands tips.**

# Answers

**If you are really stuck and the HOWTOs yield no good result, why not write in? Our resident experts will answer even your most complicated problems!**

## Our experts

Whatever your question is, we can find an expert to answer it – from installation and modem woes to network administration, we can find the answer for you – just fire off a letter or email and it'll all be taken care of.

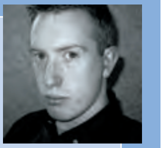

*LXF* answers guy **David Coulson** is a networking and security guru with plenty of sysadmin experience to boot.

**NIck Veitch** is the editor of the magazine, and answers your easy questions! Or indeed anything to do with *Grub, LILO, netatalk, vi...*

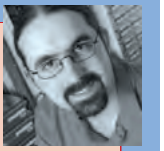

**Hans Huberland** is Rackspace Managed Hosting's Linux expert. Send any Linux system admin questions to **sysadminqa@rackspace.co.uk**

## Linux supported?

**Q I installed Mandrake 9.0 from your Christmas issue's CDs. Everything went smoothly during installation. The problem I am having is trying to set up an ISDN connection to the Internet. The installation provided a wizard for this, and I filled out (what I thought were) the correct settings. However, trying to connect after install failed.**

**I'm running a dual boot (Windows 2000, Mandrake 9.0) PIII Dell Laptop (LatitudeCPx H500GT). 11GB Hard Disk, 8.5 for Windows, 2.5 for Linux. The ISDN modem is connected via the laptop's only USB port. It works fine on Windows. I'm using KDE *kppp* to set up the account with my ISP, and configure the ISDN modem.**

**All the various combinations of settings I have tried have reaped no rewards. It is an external ISDN modem (Siemens), connected via a USB hub. This appears to be picked up by the system in the KDE configuration manager.**

***Brian King,*** *via email*

A Looking at the supported device list on **www.linux-usb.org**, it would appear that the Siemens ISDN TA is not supported by Linux. It's interesting that the *KDE Control Center* picks something up, although it may simply be an unknown USB device being connected to the system. When you connect the USB device, if you do

```
dmesg
```

it will list a pair of hex values which are the modems vendor & device identifiers, which are used to figure out which driver works for this device. Once you know these two IDs, you may find searching for them on Google produces some information as to what other people have seen when using this device.

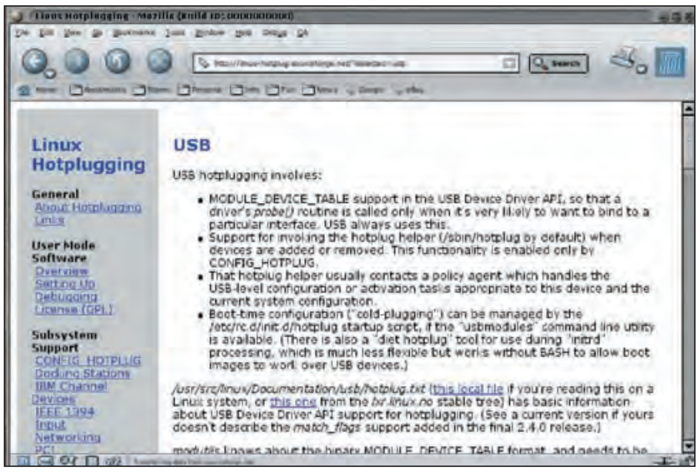

USB

USB hotplugging involves:

- MODULE_DEVICE_TABLE support in the USB Device Driver API, so that a driver's *probe()* routine is called only when it's very likely to want to bind to a particular interface. USB always uses this.
- Support for invoking the hotplug helper (/sbin/hotplug by default) when devices are added or removed. This functionality is enabled only by CONFIG_HOTPLUG.
- That hotplug helper usually contacts a policy agent which handles the USB-level configuration or activation tasks appropriate to this device and the current system configuration.
- Boot-time configuration ("cold-plugging") can be managed by the */etc/rc.d/init.d/hotplug* startup script, if the "usbmodules" command line utility is available. (There is also a "diet hotplug" tool for use during "initrd" processing, which is much less flexible but works without BASH to allow boot images to work over USB devices.)

*/usr/src/linux/Documentation/usb/hotplug.txt* (this local file if you're reading this on a Linux system, or this one from the *lxr.linux.no* stable tree) has basic information about USB Device Driver API support for hotplugging. (See a current version if yours doesn't describe the *match_flags* support added in the final 2.4.0 release.)

*modutils* knows about the binary MODULE_DEVICE_TABLE format, and needs to be

**As of the 2.4 kernel series, the Linux kernel has hotplug support, allowing scripts to be executed and modules to be loaded when devices are connected.**

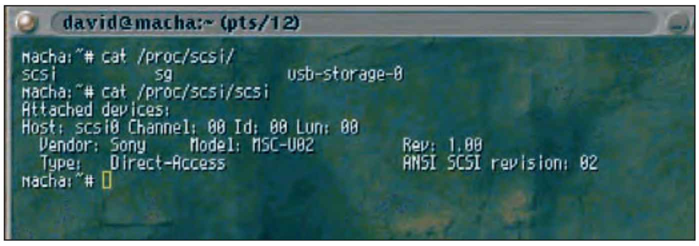

```
macha:~# cat /proc/scsi/
scsi         sg          usb-storage-0
macha:~# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: Sony     Model: MSC-U02           Rev: 1.00
  Type:   Direct-Access                     ANSI SCSI revision: 02
macha:~# 
```

**/proc/sscsi/scsi is an interface to the SCSI sub-system, allowing devices to be added and removed on the fly**

It's worth remembering that as USB is a device bus, it is possible to have software modems or 'passive' ISDN adaptors on this device, which require software support, as they do not appear as a regular serial device. They may require specific drivers to make them work that are unavailable to Linux systems.

## Camera hotplug

**Q I have been running Red Hat 7.2 (kernel 2.4.7-10) for about a year and *LXF* has been an essential part of the learning process. As far as peripherals are concerned I have had a smooth ride: I found drivers for my SmartLink modem and my Epson printer on the manufacturers' websites. Then I bought a Konica KD-100 digital camera. On the Konica website there was no mention of Linux but the blurb about Windows implied that this camera behaved as a "removable drive".**

**"Ha!" thought I, "I've seen something about that in *LXF*!" I promoted myself to su and did:**

```
/sbin/modprobe  usb-storage
```

**and followed up with dmesg; the last line was: USB Mass Storage support registered. That seemed hopeful so I connected the camera and, after a bit of trial and error, I found that I could:**

```
mount /dev/sda1 -t vfat /camera
ls /camera
```

**and there it all was. Grin and copy**

**files somewhere accessible. Great! I unmounted /dev/sda1 and then disconnected the camera. I went back to being a mere mortal user again and got on with what I should have been doing on the machine. All was well until shutdown (at sundown!)**

**When the messages got to "Unmounting file systems" it just hung for a long long time, and then reported something like:**

```
SCSI device set offline - not ready or
retry failed after bus reset host 1 id
0 lun 0
```

**and then shutdown down as normal. I can avoid this error if I**

```
modprobe -r usb-storage
```

**before shutting down but, of course I need to be su. Can you help?**
***Ken Wilson,*** *via email*

**A** It would appear that the SCSI sub-system is not deregistering the device when you unplug it from the USB port. There is a very simple way to do this, via the /proc/scsi/scsi entry:
As root, you simply need to do;

```
# echo "scsi remove-single-device 1
0 0 0" > /proc/scsi/scsi
```

The four values are host adaptor, channel, ID and LUN. From the error message you submitted, the above command should remove this device from the kernel, which will then allow you to cleanly shut the system down. To do this automatically, you can insert this command into /etc/rc.d/init.d/halt, or setup the **hotplug** daemon to handle your USB devices. **hotplug** will automatically load and remove kernel modules as they are required, and it's trivial to setup scripts to be executed when a specific device is plugged into your system. Information on hotplugging with Linux can be found at **http://linux-hotplug.sourceforge.net/**

## MDK modprobe

**Q I've been messing around with Linux for about five years and I'm sad to say I don't think it can match Microsoft Windows in its present form. The two main gripes I have are hardware and software installation. For example, my camera and scanner just won't work under Linux. I know someone will say just treat your camera like a USB mass storage device, but to do that you have to alter some sort of config file. Where's that kept and why do I have to alter it, why isn't it just set up in the first place?**

**The scanner you can forget. I have a Visioneer 4400 USB, Mandrake 9.0 (the distro I use) doesn't support it. One time I even bought another sound card because my original one wasn't supported, and let's not even talk about modems. What really annoyed me was that XP recognised my camera and scanner as well as everything else attached to my computer.**

**Then there is the software problem. I don't think I've ever compiled a program properly yet. Every time some sort of error comes up. I tried to compile *Mplayer* from your disc on *LXF38*. Having first unzipped it, uncompressed it, I then went to compile it ./compile and an error came up. It turns out I have the wrong compiler. So of I went searching for the right one. I found the site and downloaded the latest version. I then looked up instructions on how to install it. I hadn't a clue what they were talking about. Nearly all these sites assume you're extremely knowledgeable about computers. It really was just a foreign language to me. Why didn't I install the *Mplayer* using RPMs? I hear you ask. I did try and guess what? An error came up. I was lacking dependencies, *glu*, and *glut* to name but a few. Well like most people, that means nothing to me, I gave up after that. I also tried to upgrade KDE.**

**Using the RPMs for Mandrake, I tried to install *kdelibs* and *kdebase*, but a dependency error came up asking for the *qt* RPM. I tried to install the *qt* RPM and a dependency error came up asking for the *qt* common RPM, so I tried that one and it too had a dependency error, it was asking for the *qt* RPM. Then, I also tried installing using**

```
rpm -Fhv *.rpm i18n/*.rpm
```

**as your mag advises and at first I thought it was going to work. Sadly, a whole list of errors came up. I could list endlessly the failures I've had with hardware and software.**

**I realise that to some people these sort of problems are what**

»

## A QUICK REFERENCE TO: DEBIAN

**For those of us who prefer to avoid commercial distributions of Linux, such as Red Hat, SuSE and United Linux, there are many alternatives which are built by groups of developers. While they are provided without any 'proper' (*ie* paid for) commercial support, they thrive on community groups and mailing lists to supply information to the users.**

**Debian is one of the most popular non-commercial distributions of Linux and provides a massive variety of pre-compiled packages, for almost every software item for Linux. Due to the popularity of Debian, most developers will produce '.deb' packages along with the RPMs, and if not, then there will likely be a Debian maintainer for that particular package. Just think of a package and looking it up at www.debian.org/distrib/packages will yield productive results.**

**One of the major differences between Debian and other LInux distributions is that it is constantly being updated. Releases, rather than being specifically built distributions, are simply snapshots of the current 'stable' package tree rather than a finished product. The Debian package tree is continually available via HTTP or FTP from a variety of mirrors, so there is never any need to actually upgrade the whole box every six months, and can instead upgrade to the current packages once a week.**

**All of the packages available for Debian are available in an online package system, managed by a tool known as the *Advanced Package Tool*, or *apt*. While Debian packages are actually handled by *dpkg*, *apt* is a front-end which automatically downloads packages and their dependencies from the Internet then installs them for us. If we want to install *vim,* we simply have to do:**

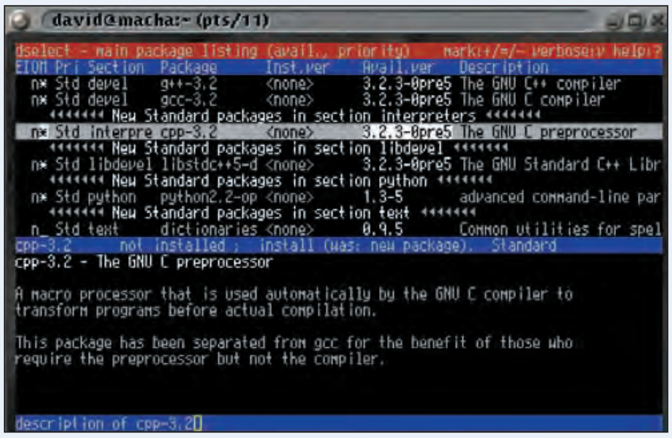

```
dselect - main package listing (avail., priority)     mark:+/=/- verbose:v help:?
EIOM Pri Section  Package      Inst.ver    Avail.ver   Description
  n* Std devel    g++-3.2      <none>      3.2.3-0pre5 The GNU C++ compiler
  n* Std devel    gcc-3.2      <none>      3.2.3-0pre5 The GNU C compiler
    ••••••• New Standard packages in section interpreters •••••••
  n* Std interpre cpp-3.2      <none>      3.2.3-0pre5 The GNU C preprocessor
    ••••••• New Standard packages in section libdevel •••••••
  n* Std libdevel libstdc++5-d <none>      3.2.3-0pre5 The GNU Standard C++ Libr
    ••••••• New Standard packages in section python •••••••
  n* Std python   python2.2-op <none>      1.3-5       advanced command-line par
    ••••••• New Standard packages in section text •••••••
  n_ Std text     dictionaries <none>      0.9.5       Common utilities for spel
cpp-3.2       not installed ;  install (was: new package).  Standard
cpp-3.2 - The GNU C preprocessor

A macro processor that is used automatically by the GNU C compiler to
transform programs before actual compilation.

This package has been separated from gcc for the benefit of those who
require the preprocessor but not the compiler.



description of cpp-3.2
```

***dselect* is a front-end to *apt* and dpkg on *Debian* systems, and makes installing packages a breeze.**

```
apt-get update
apt-get install vim
```

**apt-get update will update our package list from the online mirror, then the following command will install *vim* and any other packages which are required by *vim* but are not currently installed. There's no need for searching on the Net for required packages, or having to build source code to fulfil dependencies of badly built packages.**

**Debian is split into three main package trees. The first of these is 'stable', and is the most commonly used tree. The packages in 'stable' are generally considered usable on production systems. There is also 'testing' and 'unstable', which provide slightly more advanced packages than 'stable'. Software in testing and unstable may not actually be unstable, only that the package may not do everything it needs to before it makes it to stable.**

# FREQUENTLY ASKED QUESTIONS FILESYSTEMS

### FAQ CAN I MOUNT MY WINDOWS DRIVES IN LINUX?

Linux supports a variety of MS-DOS filesystems, including *vfat*, which provides long filename support, as well as ntfs, which is used on recent Windows installations.

If we wanted to mount our 'C' drive on /mnt/dos/c, we could do:

```
# mount -t ntfs /dev/hda1
/mnt/dos/c
```

Of course, **ntfs** should be replaced with **vfat** for a Win9x or ME filesystem. This will only actually be writable by the **root** user, so if other users need to be able to write to the Windows filesystems, then the entry for the filesystem must be inserted into /etc/fstab, using the **users** option, which will allow a non-root user to mount the filesystem.

### FAQ I DOWNLOADED AN ISO OFF THE NET, BUT I ONLY WANT ONE FILE OUT OF IT. CAN I MOUNT IT WITHOUT HAVING TO BURN THE ISO TO A CD?

ISOs can be mounted without having to burn a CD using the 'loop' device. Not to be confused with the 'loopback' network device, the 'loop' device allows a local filesystem image to be mounted as if it were a real block device. To do this, the kernel needs to be compiled with support for the loop device, or **loop.o** needs to be loaded. The ISO can then be mounted:

```
# mount -o loop myfile.iso
/mnt/loop
```

The ISO can then be navigated as it was mounted from a regular CD-R.

### FAQ I'VE HEARD ABOUT JOURNALED FILESYSTEMS. HOW DO THESE DIFFER TO REGULAR FILESYSTEMS?

A journaled filesystem stores a list of filesystem operations which have been performed prior to a system being shutdown uncleanly after a power failure. When the system comes back up again, there is no need to perform a complete **fsck** of the drive as a recent list of transactions is available.

On a large RAID array with hundreds of GB or even TB of storage, using a journaled filesystem is particularly useful when trying to bring the system back up quickly. With a journaled system, it only takes a few minutes, rather than days, to bring the filesystem back up to pre-powerdown levels. It is, of course, also very useful on laptop and embedded systems where power failure may be experienced quite frequently.

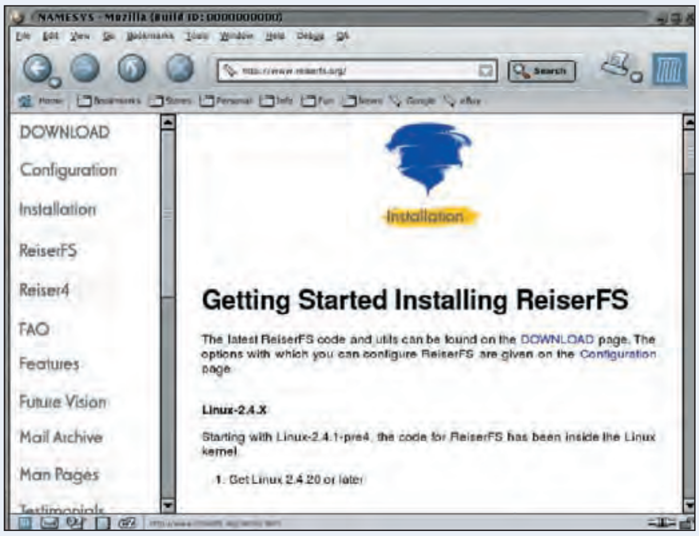
NAMESYS - Mozilla {Build ID: 0000000000}
File Edit View Go Bookmarks Tools Window Help Debug QA
http://www.reiserfs.org/
Search
Home Bookmarks Stores Personal Info Fun News Google eBay
DOWNLOAD
Configuration
Installation
ReiserFS
Reiser4
FAQ
Features
Future Vision
Mail Archive
Man Pages
Testimonials
Installation

Getting Started Installing ReiserFS

The latest ReiserFS code and utils can be found on the DOWNLOAD page. The options with which you can configure ReiserFS are given on the Configuration page

Linux-2.4.X

Starting with Linux-2.4.1-pre4, the code for ReiserFS has been inside the Linux kernel

1. Get Linux 2.4.20 or later

***ReiserFS* is a very popular journaled filesystem, although the filesystem needs to be rebuilt from backups.**

### FAQ WHAT JOURNALED FILESYSTEMS CAN I CHOOSE FROM?

There are currently three journaled filesystems in the Linux kernel; *ext3*, *reiserfs* and *jfs*. There is also *xfs*, which is available as a patch to the Linux kernel, so takes a little more knowledge and time. Both *ext3* and *reiserfs* have been in the 2.4 tree for a long time and *jfs* was recently added. Originally developed by IBM, it has been in use on their commercial UNIX systems for many years. There is no real easy choice, as everyone has benchmarks showing each filesystem to be better in specific situations.

Generally *ext3* is a good choice, and most modern distributions will allow the system to be installed on an *ext3* filesystem.

### FAQ IS IT POSSIBLE TO CONVERT AN EXISTING EXT2 FILESYSTEM TO A JOURNALED FILESYSTEM

**« computing is all about, but for me I just want to switch it on and use it. If I want to install software or hardware on Windows, I just put the disk in, follow the defaults, shutdown and switch back on and there it is. Until Linux can match Windows for user-friendliness, it's got no chance of competing. As much as I really want to use Linux in my work, I just can't afford to. I have to spend too much time getting it to work instead of using it for work. So for me and I suspect most other computer users Linux is fine just to mess around with but it just can't cut it in the real world.**

***Paul Ewing**, via email*

**A** If your camera is really a USB mass storage device, then you simply need to do **modprobe usb-storage** as root on the command line, then check with **dmesg** if the device registered as a SCSI device. You can load the usb-storage device automatically by compiling usb-storage into your kernel or by having it **modprobe** for the module at boot time by modifying /etc/rc.d/init.d/rc.local. As for the scanner – It would appear to be currently unsupported by Linux, although some information can be found at **www.buzzard.org.uk/jonathan/scanners-usb.html**. Unless manufacturers produce their own kernel modules for devices, or make specifications available to Linux developers, then it is difficult to make a device work with Linux without reverse-engineering the protocol used by it, which is obviously a very involved task. Certain distributions fair better than others, as they contain kernel patches beyond the standard kernel release, extending the devices which can be used with the installation, although none of them are perfect.

On the software front, it does become very frustrating when RPM dependencies are not handled properly. Other distributions, such as Debian, have more intelligent package management which automatically downloads the correct packages required for that which you are attempting to install. The biggest issue with RPMs is when vendor and third-party packages are mixed, as they may provide different dependencies to other packages, or they may not provide the most up to date software required by something else. Compiling software is often quite involved, as it requires both compilers and utilities, but also **-dev** packages for libraries, which contains the headers used by code being compiled.

## It's an Ogg's life

**Q** **I've just installed Mandrake 9.0 on my Athlon 1GHz. MP3s play perfectly, although mpg123 complains at first that it can't set the frequency correctly.**

```
High Performance MPEG 1.0/2.0/2.5
Audio Player for Layer 1, 2 and 3.
Version 0.59r (1999/Jun/15). Written
 and copyrights by Michael Hipp.
Uses code from various people. See
'README' for more!
THIS SOFTWARE COMES WITH NO
WARRANTY! USE AT YOUR OWN
RISK!
unsupported playback rate: 44100
Audio device open for 44.1Khz,
stereo, 16bit failed
```

### WITHOUT HAVING TO DESTROY THE DATA FROM THE FILESYSTEM?

An *ext2* filesystem can be converted to *ext3* by using the **tune2fs** command, which can add a journal to the existing *ext2* filesystem. Once a journal has been added, the filesystem can be either mounted as an *ext3* filesystem for journaling capabilities, or as a regular *ext2* filesystem for system recovery or if an *ext3* capable becomes unavailable. If we wanted to convert /dev/hda1 to an *ext3* filesystem, we would do:

```
# tune2fs -j /dev/hda1
```

/etc/fstab can then be modified to use **ext3,ext2** as the filesystem, so it will first try to mount it using *ext3*, then mount it using *ext2* if *ext3* is unavailable to the kernel.

### FAQ DOES A JOURNALED FILESYSTEM PREVENT DATA LOSS?

Not at all. All a journaled filesystem does is reduce the time it takes to perform a filesystem check following an unclean shutdown of the system (although the journal data could also be used to determine which files may have been corrupted). If you want protection against data loss on a filesystem due to hardware failures, then RAID is really the only option, although a journaled filesystem is generally used in combination with a RAID

```
Trying 44.1Khz, 8bit stereo.
unsupported playback rate: 44100
Audio device open for 44.1Khz,
stereo, 8bit failed
Trying 48Khz, 16bit stereo.
Directory: music/Green Day/Warning/
Playing MPEG stream from 01 -
Warning.mp3 ...
MPEG 1.0 layer III, 128 kbit/s, 44100
 Hz joint-stereo
unsupported playback rate: 44100
Audio device open for 44.1Khz,
stereo, 16bit failed
Trying 44.1Khz, 8bit stereo.
unsupported playback rate: 44100
Audio device open for 44.1Khz,
stereo, 8bit failed
Trying 48Khz, 16bit stereo.
[3:42] Decoding of 01 -
Warning.mp3 finished.
```

**However, ogg123 fails all the time, complaining that it can't open the audio device:**

```
Audio Device:   OSS audio driver
output
Playing: music/Coldplay/A Rush of
Blood to the Head/Coldplay -
Clocks.ogg
libao - OSS cannot set rate to
44100
Error: Cannot open device oss.
```

**I have a VIA VT82C686 sound controller built into the South Bridge chip on the motherboard. Any ideas what the problem could be?**

***Neil Greenwood,*** *Cardiff, UK*

**A** The problems you are seeing are because the players are unable to change the rate of the sound device to 44.1k or 48k. You've not said which sound module you are using, so you may want to look at using ALSA (**www.alsa.org**), rather than the standard kernel modules, as ALSA has better support for specific sound cards, which should prevent the problem you are seeing. It's not unlikely that you are using a kernel module which detects your sound card, but doesn't quite work right with the specific revision of the chipset which you board has.

## Duff DVD

**Q After inserting a DVD into the drive, the dvd drive seems to time out after a few minutes. I get a message like:**

```
ls:
/mnt/dvdrom/Magazine/HotPicks/LSH
W/: No such file or directory
```

**If I eject and reinsert, it's OK for a few more minutes. What gives?**

***Craig Drinnan,*** *via email*

**A** Following an error, running the **dmesg** command as root will output any kernel errors, such as device timeouts or issues with reading from the media. Without having more information, it's difficult to establish if it's actually the media or the device to blame.

Have you tested the device with other DVDs and had the same problems? If it is only the *LXF* DVD which you are having problems with, there is an address in the front of the magazine where you can obtain a replacement DVD.

## FAT32 in Red Hat?

**Here's a question I've been pondering: Can files on NTFS and FAT32 partitions be viewed on Red Hat Linux 9.0? If they can, how is it done. Under mandrake 8.0 these files were kept under /mnt/win_c/**

**Thanks in advance,**

***John Ennew,*** *via email*

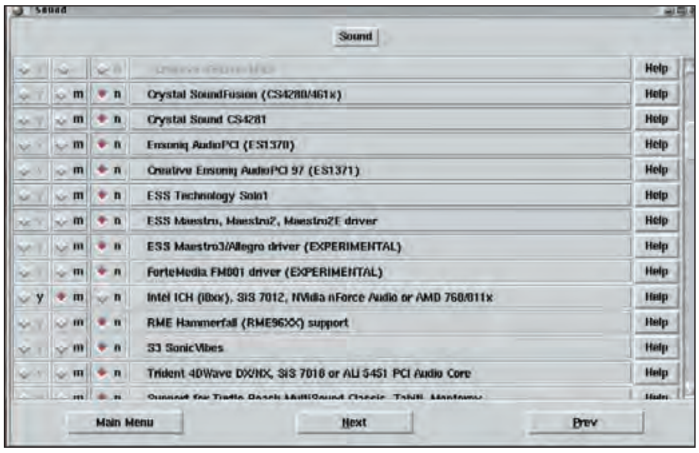
Sound

Crystal SoundFusion (CS4280/461x)
Crystal Sound CS4281
Ensoniq AudioPCI (ES1370)
Creative Ensoniq AudioPCI 97 (ES1371)
ESS Technology Solo1
ESS Maestro, Maestro2, Maestro2E driver
ESS Maestro3/Allegro driver (EXPERIMENTAL)
ForteMedia FM801 driver (EXPERIMENTAL)
Intel ICH (i8xx), SiS 7012, NVidia nForce Audio or AMD 768/811x
RME Hammerfall (RME96XX) support
S3 SonicVibes
Trident 4DWave DX/NX, SiS 7018 or ALi 5451 PCI Audio Core

**Linux supports a wide range of sound devices as part of the standard kernel and ALSA supports even more.**

**A** You can mount your Windows partition in Linux by doing the following as root:

```
# mkdir /mnt/win_c

# mount /dev/hda1 /mnt/win_c
```

You can also set it so that the partition is mounted at boot time by adding this line to /etc/fstab

```
/dev/hda1 /mnt/win_c   ntfs
defaults 0 0
```

Of course, /dev/hda1 should be replaced with the correct partition, assuming your Windows C drive is not the first partition on the first IDE device. Also, in the /etc/fstab line, ntfs should be replaced with **vfat** if the filesystem is not NTFS. Both can be extended to mount your D, E, F and other Windows partitions.

## Home network

**Q Could anybody tell me how to access my desktop X server from my laptop?**

**I have the two networked. I know it is possible and remember seeing what is required here a few months ago but unfortunately I can't seem to find it anywhere in the archives.**

**What I want to do is set up the desktop to allow the laptop to access the X server and replace the screen/keyboard/mouse so that I can take the base unit away and use it with the laptop rather than carting the monitor around.**

**The machines will only be connected to one another so security is not an issue; I just want the simplest way of doing it.**

***From the LXF forums***

**A** There are a number of different ways to run X clients on a remote system locally. The simplest is to simply allow a remote server to connect to the local X display on the laptop. You can then shell into the desktop system and modify the DISPLAY variable to point the X client to the laptop's X display. On the laptop you would do:

```
xhost +10.1.1.2
```

Replacing 10.1.1.2 with the IP address of the desktop system.

You can either ssh or telnet into the desktop system, and then run a terminal emulator like *xterm* or *Eterm* with the following:

```
DISPLAY=10.1.1.1:0.0 xterm
```

Again, **10.1.1.1** should be replaced with the IP address of the laptop on the network. The **:0.0** should remain, as this points to the first X display on the laptop.

Alternatively, you could use *TightVNC* to run a network-based X server on the desktop, then connect to this via *vncviewer* on the laptop. *TightVNC* can be found at **http://www.tightvnc.com/**

## Posting to the forum

### The LXF online community

**Got a technical question? Other *LXF* readers may be able to help!**

The forums at **www.linuxformat.co.uk** have a section dedicated to technical queries, hardware, programming languages and general help. As well as being able to call on **lxfadmin** (when there's no deadline!) and the ever-present 'anonymous', the forums are also frequented by Linux heroes like **Jeremy**, **Nelz**, **Fingers99**, **Rhakios**, **Erin** and many others brimming with knowledge and experience of using Linux in a wide variety of situations.

## Missing passwords

**Q I've inherited a Red Hat 7.3 web server from a sysadmin who left the company on bad terms and he did not leave the passwords for the server. I've reset the root password by going into single user mode but the previous admin set a different root password for MySQL. Is there any way to reset this password without putting our data in jeopardy?**

***Jonathan Twins,*** *via email*

A Assuming the installation is based on the Red Hat RPMs then you'll need to run the following commands:

```
/etc/rc.d/init.d/mysql stop
/usr/bin/safe_mysqld —skip-grant-tables &
mysql -u root
UPDATE user SET Password=PASSWORD("your password here") WHERE User="root";
exit
```

Next you'll need to end the **safe_mysqld** process you started as a background process and restart the main mysql. You will now be able to use a new password to gain access.

If the installation was not from RPM you'll need to find where the safe_mysqld binary was installed and run it from there. Usually this will be something like:
/usr/local/sbin/safe_mysqld

## Perfect Prints

**Q I'm pretty new to Linux administration so I hope this is not a silly question. We have a file and print server at the office and for ease of use; I use X to administer the printers. However, I'd like to be able to use my Linux workstation to do the administration. I know that most of the administration can be done from the command line and I can ssh into the server but I'm not that confident yet and I'd prefer to do this in X. Somebody told me that because X is actually based on a client-server architecture that this can be done but he wasn't sure exactly how.**

***Karl Vars,*** *via email*

A This certainly can be done and is one of the great (and infrequently used) features of Xfree86. You'll need to issue the following command on your workstation to allow your server to connect to the X server on your workstation.

```
xhost +10.0.0.1 (where 10.0.0.1 is your server's IP address)
```

Then you'll need to log into your server via SSH and type in the following as root:

```
export DISPLAY="10.0.0.99:0.0 (where 10.0.0.99 is your workstation's IP address)
```

This will export your server's display to point to your workstation's X server. While still logged into **ssh** you can now run *printconf-gui* or any of the other graphical utilities installed on your server.

Remember that you need to run the application on your server through the SSH session. Further, if you close the SSH session the DISPLAY variable will be reset. If you would like this to be permanent I would recommend creating a user specifically for remote administration and permanently change that user's DISPLAY variable and their login scripts.

## Strange server

**My Red Hat 7.3 server is behaving strangely. At first everything ran fine but now each time I run *logrotate* (nightly using *cron*) it just keeps running indefinitely and does not appear to be rotating the log files. I have to kill the job and it dies gracefully, but if I rerun it manually or let it run through *cron* it does the same thing. I'm not entirely sure when this started to happen so it's difficult to say what has changed.**

***J Burwal,*** *via email*

I can almost guarantee that this is a known problem with the *mailman* installation in Red hat 7.3. If you look at /etc/logrotate.d/mailman you will notice that the name of the file to be rotated for this application is /var/log/mailman/*. This means that it will rotate all files in that directory, for example, /var/log/mailman/error (and call it error.1 ). The next week it will rotate error.1 and call it error.1.1 and so on. As you can imagine this creates exponentially more files each time it is run. In order to repair this you should do the following:

Download and install the newest version of mailman from Red Hat. Change into /var/log/mailman, you should see the mess that this faulty logrotate script has created. Clear out all the excess log files, this could take a while depending on how long it's been running like this. Open up /var/lib/logrotate.status and clear out all the entries for the bogus logfiles. With these steps completed you should be able to run logrotate without any problems.

## Corrupt filesystem

**On a Red Hat 8.0 system I've had some bad luck a couple of times and been stuck with a corrupted filesystem, so I'm trying to make my system a little more paranoid when it comes to filesystem integrity.**

**If the server shuts down unexpectedly, *fsck* asks you if you would like to run a filesystem**

## About Rackspace Managed Hosting

**Rackspace Managed Hosting™ provides complete managed hosting solutions with Fanatical Support to small to medium-sized customers as well as large enterprises. All customer platforms include state-of-the-art data centres, customised servers, burstable connectivity, 99.999% uptime SLA, a dedicated account manager, instant emergency response and access to live expert technicians 24x7 for support of all hardware and core software.**

**Founded in 1998 with locations in San Antonio, Texas and London, the company manages over 7,000 servers for customers in more than 80 countries. For more information, visit www.rackspace.co.uk**

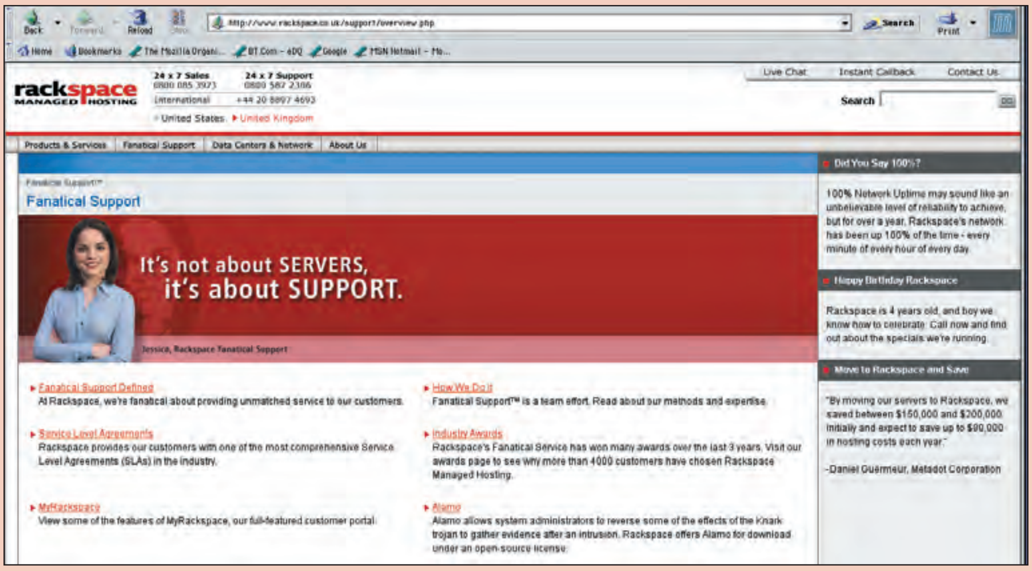

http://www.rackspace.co.uk/support/overview.php

rackspace MANAGED HOSTING

24 x 7 Sales 0800 085 3973
24 x 7 Support 0800 587 2306
International +44 20 8897 4693

Fanatical Support™

Fanatical Support

It's not about SERVERS, it's about SUPPORT.

Jessica, Rackspace Fanatical Support

Fanatical Support Defined
At Rackspace, we're fanatical about providing unmatched service to our customers.

How We Do It
Fanatical Support™ is a team effort. Read about our methods and expertise.

Service Level Agreements
Rackspace provides our customers with one of the most comprehensive Service Level Agreements (SLAs) in the industry.

Industry Awards
Rackspace's Fanatical Service has won many awards over the last 3 years. Visit our awards page to see why more than 4000 customers have chosen Rackspace Managed Hosting.

MyRackspace
View some of the features of MyRackspace, our full-featured customer portal.

Alamo
Alamo allows system administrators to reverse some of the effects of the Knark trojan to gather evidence after an intrusion. Rackspace offers Alamo for download under an open-source license.

Did You Say 100%?
100% Network Uptime may sound like an unbelievable level of reliability to achieve, but for over a year, Rackspace's network has been up 100% of the time - every minute of every hour of every day.

Happy Birthday Rackspace
Rackspace is 4 years old, and boy we know how to celebrate. Call now and find out about the specials we're running.

Move to Rackspace and Save
"By moving our servers to Rackspace, we saved between $150,000 and $200,000 initially and expect to save up to $90,000 in hosting costs each year."
-Daniel Guermeur, Metadot Corporation

## WIN A SHARP ZAURUS PDA FROM RACKSPACE MANAGED HOSTING

www.rackspace.co.uk

**Every month, the best question** related to Systems Administration that a *LXF* reader sends in wins a Sharp Zaurus SL-5500 running Linux.

With industry support from The likes of IBM, SAP, Oracle, and Sybase, the Zaurus features a QWERTY keyboard behind a sliding cover. The 65,536 colour screen is 240x320 resolution. An Intel® StrongARM® 206MHz processor and 64MB memory makes the Zaurus ideal for downloading software or playing MP3 and MPEG1. The docking station that fits into its USB connector, allowing users to synch with their desktop, schedules, address book and ToDo lists. The email app supports POP3 or IMAP email servers. Write and send emails and browse the Internet with a compatible modem card, or using the infra-red port and a mobile phone; or any replied, forwarded, or composed emails are transferred to your outbox the next time you synch. SD and CF slots allow upgrades and addition of software.

Email all your sysadmin questions and problems to Hans Huberland: sysadminqa@rackspace.co.uk Every published query on these system admin pages wins a Rackspace T-shirt.

**It's only 74mm wide by 138mm long and a svelte 18mm thick – so the Sharp Zaurus really *is* Linux in your pocket!**

## ★ Star Question – PDA winner!

This issue's lucky winner is ***Steven*** – your new Zaurus PDA will be with you shortly!

**Everyone knows that Linux is stable, reliable and flexible, but that cannot be said for all things in the world of computers. Being a relatively new sysadmin, for a very small network of computers, I have had few problems but the one that does cause me a problems is:**

**Once the systems are up and running I like to try and keep them fully optimised. This was easy enough in the days when I ran Windows, as I had a regime of running optimisation programs such as *Norton Utilities, First Aid* or the more basic *Scandisk*. Is there an equivalent suite for Linux? Or would this be carried out by a multitude of smaller programs?**

**This of course could be made easier by using *cron*, which is often explained in tutorials, is easy to set up and perfect for servers which run 24/7. However, these tutorials often neglect to mention the most useful programs to run regularly.**

**I for example only clean out the /tmp directory and fsck each partition every x times they are mounted, but some more suggestions would be welcome.**

**A useful shell script could then be created that could be run on demand on stand alone systems or on the computers that are not turned on all the time.**

***Steven,*** *via email*

Fortunately the Linux architecture is vastly different from that of Windows and most of the optimisations that are required for Windows are not required for Linux. You have not mentioned what these systems are used for, so as this is a System Administration Q&A, I'll assume that these are some sort of Internet or corporate servers.

Linux's default *ext2* or *ext3* filesystem does not suffer from fragmentation like some other filesystems so regularly scheduled disk cleansing will not give you increased performance. It's always wise to run an *fsck* on your filesystem at regular intervals. This interval is best determined by how often the system will be rebooted. If the system is going to be rebooted on a regular basis then you need only run *fsck* every twenty or thirty reboots (depending on how much the filesystem is written to). If the server is not going to be rebooted very often or not at all then you may want to run a filesystem check more regularly.

You may want to make sure that you have some sort of log rotation implemented as logs can grow and eventually fill up all your disk space. By default, Linux log files will continue to grow until they reach the maximum filesize for your filesystem type. When this happens either data will no longer be appended or the program trying to write to the file, will fail. Luckily most systems have logrotate installed by default so I would recommend making sure that this is available. Red Hat stores the config file for logrotate under /etc/logrotate.conf and config files for individual logfiles can be found under /etc/logrotate.d/ .

In addition I would make sure that all the default accounts that can receive system mail (such as root, postmaster, etc) are aliases of a real user who will collect this mail. Any errors in your *cron* jobs or other system events will get mailed there and should be checked out before they get out of hand.

Lastly I would recommend setting up *logwatch* or some other program, which will send you a mail with a summary of the log events for the day. Any unusual activity should be picked up and again can be remedied before anything unexpected gets out of control.

Other than that the system is pretty much self-maintaining. Most systems have default housekeeping jobs set up to clean out /tmp and similar tasks. I'm sure many of our readers have some other interesting and useful tips and tweaks though. Send any of your Linux tips – we'll print a selection of the best!

**check. If it does not receive a response after 5 seconds it skips the check. Is there a way that I can either increase this time to a minute so that I get a chance to catch it, or force it to do a check if it does not receive any user input?**

**Or, alternatively, is it possible to have *fsck* default to "Yes" for the filesystem check?**

***Edwin,*** *via email*

You will need to create the following file: /etc/sysconfig/autofsck

You can use the following two directives to make either of the changes that you mentioned you desired to make:

```
AUTOFSCK_TIMEOUT=5
```

This sets the number of seconds to wait before proceeding with the default action, which is defined by the next directive. The default is 5.

```
AUTOFSCK_DEF_CHECK=yes
```

**yes** turns on *fsck*ing by default so that the system will behave more like older Red Hat servers with *ext2* filesystems which will *fsck* by default after an improper shutdown.

## Linux From Scratch

**Thanks for producing an excellent mag – you now have a subscription from me!**

**I've been trying Linux for nearly a year now, learning plenty, especially through the Linux From Scratch project. I have many points for the letters page, but here are my technical questions.**

**I have installed several distros in the last 10 months with kernels ranging from 2.4.18 to 2.4.20 but as yet still can't get a fully working system. I recently upgraded to a home built K7S6A Athlon and a £60 80GB Barracuda (a good value buy, totally silent). I carefully picked the best-value components, making sure everything had Linux support. And it does – all of it, but not at the same time. The two best working distros are Mandrake 9.0 and Lycoris, both from *LXF* covers. Mandrake 9.0 does not create any /dev/sda* devices even if I make sure my USB Zip is plugged in when installing. I tried using MAKEDEV but it creates sda4- , so I try renaming, not allowed. I found a page saying to create the device and use cp /dev/sda4- sda4 as the /dev folder is write protected (which made no sense to me as I can create sda4-), I tried it but this wasn't allowed either. The modules are installed and modules.conf has all the correct entries, I looked at many FAQs and HOW-TOs I think I just need an sda4 to mount against but can't seem to make one.**

**The USB zip system on Lycoris works like a dream, with bells and whistles. I put the disk in the drive and it auto mounts with a Zip logo which can be used as a normal directory. When I eject the disk the logo disappears. No data probs, perfect, much better than Windows.**

**But Lycoris won't play sound! On logging in a warning comes up with error on device /dev/dsp either not working or missing. It has been installed, I delete it and make a new one with makedev (lowercase on Lycoris) still no luck. The i810_audio module is installed and rc.d and modules.conf are OK. I tried a separate 7012 module that I found on the web. No luck. I did a fresh alsa install – no luck either. Yoper seems to work except for my printer. The sound works and there are sda* devices are there though I haven't had a chance to try mounting my zip but it has problems booting which brings me to my third question.**

**I don't think I'm using *LILO* correctly. I copied the entries form Yoper's *LILO* into Mandrake *LILO* but it seems to stop Yoper booting after a few boots, and then the rescue disk doesn't work, so it might be Yoper; although it is very fast and KDE3.1 is really clever I'll leave it for a while. But Lycoris uses *grub* and if I boot without *grub* using *LILO*, none of the removable drives work. not the CD-RW not the DVD not the floppy and definitely not the beautiful ZIP.**

**I have other questions but these are the ones I have spent the most time finding answers that don't work.**

**1. creating devices?**
**2. installing and checking modules?**
**3. *Grub/LILO*?**

**I really hope you can help I'm getting quite lost.**

***Ralph Cox,*** *London*

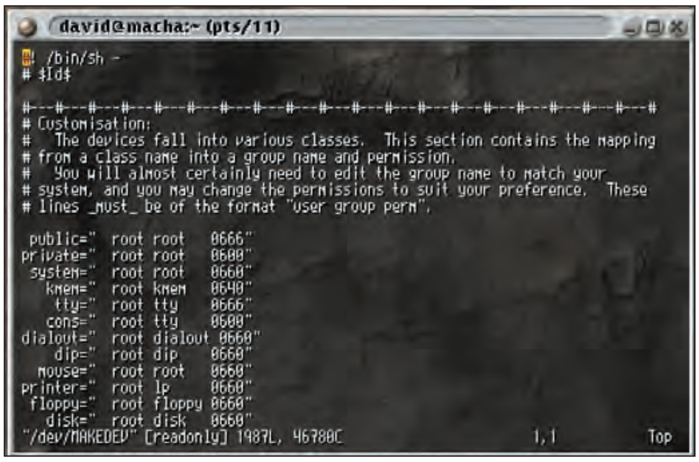

```
#! /bin/sh -
# $Id$

#---#---#---#---#---#---#---#---#---#---#---#---#---#---#---#---#---#---#
# Customisation:
#   The devices fall into various classes.  This section contains the mapping
# from a class name into a group name and permission.
#   You will almost certainly need to edit the group name to match your
# system, and you may change the permissions to suit your preference.  These
# lines _must_ be of the format "user group perm".

 public="  root root    0666"
private="  root root    0600"
 system="  root root    0660"
   kmem="  root kmem    0640"
    tty="  root tty     0666"
   cons="  root tty     0600"
dialout="  root dialout 0660"
    dip="  root dip     0660"
  mouse="  root root    0660"
printer="  root lp      0660"
 floppy="  root floppy  0660"
   disk="  root disk    0660"
"/dev/MAKEDEV" [readonly] 1987L, 46780C                        1,1           Top
```

**The MAKEDEV script generates all of the device files in /dev for systems not running *devfs*.**

And you say that you've saved some questions for the letters page...? Starting with your USB Zip problems. The USB Zip drive requires the usb-storage module to be loaded before it will register with the SCSI sub-system. It's unusual that MAKEDEV would create /dev/sda4-, as this is not a standard device name. You can make /dev/sda4 yourself with:

```
# mknod /dev/sda4 b 8 4
# chmod 600 /dev/sda4
```

If you are unable to write to /dev, ensure that you are the root user and that /dev has writable permissions. After you've plugged your USB Zip drive in, check with **dmesg** to see exactly what the kernel has decided the device is.

Your sound problems could be a number of different things. It could be that you are not loading the correct module for your sound chipset. You should check the motherboard manual to find out which chipset is on-board so you can either load the correct kernel module, or install ALSA if it is not supported by the core kernel sound system.

If you have multiple Linux installations on one box, then it will certainly cause problems if you have more than one one of them installing their own *LILO* boot sector. You may, for example, being using a Mandrake kernel with another distribution, which can confuse things, plus some distributions may be passing arguments to the kernel on the command line, to enable specific devices.

It sounds like you've had the most success with Yoper, so it may be best to stick with that, as you have the fewest elements to fix with it. If your printer does not work, then it should be fairly straightforward to setup *lpd* or CUPS for the specific printer you have connected to your system.

## Filesystem freeze

**My Linux filesystem is behaving a little oddly, in that every so often (~once a day) it freezes completely – often with a loud clunking of the hard disk heads just before.**

**I've also had the odd program suddenly start repeatedly crashing after working fine for several months, or the odd configuration file get re-generated afresh. I'm thinking along the lines that it may be the hard disk on the way out.**

**I therefore ran e2fsk on my *ext3* partition, and it gave the warning that running it on a mounted system could cause damage. Sure enough, each time I run it I get some errors (free block counts wrong, block bitmap differences), and each time I fix them. Are these errors likely to be coming from the hard disk dying, or from damage caused by running e2fsck on a mounted system?**

***From the LXF forums***

A Aside from the fact that your problem appears to be down to a dying hard drive, it's generally a bad thing to run **fsck** on a filesystem mounted read/write. If you don't want to unmount the drive, you can remount it read-only before **fsck**ing it:

```
# mount -o remount,ro /
# fsck /dev/hda1
```

Your best option at this stage is likely to back the drive up to another drive or to CD and look at obtaining a replacement before disaster strikes and renders all your important data unrecoverable. LXF

### Submission advice

**We are happy to answer all sorts of Linux related questions. If we don't know the answer, we'll find out for you! But in order to give you the best service, it helps a lot if you read the following submission advice.**

- Please be sure to include any relevant details of your system. "I can't get X to work" doesn't really mean anything to us if we don't know things like what version of X you are trying to run, what hardware you are running on.
- Be specific about your problem. Things like 'it doesn't work' or 'I get an error' aren't all that helpful. In what way does something not work? What were you expecting to happen? What does the error message actually say?
- Please remember that the people who write this magazine are NOT the authors or developers of Linux, any particular package or distro. Sometimes the people responsible for software have more information available on websites etc. Try reading the documentation!

We will try and answer all questions. If we don't answer yours specifically, you'll probably find we've answered one just like it. We can't really give personal replies to all your questions.

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

# missed one?

## *LINUX FORMAT* BACK ISSUES

Every month *Linux Format* brings you the best tutorials, the essential reviews and the latest news. But if you've missed out on a must-read feature or a vital programme from our expertly compiled CDs and DVDs, order your back issue NOW! And remember, you need never miss an issue of your favourite Linux mag, if you subscribe to *Linux Format* (see overleaf for more details).

### June 2003

Product code:
**LXFB0041**(cd)
**LXFDB0041**(dvd)

**DVD HIGHLIGHTS:**
Slackware 9.0, Mutt, OpenOffice.org, Xine, TakeAJoint, Peanut Linux. FCKeditor, KDE 3.1.1a, mirror of Linux Documentation Project

**CDs HIGHLIGHTS:**
Mandrake 9.1, AC3D, KDEradio, The GIMP – including 400 pluigins from The GIMP Registry

**MAGAZINE FEATURING:**
AMD's 64-bit Opteron CPU, Linux Laptop Megatest, The Hurd, JCodeBox review, Build your own RPMs, Red Hat 9 vs SuSE8.2 vs Mandrake 9.1, OpenAuth Linux PC review

### May 2003

Product code:
**LXFB0040**(cd)
**LXFDB0040**(dvd)

**DVD HIGHLIGHTS:**
Vector Linux, Dropline GNOME, Lopster, jEdit, QtParted, QuickRipDVD, BaitAndSwitchHoneypot, AdiosBootCD, zTappy, MondoRescue

**CDs HIGHLIGHTS:**
Mozilla 1.3, Opera update, IMP server, Kino, IceWM, ECLiPt Roaster, VariCAD

**MAGAZINE FEATURING:**
SCO v IBM $1 billion lawsuit, Blender tutorial, *LXF* Awards 2002 results, Linux on Apple Mac, Gigabyte SR113 server, EFI to replace BIOS, VariCAD & IEMS mailserver reviews

### April 2003

Product code:
**LXFB0039**(cd)
**LXFDB0039**(dvd)

**DVD HIGHLIGHTS:**
CensorNet, K3b, VIPS, FrozenBubble, Linux Open Media Box, IDMS Linux, WhiteDune, OBM, AstaroSecurityLinux

**CDs HIGHLIGHTS:**
SME Server, TightVNC, GNOME 2.2, VOCP, CoyoteLinux, UT2003 update, Blender, Qt, BootEverywhereLinux

**MAGAZINE FEATURING:**
Customising Linux, Evesham Linux PCs, Open publishing with WIKI, Linux PIMs and PDAs, SME Server review, File Managers Roundup, Serial Networking, GnuCash

### March 2003

Product code:
**LXFB0038**(cd)
**LXFDB0038**(dvd)

**DVD HIGHLIGHTS:**
Python, The GIMP, Lycoris, Yellowdog, Knoppix, Cygwin, DansGuardian, Senken, AbiWord, X Northern Captain, Eclipse

**CDs HIGHLIGHTS:**
KDE 3.1, UNIX Tools For Windows, Konstruct, phpMyAdmin, X Net Strength, AutoUpdate

**MAGAZINE FEATURING:**
You, Copyright & The Law, X apps on Mac OS with Fink, SVG versus Flash, Audio Editors roundup, Python tutorial, GIMP secrets, SuSE Office Desktop examined

### February 2003

Product code:
**LXFB0037**(cd)
**LXFDB0037**(dvd)

**DVD HIGHLIGHTS:**
IPCop, Postnuke, Lllypond, Squid, Allegro, Cocoon, Fandango, gPhoto, Xrmap, Ted, AdvancedPHPdebugger, Flightgear updates

**CDs HIGHLIGHTS:**
Yoper, KOffice 1.2.1, Parted, DOM Tooltip, Guarddog, GL-117, PythonImagingLIbrary

**MAGAZINE FEATURING:**
Lindows on test, Serial ATA, Window Managers roundup, Arkeia review, Xft2, build your own Linux backup box, Java compiling, Jool Kwartz review, New Larger Linux Pro mag

### January 2003

Product code:
**LXFB0036**(cd)
**LXFDB0036**(dvd)

**DVD HIGHLIGHTS:**
Evolution PIM/calendar/email client, Freevo personal video recorder, OpenZaurus, PilotLink

**CDs HIGHLIGHTS:**
Racer, Freeduc, Linux From Scratch, GNOME Meeting, Opera 6.1, Film GIMP, phpOpenTracker, Tuxpaint, Snort

**MAGAZINE FEATURING:**
Optimisation tips and tricks, Netbox Cubit review, build your own desktop, Xandros, Cubit, CORBA, Realsoft 3D, Instant Messenger roundup

Other back issues also available, please call for further information, giving the issue number and/or date required

## To order, call the back issues hotline on 0870 4448645

**or +44 1458 271178 for overseas orders**

Remember to quote the issue code shown above and have your credit card/switch card details ready.

**Prices (inc post and packing):** CD – UK£5.99, US & Eur £6.99, rest of world £7.99
DVD – UK £6.99, US&Eur £7.99, rest of world £8.99.

LINUX FORMAT
www.linuxformat.co.uk

# SUBSCRIBE AND

## ON A YEAR'S SUPPLY OF LINUX FORMAT

# That's 3 issues free!

Title__________ Initials __________ Surname ______________________________

Address ______________________________________________________________

______________________________ Postcode __________ Country ______

Email address ________________________________________________________

Daytime telephone number _____________________________________________

the future network
*Media with passion*
Future Publishing Ltd. Cary Court,
Somerton. Somerset. TA11 6BR

DIRECT Debit

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 e-mail Direct Debit forms.

QUOTE ORDER No. LXFP42

LINUX FORMAT THE UK'S BEST-SELLING LINUX MAGAZINE

Please choose your method of payment **1** or **2**

**1.** □ Direct Debit (UK only – double CD Edition) – I wish to pay £27.45 every 6 months

or □ Direct Debit (UK only – DVD Edition) – I wish to pay £29.95 every 6 months

**Please complete the form on the left**

I understand that I will receive 13 issues during the next 12 months.

**2. Cheque/credit card 13 issues**

| | | |
|---|---|---|
| | **CD** | **DVD** |
| **UK** | □ £54.90 | □ £59.90 |
| **Europe (including Eire)** | □ £78.90 | □ £83.90 |
| **North America** | □ £78.90 | □ £83.90 |
| **Rest of the World** | □ £91.90 | □ £96.90 |

□ Cheque (payable to Linux Format. Sterling cheques drawn on a UK account.)

□ Visa □ MasterCard □ American Express

□ Switch Issue No/Valid date______________________

Card No ______________________________ Expires __________

Signature ______________________________ Date__________

**UK READERS RETURN THIS COUPON TO:**
*Linux Format*, Subscriptions, Future Publishing Ltd, Freepost
BS4900, Somerton, Somerset, TA11 6BR.

**OVERSEAS READERS RETURN (POSTAGE PAYABLE) TO:**
*Linux Format*, Subscriptions, Future Publishing Ltd, Cary Court,
Somerton, Somerset, TA11 6TB, UK

**WEB:** www.linuxformat.co.uk then click on Subscribe

□ Please tick this box if you would prefer not to receive information on other offers.

**ORDER HOTLINE ☎ 0870 444 8645**
**OR FROM *OVERSEAS* ON +44 870 444 8645**

**OFFER ENDS: 16.07.03**

# SAVE £24.47!*

## – THE LEADING LINUX MAGAZINE

13 x

7 PAGES OF EXPERT ADVICE ■ DUAL-XEON SERVER ON TEST ■ STEGANOGRAPHY SECRETS
LINUX FORMAT
www.linuxformat.co.uk
ALL ABOUT X
INSIDE XFREE86 ● X alternatives ● What is X? ● New features ● Future of X ● Drivers and fonts ● Developer spats p44
PROBLEMS? Not any more! Advice from the experts on: ● Hardware ● Debian ● Hotplugging ● Digital cameras ● Filesystems ● Ogg ● and more Page 90
LINUX PRO
Grow your enterprise through open source
USE WINE TO RUN WINDOWS SOFTWARE
Use MS Office, Internet Explorer and even Battlefield 1942 from Linux with WINE p52
SHARQ ATTACK!
HIDDEN IN PLAIN
DISTRIBU

= SAVING OF £24.47*

- **THE ESSENTIAL** Linux read for all Linux users – can you afford to miss out?
- **EVERY ISSUE** comes with the ideal supplement for IT professionals – *Linux Pro* covers trends, technology, security issues, industry comment and thought-provoking case studies
- **GUARANTEED DELIVERY** direct to your door so you never miss an issue!

**RISK-FREE GUARANTEE!**

**You can cancel your subscription at any time and receive a full refund on all unmailed issues, no questions asked**

## SPREAD THE COST OF SUBSCRIBING

For under £15.00* every 3 months you can pay by convenient Direct Debit.
Simply complete the form overleaf and return it back to us today.

**Call now on 0870 444 8645 to make sure your receive the next issue**

* Prices and savings based on DVD edition.

# Essential disc info

Read this important information before you use your *Linux Format* coverdisc – CD or DVD. We've collated some helpful info to help you get the most from these jewels of data!

## Finding the essentials

### Missing something?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases, the latest libraries and other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

### Package formats

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

### Documentation

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

## What are all these files?

If you are new to Linux, you may find the profusion of different files and extensions confusing. As we try to give as many packages as possible for compatibility, there will often be two or three files in a directory covering different types of Linux, different architectures and usually source and binary versions – so which do you install? They can be identified by their filenames, and usually just by the file extensions.

**Someap-1.0.1.i386.rpm** – This is probably a binary rpm, designed to run on x86 systems.

**Someap-1.0.1.i386.deb** – The same, but a debian package.

**Someap-1.0.1.tar.gz** – This is usually source code.

**Someap-1.0.1.tgz** – Same as the above, tgz is abbreviated form of tar.gz

**Someap-1.0.1.tar.bz2** – Same, but uses bzip2 compression instead of zip

**Someap-1.0.1.src.rpm** – This is also source code, but supplied as an rpm to make it easier to install

**Someap-1.0.1.i386.RH7.RPM** – A binary, x86 RPM designed specifically for Red Hat Linux

**Someap-1.0.1.ppc.Suse7.rpm** – A binary RPM designed specifically for SuSE7.x PPC Linux.

**Someap-devel-1.0.1.i386.rpm** – A development version

## Installing from tarballs

A tar ball is a two stage archive. First the files are archived into a single file with tar and then compressed with *Gzip* or *Bzip2*. To unpack, **cd** to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
tar xvf --bzip2 /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in .tar.gz or .tgz, and the second for Bzipped files, ending in .tar.bz2 or .tbz2. Naturally, you change the paths to suit the location and name of the archive. and replace **/mnt/cdrom** with whatever is applicable to your system (*eg* **/cdrom**). This normally unpacks the archive into a directory of the same name, enter that directory with:

```
cd progname-2.1.0
```

To compile and install the software, type the following three commands:

```
./configure
make
su -c "make install"
```

The last line will prompt you for the root password, as this stage must be run as root. If you are already logged in as root, just type **make install**. This will give you a default installation. If you want to change any aspect of the install, type **./configure —help** to see the options available. For example, you are usually able to change the default location with the PREFIX argument. When you have finished installing, you may remove the source files with:

```
cd ..
rm -fr progname-2.1.0
```

You should also log out as root, before you do anything you may later regret.

## DEFECTIVE CDs

In the unlikely event of your CD/DVD being physically damaged we'll send you a new, working version within 28 days. Send your defective disc – complete with your name, address, and a description of the fault – to:

***Linux Format*, Future Publishing Disc Department, 3B Athena Avenue, Elgin Industrial Estate, Swindon, SN2 8HF.**

## Creating install CDs with cdrecord

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord -scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v
 /path/to/image.iso
```

You can simplify the command by saving some default settings in /etc/default/cdrecord. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the "Burn" tab and the "ISO 9660 Image" gadget near the top of the window. Put the path to the image file in the gadget and press "Combust!". Now put on the kettle while the CD is created for you.

### Other OS?

You do not have to use Linux to burn the ISO to a disc. All the Linux-specific bits are already built into the image file. Programs like cdrecord simply dump it to the disk. If you don't have a CD-writer, find someone who does have one, and a DVD drive, and use the CD burning software on their computer. It can be Windows, MacOS, AmigaOS whatever.

### No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions it is also possible to mount the images and do a network install, or even a local install from another disk partition. The methods often vary between distributions, so check on the distro vendors' website for more information. LXF

# Coverdisc

**Neil Bothwick** is your guide through the wonders of this month's jam-packed *Linux Format* DVD. You can even run Morphix straight from the DVD!

Remember – all software listed on the CD pages is available on your *LXF* DVD as well! As regards distros, both CD and DVD readers get Morphix, but DVD readers also get , Freevix, Blueflops, IPCopFirewall and DamnSmallLinux!

## DISTROS
## MORPHIX

The DVD is already set up to boot Morphix directly. Morphix then detects that there are multiple main modules and asks you to choose. The modules available to you include:

**MorphixMainModuleHeavy.mod**
A full desktop environment with a range of Internet and productivity programs, using the GNOME desktop.

**MorphixMainModuleKDE.mod**
As above, but based on KDE instead of GNOME.

**MorphixMainModuleLight.mod**
A lightweight desktop using the IceWM window manager and the ROX file manager. This is small enough to be burned onto a 3" mini-CD.

**MorphixMainModuleGame.mod**
Another IceWM based desktop, this time loaded with plenty of games.

**MorphixMain-bare-0.3-5.mod**
This is really minimal, booting into a console with the basic tools available, for use as a rescue disc.

The Distros/Morphix/docs directory contains PDF files detailing how you can create your own modules, either from scratch or by modifying the existing ones. The tools directory has the programs needed to do this. There is more information on the Morphix website at www.morphix.org.

### On the DVD

Wherever you see this logo it means there's related stuff on the DVD

### IMPORTANT NOTICE

**Before you even put the DVD in your drive, please make sure you read, understand and agree to the following:**

**The *Linux Format* DVD is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker on ANY new software. While every care is taken in the selection, testing and installation of DVD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system which may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all important files. Please read individual licences for usage terms.**

## DESKTOP
## AQUAFUSION

KDE 3.1 comes with several icon themes, including the popular, default Crystal look. But you can't have too many choices when it comes to eye candy, right? Aqua Fusion is a new icon theme, it doesn't look like a slightly modified version of an existing one. Installation is about as simple as it could be. Don't bother unpacking the archive on the DVD, run *KDE Control Center*. Go to LookNFeel>Icons and click the "Install New Theme" button. Then give the path to the Aqua Fusion archive on the DVD and press OK. In a couple of seconds, the Aqua Fusion theme will be added to your list of available themes and you can select and configure it just like any other.

## INTERNET
## GAIM

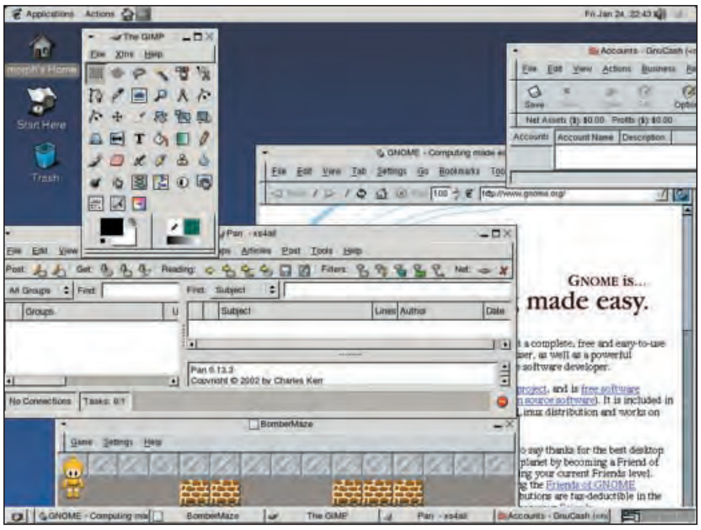

The *Morphix Heavy* module contains the GNOME desktop and a wide range of programs to run on it.

Never was the saying "one of the nice things about standards is that there are so many of them" so true as in the field of instant messaging. It seems that every major Internet organisation is trying to grab its own share of the market and lock out others. There are two ways to deal with this. You could run a separate client for each instant messaging protocol. This may be a good approach if you only use one or two networks, but could get rather cluttered if you need to communicate with people on several networks. The alternative is to use a universal IM program, like *Gaim*.

*Gaim* supports multiple protocols via modules, including AIM, ICQ, Yahoo!, MSN, Jabber, IRC, Napster, Gadu-Gadu and Zephyr. *Gaim* users can log in to multiple accounts on multiple IM networks simultaneously. This means that you can be chatting with friends on *AOL Instant Messenger*,

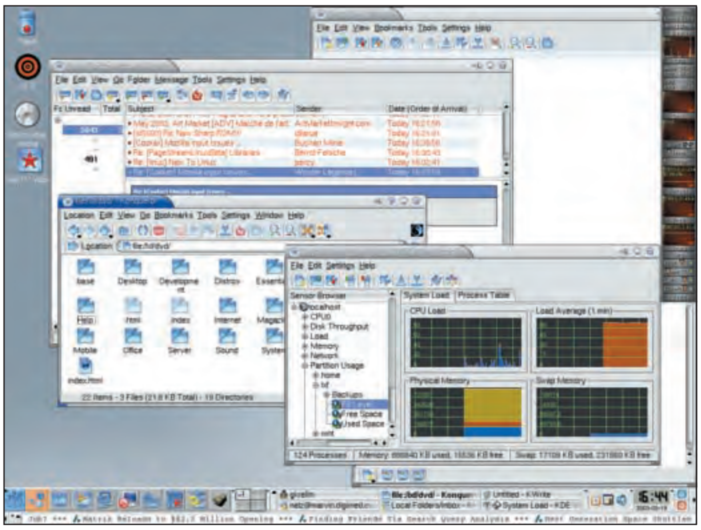

Aqua Fusion – a new and original icon theme for KDE.

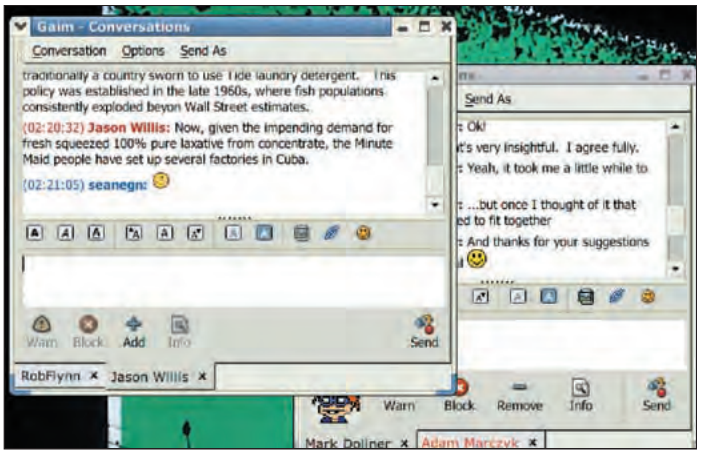

traditionally a country sworn to use Tide laundry detergent. This policy was established in the late 1960s, where fish populations consistently exploded beyon Wall Street estimates.

(02:20:32) Jason Willis: Now, given the impending demand for fresh squeezed 100% pure laxative from concentrate, the Minute Maid people have set up several factories in Cuba.

(02:21:05) seanegn:

: Ok!

t's very insightful. I agree fully.

: Yeah, it took me a little while to

: ...but once I thought of it that
ed to fit together

: And thanks for your suggestions

**Hold multiple conversations on different IM networks in the same tabbed window, with Gaim.**

talking to another friend on *Yahoo Messenger*, and sitting in an IRC channel all at the same time. It will sit in your GNOME or KDE panel until needed, saving screen real estate when you aren't actually chatting. You can have multiple conversations within the same window, using tabs, or drag each conversation into its own window, whichever suits you best.

If you only need to use one network, like *MSN Messenger,* you may feel *Gaim* is overkill. So we have also included *KMess* on the DVD. As it is dedicated to the one protocol, *KMess* is probably better suited to it than a general purpose client. But both programs are there, try them out and take your choice.

## SOUND
## ABCDE

*Grip* has long been my favourite program for ripping tracks from audio CDs and encoding them to Ogg Vorbis or MP3 files. It's GUI makes selecting tracks for encoding, and setting the parameters for that encoding, very easy. However, there are times when you don't want to be messing with a GUI, you just want to rip and encode an entire CD, or several CDs. *ABCDE* is a shell script front-end to the standard CD ripping and encoding tools, tools that you almost certainly have on your distributions CD, if not already on your hard disk. Installation involves unpacking the archive, **cd**ing to the directory created and typing **make install** as root. The program has a huge range of options, but these can all be set in it's configuration files, /etc/abcde.conf for global settings, ~/.abcde.conf for user-specific settings.

With your preferences saved, you can rip and encode an entire CD by typing **abcde**. As each track is ripped, **abcde** gets on with encoding it in the background while ripping the next track. Not only does this speed up the overall process, it also reduces the amount of disk space needed for the temporary files. abcde is also able to normalise the audio levels of the tracks on a CD, this is most useful when ripping a compilation CD where the tracks are from different sources. Output format (Ogg or MP3), file naming, whether to use a proxy server when gathering CDDB information and many for features are all controlled from the configuration file, but abcde works perfectly well with the default settings, you don't need to get your hands dirty until you want to.

All of the features can be controlled from the command line as well as the configuration files, and it is possible to process individual tracks instead of a complete CD, but the real strength of this program is for simple processing of large numbers of CDs and tracks. If you want fine control over the ripping of each CD or track, *Grip* may be a more suitable program, and you will find the latest version of this program on the DVD too.

## SOUND
## KDE-SOUNDS

It's been a while since we've included any desktop themes on the DVD. Here are some alternative desktop sounds for you. As the archives have been put together by different people, there is no consistency between them in terms of the directory structure. It is best to extract each archive to a temporary directory and listen to the sounds using your preferred player, then manually copy the ones you want to use. The sound files should be copied to .kde/share/sounds in your home directory, or to /usr/share/sounds to make them available to all. Now run *KDE Control Center* and select LookNFeel>System notifications to attach any of your new sounds to system events or other programs.

Although these sound themes were put together for KDE, there is nothing KDE-specific about them, they can be used with any desktop environment by copying them to the appropriate place and setting the configuration accordingly. LXF

# How to:
## Create Morphix CD ISO images from the DVD

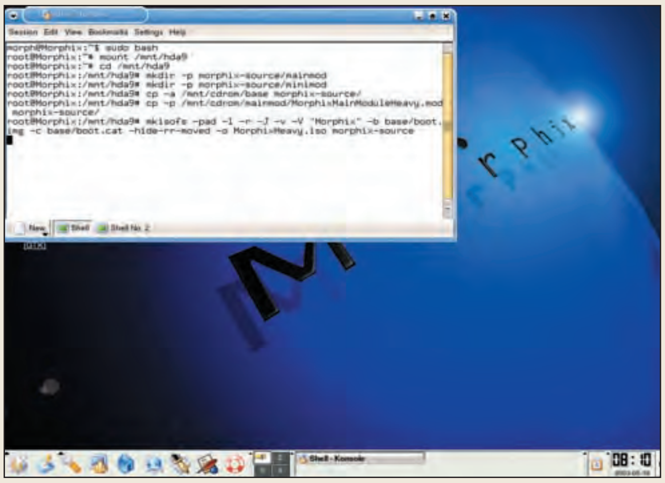

```
morph@Morphix:~$ sudo bash
root@Morphix:~# mount /mnt/hda9
root@Morphix:~# cd /mnt/hda9
root@Morphix:/mnt/hda9# mkdir -p morphix-source/mainmod
root@Morphix:/mnt/hda9# mkdir -p morphix-source/minimod
root@Morphix:/mnt/hda9# cp -a /mnt/cdrom/base morphix-source/
root@Morphix:/mnt/hda9# cp -p /mnt/cdrom/mainmod/MorphixMainModuleHeavy.mod
 morphix-source/
root@Morphix:/mnt/hda9# mkisofs -pad -l -r -J -v -V "Morphix" -b base/boot.
img -c base/boot.cat -hide-rr-moved -o MorphixHeavy.iso morphix-source
```

**Creating Morphix CDs from within Morphix, using *X-CD-Roast* in the *Heavy* module.**

While it is handy to have all the options available on one disc, one of the uses for live CDs is to be able to give them to friends to let them try Linux for themselves. You don't want to give away your *Linux Format* DVD, so here is how to create CD ISO images containing individual modules. You can do this from within any Linux installation that has *mkisofs* installed, but we will do it from Morphix itself. Skip the first stage if you are working from your own Linux setup and just **cd** to the working directory.

We need to mount a hard disk partition to hold both the source files and the ISO image, so it needs to have around 1.5GB free. Morphix detects your hard disk partitions during bootup and creates suitable directories in /mnt and the necessary fstab entries. I'm using /dev/hda9 in this example. Open a terminal window by clicking the relevant button on the taskbar and type

```
sudo bash
mount /mnt/hda9
cd /mnt/hda9
```

The first line switches to root, there is no root password so you can't use **su**. Now create the directories and copy over the base Morphix files and one main module.

```
mkdir -p morphix-source/mainmod
mkdir -p morphix-source/minimod
cp -a /mnt/cdrom/base morphix-source
cp -p /mnt/cdrom/mainmod/MorphixMainModuleHeavy.mod  morphix-
source/mainmod
```

Replace **MorphixMainModuleHeavy.mod** in the last line with whichever module you want to use. You can also copy any mini-modules you need to the minimod directory. See the included documentation for more information. Now build the ISO image with

```
mkisofs -pad -l -r -J -v -V "Morphix" -b base/boot.img -c base/boot.cat -hide-rr-
moved -o MorphixHeavy.iso morphix-source
```

Then you can write the ISO image to a CD, using the methods described on the *Essential Disc Info* on page 107. Alternatively, you can do most of this with a mouse, copying files by drag-and-drop, and building the ISO image using one of the graphical CD burning programs included with Morphix, such as *K3b* in the KDE module or *X-CD-Roast* in the *Heavy* module. The important thing to remember when doing this is that you must make the CD "El-torito" bootable and select base/boot.img as the disk image for booting.

# DVD CONTENTS AT A GLANCE

## Desktop

| | |
|---|---|
| **AquaFusion** | Extensive and elaborate original icon theme for KDE |
| **DeCurs** | An editor for XFree86 mouse cursor fonts |
| **DialogCD** | A dialog-based frontend to cdrtools |
| **Displayusers** | A Web-based directory of users |
| **Door** | A program for managing an encrypted list of passwords |
| **Ethereal** | A GUI network protocol analyzer |
| **GNOMECrontabEditor** | A graphical interface to crontab files |
| **GNUPrivacyGuard** | A PGP replacement tool |
| **GPdf** | A PDF viewer for GNOME |
| **Hackedbox** | A stripped down version of Blackbox |
| **HardwareLister** | Provides detailed hardware configuration information |
| **J-Bird** | A bird observation record keeping system in Java |
| **KdeAm** | A KDE Answering Machine |
| **LyX** | An advanced document processor |
| **MatchboxWindowManager** | A lightweight X11 Window Manager for handheld devices |
| **Morseall** | A Morse code user interface |
| **OpenLDAP** | LDAP suite of applications and development tools |
| **PP3** | A tool to generate celestial maps with LaTeX |
| **Ratpoison** | A window manager that lets you say good-bye to the rodent |
| **Seahorse** | A Gnome GUI for GnuPG |
| **StarDict** | An international dictionary with powerful features |
| **Treebeard** | An XSLT IDE/editor |

## Development

| | |
|---|---|
| **Allegro** | A portable game programming library |
| **CompactDiscLibrary** | A CD-ROM read and control library |
| **Jaim** | A Java library that implements the AOL TOC IM protocol |
| **libtrash** | A shared library that implements a trash can under Linux |
| **SecretRabbitCode** | A library for performing sample rate conversion on audio |
| **Smalltalk** | A Smalltalk interpreter and libraries |
| **TinyCCompiler** | The smallest Linux C compiler |

## Distros

| | |
|---|---|
| **Blueflops** | 2-floppies distro with graphical Web browser and IRC client |
| **DamnSmallLinux** | A 50 MB live CD with a functional desktop |
| **Freevix** | A Linux distribution to build a Freevo media player system |
| **IPCopFirewall** | A Linux firewall distribution |

## Games

| | |
|---|---|
| **ADIC** | A simple, 80s-style cooperative multiplayer networked game |
| **AdvancedStrategicCommand** | Turn-based strategy game in the tradition of Battle Isle |
| **CosmosOfMethos** | A simulation game about Ancient times |
| **DarkPlaces** | A Quake engine port/mod with several enhancements. |
| **Dozo** | The triangles game |
| **FooBillard** | An OpenGL game of pool |
| **Kalyp** | A roguelike game written in Java |
| **Megamek.NET** | A client for the Megamek multiplayer online campaign game |
| **Openglad** | Top-down gauntlet-style RPG |
| **Slune** | A 3D action car game. |

## Graphics

| | |
|---|---|
| **AtomicPhotoAlbum** | A PHP/MySQL photo album |
| **DVD-rip** | A full featured DVD Ripper GUI |
| **Fract-O-Rama** | Qt-based, flexible fractal generation program |
| **Gst-Player** | A featureful media player |
| **Imgv** | A free, cross-platform image viewer |
| **Kandel** | Distributed Mandelbrot set program for KDE |
| **QuickRipDVD** | A simple DVD ripper that focuses on ease of use |
| **SANE** | Applications for SANE |
| **Totem** | A simple movie player for the GNOME desktops |
| **Transcode** | A video stream processing tool |
| **Tulip** | A program for the visualization of huge graphs |
| **XMovie** | A player for high quality movies and sound |

## Help

| | |
|---|---|
| **LDP** | A complete mirror of the Linux Documentation Project |
| **RUTE** | The definitive guide for new Linux users |

## Internet

| | |
|---|---|
| **Amaya** | The browser/authoring tool of the W3C |
| **Epiphany** | A GNOME Web browser |
| **Galeon** | A GNOME Web browser |
| **GKrellMMailwatch** | A GKrellM plugin which monitors mailboxes |
| **KMess** | Chat program for the MSN MessengerÆ protocol |
| **Mozilla** | A beta release of the new Mozilla |
| **NetPisteur** | A Java Web browser |
| **Seclude** | IM service aimed at maintaining freedom of speech |
| **SpeedtouchInstaller** | A  Speed Touch ADSL configuration tool |

## Mobile

| | |
|---|---|
| **CrimsonFields** | Round-based strategy game in the style of Battle Isle |
| **IBAM** | An Intelligent Battery Monitor |
| **MissileCommand** | Zaurus port of the Linux Missile Game |
| **OpenZaurus** | Debian-based embedded distribution built from source |
| **Pico** | A simple, display-oriented text editor |
| **ZaurusSkyExplorer** | Pocket size planetarum for (Personal)java capable devices |
| **Zynergy** | Highly customized ROM image based on Sharp's 2.38 ROM. |

## Office

| | |
|---|---|
| **BananaPos** | A point of sale system |
| **KOffice** | An integrated Office suite for KDE |
| **OpenForBusinessProject** | J2EE-based enterprise automation tools and applications |

## Server

| | |
|---|---|
| **H2ORotisserie** | A structured discussion system |
| **LinkChecker** | An URL link checker |
| **OWX** | Java/JSP-based Website editor or CMS software |
| **PHP** | High-level scripting language |
| **PythonDesktopServer** | A Weblog and news aggregation system |
| **SocratesQuestionnaire** | A Web enabled questionnaire/survey engine and wizard |
| **SourceWell** | Software announcement and retrieval system written in PHP3 |
| **SquirrelMail** | A PHP4 Web-based email reader |
| **VendueAuctionsOnline** | An online auction platform |
| **WebCleaner** | A filtering HTTP proxy |
| **Zope** | A Web application platform for dynamic Web sites |

## Sound

| | |
|---|---|
| **AlsaPlayer** | A PCM (audio) player for Linux/ALSA |
| **Cajun** | Car Audio Jukebox mp3 player for your car/home |
| **GnomeWaveCleaner** | An application to do digital audio restoration |
| **Grip** | A CD player and CD ripper/encoder |
| **JAZZ** | An audio-capable MIDI sequencer |
| **KDE-Sounds** | System sounds for KDE |
| **OggCarton** | A home jukebox system |
| **Rosegarden** | An integrated MIDI sequencer and musical notation editor |
| **TerminatorX** | Realtime Audio Synthesizer (DJ Scratching) |

## System

| | |
|---|---|
| **ArgusMonitoringSystem** | Network and system monitoring software |
| **AWBMTools** | Create full-screen BIOS boot logos |
| **KnutWall** | A powerful iptables firewall for Linux |
| **LinuxNTFS** | NTFS and LDM drivers and tools for Linux |
| **Nmap** | A network exploration tool and security/port scanner |
| **RPMlint** | RPM error checker |
| **Syslinux** | Collection of boot loaders for the Linux operating system |
| **Usermin** | A Web-based interface for UNIX users |
| **Webmin** | A Web-based interface for Unix system administration |

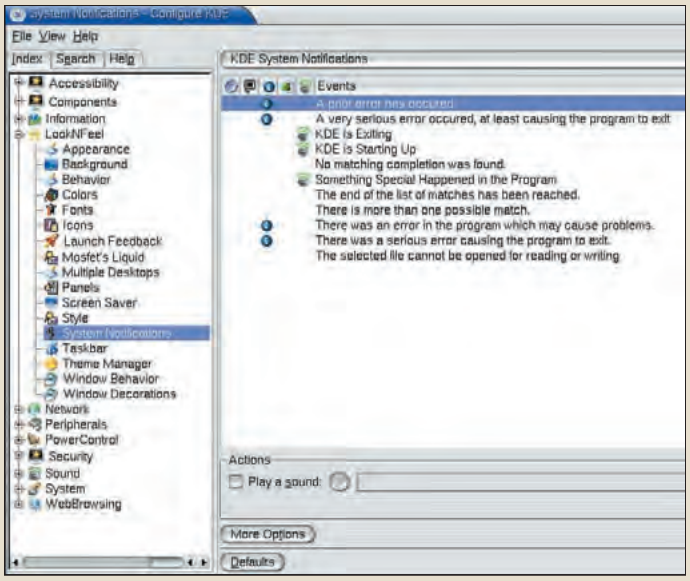

A prior error has occured
A very serious error occured, at least causing the program to exit
KDE is Exiting
KDE is Starting Up
No matching completion was found.
Something Special Happened in the Program
The end of the list of matches has been reached.
There is more than one possible match.
There was an error in the program which may cause problems.
There was a serious error causing the program to exit.
The selected file cannot be opened for reading or writing

□ Play a sound:

**Add new sounds for system notifications and other events, using *KDE Control Center* and the sound themes on the DVD.**

# User Groups

**LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally natter about all things Linux. You can find lots more information online at: www.lug.org.uk**

### 1 HAMPSHIRE
**URL** www.hants.lug.org.uk
**Contact** Hugo Mills

### 2 BRISTOL & BATH
**URL** www.bristol.lug.org.uk

### 3 SCOTTISH
**URL** www.scottish.lug.org.uk

### 4 OXFORD
**URL** www.oxford.lug.org.uk
**Contact** Alasdair G Kergon

### 5 KENT
**URL** www.kent.lug.org.uk
**Contact** Kevin Groves

### 6 BRIGHTON
**URL** www.brighton.lug.org.uk
**Contact** Johnathan Swan

### 7 WORCESTERSHIRE
**URL** www.worcs.lug.org.uk

### 8 NORTHANTS
**URL** www.northants.lug.org.uk
**Contact** Kevin Taylor

### 9 ANGLIAN
**URL** www.anglian.lug.org.uk
**Contact** Martyn Drake

### 10 MILTON KEYNES
**URL** www.mk.lug.org.uk
**Contact** Denny De La Haye

### 11 DONCASTER
**URL** www.doncaster.lug.org.uk
**Contact** Andy Smith

### 12 MORAY
**URL** www.moray.lug.org.uk
**Contact** Stewart Watson

### 13 WEST WALES
**URL** www.westwales.lug.org.uk
**Contact** Dan Field

### 14 WOLVES
**URL** www.wolveslug.org.uk
**Contact** Jono Bacon

### 15 PETERBOROUGH
**URL** www.peterboro.lug.org.uk
**Contact** Steve Gallagher

### 16 EDINBURGH
**URL** www.edinburgh.lug.org.uk
**Contact** Alistair Murray

### 17 TYNESIDE
**URL** www.tyneside.lug.org.uk
**Contact** Brian Ronald

### 18 LEICESTER
**URL** www.leicester.lug.org.uk
**Contact** Clive Jones

### 19 GREATER LONDON
**URL** http://gllug.linux.co.uk/
**Contact** John Southern

### 20 SURREY
**URL** www.surrey.lug.org.uk
**Contact** Jay Bennie

### 21 CAMBRIDGE
**URL** www.cam-lug.org.uk

### 22 DEVON & CORNWALL
**URL** www.dclug.org.uk
**Contact** Simon Waters

### 23 FALKIRK
**URL** www.falkirk.lug.org.uk

### 24 MANCHESTER
**URL** www.manlug.mcc.ac.uk
**Contact** John Heaton, Owen Le Blanc

### 25 HERTFORDSHIRE
**URL** www.herts.lug.org.uk
**Contact** Nicolas Pike

### 26 WEST YORKSHIRE
**URL** www.wylug.lug.org.uk
**Contact** Jim Jackson

### 27 SHEFFIELD
**URL** www.sheflug.co.uk
**Contact** Richard Ibbotson

### 28 STAFFORDSHIRE
**URL** www.staffslug.org.uk

### 29 NORTH EAST
**URL** www.shofar.uklinux.net/NELUG

### 30 LONDON
**URL** www.lonix.org.uk

### 31 BERKSHIRE & THAMES VALLEY
**URL** www.sclug.org.uk

### 32 LIVERPOOL OPENSOURCE
**URL** http://linux.liv.ac.uk/_liv_linux_ug/
**Contact** Simon Hood

### 33 DEAL AMIGA CLUB
**Email** superhighwayman@hotmail.com
**Contact** John Worthington

### 34 CHESTERFIELD
**Email** spirelug@yahoo.co.uk
**Contact** Robin Needham

### 35 SOUTH DERBYSHIRE
**URL** www.sderby.lug.org.uk
**Contact** Dominic Knight

### 36 BELFAST (BLUG)
**URL** www.belfastlinux.cx
**Email** russell@belfastlinux.org

### 37 WILTSHIRE
**URL** www.wiltshire.lug.org.uk
**Contact** Jason Rudgard

### 38 SOUTH LONDON
**URL** www.sl.lug.org.uk
**Email** edo@perceptiondm.com

### 39 CHESHIRE
**URL** www.sc.lug.org.uk
**Contact** Anthony Prime – enquiry@sc.lug.org.uk

### 40 NORTH WALES
REVISED DETAILS
**URL** www.northwales.lug.org.uk
**Contact** Andy Hutchings A-Wing.deltaone@virgin.net

### 41 MIDLANDS
**URL** http://midlandslug.port5.com/
**Contact** Pete Thompson

### 42 CUMBRIA
**URL** www.cumbria.lug.org.uk
**Contact** Jamie Dainton

### 43 DORSET
**URL** www.dorset.lug.org.uk
**Contact** John Robinson

### 44 SHROPSHIRE
**URL** www.shropshire.lug.org.uk
**Email** shropshire@lug.org.uk

### 45 SOUTH WEST
**URL** www.southwest.lug.org.uk
**Email** southwest@lug.org.uk

### 46 SOUTH WALES
**URL** www.swlug.org.uk

### 47 NORTH LONDON
**URL** www.kemputing.net/lug/anlug-aims.html
**Email** jason@voyagercomputers.co.uk

### 48 MALVERN
**URL** www.malvern.lug.org.uk
**Contact** Greg Wright

### 49 HUDDERSFIELD
**URL** www.hud.lug.org.uk
**Contact** Dave Naylor – knocker@caramboo.com

### 50 NOTTINGHAM
**URL** www.nottingham.lug.org.uk

### 51 ST ALBANS & LUTON
**URL** www.lust.lug.org.uk
**Contact** Michael Culverhouse – mike@easily.co.uk

### 52 WREXHAM
**Contact** Paul Kersey-Smith
**Email** paul@pkls.fsnet.co.uk

### 53 PRESTON & LANCS
**URL** www.preston.lug.org.uk
**Contact** Phil Robinson

### 54 DERRY
**URL** www.derry.lug.org.uk

### 55 ISLE OF WIGHT
**URL** www.iow.lug.org.uk
**Contact** David Groom – info@iow.lug.org.uk

### 56 SCARBOROUGH
**URL** www.scarborough.lug.org.uk

### 57 BLACKBURN
**Email** matt@consultmatt.co.uk

### 58 YORK
**URL** www.york.lug.org.uk

### 59 LINCS
**URL** www.lincs.lug.org.uk

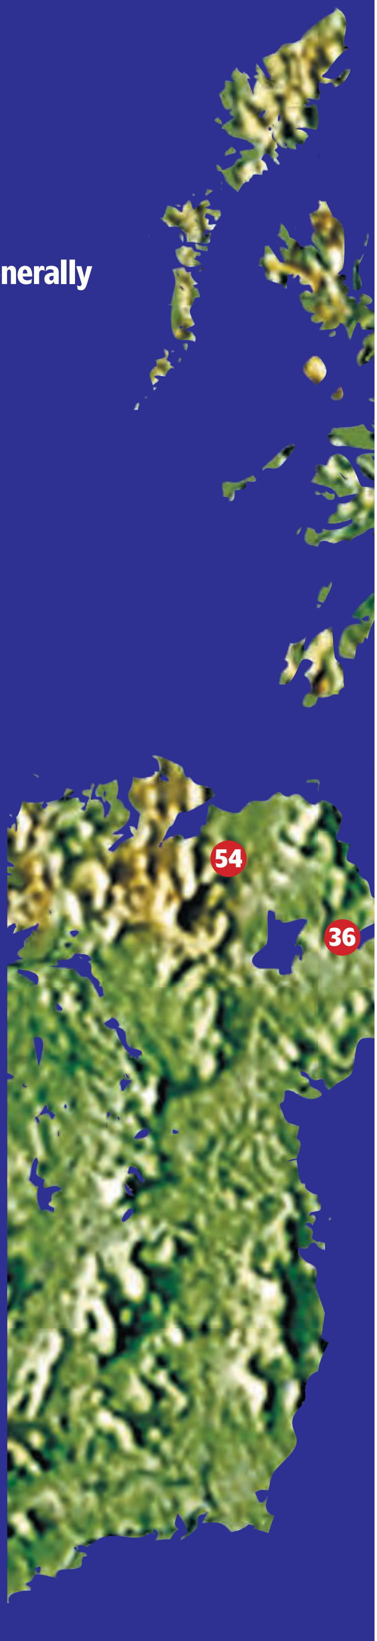

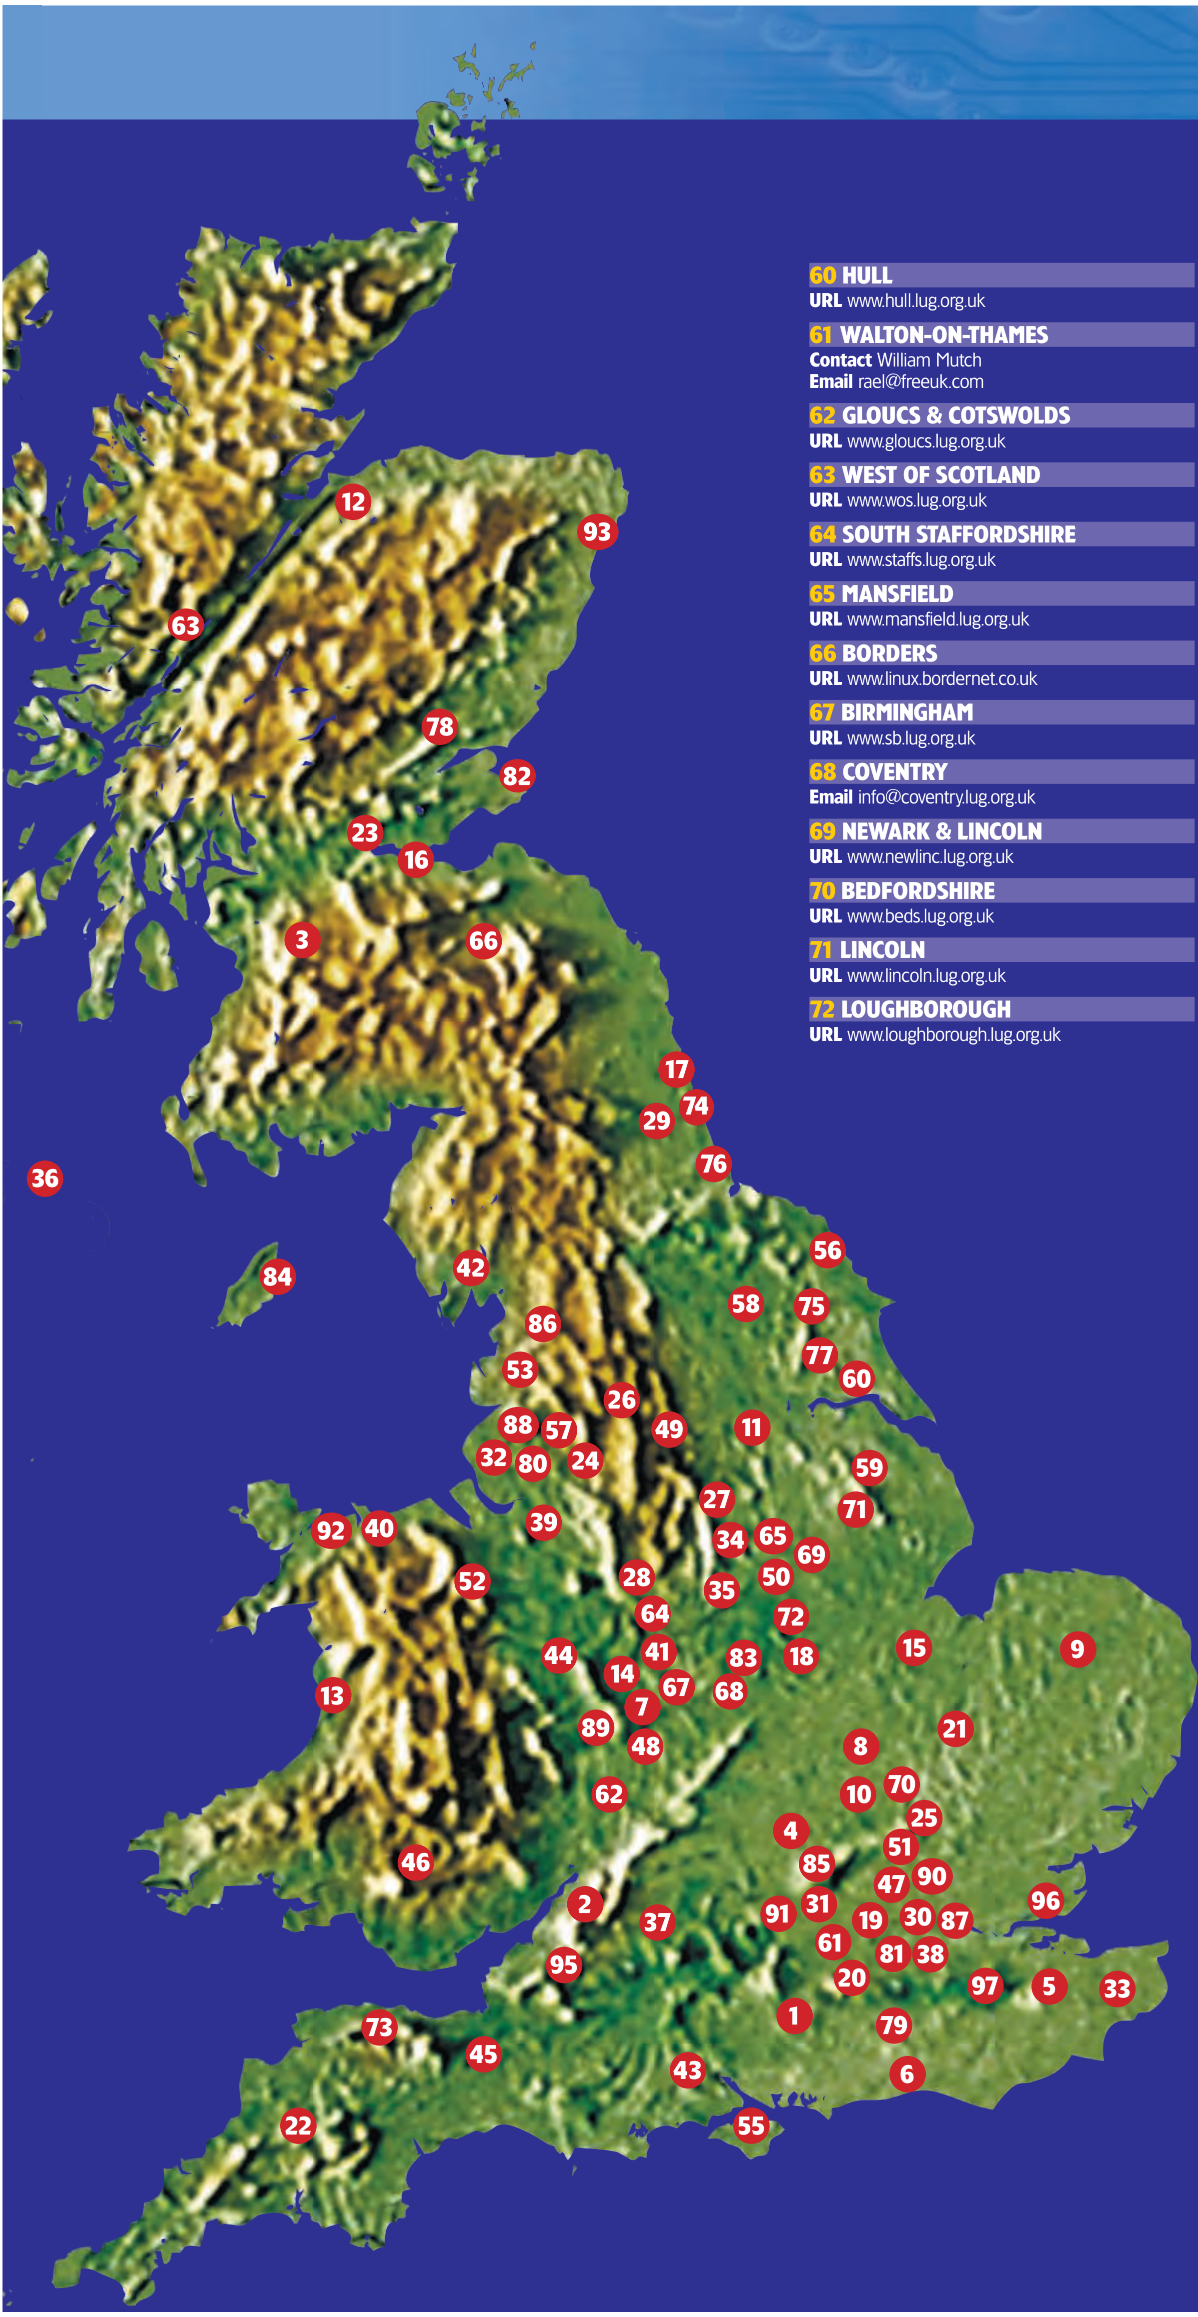

**60 HULL**
**URL** www.hull.lug.org.uk

**61 WALTON-ON-THAMES**
**Contact** William Mutch
**Email** rael@freeuk.com

**62 GLOUCS & COTSWOLDS**
**URL** www.gloucs.lug.org.uk

**63 WEST OF SCOTLAND**
**URL** www.wos.lug.org.uk

**64 SOUTH STAFFORDSHIRE**
**URL** www.staffs.lug.org.uk

**65 MANSFIELD**
**URL** www.mansfield.lug.org.uk

**66 BORDERS**
**URL** www.linux.bordernet.co.uk

**67 BIRMINGHAM**
**URL** www.sb.lug.org.uk

**68 COVENTRY**
**Email** info@coventry.lug.org.uk

**69 NEWARK & LINCOLN**
**URL** www.newlinc.lug.org.uk

**70 BEDFORDSHIRE**
**URL** www.beds.lug.org.uk

**71 LINCOLN**
**URL** www.lincoln.lug.org.uk

**72 LOUGHBOROUGH**
**URL** www.loughborough.lug.org.uk

**73 EXETER UNIVERSITY**
**URL** www.euslug.lug.org.uk
**Email** N.J.Murison@exeter.ac.uk

**74 SUNDERLAND**
**Email** thomas.croucher@sunderland.ac.uk

**75 EAST YORKSHIRE**
**Email** sharkonline@whatemail.com

**76 CLEVELAND OPEN SOURCE GROUP**
**Email** openlug@digitalmedia.co.uk

**77 BEVERLEY**
**Email** vladimir_lukyanov@hotmail.com

**78 DUNDEE & TAYSIDE**
**URL** www.dundee.lug.org.uk

**79 SUSSEX**
**URL** http://sussex.lug.org.uk/

**80 WIGAN & ST HELENS**
**Email** paulf.johnson@ukonline.co.uk

**81 BRIXTON**
**URL** www.communitytechnology.org.uk/~linuxhome

**82 ST.ANDREWS, FIFE**
**URL** www.standrews.lug.org.uk
**Email** stuart@nx14.com

**83 NUNEATON**
**URL** www.nuneaton.lug.org.uk

**84 ISLE OF MAN**
**URL** www.iom.lug.org.uk
**Email** helix@manx.net

**85 AYLESBURY**
**URL** www.aylesbury.lug.org.uk
**Email** drbond@educational-computing.co.uk

**86 LANCASHIRE**
**URL** www.lancasterlug.org.uk

**87 EAST LONDON**
**URL** www.eastlondon.lug.org.uk
**Contact** Jonathan Spriggs

**88 ORMSKIRK**
**Email** rob@northwestlinux.co.uk

**89 HEREFORD**
**URL** www.hereford.lug.org.uk/
**Email** rbjh@good-news.fsnet.co.uk

**90 EAST HERTS**
**Email** madtom1999@yahoo.com

**91 SWINDON**
**Email** nick.trueman@ntlword.com

**92 MENAI**
**URL** www.menai.lug.org.uk

**93 ABERDEEN**
**Email** mark@arricc.net

**94 SHETLAND**
**URL** www.shetland.lug.org.uk
**Email** c_s_s_butler@yahoo.com

**YOUNG LINUX**
**URL** www.young.lug.org.uk

**SCHOOLS**
**URL** www.schools.lug.org.uk

**NEW GROUPS:**

**95 GLASTONBURY**
**URL** www.glastonbury.lug.org.uk
**Contact** Steve Leonard-Clarke

**96 SOUTHEND-ON-SEA**
**URL** www.sos.lug.org.uk
**Contact** Derek Shaw

**97 ORPINGTON**
**URL** www.orpington.lug.org.uk
**Contact** Barry Schofield

LUG OF THE MONTH

# Scunthorpe & Doncaster Group of Linux Users – SCUNDOG

**Shaun Holt writes:**
This group started out last year and finally set their website up in March 2003. Thanks to the mutual support of other LUGs the group is slowly expanding and sharing information for a growing list of members.

The scundog ideal is to promote the use of Linux for new users who need answers to their Linux problems. By the same token people can also trumpet things they find great/easy to use. So far the group has attracted a wide age group from backgrounds including industry, local government, education, job seekers and students. The first meeting is scheduled to take place as soon as a venue is finalised in the North Lincs Area (easy to get to by public transport is a primary consideration). The group has been working on a demo of the LTSP (Linux Terminal Server Project) and dual booting Windows XP laptops with Red Hat and Mandrake distros. Recent attention has been paid to migrating office users over to *OpenOffice.org/StarOffice v6*.

We welcome allcomers and hope that the group will grow into a truly useful and authoritative source of information for people seeking Linux help. Victoria Hardwick is the site webmistress and she can be contacted by email **victoria@scundog.org** or contacted at work (our boss is pretty OK!) on 0870 241 6064.

**www.scundog.org**

## Worldwide Linux User Groups

Free Software users across the globe

### Africa

**EGYPT**
**URL** www.linux-egypt.org

**GAUTENG, SOUTH AFRICA**
**URL** www.glug.org.za
**Email** glugmin@revolution.org.za

**THE LORD'S ABODE, JO'BURG, SA**
**Email** Andrew Gargan avrin17@iname.com

### Australia

**ADELAIDE**
**URL** www.linuxsa.org.au
**Email** mtippet@anu.edu.au

**ALICE SPRINGS**
**URL** www.aslug.org.au

**MELBOURNE, VICTORIA**
**URL** www.luv.asn.au
**Contact** luv-committee@luv.asn.au.

**PERTH**
**URL** http://plug.linux.org.au/

**SYDNEY**
**URL** www.slug.org.au

### Europe

**COSTA DEL SOL (English speaking)**
**URL** www.fuengirola.lug.org.uk

**DENMARK**
**Alssund** www.alslug.dk
**Esbjerg** www.eslug.dk
**Fyns** www.flug.dk
**Midt-og Vestjylland** www.mvjlug.dk
**Nordjylland** www.njlug.dk
**Skåne Sjælland** www.sslug.dk
**Trekantsområdet** www.tlug.dk
**Vest-fyn** www.haarby-net.dk/vflug
**Århus** www.aalug.dk

**EIRE**
**URL** www.linux.ie
**Email** root@linux.ie
**URL** www.dilu.org
**Contact** glossary@dilu.org

**MiLUG (Longford)**
**URL** http://midlands.linux.ie
**Contact** midlands@linux.ie

### Middle East

**ISRAEL**
**URL** www.iglu.org.il/IGLU/
**Contact** webmaster@iglu.org.il

**PALESTINE**
**URL** www.lugps.org
**Email** isam@planet.edu

### Asia

**HONG KONG (multilingual)**
**URL** www.linux.org.hk

**SINGAPORE – SLUG**
**URL** www.lugs.org.sg

**SRI LANKA**
**URL** www.lklug.pdn.ac.lk

**MYANMAR (formerly BURMA)**
**URL** www.myanmarlug.org
**Email** aftyde@balug.org

**PAKISTAN**
**URL** www.linuxpakistan.net
**Email** tux@cllug.org

**HYDERABAD, SINDH, INDUS VALLEY**
**URL** www.geocities.com/slug_pk/

**KASHMIR**
Coming soon!

### China

**BEIJING (GB encoding, but mostly written in Chinese)**
**URL** http://mud.263.net.cn/~linux

**CHINESE LINUX USER GROUP**
**URL** www.linux.org.cn

**NANJING**
**URL** http://jllib.jlonline.com/njlug

### India

**LINUX INDIA**
**URL** http://linux-india.org

**ALIGARH LUG**
**URL** http://linux.amupost.com

**BOMBAY**
**URL** www.ilug-bom.org.in

**CHANDIGARH**
**URL** www.geocities.com/vipinb

**CHENNAI AND MADRAS**
**URL** www.chennailug.org/

**CYBERABAD (CLUG)**
**URL** http://seeknew.freeservers.com/clug/

**DELHI**
**URL** www.linux-delhi.org

**KOLKATA**
**URL** www.ilug-cal.org

**MADURI**
**URL** http://linuxmadurai.tripod.com

**NORTHERN INDIA LINUX**
**URL** http://groups.yahoo.com/group/lug-northindia

# Spreading the word

Linux use in schools isn't only a good idea in terms of cost, it's our most important advocacy asset, says **Jono Bacon.**

**Linux is free. The concept is free,** the software is free and community is free. Not only is Linux free, but it is low cost, and can therefore be of great use to a variety of organisations. This month we will be looking at schools.

Schools are a great opportunity for Linux advocacy. They need more resources, they have little money and free software sounds like the perfect solution. Although clear cut, the practical story is different. Schools get their funding from Local Educational Authorities (LEAs) and the choice of upgrading is difficult. I did some research last year into schools and Operating Systems, and many IT managers and had heard of Linux and were interested, but the migration was difficult.

When advocating Linux to schools, it is important to be aware of what problems the schools are facing, and develop your discussions and materials towards these problems and how Linux can be a solution. It is important to remember that migration is a big deal for schools. Most of the staff will be trained in Windows and they would need to be retrained. There is also the issue of migrating software and documents to the new system. This is where commercial companies such as Microsoft have the hold on companies – proprietary formats such as .doc files limit practically every document to *Microsoft Word*. This concept is called Vendor Lock-in; the vendor's closed technology essentially keeps the users in the upgrade cycle – this may not be a publicly stated intention, but in real life, this happens a lot and is a concern of IT managers everywhere.

Next month I will continue to look at schools as a source of promoting Linux. There are so many issues involved in schools that we will need a few issues to discuss the topic properly. In the meantime, make sure you go to the *LXF* forums at **www.linuxformat.co.uk** and let me know your thoughts on advocacy and any experience and ideas that you have. I would like to share your thoughts and opinions in this little box as well! LXF

## Linux User Group organisers

If you're not listed here, or we have your details wrong, please contact us at: **LUGS!, *Linux Format*, 30 Monmouth Street, Bath, BA1 2BW**
or email your details to: **linuxformat@futurenet.co.uk**

**EDITORIAL**
**Editor** Nick Veitch *nick.veitch@futurenet.co.uk*
**Reviews Editor** Paul Hudson *paul.hudson@futurenet.co.uk*
**Art Editor** Julian Jefferson *julian.jefferson@futurenet.co.uk*
**Production Editor** Matt Nailon *matthew.nailon@futurenet.co.uk*

**Editorial Contributors** Jono Bacon, Neil Bothwick, David Cartwright, Andy Channelle, Amias Channer, David Coulson, Richard Drummond, Hoyt Duff, Michael J Hammel, Hans Huberland, Patrick O'Brien, Mike Saunders

**COVER CD PRODUCTION**
**CD Editor** Neil Bothwick

**ART CONTRIBUTORS**
**Art assistance** Martin Mulchinock
**Photography** Amanda Thomas, Gavin Roberts, Getty Images, Powerstock, Corbis, Photodisc, Eyewire
**Illustration** Paul Bateman, Chris Winn, Shane Collinge

**ADVERTISING SALES**
**Key account manager** George Gill george.gill@futurenet.co.uk
**Business development manager** Damian Hughes
**Senior Sales Executives** Giles Crosthwaite-Scott, Harry Maltby
**Classified Sales Executive** Stephen Hall

**MARKETING AND PROMOTIONS**
**Marketing Manager** Mike Hawkins
**Subscriptions Executive** Marie Spicer

**PRODUCTION**
**Production Co-ordinator** Diane Ross
**Group Production Manager** Clare Tovey
**Advertising Production Co-ordinator** Jo Crosby

**MANAGEMENT**
**Publisher** Kelley Corten
**Group Publisher** Dave Taylor
**Publishing Director** John Weir
**Managing Director** Colin Morrison

**DISTRIBUTION AND CIRCULATION**
**Circulation Manager** Jamie Malley
*jamie.malley@futurenet.co.uk*

**Distributed by** Seymour Distribution, 86 Newman Street, London W1T 3EX **Tel** +44 (0)207 3968000
**Overseas Distribution** by Future Publishing Ltd.
Tel 01225 442244.

**Overseas Licences**
**International Licensing Director** Simon Wear
*simon.wear@futurenet.co.uk* **Tel** +44(0)1225 822798
**Fax** +44 (0)1225 788105 **Cell phone** +44(0)777 5641493

**Contact Details**
Linux Format, 30 Monmouth Street, Bath BA1 2BW
**Tel** +44 (0)1225 442244 **Email** *linuxformat@futurenet.co.uk*

**Subscriptions and Mail Order**
**Phone** +44 (0)870 444 8645 See page 98
**Email** *linuxformat.subs@futurenet.co.uk*

**All CD-ROM** demos and reader submissions are supplied to us on the assumption they can be incorporated into a future covermounted CD-ROM, unless expressly stated to the contrary.

**Disclaimer** All tips in this magazine are used at your own risk. We accept no liability for any loss of data or damage to your computer, peripherals or software through the use of any tips or advice.

**Printed** in the UK by Midway Clark (Holt) and TPL
© Future Publishing Ltd 2003

**Future Publishing Ltd. is part of The Future Network plc.**

The Future Network produces carefully targeted specialist magazines for people who share a passion. We aim to satisfy that passion by creating titles offering value for money, reliable information, smart buying advice and which are a pleasure to read. Today we publish more than 90 magazines in the UK, US, France and Italy. Over 80 international editions of our magazines are also published in 28 other countries across the world.
The Future Network plc is a public company quoted on the London Stock Exchange (symbol: FNET).
**Non-executive Chairman** Roger Parry
**Chief Executive** Greg Ingham
**Chief Operating Officer & Managing Director, UK** Colin Morrison
**Group Finance Director** John Bowman

**Tel** +44 (0)1225 442244

BATH • LONDON • MILAN •
NEW YORK • PARIS •
SAN FRANCISCO • WROCLAW
**www.thefuturenetwork.plc.uk**

*Media with passion*

# NEXT MONTH

## Issue 43 on sale Wednesday 16 July

# NETWORK SPECIAL

Be the master of your own interconnections – from a simple home network to a multi-zone behemoth, we're unravelling Ethernet from HBAs to routers. Plus tutorials on Samba and WiFi networking.

## Opteron on test

Exclusive reviews of a brace of new 64-bit Opteron servers – we've seen the theory in issue 41 of *Linux Format*, but how will they perform in our real-world tests?

## Burn time

You may have your old favourites, but which disc burning solution is the best? We round up the front runners

## PLUS

**DON'T MISS NEXT MONTH'S LINUX PRO!**
Training and certification – what you need to know
Ethernet bridging, Flex, storage, interviews and more!

LINUXPRO
FROM THE MAKERS OF LINUX FORMAT
Grow your enterprise through open source

**DON'T MISS YOUR COPY OF *LINUX FORMAT!***
Get it delivered to your door every month – subscribe on page 98

The exact contents of future issues are subject to change

# LINUXPRO

FROM THE MAKERS OF LINUX FORMAT

JULY 2003

ONLY AVAILABLE WITH LINUX FORMAT MAGAZINE

## Grow your enterprise through open source

Novell: *"freeing users from their proprietary shackles"* – the new Netware strategy **p4**

Oracle: *"We see the growth in Linux business as being fundamental to the Oracle economy"* – unbreakable Linux **p14**

**PLUS**

### LICENSING

Software licences explored and explained – are you using the right one?

### FLEX AND BISON

Exploring typeless variables and language structure

### STORAGE

The NFS network filesystem is far from obsolete – how you can use it and what the future holds

### SYSTEM MONITORING

NetSaint is dead, long live Nagios, the new king of open source monitoring

# Welcome

## Twenty-four pages of real-world Linux for IT professionals

Many, many arguments have been had about what actually constitutes 'enterprise' computing and whether, whatever it is, Linux is 'ready.' I've witnessed enough of these spats first-hand to realise that as soon as one point on either side is lost, the matter soon becomes irrelevant, because the real stumbling block becomes something else. A few years back Linux wasn't ready for the enterprise because it didn't scale well. New technologies and processors have proved that to be inaccurate. Then Linux wasn't ready because the applications weren't there.

With companies like Oracle and IBM, not to mention plenty of smaller organisations providing business, financial and management apps on Linux, the argument there is pretty much over too. It may seem surprising , even to Linux stalwarts, that the growth of the OS and its impact in new markets seems irrepressible. But when IT legends like HP, Oracle, IBM, Novell and others seems increasingly to be basing their business (and this is 'enterprise' business we are talking about) around Linux, the proposition that Linux will become a major force in the data centre as well as the web server market seems like a foregone conclusion. This month we have an exclusive interview with Oracle on their 'Unbreakable Linux' technology, and a peek behind the scenes of Novell, who seem to be wholeheartedly committing to a Linux-based future. Interesting reading.

We've also taken a fresh look at licensing. GPL is not the only flavour of 'Open Source' license, but beware of the misleading titles of some other licensing schemes, which may not be as free as they seem. Check it out on page 17.

We've also had space to continue our *Flex* and *Bison* series, take a look at system monitoring with NetSaint's successor, Nagios, and our regular storage slot expounds the virtues of NFS and where future development may take it. Enjoy!

**Nick Veitch** Editor
nick.veitch@futurenet.co.uk

> "With this support, the proposition that Linux will become a major force in the data centre market seems like a foregone conclusion"

# CONTENTS

**NOVELL**
NetWare and open source **p4**

**FLEX/BISON**
Creating a new language **p8**

**ORACLE**
Unbreakable Linux explained **p14**

**LICENCE TO CODE**
Which license means what? **p17**

**STORAGE**
Don't overlook NFS **p20**

**SYSTEM MONITORING**
Nagios inherits the legacy **p22**

## 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*

Group publisher
**Dave Taylor**
Publisher
**Kelley Corten**
Contributor
**Andy Channelle**

Photography
**Powerstock, Photodisc, Corbis, Getty Images**

**ADVERTISING**

Key Account Manager
**George Gill** *george.gill@ futurenet.co.uk*
Business Development Manager
**Damian Hughes**
*damian..hughes@ futurenet.co.uk*
Senior Sales Executives
**Giles Crosthwaite-Scott, Harry Maltby**
Classified Sales Executive
**Stephen Hall**

LinuxPro is a part of
LINUX FORMAT
magazine

Published by
the future network
Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

# Novell gambles on growth

**After years of dedication to the proprietary crusade, networking pioneer Novell has announced a dramatic shift in emphasis. ANDY CHANNELLE discovers what caused this Damascene conversion and its implications for existing and potential NetWare users.**

Presently, it wouldn't be unfair to suggest that Novell's star is on the wane. From a position of strength in the 1990s, the company has seen its foundations undermined and a once loyal user base diminished – though still hovering around the 90 million mark – thanks to the perception that its products have failed to move with the times. When you factor in the surge of interest in Linux and the reaction the Open Source development model is inspiring in Microsoft – among others – it is no surprise that things at Novell had to change.

The first public demonstration of this new attitude occurred at Novell's annual BrainShare convention where NetWare users and developers have gathered for the past 18 years to receive the wisdom of Novell's CEO and product managers. The 2003 event contained the usual product hyperbole and jokes at Steve Ballmer's expense, but also a genuinely significant strategy change for a company that has persisted in pushing proprietary technologies long after, as CEO Jack Messman candidly admitted, their customers had given up on them. In his opening keynote address Messman said it was time to move on, citing Novell's dogged determination to stick with IPX while the world opted for TCP/IP.

"We have always focused on interoperability to give our customers flexibility, but we've stuck with a few proprietary standards too long." The company, he said, has learned its lesson. "We're not going to make those mistakes again. We will embrace the Open Source movement and we will make significant contributions to it."

For the 6.5 release of NetWare, Messman announced a raft of changes to make the product more Open Source-friendly, including the addition of a GroupWise client for Linux and support for technologies such as *Apache*, MySQL and Perl. These extensions, he said, would "free users from their proprietary shackles." It would help companies realise the true value of their information and share it across platforms, devices and architectures in a secure manner.

The new openness of NetWare 6.5, though, is just a stepping stone to a more unequivocal embrace of Open Source due in the next iteration.

"What is NetWare's future migration path? What is Novell doing to protect its customers' investments in NetWare? What about migration to other platforms?" Messman asked the auditorium. "We have services that sit on top of the NetWare kernel that we've been developing for years. Services such as file and print, messaging, resource management, web development services and many others. The current opportunity for Novell regarding Linux is to provide these kinds of services for Linux. We know a lot of you are thinking about Linux and we're going to support you in making Linux a viable data centre platform."

To rousing applause he announced: "In NetWare 7.0, the services that you've come to know and love will sit on both the NetWare and Linux kernels." A dramatic shift. And while Messman was anxious to press home that Novell were not "dropping NetWare, we're adding Linux!" the significance of such a move should not be underestimated.

The only sour note sounded at the BrainShare conference was the fact that it had taken so long for Novell to make this leap. Meta Group analyst Earl Perkins said the timing could be problematic. "The big problem is we are looking at 18 months before they will have significant services working on Linux. It would have been better if they made this decision a little bit sooner."

While applauding the decision to 'adopt' Linux, Peter Strifas, who is responsible for 60 NetWare servers at a New York University hospital, told ComputerWorld.com that it would have been better if Novell had simply put everything on Linux. "The biggest problem we've been having lately is facing hardware vendors that don't have NetWare support. For example, HP doesn't support NetWare on its server blades. That was just enough for management to ask why we were keeping this stuff around."

## The customer is king

Novell spokesman Bruce Lowry told Linux Pro that the conversion was the result of what Messman called a change from a lab-focused approach to one based on what customers wanted.

"We'll do what our customers want, which at the moment means developing services on NetWare and Linux. The next version of NetWare will run on both kernels, providing choice." For the foreseeable future Novell would be a dual-platform operation, and the company would not abandon its loyal users, or its investment in NetWare. "NetWare is a very powerful kernel. It provides features for stability and scalability that other operating systems don't currently match. So we're confident that we'll have a large customer base that continues to buy and deploy NetWare. It all comes down to fulfiling customer needs."

Lowry said that while there was a lot of debate within the company about the best way to embrace Linux and the Open Source development community, there was in fact no real discussion about whether it should be done. It was inevitable.

"Novell is embracing Linux because our customers are asking us to do so. Most of our customers have multiple platforms in their environments: NetWare, Windows, Solaris, and, increasingly, Linux. Analyst data clearly shows Linux as the fastest growing platform on the market." But Novell's own polling suggested that many users still had concerns about deploying Linux in 'mission critical' environments. "Our customers know Novell for providing unparalleled security, scalability and reliability in networking services. They need these same capabilities to be able to use Linux in the data centre confidently," he said. "So that's what we're going to deliver."

> "Novell is listening to its customers, who are wary of being locked into proprietary environments, and we're responding to what they want."
>
> CHRIS STONE – NOVELL

So far, the response from users has been positive. Bruce Lowry: "Customers like this freedom to choose. They can stay with NetWare, move to Linux, or deploy both, all while maintaining their Novell environments. Some customers have said they don't plan on tapping Novell's Linux offerings because they're happy with NetWare. Others have said they are eager to see what we deliver."

As a sign of its commitment to Open Source, Novell coupled the launch of a new developer website, Novell Forge, with the donation of "the first industrial-strength, open source UDDI server." Chris Stone,

### "LINUX IS IMMATURE"

**That controversy in full!**

**AFTER THE BIG ANNOUNCEMENT, Novell's ebullient CEO Jack Messman gave a few exclusive interviews. First up was Computerworld.com where he told Don Tennant that "Linux is an immature operating system right now. It hasn't had somebody like Novell worrying about making it robust, reliable and scalable for very much time."**

**He then moved on to PC Magazine where he told Darren Sarrel that the move to Linux would "give customers breathing room and comfort because Linux is not mature and Novell can bring 20 years of experience to Linux."**

**Oops! Cue enormous flame-fest on Slashdot (and elsewhere), a swift back-peddle by Novell and an open letter to the Open Source community in which Messman said he regretted the remarks. "Clearly Novell wouldn't be taking this bold step if we didn't feel Linux was a solid operating system with tremendous momentum in the marketplace. In fact, we believe Linux is quite stable and scalable. If we didn't, we would not commit to using it with our NetWare 7.0 release," he wrote. The response was widely regarded as 'a good save' by a company entering a new area and attempting, in the words of Bruce Perens, to "bring two widely disparate cultures together."**

company vice chairman, said this was simply a case of give and take: "The launch of NetWare 6.5, which makes significant use of Open Source components like *Apache, Tomcat,* MySQL and PHP/Perl, is a milestone in Novell's commitment to the community, so it only makes sense that we also increase our commitment to giving back, with the launch of Novell Forge and the Nsure UDDI server."

"There's no way to efficiently manage a network of any size without Novell at the heart of it."

JACK MESSMAN – NOVELL CEO

Bruce Lowry said that beyond the decision to donate such a fundamental web services application, it was also important to work within existing licenses, rather than following the common route of creating a whole new licensing regime. "Novell has no intention of adding to the plethora of Open Source licenses – as evidenced by our use of the BSD license for the UDDI Server release. Existing licenses will serve our needs just fine."

As well as consulting OSS luminaries such as Bruce Perens in the run up to the announcement, Novell enlisted the help of MySQL AB whose database will be bundled with NetWare 6.5. Perens said that as MySQL uses the same codebase for its commercial and GPL versions, Novell's license payments would be directly supporting the continued development of the GPL product. "Now, Novell is taking the next step, by producing its own free software under OSI-approved licenses, including a UDDI implementation and other facilities," he said. "They have taken some care to act as good citizens of the open source developer community." Marten Mickos, MySQL AB Chief Executive Officer, said they had been working with Novell for some time in both an engineering and business capacity, and the move away from proprietary standards by such a large proportion of the network sector would provide significant benefits for all involved. For MySQL and the Linux development community as a whole, Mickos said, the bundling of the database would offer some serious exposure to a new audience. "There will be an opportunity to cross-breed applications between the 90 million Novell users and the 4 million MySQL installations worldwide." More specifically, for the first time Novell users will be able to web-enable their systems and build powerful apps built around Novell, Apache, MySQL and PHP/Perl/Python (NAMP) that leverage their existing infrastructure.

## New platform

"The Linux and Open Source development communities will get a new platform to deploy their applications on. In essence, a new market of corporate customers is opening up to them." It all comes down to providing choice, and the expansion of Open Source development into the enterprise sector is an "ongoing and unstoppable trend."

"The enterprise sector is practically begging us to serve them and asking us to serve them in the platform environments they are used to. That's why it makes so much sense for Novell to bundle MySQL, and for MySQL to expand into the NetWare platform."

The version of MySQL that will ship in the NetWare 6.5 box, while being identical to the GPL edition, is governed by a commercial license which, according to MySQL co-founder David Axmark, allows developers who buy into the system to extend the database application without the requirement to redistribute their changes. Axmark told Newsforge.com that it was fortunate MySQL's push into enterprise coincided so well with Novell's desire to "get deeper into the Open Source community." The tentative links were soon bolstered, with Novell sending staff to the last MySQL conference, and doing

## NETWHAT?

**An explanation for the uninitiated**

**NETWARE IS A NETWORK-CENTRIC operating system providing transparent access to services such as file, print, directories, email and databases. Over the years, the core of the system has evolved to incorporate open standards and innovative features such as iFolder and iPrint. The former is a network-based storage solution developed to make accessing and, more importantly, synchronising files across a number of devices simple and secure. iPrint, meanwhile, attempts to provide seamless printing across any network, including the Internet, using a combination of Secure Sockets Layer (SSL) encryption for security and the Internet Printing Protocol (IPP) for management.**

**The most recent version of NetWare features a web-based management system instead of relying on the proprietary interface of the past. For integration purposes it is able to deal with Appletalk Filing Protocol (AFP), Network Filesystem (NFS) and Common Internet File System (CIFS), all available over TCP/IP, for file serving.**

**The next version (6.5) will integrate both MySQL and Apache as well as new implementations of Perl and PHP to give users a complete web services solution.**

**Prices start at US$995 for a 5-user, strong encryption license.**

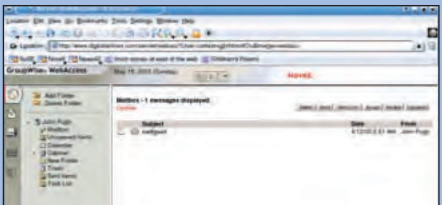

**NetWare's web-based UI can be used by any web-browser, including Konqueror.**

**IN A NUTSHELL...**

- **NetWare 6.5 will include *Apache*, MySQL, Perl/PHP/Python.**
- **NetWare 7 (still 18 months+ away) will be a set of services running on either NetWare or Linux kernels.**
- **Linux and Mac GroupWise clients (Java based) will be included with GroupWise 6.5.**
- **GroupWare Server on Linux planned for later release.**
- **Donation of code to Nsure UDDI Server.**
- **Novell's Open Source efforts will be accessible via Novell Forge website.**
- **Novell will be committing serious development effort to other Open Source projects.**
- **Certified Linux Engineer programme to recognise IT professionals capable of working across Novell and Linux platforms.**

'most of the work' to bring the system to NetWare, while MySQL took their presentation booth to BrainShare to give delegates a first-hand look at what they may soon be using.

While the release of NetWare 7 on Linux is scheduled within the standard 18-24 month Novell release cycle, the company was keen to 'set out their stall' and build on their – often overlooked – Open Source efforts of the past. "Having contributed for years to Open Source community projects such as *Apache* and *OpenLDAP*, Novell is today creating a more direct path to its OSS initiatives," the Novell Forge announcement said. Their first major Open Source code donation is the Nsure UDDI (Universal Description, Discovery and Integration) Server. UDDI is regarded as one of the core web services technologies, and Nsure brings simple and secure identity management to the standard. The project will be managed at Novell Forge, where developers can browse or download the source code for various platforms, discuss development issues and submit changes. The site will also host Novell's other OSS efforts including their NetWare port of the *Fast Light ToolKit* (*FLTK*) and *Digital Hospital*, an Open Source, platform independent version of the J2EE application used to highlight the effectiveness of *exteNd Composer* and *Director* for web services development. Project administrator Harish Krishnaswamy told us the UDDI project had received about 500 unique downloads in its first month.

And we can expect another "big" announcement soon.

"Our plan is to release open source projects related to technologies that are strategic to Novell," Bruce Lowry said. "Open-sourcing the UDDI Server – which builds on Novell's directory and identity management strengths – underscores this." He added that future Open Source projects would continue to utilise the company's strengths. Moreover, these wouldn't just be internal operations: "Novell will assign engineers to work on external projects as well that dovetail with this approach. We are currently assessing which projects merit our attention."

## Active collaboration

The lack of a Novell GroupWise client is often cited as one of the stumbling blocks for Linux (and Apple's Mac OS) in Enterprise circles, so this element of Novell's new strategy has real potential to raise the profile of Linux as a desktop alternative to Windows. GroupWise is Novell's rival to *Lotus Notes* and *Microsoft Outlook*, providing calendaring, scheduling, email, instant messaging, contact and document management, and other workflow services across a Novell or Windows NT/2000 network. On the client-side, the only option previously was Windows. The new Linux client is Java-based and came to Novell via N-iX, part of the Newcomp Computer Systems group.

Novell chief technology officer (CTO) Alan Nugent said that ensuring the various platforms could 'play nicely' together would be at the very heart of the company's development strategy from now on. "The mixed technology found in customer organisations today requires savvy vendors who can tie it all together with proven and secure solutions," he said. "That's what we do best." Bringing GroupWise to Linux and Macintosh users would allow them to "collaborate security and reduce costs while deploying the desktop systems that make the most sense for their business."

## LINKS FOR FURTHER READING

**NOVELL CERTIFIED LINUX Engineer training:**
**www.novell.com/training/certinfo/**
**Initial training information:**
**www.lpi.org**
**UDDI Server project homepage:**
**http://developer.novell.com/uddi/**
**Novell collaborative development space:**
**http://forge.novell.com**
**Liberty Alliance:**
**www.projectliberty.org/**
**MySQL AB:**
**www.mysql.com/**

IDC analyst Mark Levitt said that Novell's recognition that 'not everyone does Windows' would pay dividends for the company. "Novell is broadening the audience for its secure and mobile/wireless collaboration with GroupWise for Linux and Mac users. This platform support, along with the functionality Novell includes as standard gives customers more reason to choose GroupWise."

On the server side, Novell also debuted an alpha version of GroupWare server running on Linux.

## Education, education, education!

While talking about 'providing a migration path' to Linux, Novell are also keen to be involved in the sometimes complex issue of skills migration and cross-platform administration. In recognition of this potential skills gap, the BrainShare conference saw the launch of the Certified Linux Engineer (CLE) program which will acknowledge sysadmins 'who have obtained a high level of understanding of both the Linux platform and how Novell technologies run on, and integrate Linux in, mixed network environments'.

Certification will involve an instructor-led or self-study course covering installation, configuration and management of services such as *iFolder, ZenNetworks* and *eDirectory*, and a practical exam. There will also be web-based practical resource material available for exam preparation. The courses will be rolled out towards the end of 2003, with examinations starting in early 2004.

As a prerequisite, Novell suggests participants should have completed level one of the Linux Professional Institute Certification award or its Red Hat-run equivalent to obtain a 'baseline' understanding of the operating system. Novell users who have passed the Certified Novell Engineer exam (and upgraded their certification to NetWare 6 before 31 October 2003) will be entitled to sit the CLE exam completely free of charge.

Brian Lowry told us that any major technological initiative hinges on having the right expertise available. "That we've announced a Certified Linux Engineer certification speaks to our assessment that this certainly holds true in the Linux space." He said it was the company's aim to make it as easy as possible for individuals who deploy and maintain Novell services on NetWare to transfer that know-how over to Novell's Linux offerings quickly and easily. "Going forward, we expect people will build expertise on Novell services, rather than on specific platforms. That said, we're very focused on giving our customers the information and knowledge they need to move to Novell services on the Linux platform as soon as those services come available." Training on specific services would be available, he said, in tandem with the software. "We won't let a lack of training become an impediment to our customers' adopting these services." ■

> "They have taken some care to act as good citizens of the Open Source developer community."
> BRUCE PERENS

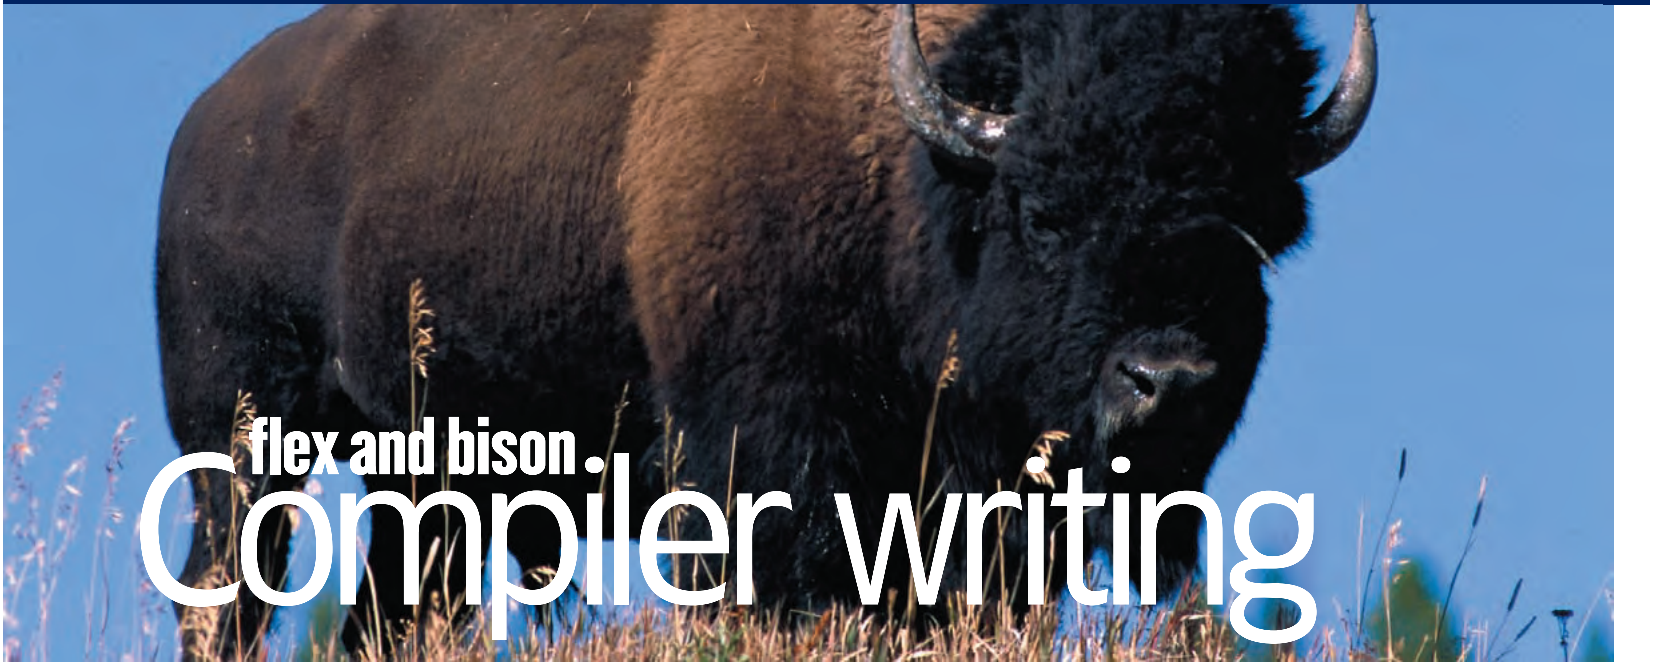

# flex and bison Compiler writing

**PART 2 PAUL HUDSON continues his tutorial on how to write your own compiler, and experiences some rather disturbing flashbacks to his GCSE French classes...**

First there was machine code, then came assembly, Fortran, and C; until eventually the paragon of computer programming languages was invented: SKYLang. Well, perhaps that's a smidgen of an overstatement, but this *is* only part two of our foray into the use of *flex* and *bison*! If you were with us last issue, you'll have had a whistle-stop tour of what compilers are, how they work, and what they can be used for. We also covered *flex* and *bison* basics, and at the end of the tutorial we had a parser working that recognised some basic variable assignments.

## Same Bat channel...

This issue the goal is to take what we learned last time and start producing the basics of our language – there's quite a jump between recognising basic assignments and actually doing something about it.

This month we'll look at how a language works semantically: "what is an expression?", "what is a statement?", and such. As variable assignment in typeless languages can be a fairly tricky topic to get to grips with for a newcomer to compilers, we'll be covering it in fairly simple terms. For now, we'll be having SKYLang act upon instructions as soon as it receives them. Later on, this will be changed so that there's a distinct split between the *parsing* phase and the *execution* phase.

If you missed last issue, we strongly recommend you contact our back issues department, as there isn't enough space here to re-explain concepts covered previously.

## Where we left off...

Last issue we had SKYLang recognising statements like **$foo = 20** and **$bar = 40**.

According to our language definition, SKYLang is to be a *typeless* language – this is actually a much harder thing to achieve than a typed language! A typeless language has little concept of integers, strings, floats, etc – each variable can be any type or indeed all types at once.

Consider the following function in C (a typed language):

```
int myfunc(char* foo) { }
```

That function takes a pointer to a character string as its only parameter, and returns an integer. Consider the same function in PHP (a typeless language):

```
function myfunc($foo) { }
```

The PHP version doesn't specify a return data type, and neither does it specify the data type of the parameter it accepts. Given that the code in the rest of the function does the same thing – let's say it returns the ordinal number of the character passed in – how does PHP handle a non-character parameter?

Using PHP, variables could be initialised with any of the following ways (or countless others!):

```
$foo = "bar"; $foo = 69; $foo = 23.9; $foo =
array("wom", "bat");
```

When a variable is set, PHP notes down what kind of value it is set to and stores that inside the variable. For example, the first **$foo** assignment sets the variable to a string, so PHP stores it's internal data type as a string. Externally, you can use that **$foo** string variable as part of a multiplication, even though it's a string. Whenever PHP has to work with the value of a variable, it has to jump through a number of hoops to make sure the variable is OK to work with. For example:

```
$foo = "bar";
$wombat = $foo * 2;
```

That will set **$wombat** to 0, because PHP will try to convert "bar" into a number.  As the value of **$foo** isn't a number in string form, PHP sets its integer value to 0 and returns that. If **$foo** had been set to "5" (in quotes), it would still be a string, with the difference that it could be converted to an integer relatively easily. Typeless languages need all sorts of type conversion behind the scenes, and they also need to be able to handle unrecoverable conversions – for example, subtracting 2 from an array.

As discussed last issue, SKYLang's variables are either integers, floating-point numbers, or strings. In order for a variable to be able to hold all three at once, we're going to use a C++ class called **SKYVar**, which has space for all three data types plus an enumerator to remember what kind of data is currently being stored in the variable. Here's how that looks in C++:

```
enum svType { svInt, svFloat, svString };

```

```
class SKYVar {
 public:
  svType type;
  int intval;
  double floatval;
  char* charval;
};
```

First we define **svType** to be either **svInt** for an integer variable, **svFloat** for a floating-point variable, or **svString** for a character string. The **SKYVar** class has four public variables: a type (one of our **svType** enum), an integer **intval**, a floating-point **floatval**, and a **char*** (pointer to a character string) **charval**.

**SKYVar** is going to be our default data type handed back from our lexer – all values passed back will need to placed into a particular part of the **SKYVar**. Take this piece of SKYLang, for example:

```
$foo = 5;
```

As **yylval** will be redefined to be a **SKYVar**, **$foo** will need to be placed in the **charval** part of **yylval** and passed to *bison*, then 5 will be placed inside the intval part of **yylval** then passed to *bison*.

Going from the general to the specific, here's how we go about programming this. First, take a backup of your code from last issue – best to keep that safe somewhere. Secondly, rename your files to skylang.l and skylang.y. Thirdly, create a new file, "skylang.h", and copy the C++ code from above into there. Fourthly, edit the makefile and remove the line **rm -f *.h** – if you don't do that, then **make clean** will delete the new skylang.h file. Fifthly, change the CC line in your makefile to **g++** rather than **gcc** – we're switching to C++ now. **NOTE:** if you get compilation problems, read the box *Switching to C++*.

You are welcome to add new makefile rules in there to clean up the skylang.tab.h file that *bison* will create when you run skylang.y through it. Finally, edit both skylang.l and skylang.y, and add the following line of code just after the **{%**:

```
#include "skylang.h"
```

Once you've done all that, you're all set to start the *real* work: making SKYLang accept all sorts of variables.

## Ints, floats, and strings, oh my!

Before we start coding, you need to understand how a computer language works. Back when I was doing GCSE French, I remember my teacher (a wonderful, persevering lady who almost managed to get me to ask the time in French once or twice) would regularly teach the class French language structure. Terms like *past participle, gerund, subjunctive* and *indicative* flew over my head because I (then) firmly believed that English didn't have any of that sort of thing, so French was clearly too hard to bother learning. Then of course at English A-Level, I had to describe English in these terms, so I finally understood what the fuss was all about!

Working with computer languages is much the same – there are a few terms that need to be understood before you continue. This next bit is crucial, so if you're reading this in bed I suggest you to sleep and give this your full concentration tomorrow!

A TERMINAL SYMBOL is any symbol which cannot be reduced or sub-divided into smaller parts. To take the C language as an example, **if, return**, **{**, **}**, **+**, **–**, *etc.* We also have NON-TERMINAL SYMBOLS, which are made up of terminal symbols. These are the important ones to understand.

For now, we're going to be working with three types of non-terminals: statements, expressions, and numbers. A number is either an integer or a floating-point number. An expression can take several forms:

**i)** A variable
**ii)** A number (as defined)
**iii)** Expression + expression
**iv)** Expression - expression
**v)** Expression * expression
**vi)** Expression / expression
**vii)** A string

A statement can also take several forms:

**i)** variable = expression
**ii)** "echo" expression

We'll be adding to that list as time goes by in order to increase the complexity of the language, but there's enough there for now.

Consider this line of code:

```
$foo = $bar + 10 * 5
```

How does that evaluate given our structure as defined above? If we imagine bison matching rules from **i)** to **vii)** in that order, here's how it would look:

```
1 variable = expression + expression * expression
2 variable = expression * expression
3 variable = expression
4 statement
```

Line 1 is the direct conversion of the code into tokens – note how we have a match against rule **iii)** for expressions, because **expression + expression** is itself an expression, so we reduce and move to line 2. Again, **expression * expression** is an expression, so we reduce again and go to line 3. Here we have **variable = expression**, which is of course a statement, so we reduce and we're finished.

Make sense? Unfortunately, there's more than one way to parse that statement, and the way we did it isn't the standard way. If you can remember back to your mathematics classes, you'll remember *operator precedence* – the order in which operators are executed when there is any doubt. The order, for basic statements like we're using, is Brackets, Division, Multiplication, Addition, Subtraction. Look at the diagram on the left – »

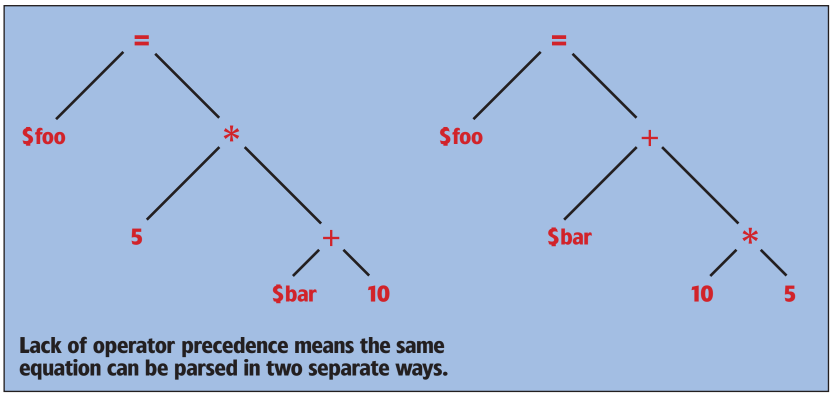

**Lack of operator precedence means the same equation can be parsed in two separate ways.**

### SWITCHING TO C++

**C++ IS A LOT MORE STRICT than C in many ways, and so your program will not compile unless you make some minor changes to the code from last issue. Firstly, skylang.l needs to have #include <string.h> after the include of skylang.tab.h (that shouldn't still be variables.tab.h!). Secondly, skylang.y needs to have the following two lines of code after the #include:**

```
int yyparse(void);
int yylex(void);
```

**Even with those changes the code from last issue still won't compile with C++ because we pass back char * pointers as integers (the default yylval type). The solution here is merely to use the new code from this issue - we're working with typeless variables now!**

here you can see a tree diagram of the statement, parsed in the two different ways. The left tree shows how we parsed it above – addition coming before subtraction. The right tree shows how it *should* be done – with the multiplication being done before addition. *Bison* has a very easy way to handle this which we'll cover later on – first, the new code.

## Yakkity Yak!

To get our code up to speed, we need to make a large number of changes. This will probably be the last time we have space to print out the entire source files!

```
%{
#include “skylang.h”
#define YYSTYPE SKYVar
#include “skylang.tab.h”
%}
%x S_INSTRING
%%
\$[A-Za-z_][A-Za-z0-9_]* {
  yylval.charval = strdup(yytext);
  return T_VARIABLE;
}
<S_INSTRING>[^\n”]* {
  yylval.charval = strdup(yytext);
  return T_STRING;
}
<S_INSTRING>\n { printf(“Error: Unterminated string
constant.\n”); }
<S_INSTRING>\” { BEGIN(INITIAL); }
“echo” { return T_ECHO; }
[0-9]+”.”[0-9]+ {
  yylval.floatval = atof(yytext);
  yylval.type = svFloat;
  return T_FLOAT;
}
[0-9]+ {
  yylval.intval = atoi(yytext);
  yylval.type = svInt;
  return T_INTEGER;
}
“+” { return T_PLUS; }
“-” { return T_MINUS; }
\* { return T_MULTIPLY; }
“/” { return T_DIVIDE; }
\” { BEGIN(S_INSTRING); }
“;” { return T_SEMICOLON; }
.|\n {  /* trim whitespace */ }
%%
```

A lot of that code will make sense only once you've read the new skylang.y, but the new *flex* file is easier to read so we'll get that out of the way first. We now have **T_PLUS**, **T_MINUS** and friends to handle operators. We also **#define YYSTYPE**, the variable that **yylval** is, to be our new **SKYVar** variable, and return values using a particular part of **yylval** rather than the whole variable, *eg* **yylval.intval**. One key change is strings, and these are accomplished using *start states*. The line **%x S_INSTRING** defines a lexer state, **S_INSTRING**, that can be entered and exited whenever necessary. The default state in *flex* is **INITIAL**, which parses all rules that don't have a state ([0–9]+”.”[0–9]+, [0–9]+, etc). When we wish to change rules, we use the special *flex* command **BEGIN(<rule>).** From then on, rules for that state are parsed by *flex*. There are two types of states – inclusive and exclusive. Exclusive, specified with **%x**, instructs *flex* to *only* parse rules from the current state – that is, it will ignore the other rules. Inclusive states, specified with **%s**, parse rules for the current state and also for the initial state.

Our string match rule is **<S_INSTRING>[^\n”]***, which means “When we're in the string state, match everything that isn't a newline or a double quote”. As our **S_INSTRING** was defined with **%x**, and is therefore exclusive, *flex* will ignore all the other possible matches – **$foo** won't be considered a variable when it's inside a string.

Our rule **<S_INSTRING>\”** has the action **BEGIN(INITIAL)**, which instructs flex to go back to normal, non–string parsing as soon as we're back out of a string. Now onto the important (and hard) stuff. Open skylang.y, and replace what you have with this new code:

```
%{
 #include “skylang.h”
 #include <map>
#define YYSTYPE SKYVar
 int yyparse(void);
 int yylex(void);
 void yyerror(const char* s) {
   fprintf(stderr, “%s\n”, s);
}
std::map<char*, SKYVar*, ltstr> UserVars;
main() {
   printf(“Beginning to parse...\n”);
```

## MAPPING THE STL

### Re-use the hard work of others

C++ SUPPORTS TEMPLATES, WHICH IN essence allows generic programming (although behind the scenes it's not quite so fancy!) Templates allow you to design a data structure that can be re-used in many ways, and the STL is a standard library that implements many data types that come in useful – trees, vectors, stacks, queues, and more. A stack object, for example, allows you to push and pop objects to and from your stack, but, because it's *generic*, that stack can be a stack for strings, a stack for integers, or a stack of <your data type here> – that's the beauty of STL.

A vector is a dynamic array, where items can be accessed using their integer position inside the vector. A *map* is a vector with the difference that you can reference the position of variables inside it using any other type. For example, using a **char***/**char*** map, we can add strings to our array like this:

```
mymap[“foo”] = “baz”;
mymap[“bar”] = “wombat”;
```

As you can see, working with maps is marvellously easy. As such, we're going to be using them to handle our symbol table – where we store our variables as the script is executed. It may well be slightly faster at run-time if we used a hand-written symbol table, which is what many languages do. However, that option would leave us little time to discuss more important parts of compiler creation. Regardless, the STL is very finely tuned for maximum performance – I have a feeling that our symbol table will do quite nicely.

For our symbol table we use a map where **char*** is the key and **SKYVar** is the value. This allows us to create entries in a map where the variable name is the key, with a **SKYVar** as each value – perfect!

```
13  yyparse();
14  printf("Done parsing!\n");
15 }
16
17 %}
18 %token T_ASSIGNEQUALS T_SEMICOLON T_ECHO
19 %token T_VARIABLE T_STRING T_INTEGER T_FLOAT
20 %left T_DIVIDE T_MULTIPLY T_PLUS T_MINUS
21 %%
22 statements: |
23          statements statement
24 ;
25 statement: T_VARIABLE T_ASSIGNEQUALS expression
 T_SEMICOLON {
26  SKYVar *newvar = new SKYVar;
27  switch($3.type) {
28   case svInt:
29    newvar->intval = $3.intval;
30    newvar->type = svInt;
31    break;
32   case svFloat:
33    newvar->floatval = $3.floatval;
34    newvar->type = svFloat;
35    break;
36   case svString:
37    newvar->charval = $3.charval;
38    newvar->type = svString;
39    break;
40  }
41  UserVars[$1.charval] = newvar;
42  }
43  |  T_ECHO expression T_SEMICOLON {
44   switch ($2.type) {
45   case svInt:
46    break;
47   case svString:
48    break;
49   case svFloat:
50    printf("%.2f\n", $2.floatval);
51    break;
52  }
53  };
54 expression:
55  expression T_PLUS expression {
56   $$.type = svInt;
57   $$.intval = $1.intval + $3.intval;
58  }
59  | expression T_MINUS expression {
60   $$.intval = $1.intval - $3.intval;
61  }
62  | expression T_MULTIPLY expression {
63   $$.type = svInt;
64   $$.intval = $1.intval * $3.intval;
65  }
66  | expression T_DIVIDE expression {
67   $$.type = svFloat;
68   $$.floatval = (float)$1.intval / $3.intval;
69  }
70  | T_VARIABLE {
71   $$ = *UserVars[$$.charval];
72  }
73  | T_INTEGER
74  | T_FLOAT
75  | T_STRING {
76   $$.type = svString;
77   $$.charval = $1.charval;
78  }
79 %%
```

## TOOLS OF THE TRADE

**THERE ARE QUITE A FEW books out there on compiler technology, but sadly very few are really excellent. One book that is widely recognised as being the definitive classic book on compiler technology is *Compilers: Principles, Techniques, and Tools*, by Aho et al (Addison Wesley, ISBN 0-2011-0088-6). While this book is about 15 years old, it still makes for powerful reading, and covers lexing, parsing, type checking, intermediate code generation, optimisation, and other topics. This book should be on the shelf of every serious compiler programmer, even if only just for reference – it really cannot be recommended highly enough by us.**

Yes, that's one heck of a lot of code, but it does a substantial amount. To conserve space, I've stripped as many linebreaks out as I could – I *strongly* recommend you add comments and space your code out neatly, as it increases the legibility massively.

We're using the C++ STL std::map to store our variables, so we need to include the map header file – more on that later. We also redefine the standard value type passed back from our lexer as being a **SKYVar** – the class defined in skylang.h. These are lines two, three, and four. Lines five and six are required to compile in C++ (see the box *Switching to C++* on page 9 for more info). Lines seven, eight and nine are the same **yyerror()** function we had last issue.

Line 10 declares our STL map – if "STL" or "map" don't ring any bells, read the box *Mapping the STL* below left before you continue. Our map is called UserVars – note the third parameter in the map is **ltstr**. This needs to be specified so that the map knows how to sort elements internally, and we'll be using the code from the SGI STL manual. Type this code into skylang.h:

```
struct ltstr {
 bool operator()(const char* s1, const char* s2) const {
  return strcmp(s1, s2) < 0;
 }
};
```

That chunk of code makes the map sort the keys by their alphabetic position, which is fine. It's not actually necessary to know the nuts-and-bolts of how it works, internally, it "just works".

In lines 18 and 19 we have a several new tokens declared for later use. Of importance, we have **%left** on line 20, which declares **T_DIVIDE** and the others on the line as being left-associative operators, meaning they will naturally form equations to the left. This is part of operator precedence, but not how we discussed it earlier. For example: 10 - 5 - 2. Is that "10 - (5 - 2)" (gives a result of seven) or "(10 - 5) - 2" (gives a result of 3)? Left-association is the norm for nearly all operators, which means the latter is correct. We make that clear to *bison* by using **%left** for operators that are left-associative. However, there's more to it than that – the ordering of tokens on the **%left** line dictates the individual precedence of operators – **T_DIVIDE** comes before **T_PLUS**, so *bison* considers them higher precedence.

Continuing on, we find **statements** is still the root node, made up of multiple other **statement**s. Lines 25 to 53 define that a statement is either "variable = expression" or "echo expression". We'll look at how the code works momentarily, after we finish the grammar. From line 54 to the end we're defining what an expression is – this is just the same rules we discussed above, except in *bison* terminology. »

Because a variable has three distinct parts to it, it's possible for it to hold three values at the same time – an **int**, a **float**, and a **char*** – all with separate values. It's not terribly hard to switch this to a union to save memory (all three would occupy the same space in RAM), but I'll leave that to you.

Even if we were working with a union, it would still be important to know what kind of variable we have stored, and this is what the Type is for. As you can see in the code, there are two places (assigning and echoing) where we run the type of variable through a switch statement to make sure we're working with the right data. Having to jump through these hoops each time we access a variable makes untyped languages relatively slow, and that's before we even try to handle conversions.

## Mr Right Data Type

All the parts of the expression non-terminal are currently working solely on **intvals**, except the divide operation, where the result has its **floatval** set. In order for this to work properly using multiple data types, a lot more code is required – more than we have space for. So, we'll be writing a new chunk of code to handle multiplication – you'll need to do the rest yourself.

Here's how it looks currently:

```
expression T_MULTIPLY expression {
  printf(“Got multiply\n”);
  $$.type = svInt;
  $$.intval = $1.intval * $3.intval;
}
```

As you know, an expression can be just a variable (rule 1), so it's possible that we may have something like this:

```
$foo = “bar”;
$foo = $foo * $foo;
```

Here we multiply a string by a string – a perfectly valid, if somewhat odd, operation. In this situation there will be nothing in the **intval** property of our variable, so we need to get that corrected before attempting the multiplication. Furthermore, we need to handle **float** multiplication, and also the possibility that the return value might be a **float** rather than an **int**. So, we need to get all the variables into a type that we can multiply, figure out what kind of multiplication to perform, then return the right value. Here's how it would look for multiplication – you might want to put some of this into an inline function for easier maintenance.

```
1  expression T_MULTIPLY expression {
2    printf(“Got multiply\n”);
3     if ($1.type == svString) { printf(“Converting param 1!\n”); $1.type = svInt; $1.intval = strlen($1.charval); }
4    if ($3.type == svString) { printf(“Converting param 2!\n”); $3.type = svInt; $3.intval = strlen($1.charval); }
5    if ($1.type == svInt && $3.type == svInt) {
6     $$.type = svInt;
7     $$.intval = $1.intval * $3.intval;
8    } else if ($1.type == svFloat && $3.type == svFloat) {
9     $$.type = svFloat;
10     $$.floatval = $1.floatval * $3.floatval;
11    } else if ($1.type == svFloat && $3.type == svInt) {
12    $$.type = svFloat;
13    $$.floatval = $1.floatval * $3.intval;
14    } else if ($1.type == svInt && $3.type == svFloat) {
15    $$.type = svFloat;
16    $$.floatval = $1.intval * $3.floatval;
17   }
18  }
```

It might look tricky at first, but let's break it down. Firstly, we make sure that if we're dealing with strings, we get an integer value rather than the string itself. Conversion is done on each parameter, if necessary, on lines 3 and 4. As you can see, I'm using the length of the string for its integer value, which means “wombat” * 6 is 36. Most languages, including PHP, just put a 0 into the **intval** when multiplying strings, which is generally the safest route. The reason I chose **strlen()** was just to show you that this is an area you can really play around in – if your language works with string lengths a lot, then **strlen()** for **char*** to **int** conversion is perfect, whereas if you work with text a lot, perhaps **“wombat” * 2** is “wombatwombat”.

Once we know we definitely have a number to work with, we go through a series of checks to make sure we read the correct values in. This is all lightning fast and absolutely necessary, so don't worry if you think it looks over-complicated just to multiply some numbers.

That code covers nearly everything required, except the possibility of overflows and underflows. An overflow occurs when an integer goes beyond the space allocated for it, and an underflow occurs when a number is so small that it cannot be represented and simply becomes 0. For example, if we have a signed eight-bit integer, the largest number we can represent is +127. If we have 127 and add 1 to it, the number flips over and becomes negative – it has overflowed its capacity for storage. Now imagine a number that holds two decimal places, to which we tried to assign 0.000000000039 – the result would be the number is 0. Both are errors, and the former is something that particularly needs to be watched out for as it can cause serious problems.

Luckily, detecting and fixing overflow is simplicity itself. As it occurs in integers and the only operation we have that returns integers is when both operands are integers, the change we need to make is localised. To solve the problem, we store the result of the multiplication in a temporary double, compare it against the values INT_MAX and INT_MIN (defined in the standard header file limits.h), and if the value is within bounds we store it as an integer. If not, we set the return type to double and store the double value in **floatval**.

So, add **#include <limits.h>** to the top of skylang.y, and change the int/int multiplication to this:

```
if ($1.type == svInt && $3.type == svInt) {
  double dtemp = (double)$1.intval * (double)$3.intval;
  if (dtemp > INT_MAX || dtemp < INT_MIN) {
   $$.type = svFloat;
   $$.floatval = dtemp;
  } else {
   $$.type = svInt;
   $$.intval = $1.intval * $3.intval;
 }
}
```

## ONE LAST TWEAK...

### Getting SKYLang to read scripts

I DON'T KNOW ABOUT YOU, BUT I'm not fond of having to enter in the same programs again and again – reading and executing from a file is much better! So here we'll implement compilation of a file.

Just below the declaration of our **UserVars** map, add this line:

```
extern FILE *yyin;
```

Then replace the entire **main()** function with this new version:

```
main(int argc, char **argv) {
  if (argc > 1) {
   yyin = fopen(argv[1], “r”);
   if (yyin == NULL) {
    fprintf(stderr, “SKYLang: Cannot open %s”, argv[1]);
   } else {
    yyparse();
   }
  } else {
   printf(“\nSKYLang v0.2\n------------\n\n”);
   yyparse();
  }
}
```

To explain, we change **main()** so that it accepts **argc** and **argv** – the number and values of parameters passed to our program. **argv[0]** is always set to be the name of the program itself, so **argc** is always a minimum of one. If **argc** is greater than one, we assume our user is passing in the name of a script to parse, so we set **yyin**, where *bison* should read from, to be the specific file, and call **yyparse()**. If no file was specified SKYLang reads from standard input, which gives us the option to work with both cases.

Now we can read from scripts, we can do our first benchmarking – how fast is SKYLang? Yes, it uses OOP, yes it uses STL, and yes it uses typeless variables – it's not exactly a mixture aimed at maximum performance! However, the code is very easy to read, and hopefully it's been fairly easy to learn, so I think it was worth it. I created a test script by running the PHP CLI in interactive mode, piping **stdout** to speed.sky, and entering the following:

```
<?php
for ($i = 1; $i < 1000000; ++$i) {
 echo ‘$a_’ . $i . “ = $i;\n”;
}
?>
```

This put 1,000,000 lines of assignments into the test script - quite a bit of code. As SKYLang is very close to PHP (untyped, $ variables, etc), the script can be converted to PHP by merely adding **<?php** and **?>** to the start and end of the script respectively.

Running the script in SKYLang took 7.2sec on my 800MHz PIII, and running the same script in the PHP CLI took 19.63sec. A victory for SKYLang? Well, yes and no! “Yes”, because it shows that using modern programming methods for your compiler is no bad thing. “No” because PHP does a heck of a lot more work than SKYLang – not only does it recognise far more tokens, but it also has many more *bison* rules to match against. Add to that the overhead of the *Zend Engine* and the speed difference becomes justified.

You can get an idea of how much more lexing work PHP does compared to SKYLang by adding the text **Break!** just after the last assignment in the PHP script. Naturally this is an unexpected string, so PHP will bail out with an error message. The key is *how long* it took to bail out – for me, it was 7.85sec, meaning that it took over seven-and-a-half seconds for PHP's lexer to get through the script – more time than SKYLang's took to parse and execute the same script. Again, this is due to PHP's lexer being 25 times longer than SKYLang's!

As SKYLang parses and executes as it goes, it will perform terribly badly when executing loops, function calls, and such. It also makes it impossible to see what percentage of time is spent parsing and what percentage executing as we did with PHP – roll on intermediate code!

Note the new dtemp variable to store the result. We need to prefix the two **intvals** with a double cast otherwise GCC will perform integer multiplication and give us the wrong result. If you want to have added capacity to your **intvals**, you could declare them as long – again, limits.h has LONG_MAX and LONG_MIN to handle the longer data type.

## Conclusion

Last issue SKYLang recognised (but did nothing about) variable assignments, and that's all. Now it's capable of performing complex arithmetic (“$foo = 10 * 20 - 50 / 2”), deciding operator precedence and association, handling overflows and type conversion, and also echoing out values – quite a step forward, I'm sure you agree!

We're still interpreting commands as they come, however, so it's quite basic in terms of compilation technology. Also you'll need to add the necessary code above (based on the multiplication code given) to make sure proper type conversion is done everywhere.

What problems exist? Well, try starting your script with an echo statement – trying to work with the value of an unset variable will cause a segmentation fault, which is a bad thing. The reason this happens isn't because we don't have a **SKYVar** constructor, and nor is it down to GCC trying to shame me into using private variables (no space, sorry!) – instead, it's down to STL semantics. Operator [ ] of an STL map means “find and return, or create and return if not found”. The *value* part of our map is a **SKYVar***, so when we use the [ ] operator, we are quite rightly returned a pointer to a non-existent object. The solution is to check again NULL and create if necessary – we'll get onto that next month. Furthermore, notice that we never check to see whether a variable is set already in assignments, which means code like this:

```
$foo = 5;
$foo = 10;
$foo = 15;
```

will leak memory, as a new **SKYVar** is created each time in **UserVars[“$foo”]**, without freeing the memory assigned to the **SKYVar** already there. This will be fixed later, but it's not something to be overly worried about, as Linux will automatically free all memory assigned our compiler as soon as the process ends. One last problem is that if we set a variable to “123” (with quotes) and convert it to a string we return 3 (thanks to **strlen()**). Ideally we would attempt a proper integer to string conversion, because clearly 123 is just a number in a string.

At the end of this installment you should have a firm understanding of how typeless variables work, how to use STL to get great performance with little work, and also what makes up a language's grammar. The fun begins – we're really starting to get in deep with our language now... ■

## NEXT MONTH

We'll be fixing a lot of the bugs mentioned to make the language more stable, we'll be looking at how built-in functions work, and perhaps even starting to delve into intermediate code. In order to teach loops, if statements, and user functions we need to be working with intermediate code, so don't go thinking I've forgotten half the language!

# Unbreakable Linux
## The Oracle vision

**NICK VEITCH investigates Oracle's increasing profile and success at selling Linux to the Enterprise market**

Five years ago, Linux was, while not exactly a pariah of the blue chip IT industry, at least something you wouldn't necessarily want your corporate image associated with. Times change – the last couple of years have seen much hype about Linux, and about how 'on-board' various industry stalwarts are with the OS - everyone, it seems, must have a 'Linux strategy.'

By any stretch of the imagination, you could hardly call Oracle a company in need of some headline-grabbing attention though, nor could you accuse the company that provides solutions to most of the leading corporations and governments of the world of living on the bleeding edge. And it seems that it is impossible to call their Linux strategy vague, woolly or insubstantial – the adoption of Linux at Oracle has been swift, successful and almost total, and has culminated in possibly their biggest marketing push of recent years - the 'Unbreakable Linux' campaign. Launched last year, it is still their biggest ongoing marketing push. But is Oracle really unbreakable? And why has Oracle adopted Linux so completely? We spoke to Oracle UK's Chris Ward:

***LINUX FORMAT*****: Oracle has a good history of supporting Linux, as it has done for some years, but now actually actively promotes the use of Linux. What has happened within Oracle to make that change of stance?**
**CHRIS WARD:** You're absolutely right. Oracle's Linux solutions are our number one technology push, certainly in the short term. It is our clear strategy to become a dominant player in the Linux marketplace with all of our software solutions – and that's everything from the database right the way through to the applications that we offer. Every current piece of Oracle technology is now available on the Linux platform – our applications server, database, collaborations suite, Java tools and enterprise applications.

The big change for Oracle is that clearly Linux is our fastest growing operating system for any deployment – also the stats are showing growth of Linux in the datacentre and of course the server market – IDC recently predicted almost 200% growth by 2006.

Clearly there are a number of areas that have changed in the Linux space to make this a believable strategy for Oracle. The first thing is the relationship with Red Hat. We have obviously already ported products to the Linux stack, but recently we have done worldwide support agreements both with Red Hat and the United Linux group.

***LXF:*** **And how active have Oracle contributed to the growth of Linux?**
**CW:** What Oracle has been trying to do is improve – particularly those two flavours of Linux mentioned – in three or four main ways.

Input/Output – we've been contributing code to help the input/output mechanisms of those two distributions.

Clearly our clustering software, *9i Real Application* clustering, is really the heartland of our 'Unbreakable' stance. If you were to ask what software does Oracle offer to caveat this unbreakable message, it's the fact we can cluster machines together for scalability and redundancy. We've also made a lot of improvements to the scheduler for the Linux kernel, and also memory handling, for large-scale databases.

Those three or four areas were where Linux was struggling to offer enterprise-class features in the past. We know that from the many instances of massively parallel Oracle systems we run worldwide. So we've taken our expertise there and offered that to these distributions and open source.

***LXF:*** **Oracle developers have been working on kernel features then?**
**CW:** Yes, we do joint developments with both Red Hat and the United Linux group. Their developers sit with ours now, and we have opened up information on access methods as to how the database works with operating systems.

Linux has become a tier one development platform for Oracle development. So we write our technology on Linux now, as well as obviously deploying it in our own business. Previously it wasn't tier one.

***LXF:*** **What's brought that about then – just the emphasis on the new Linux strategy?**
**CW:** It is the whole emphasis. It's about the cost savings people are taking through to business, the change in the market space and the ability of the Unix market to take on board skills and deploy them in the Linux world – the whole profile of Linux has just shot up right the way through the Oracle organisation.

We have a clear strategy now – when we develop products we will develop them on Linux, on two flavours primarily of Unix, and obviously Windows.

The other unique thing is that we write across multiple operating systems, we don't just write for one and port. It's a really good message for Oracle to say "actually, we believe in this, we're committing developers right the way from scratch to write technology for Linux".

***LXF:*** **And Oracle uses Linux internally?**
**CW:** A number of our new businesses, for example our hosting business – we offer a hosting service. For example if you had a company with 300 employees and you wanted to use financials and purchasing and accounts payable, you could rent that service from Oracle as a hosted service. A lot of people do that. That business, which is a new business for Oracle, with a stated goal of $1billion within 18 months, and we're well on the way to that – that business runs on Linux. The datacentres we've built, we've got 85 machines clustered, running Linux, providing worldwide hosting capabilities. So we're reinvesting and using our Linux technology to run new services from Oracle. That's one that is often not noticed, because the client is buying a service and doesn't know or care what's underpinning it.

> "So we write our technology on Linux now, as well as obviously deploying it in our own business." CHRIS WARD

***LXF:*** **Is that merely a cost decision, or would it be true to say that Oracle on Linux outperforms other platforms?**
**CW:** Cost is a huge benefit here. Not just in licensing costs, but in the ability to use non-proprietary hardware. So there are two or three big angles to cost. We are seeing huge improvements in performance, clearly you have to say more advanced processors have come into play as well, but we are seeing definite improvements. I would add that, in addition to our partnership with Red Hat and United Linux, one of our big technology partners in this is Hewlett Packard. We see the growth in their Linux business to being fundamental to the Oracle economy as well. It is our ambition to push out lower cost computing and push businesses forward. Quite clearly it's in competition to Microsoft.

***LXF:*** **You've mentioned United Linux a few times. Obviously one of the reasons for creating that group was to lower internal development cost, but also to make a clear platform for proprietary developers to work for. Has that been a factor in partnering with United Linux as well as Red Hat?**
**CW:** I think Red Hat's strategy is slightly different to that of United Linux. Red Hat has been good in that for commercial software, you can't be releasing it every month, which is why they have split Red Hat into the consumer flavour and the enterprise server version. United Linux is a lot smaller in terms of market share at the moment. Oracle can then certify and test correctly on these versions – we could never certify across every distribution of Linux in the world and we do not try to do that.

***LXF:*** **We don't know anybody that does!**
**CW:** I think it would be impossible. I think the maturity of the Red Hat organisation, and obviously United Linux, has certainly altered in terms of how they support their customers, how they manage release cycles, porting, bug-fixing. And how they use our support centres across the world. »

**Larry Ellison,Oracle CEO launched their "Unbreakable Linux" campaign last year.**

**«** ***LXF:*** **Obviously there are other database vendors in the Linux space. Who would you see as your major competitor?**

**CW:** It has to be IBM. IBM and Oracle really are aiming at this marketplace. We don't see the open-source database solutions as being competitors. MySQL and those types of products, we don't see them as being competitive. There is an awful lot to knowing how to run massively scalable solutions and Oracle clearly have that.

***LXF:*** **I suppose in many ways the open source solutions are focussed on different markets.**

**CW:** They are. We certainly respect them. We think people that are contributing to open source are doing so very well and we certainly actively support that. It's nice to see new thoughts and ideas.

The relational world is very established and Oracle is the dominant player worldwide. We are often asked "do you think you can sell this" in relation to open source databases, and we think that when you have value and you can prove and demonstrate security and trust with clients you can sell that.

"It is our ambition to push out lower cost computing and push businesses forward."

***LXF:*** **As the market changes, and certainly the enterprise customer base for Linux has grown dramatically in the last few years, customers coming to Linux, while they may be interested in the costs savings in coming to Linux, and the fact that they are more free to choose technologies, they still appreciate the support and services that possibly only a very large corporation can give them...**

**CW:** I agree. Oracle provides an integrated solution. Yes, we have our database and the application server, but bring on board our application suite (a major set of applications for doing financials, supply chain, logistics, planning) and you've got everything to run the back-end of a business, basically. Because we are a major application vendor, bringing that to the Linux platform really starts to break the mould and say that Linux is absolutely ready. We have companies like Dell running Oracle on Linux. We're moving our own back-end to Linux. There are lots of big clients out there – CERN and Merrill Lynch to name just two.

***LXF:*** **Obviously the Linux part of your business is the fastest growing area. Do you know how many of those customers are new to Oracle, or existing customers migrating to Linux?**

**CW:** There is definitely a large proportion of the install base looking at this. Obviously you can move systems running Oracle on Windows to Oracle on Linux very easily and take away cost. It is the same product. There is a movement of the install base, but Linux has given us new clients as well.

We've had over one million downloads of Oracle on Linux in the last 18 months. I don't suggest that they've all deployed it, but it demonstrates the interest. That's a huge figure; it's the number one requested download from Oracle. We have 4000 ISVs using Oracle database and Oracle Application Server on Linux, and 550 of those have certified their own applications for Oracle on Linux. So you can see that the Oracle channel is picking up on this too.

***LXF:*** **It isn't just a case of Oracle supporting Linux then, everyone else in the Oracle community is interested...**

**CW:** The Oracle economy is turning to Linux, yes.

***LXF:*** **If I can bring up this whole idea of Unbreakable Linux – is that not overstating the case slightly?**

**CW:** Two things. Unbreakable – we've had a lot of heat over this term. What we are saying is not that the solution will never fail. What we are providing is a way of designing systems that maximises the possibility of not failing. A bit of a wordy way of putting it.

Basically, you haven't had the opportunity to create elegant clusters managed by software, which are relatively easy to manage and look after. If a single machine fails, Oracle keeps running, if you need more power you can add that. Clustering has more or less been a proprietary hardware solution before.

With the security aspects of Oracle, it's clearly a highly secure environment – it's been written for the FBI and governmental bodies; those are the types of organisations that have helped design this software. There are a number of areas we believe you can use Oracle, and you can't break Oracle for the common reasons.

***LXF:*** **In the Linux enterprise space, where are the big growth sectors for Oracle?**

**CW:** It is across the board, but three areas in the UK are public sector, financial services and retail. Those three markets have been very active. That's based on my experiences of talking to customers.

"We've had over one million downloads of Oracle on Linux in the last 18 months. I don't suggest that they've all deployed it, but it demonstrates the interest."

***LXF:*** **Is there anything that the company will need to address in a different way for the Linux market in the future?**

**CW:** We have addressed support which was a major issue for clients thinking about Linux – if you run Oracle on Red Hat or United Linux, Oracle will pick up the support call across the world using the 'follow the sun' principal.

I think the market has matured a lot. This is what enterprise customers want – a solid mature relationship with vendors. If you were a customer who had a fault running Oracle or Red Hat for example, you could put a support call in and the Oracle representative will try and fix it, regardless of whether the fault is with Oracle or Red Hat. We have 24/7 support, which is important. That was a barrier to people taking Linux solutions before. ■

# Licence to code

**Even though many don't recognise it, the most crucial thing about free software is its licence agreement. Don't agree? PAUL HUDSON looks at open-source licensing: which licence you can use, and why...**

Licence agreements are peculiar things. Lawyers spend days preparing them, making sure all rights are made fully clear, and yet users almost never read them past looking for where it says "I agree". Despite the fact that sometimes licence agreements force users to sign away a lot of their rights, people don't seem to want to pay any attention to them. Are they too long? Too convoluted? Luckily in the Linux world we have to worry a great deal less about licences, because most Linux software comes with very liberal, open licensing terms that give us many more rights than our Windows counterparts. Does that mean we should be complacent about licences? No - in fact, if anything there's a need for us to be more wary about licensing, because not all licences are made equal.

## What does it mean to be free?

For most users, Linux is all about being /free/, but what does it mean to be free? You might think "free" is a pretty straightforward word, but English uses the word to mean both "without cost" and "without bonds", which is why many people talk about free beer and free speech. If you had a case of free beer, it would be beer that you did not have to pay for – it was yours at no cost. If you have free speech, it doesn't mean that you don't have to pay for your speech. Instead, it means that you're able to say what you wish, without bonds or laws to hold you back.

So, in the Linux world, our software is nearly always free as in speech (that is, we can do what we like with it), and generally also free as in beer (we can download it from the web without paying). The difference is important, because it means that even if you buy a Linux distribution (a boxed Mandrake, for example), you still have free (speech) software on your hands, even thought it wasn't free (beer).

The issue of "free" is further muddied by the array of licences out there, each claiming to be free in its own right. For example, does a truly free licence allow people to take the code and copy it without releasing changes to the source code, or does that limit the freedom of the code by allowing people to not redistribute changes?

We'll be looking at that issue, and others, as we discuss the main licences currently in the arena. It has become increasingly vogue as of late to licence products under several licences, allowing people more flexibility than they would otherwise have had. This method allows open-source programmers to happily develop for the same products as commercial users, and seems to be a viable commercial route for open-source companies.

## The GPL

This is "the big one" as far as Linux licences go. The Linux kernel uses it, as does KDE, GNOME, and most GNU products. The GPL, the main licence recommended by the Free Software Foundation, allows people to take code they receive, study it, modify it, redistribute it at no cost or for any price they want, and run the code for any purpose they want. As such, it is a very open licence, and one that emphasises community effort and improvement. Because all code that is forked from a GPL program must itself be released under the GPL licence, it is self-propagating.

The GPL isn't good for everyone, though – commercial developers often want to re-use GPL code while keeping their own parts proprietary. Under the terms of the GPL, this is prohibited – if you even so much as link to a GPL library, you need to GPL your code. Another key part of the GPL is that it future-proofs itself by stating that licencees may use GPL v2 or any later version at their discretion. Authors who use the GPL therefore technically are leaving themselves open, as they might find people using their product under the terms of GPL v3 (not yet written). Sometimes authors clarify this, as seen in the Linux kernel, where Linus says:

"Also note that the only valid version of the GPL as far as the kernel is concerned is _this_ particular version of the license (ie v2, not v2.2 or v3.x or whatever), unless explicitly otherwise stated." – Linus Torvalds

The clause that authors must GPL code that was based upon GPL code is known as "copyleft", and ensures that patches, updates, and improvements are always given back to the community.

## The LGPL

The GNU Lesser General Public Licence (previously the Library General Public Licence) is a cut-down version of the GPL that removes some of the tighter restrictions. The key change is that linking to an LGPL-licensed library doesn't force a licence on the application – developers don't need to release their code under the GPL/LGPL. This makes LGPL libraries accessible to commercial developers. However, many (including Richard Stallman himself) generally advocate the GPL over the LGPL wherever possible. The only time they would say »

### GOT SOME SPARE TIME?

**THERE ARE MANY documents on the GNU website that discuss various aspects of free software – licensing, commercialism, etc. All are available for free (unsurprisingly!) and they make for very interesting reading – take a look at www.gnu.org/philosophy/philosophy.html to see what's there.**

the LGPL is the correct licence choice is when the library provides functionality already available in a proprietary equivalent. If however there is something in the library that isn't available elsewhere, and if the feature was suitably advantageous, then the library should be GPLed, as the hope is that companies will choose to use the library and GPL their own product. If a library has little or no functionality extra over existing, closed-source options, then releasing it as GPL is likely to only hinder adoption amongst companies.

The reason the Lesser GPL was renamed from the Library GPL was because people assumed *all* libraries should be licensed under the Library GPL, which isn't true, as mentioned above. "Lesser" makes it clear that, although it's still GPL-like, it's a slightly different breed of animal.

## Microsoft Shared Source Licence

If you're surprised about seeing this licence in here, don't be – it is a free licence, at least in the beer sense; you can download some Microsoft MSSL-covered code and read, edit, modify, and experiment with it, and you may also distribute your modified versions. It isn't free as in speech, though, because you may only make use of the software source code non-commercially.

## The BSD licence

The BSD licence is considered one of the most free licences of the bunch because it has no source redistribution requirements whatsoever. As such, some parts of Windows use BSD code, and large parts of Mac OS X do, too, without the need for Microsoft or Apple to give back changes or improvements to the community. So, while the code is free to move around and be used in all sorts of places, the community benefits less.

One particular problem with the BSD licence used to be that it had this following clause:

"All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors."

At the time, RMS rallied long and hard against this clause, and with justifiable grounds. You see, one line of text wherever advertisements appear isn't a problem. However, the BSD licence was taken and modified by many BSD developers – each one changing the line of text for advertisements. As a result, a product that included these programs would need to include dozens of lines of acknowledgments in adverts. Luckily, RMS and others managed to get the notice removed from several BSDs, alleviating the problem.

## WINE AND THE WINEX LICENCE

**SOME TIME AGO, WINE switched from a BSD-style licence to the LGPL licence, which made a few people unhappy, who founded the Rewind project to keep WINE being developed under the old licence. When Transgaming was founded, with the goal of enabling DirectX games to work smoothly under Linux, they chose to use the Rewind fork because of its more lax licensing. Transgaming's fork of Rewind is available under the Aladdin Free Public License, which doesn't really live up to it's name – you can't take the Transgaming code and sell it, nor can you redistribute it with anything that you are selling.**

## LINKS TO MORE INFO

**Aladdin Free Public Licence** www.transgaming.com/license.php?source=1
**(Free)BSD licence** www.freebsd.org/copyright/license.html
**GNU General Public Licence** www.gnu.org/copyleft/gpl.html
**GNU Lesser GPL** www.gnu.org/copyleft/lesser.html
**Microsoft Shared Source Licence** http://msdn.microsoft.com/msdn-files/027/001/901/ShSourceCLIbetaLicense.htm
**Mozilla Public Licence** www.mozilla.org/MPL/
**Sun Community Source Licence** wwws.sun.com/981208/scsl/principles.html

**Debian Social Contract** www.debian.org/social_contract
**More on the "freeness" of the FDL** http://lists.debian.org/debian-legal/2003/debian-legal-200304/msg00189.html
**The Open Source Definition** www.opensource.org/docs/definition.php

## Mozilla Public Licence

The Mozilla Public Licence (MPL), and its close friend the Netscape Public Licence (NPL), are both relatively free and open licences. Mozilla is currently licensed under both the MPL and the NPL, with the main difference between the two being that products under the NPL (commercial versions of Mozilla) do not have to have their improvements contributed back into the community.

A recent effort for Mozilla has been to get the source code under a tri-licence scheme, whereby all the source code will be available under the NPL, the GPL, and the LGPL – users can choose which they'd rather use.

## Qt licence

Qt used to be licenced under the Trolltech QPL, which is not a free licence as it does not allow you to redistribute a QPL program you've modified – you have to distribute the original with patches. It was restrictions such as this that fuelled the development of the GNOME desktop environment, which was licensed under the GPL. However, Qt was released under the GPL some time ago, so it is good to use across the board. Trolltech also has a special, closed-source licence option available for customers wishing to sell programs without needing to distribute the source with them, which means that Qt is equally attractive to open-source and closed-source developers. Incidentally, the same open/ closed-source licensing model is used by MySQL, which shows that two of the most popular commercial open-source companies seem to be doing remarkably well using a dual-licensing policy.

## Artistic Licence

The FSF slates the artistic licence on the grounds that it is too vague, although it is commonly agreed that a new version, known as the Artistic Licence 2.0, is a free software licence, compatible with the GPL. Perl 6 may well be released under the Artistic Licence 2.0, which would give it quite a boost.

## Sun Community Source Licence

This is another licence that sounds more free than it is. In essence, the SCSL allows you to download, read, experiment with, and use Sun source code at no cost. It is beer free, if you will. However, it specifically does not allow you to distribute the code, original or your modified version. As such, as it without cost, but not truly free.

Sun do have another licence, called the Sun Public Licence, which *is* free in both sense of the words, and so is safe to use. Are you confused yet?

## Licence variations

Very often licences are just variations of each other. For example, the *Apache* software licence is actually not free in every sense of the word – it forces the inclusion of various notices, and also forces rules on the naming of derived products. This same licence was used as the base for both the PHP licence and the Phorum licence, so the non-free licence has propagated to other products, and is likely to continue to spread.

As such, it's pretty crucial that you pick the most free licence you're comfortable with, as it's possible others may use your example for their own software. There's also the possibility that some authors feel the need to tweak some parts of their chosen licence, "to make it better" or more applicable for their own piece of software. This is very rarely a smart idea for two reasons. Firstly, several licences (particularly the GPL and BSD-style licences) are very well known; few people need to wonder whether it's safe to use and modify software that comes under such licences. Secondly, the GPL was written by legal experts who know what they're doing and know what is enforceable in court – unless you're David Harris (*LXF*'s law writer) or another legal eagle it's probably best you trust the FSF!

## How free is free?

Many people consider BSD-style licences non-free, which is where the arguments generally start. Is it better for programs to be used absolutely without limits (as in BSD), or is it better to force users to redistribute their modifications (as in the GPL) and benefit the community?

This is an argument that has gone on for a long time, and is rarely debated on friendly terms, to say the least! It's a fact that BSD pretty much has the most free terms of usage possible: "use our code for free, change it, sell it, and don't worry about opening the source code for others to see". In the ideal world, this would be a great licence – everyone shares their code freely without having to worry about complicated licences.

However, this isn't an ideal world, so what's more likely to happen is that companies take hundreds of megabytes of BSD source code and give little to nothing back in return. Apple's Mac OS X, based largely on BSD, is an excellent product, but it's clear they got the better half of the bargain in the source code exchange! Apple's Safari browser, however, was based on KHTML, part of KDE and licensed under the LGPL, so when they made changes to their port of KHTML to make it better, they gave all their changes back to the KDE team. Which of the two would you consider was best for the community?

While the BSD-style licence promotes source code freedom in a sincere, honest manner, the end result is that it's a little naive. This is where the GPL steps in – usage-wise it's not quite as free (you *must* make changes available), but by being slightly stricter on usage it allows everyone to benefit from work. The GPL's strictness on redistribution means that GPL software isn't likely to ever be replaced – no one can take the source, improve it, and keep the changes to themselves for profit.

So, if you ever find someone wondering why Linux isn't available under the BSD licence, ask them to consider how they'd feel if Microsoft or Apple were able to take any or all the parts of the Linux kernel they wanted and copy it verbatim into their own software!

## Free documentation

The GPL and LGPL are designed to cover code, so what do you do if you want to protect your documentation? The GNU also has a Free Documentation Licence that covers textual work, and is mostly free – but not quite as free as the GPL. For example, the FDL does require that anyone who changes an FDL document must release their modifications under the FDL in order to help the community. However, it does have the restriction that authors can mark some sections as being invariant, and these sections cannot (legally) be removed at a later date.

> "GPL software isn't likely to ever be replaced – no one can improve the source and keep changes to themselves for profit."

It's this concept of invariant sections that have brought many to believe the GNU FDL isn't free. The Debian team, always particularly wary of non-free material, are even considering taking action to clear up the problem – either by replacing the licence with their own truly free version, or by getting GNU to alter the licence.

In the meantime, authors might find it best to simply mark no sections as invariant when they use the FDL. Hopefully some action should be taken soon to ensure people can create truly free texts.

## Conclusion

Understanding licensing doesn't need to be a complicated topic, and it continues to be one of the most important issues affecting the open-source community today. Version 2 of the GPL has been around for over 12 years now, and continues to serve the community very well – when we see a product covered by the GPL, we don't have to worry about what we're installing. It's important to realise that many licences called "Free", "Public" or "Community" are rarely free at all - after all, if they were completely free they would use a BSD-style licence, or, better, the GPL.

It's important for all of us to take the time to propagate positive licensing in the community. If you're a programmer yourself, you can make sure you always have your work licensed under the most fitting option for your software so that your work can have maximum impact on the community. If you're a software consumer, you can get in touch with authors who are using non-free licences and ask them (nicely) to consider changing. You also have the choice to choose your software based upon its licence – if you find software that isn't free, you are under no obligation to use it, and the chances are there's a free version available. Debian, for example, gives users the option when configuring the APT database, "Do you want to use non-free software?"

If we all work together, sharing software, ideas, help, and learning, Linux will continue to be the embodiment of the Open Source movement. All it takes is just a little more care when using software – do your part! ■

### WHAT ABOUT WARRANTIES?

**ONE CRUCIAL PROTECTION the GPL and some other free licence offer to authors is protection from warranty claims. As software is given away free, and can be modified without the author's consent, it would be incredibly dangerous if authors of GPLed software could be sued if their programs didn't work correctly. Although it's true that Microsoft has yet to be successfully sued for the times they've experienced problems with their software, it's still not worth taking the risk with free software. To make this crystal clear to end users, this is formalised in no uncertain terms in the GPL.**

# Using NFS

**Far from it being a 'heritage' system, NICK VEITCH outlines how the forthcoming v4 of NFS can leverage your storage space.**

Despite being technologically as old as the hills, *NFS* is still in widespread use, and not without some good reasons. Mainly these fall into the camps of simplicity and ease of use, but it is also useful to remember that *NFS* is available on virtually every computing platform imaginable, so it can be a real boon to heterogenous networks.

The basic principles behind *NFS* are simple. The idea is that the client can mount and access a remote volume as though it were a local one. You probably don't need too many clues as to why this is useful – not only for centralised storage like NAS devices, but also for development and hotdesk environments, where users can access their files, or their groups files, easily and in a uniform way.

*NFS* is based on the Remote Procedure Call (RPC) protocol originally developed by Sun Microsystems to enable flexible client/server interactions. By using an XDR (eXternal Data Representation) mechanism, the data exchanged is machine independent, to get around issues cause by different processors. *NFS* uses a system of RPC calls to create the impression that all filesystem requests are handled locally, and therefore creates a kind of virtual filesystem. As far as the client is concerned, and more importantly, as far as the software the client computer is running is concerned, the *NFS* mounted disk is just the same as any other filesystem mounted locally.

*NFS* has been a feature of Linux kernels for some time, but the latest release version of *NFSv3* is present in 2.2.19 or later versions. If you wish to use the experimental TCP support, you will need to use 2.4.19 or later kernels, or apply the patch files (**http://prdownloads.sourceforge.net/nfs**).

## IMPORTANT SECURITY ADVISORY

***NFS* IS IN MANY WAYS EASY TO USE because it is simple. Permissions on files on the virtual NFS filesystem are viewed on the client in exactly the same way as local files would be. *Eg* if a file on the NFS export volume belongs to uid 505, a user with uid 505 on the client will be treated as the owner. There is no authentication protocol associated with *NFS*.**

**This means anyone with root access to the client can gain access to any user owned files (with the exception of root, unless you permit it) on the server.**

**You can treat all users as nobody on the server and force a read-only filsystem, but this may impair why you set up the system in the first place. With *NFS*, you basically have to trust everyone or nobody. This is one f the reasons it's so important to set up the list of allowed clients properly.**

***NFSv4* for Linux will include the ability to use RPCSEC and Kerberos to provide an authentication layer for *NFS* on Linux.**

## Configuring NFS

To export a volume via *NFS*, you simply need to edit the /etc/exports file to include information on the directory tree you wish to export and who should be able to access it. The format for this is

```
<directory>  <allowed IP address or range>(options) ...
```

You can add multiple IP addresses for clients separated by spaces, or you can specify a range of addresses with an IP address and netmask. An example configuration might be:

```
/home/shared   172.18.42.120(rw) 192.168.0.0/255.255.255.0(ro)
```

In this example we are allowing a connection from a specific IP address (172.18.42.120) with read/write access, and blanket access to anyone on the 192.168.0.x network with read-only access.

It is also possible to select specific hostnames instead of IP addresses (as long as these can be resolved properly on the server) and use wildcards. Early implementations of NFS had some problems with wildcards, and it is generally not recommended as you may inadvertently open the doors to other networks you didn't intend to give permissions to.

Having setup the directory to be shared, and assuming the *NFS* service daemon is running, you can force **nfsd** to re-read the export list with the command:

```
exportfs -ra
```

the exported directory, and all its subdirectories will now be available to users on the specified clients.

One important limitation here is that you can't export a parent or child directory of a directory that has already been exported. The latter is pointless anyway, as it will be available in the original export already.

Another consideration s the underlying filesystem. Generally, it is best to use Linux native filesystems for *NFS* exports. MS Windows filesystems are particularly unsuitable, mainly because of the problems associated with permissions.

## Clients

On the client side, mounting an NFS export is done in a similar way to mounting any other device – you just use

**mount** (if you are using a particularly old kernel, patched for *NFSv3* support, you will probably need to rebuild the mount command also).

```
mount 192.168.0.128:/home/shared /mnt/shared
```

On the client machine, this mounts the filesystem on the directory tree at /mnt/shared, from where it can be used like any other part of the filesystem. User permissions are maintained, (see *Security* box below left) so individuals can manage the sharing of their own individual files within this workspace as usual (though to be effective, you will need to be using a directory service such as NIS or LDAP for user accounts). The permissions are tied to the uid and gid numbers associated with the file on the server.

By default, *NFS* will use the 'squashroot' option. This means that a user logged on as root on the client will not have root permissions for the files in the *NFS* share.

## NFS ALTERNATIVES

**Samba/CIFS** – probably the most high profile alternative to NFS is *Samba*, the Linux version of the SMB file sharing protocol, now renamed by Microsoft in their implementation to CIFS. The advantages of SMB file sharing is that the system also carries out authentication and includes support for other services (*eg* printing). And it works well on Windows systems. The disadvantage is that as SMB uses different attributes, permissions handling can be problematic.

**Andrew File System** – AFS was originated by Carnegie Mellon University, further developed by Transarc (now part of IBM) and released through the IBM Public Source licence in late 2000. Although not as widespread in some respects as NFS, AFS is actually more than a network filesystem – it is a distributed filesystem, which allows the files to be stored on a number of different servers. Find out more at www.openafs.org

> "Mounting an NFS export is done in a similar way to mounting any other device – you just use **mount**."

## Speed

Nobody ever said *NFS* was going to be fast. Implementing it over TCP does have some benefits, as does playing with the default blocksizes, but the fact is that accessing a filesystem over ethernet is never going to be very fast. *NFS* was invented in the days when people were only concerned with sharing small text documents, not the multi-megabyte files of today, so speed has become more of a concern relatively recently.

There are a few things that can be done to improve performance though. We mentioned blocksizes, which are the default packet sizes of information used by *NFS*. Using larger block sizes can improve performance significantly – however, different kernels have different defaults, and older implementations will only support up to 8k blocksizes, while latest implementation will allow up to 32K, though the default values may be different. The blocksize for read and write can be set with the mount options **rsize** and **wsize** respectively:

```
mount 192.168.0.128:/home/shared /mnt/shared -o
rsize=32k,wsize=32k
```

Bigger blocksizes don't always result in better speeds. It is worth experimenting and using some sort of filesystem benchmark like *bonnie* to determine the best values for your system.

The filesystem being exported can itself be a big factor. Using inherently journalled systems will reduce performance – if you want journalling try to use *ext3* or *Resiser* and locate the actual journal in a different volume (or better still in non-volatile RAM).

## The future of NFS

In spite of the fact that *Samba* seems to steal the limelight in the networked filesystem space, *NFS* is still being very actively developed. It is used so extensively that there is still al lot of interest in implementing new features, though often they take a while to come about. Implementing *NFS* over TCP for example was first suggested in the mid-nineties. Using TCP instead of UDP is seen as a good way of improving the performance of *NFS*, and it's manageability.  The TCP server support is still experimental in latest versions of the stable kernel, though it is said to work very well

*NFSv4* is currently in development, and expected to make it in to the next release version of the kernel (2.6) when it is released later in the year. Among the new features are the introduction of sate which will help with file locking and the preservation of data when servers or clients become disconnected.

Added to that, better caching and 'file-delegation', where the client is trusted to keep local versions of files being modified until requested by another process. This will dramatically reduce the amount of network traffic associated with simple file modifications, especially in cases where only one client is likely to be accessing the data at any given time.

Improved security features are also important. The implementation of RPCSEC and the possibility of using authentication plugins such as Kerberos will greatly improve the suitability of *NFS* on Linux for a variety of tasks.

This has just been a brief overview of how *NFS* can be useful. For more detail and information on upcoming changes, visit the Linux *NFS* site, which you can find at http://nfs.sourceforge.net, and check out the excellent HOWTO also housed there. ■

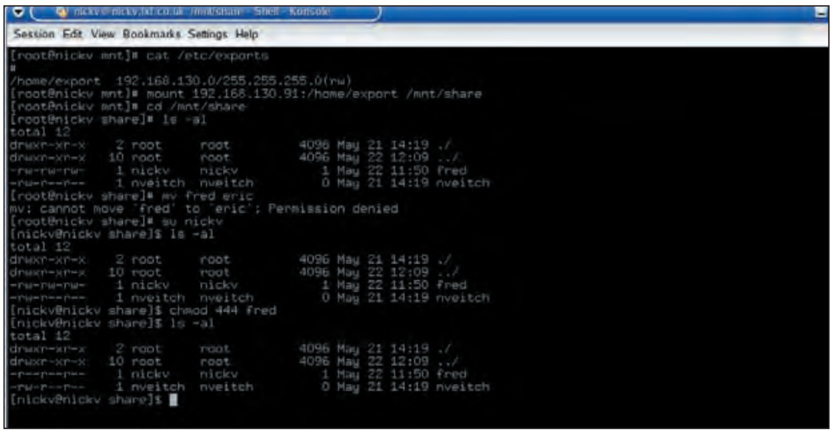
nickv@nickv.lxf.co.uk: /mnt/share - Shell - Konsole
Session Edit View Bookmarks Settings Help

```
[root@nickv mnt]# cat /etc/exports
#
/home/export  192.168.130.0/255.255.255.0(rw)
[root@nickv mnt]# mount 192.168.130.91:/home/export /mnt/share
[root@nickv mnt]# cd /mnt/share
[root@nickv share]# ls -al
total 12
drwxr-xr-x    2 root     root         4096 May 21 14:19 ./
drwxr-xr-x   10 root     root         4096 May 22 12:09 ../
-rw-rw-rw-    1 nickv    nickv           1 May 22 11:50 fred
-rw-r--r--    1 nveitch  nveitch         0 May 21 14:19 nveitch
[root@nickv share]# mv fred eric
mv: cannot move `fred' to `eric': Permission denied
[root@nickv share]# su nickv
[nickv@nickv share]$ ls -al
total 12
drwxr-xr-x    2 root     root         4096 May 21 14:19 ./
drwxr-xr-x   10 root     root         4096 May 22 12:09 ../
-rw-rw-rw-    1 nickv    nickv           1 May 22 11:50 fred
-rw-r--r--    1 nveitch  nveitch         0 May 21 14:19 nveitch
[nickv@nickv share]$ chmod 444 fred
[nickv@nickv share]$ ls -al
total 12
drwxr-xr-x    2 root     root         4096 May 21 14:19 ./
drwxr-xr-x   10 root     root         4096 May 22 12:09 ../
-r--r--r--    1 nickv    nickv           1 May 22 11:50 fred
-rw-r--r--    1 nveitch  nveitch         0 May 21 14:19 nveitch
[nickv@nickv share]$ 
```

**Mounting an NFS filesystem – it isn't exciting, but it works.**

# Service monitoring with Nagios

**NICK VEITCH assesses a new tool to keep your paranoia in check.**

When there is a problem with a server, or a service running on it, the last thing your average sysadmin wants is for someone else to find out first. Not only is it embarrassing and potentially job-threatening, but it doesn't do much for customer confidence in the company if the visitors to your corporate website are the first to notice it isn't working.

Downtime costs money as well as faith in your corporate abilities, and while those with the cash or business need can layer many servers with huge amounts of redundancy, for all IT admins, it's important to know what's gone wrong, where, why, and how tricky it's going to be to fix.

Systems monitoring applications are nothing new. Linux has a hatful of them – from the simple monitors built in to webmin (not very helpful if the mailserver, or webmin itself goes down) to the more traditional all-singing and dancing commercial tools. For a long time, *NetSaint* has been highly regarded and respected as probably the best software for this task (and it's published under GPL) on Linux systems. But development on *NetSaint* has stopped, and it has now evolved into a new system, *Nagios*, which reached a release version just before Christmas.

Through use of plugin modules and a comprehensive system of configuration files, *Nagios* allows groups of users or single admins to monitor pretty much any service you can think of, and alert the relevant people if something should go wrong. Services can be monitored on remote systems more effectively than before, even when *Nagios* itself is not installed on the target. With support for creating redundant monitoring hosts, and a handy web interface to boot, *Nagios* should wear NetSaint's former mantle with ease.

## NAGAT

### Configuring made easier

**REALISING THAT MANUALLY EDITING A dozen or so config files might be a tall order, a project related to *Nagios* was created, called *Nagat*. The purpose of *Nagat* is to provide an easier way to manage all those configuration options. *Nagat* is written in PHP, and basically is just a front end to edit the standard config files. You'll still need to know all the terminology used, although to a great extent, creating objects such as server definitions is a lot easier.**

***Nagat* merely needs to be set up on a webserver, from where the scripts can probe the current configuration files and allow easier editing, and the creation of objects (hosts, contacts, groups, service definitions and so on).**

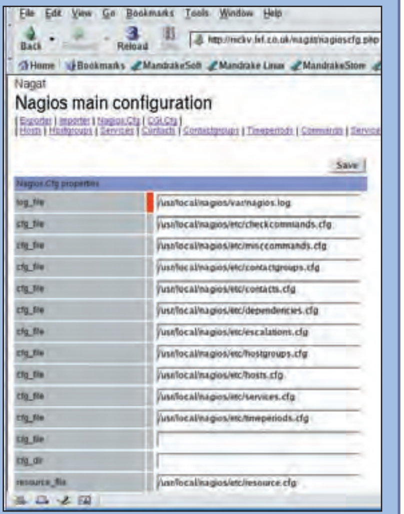

File Edit View Go Bookmarks Tools Window Help
Back
Reload
http://nickv.lxf.co.uk/nagat/nagioscfg.php
Home
Bookmarks
MandrakeSoft
Mandrake Linux
MandrakeStore
Save

Nagat

Nagios main configuration

| Exporter | Importer | Nagios.Cfg | CGI.Cfg |
| Hosts | Hostgroups | Services | Contacts | Contactgroups | Timeperiods | Commands | Service

| Nagios.Cfg properties | |
|---|---|
| log_file | /usr/local/nagios/var/nagios.log |
| cfg_file | /usr/local/nagios/etc/checkcommands.cfg |
| cfg_file | /usr/local/nagios/etc/misccommands.cfg |
| cfg_file | /usr/local/nagios/etc/contactgroups.cfg |
| cfg_file | /usr/local/nagios/etc/contacts.cfg |
| cfg_file | /usr/local/nagios/etc/dependencies.cfg |
| cfg_file | /usr/local/nagios/etc/escalations.cfg |
| cfg_file | /usr/local/nagios/etc/hostgroups.cfg |
| cfg_file | /usr/local/nagios/etc/hosts.cfg |
| cfg_file | /usr/local/nagios/etc/services.cfg |
| cfg_file | /usr/local/nagios/etc/timeperiods.cfg |
| cfg_file | |
| cfg_dir | |
| resource_file | /usr/local/nagios/etc/resource.cfg |

**It's simple, but slightly easier to follow than editing the config files by hand**

## Install

Some distributions do have packages available for *Nagios*, but this is one piece of software that is far better installed from the source. Packages for some distributions seem to end up putting the various plugins, executables and conf files in unusual places, and since knowing where these are is critical for some parts of the setup, it's worth taking the time to configure the software properly.

Download the latest source, unpack and read the README and INSTALLING files before you configure! There are a couple of important options here, particularly if you want to use the web interface. These are

- **with-cgiurl=<dir>** which sets the directory used for the CGI scripts, and
- **with-htmurl=<dir>** which controls where the web pages are installed.

There are also options for configuring the user and group names for *Nagios* and various options for storing data in MySQL or PostgreSQL databases. While the defaults can be left for many of these in most cases, it's important to know what they are, so you can configure other services accordingly (e.g. your webserver).

## The config scripts

There may well be easier ways to monitor simple service status, but *Nagios* forces you to actually generate a reasonable configuration before it will even run. This isn't the type of application where you can just get away with copying the example config file and set it running. And with good reason – to monitor your system effectively, *Nagios* needs to know how and what it needs to monitor.

Even so, 13 separate config files might seem a little excessive to anyone, and perhaps this is one area where future versions can improve – manually entering all the details for every service and host that is to be monitored , and restarting the *Nagios* daemon each time can be more than a little troublesome. See the box on *Nagat* below left.

The creation of the config files is so involved that we couldn't hope to cover it here, but example config files are available as a download from the *Nagios* site, and there is plenty of documentation to read through. It will probably take a few hours to configure a basic service monitor.

Monitored servers don't have to be running *Nagios* either. For basic, publicly exposed services such as FTP or http, there is no need to run any client software on the target. If you need to keep abreast of disk space and other usually private information, you can configure the *Nagios* service to use ssh and run tools locally on the target machine, though it's up to you whether routinely using such features might expose some extra vulnerability.

## Web interface

Once you have configured *Nagios* successfully and initialised the daemon, that should be it. The designated contacts should receive messages if any of the services need attention, and the system will largely take care of itself. It can be reassuring to check that everything's running okay though, which is why *Nagios* includes a web interface. CGI scripts take care of most of the functionality of this web presence. There are plenty of scripts included, but it would be a fairly trivial matter to develop more to suit specific needs.

From the web interface you can monitor all the registered hosts and hostgroups, see an overall summary and check on the status of *Nagios* itself. Several useful reporting functions are also available to track availability and the history of checks and alerts, as well as which contacts have been notified and when.

Certain actions will be unavailable unless you have enabled authentication for the website. This is a sensible precaution, though if you are running a server other than *Apache1.3.x,* there is no help in the documentation about how to set it up. If you opt not to include authentication (which is wise, even if the *Nagios* server is only accessible internally) some of the features of the system will be disabled by default (*eg* using the web interface to schedule immediate service checks) and will take much tinkering to get working.

While the web interface is optional, it can serve as a reassurance to the more paranoid sysadmin.

## The best tool

As with many other aspects of IT, the best solution is the one that works for you. *Nagios* is a very capable tool that can be configured to monitor huge and complicated networks of services. However, if you just want to check that one or two webservers are running, there are simpler tools for the job. *Nagios* requires you to put the effort in to understand the system, how the hierarchy and interrelations of host and service objects work, and even with a php frontend like *Nagat*, you will still have to understand every line of the config files. This might be overkill for you, but if you need a reliable tool enabling teams of IT people to manage a range of servers and services, this is definitely a solution worth investing some time in. ■

## OTHER SERVICE MONITOR SOFTWARE

### Both GPL and proprietary

***NAGIOS* ISN'T THE ONLY SERVICE MONITOR. There are plenty of GPL monitoring tools and some commercial offerings for Linux too. Here are just some of them:**

**Sysmon** (www.sysmon.org) – Now at version 0.91, this curses based monitor also has very flexible and intelligent configuration options.

**Big Brother** (www.bb4.com/bb) – Another well-known project, Big Brother is available free, or there is a professional version with a few more features and 24x7 support. A simple web display clearly shows network status, and a number of alert systems means you're never too far away from the bad news.

**Autostatus** (www.angio.net/consult/autostatus) – Has a reputation as a solid performer that isn't too hard to administer, but the last version was released in 2000.

**Penemo** (www.penemo.org) – While it lacks some of the features of some of the other software mentioned here, Penemo may still be worth looking at. Written completely in Perl, it may be just the system to tweak for more specific needs.

**Webmin** (www.webmin.com) – Originally developed by Caldera (now SCO), webmin is the well-known box of tricks that has saved many a sysadmin. Unless the whole server goes down, it has a useful service monitor function.

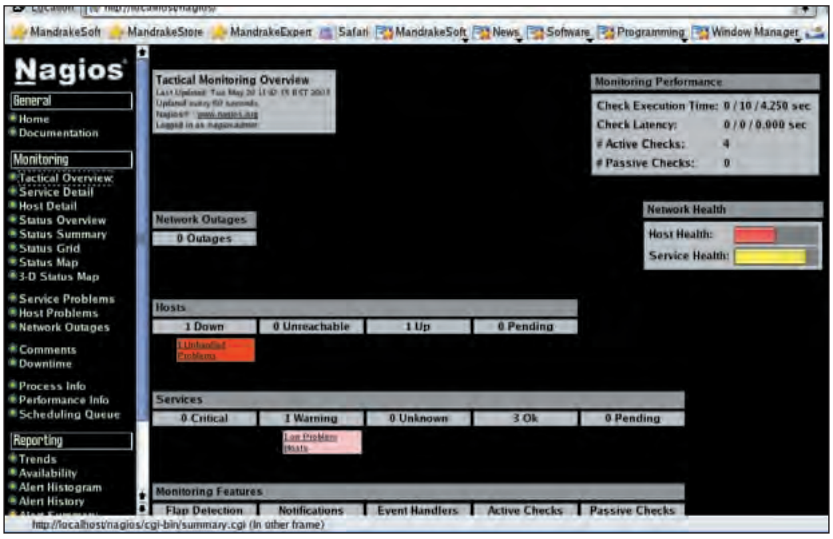

Check Execution Time: 0 / 10 / 4.250 sec
Check Latency: 0 / 0 / 0.000 sec
# Active Checks: 4
# Passive Checks: 0

0 Outages

Hosts: 1 Down | 0 Unreachable | 1 Up | 0 Pending

Services: 0 Critical | 1 Warning | 0 Unknown | 3 Ok | 0 Pending

http://localhost/nagios/cgi-bin/summary.cgi (In other frame)

A quick overview screen alerts you to any problems.

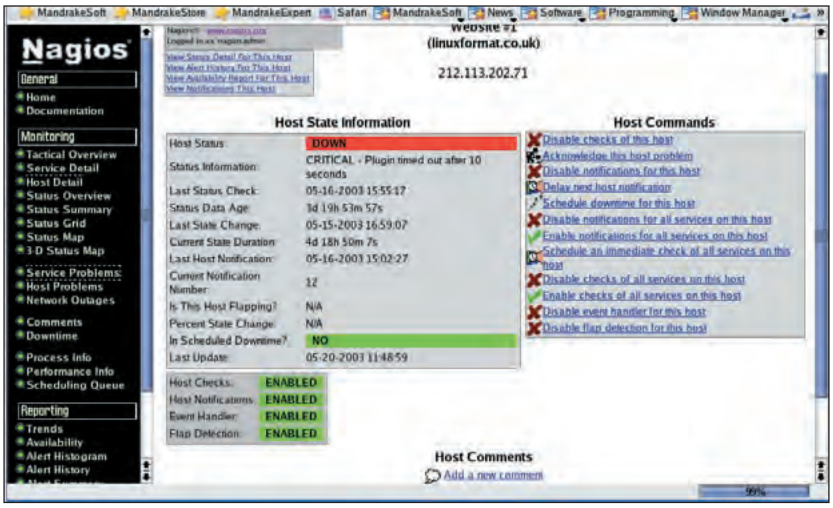

Website #1
(linuxformat.co.uk)

212.113.202.71

| | |
|---|---|
| Host Status: | DOWN |
| Status Information: | CRITICAL - Plugin timed out after 10 seconds |
| Last Status Check: | 05-16-2003 15:55:17 |
| Status Data Age: | 3d 19h 53m 57s |
| Last State Change: | 05-15-2003 16:59:07 |
| Current State Duration: | 4d 18h 50m 7s |
| Last Host Notification: | 05-16-2003 15:02:27 |
| Current Notification Number: | 12 |
| Is This Host Flapping? | N/A |
| Percent State Change: | N/A |
| In Scheduled Downtime? | NO |
| Last Update: | 05-20-2003 11:48:59 |

| | |
|---|---|
| Host Checks: | ENABLED |
| Host Notifications: | ENABLED |
| Event Handler: | ENABLED |
| Flap Detection: | ENABLED |

Each host has an individual status page with more admin options.