

# news@UK

*The Newsletter of the UK UNIX Users Group*

Published electronically at <http://www.ukuug.org/newsletter/>

---

Volume 13, Number 1

ISSN 0965-9412

March 2004

---

## Contents

News from the Secretariat	3
Tutorial: Advanced Topics in DNS Administration - by Jim Reid	3
Emergency General Meeting	4
Open Source Award 2004	4
Announcement: AUUG 2004 Call for Papers	5
Announcement: Linux User and Developer Expo 2004	5
Announcement: LISA '04 - call for papers	5
OSWC-2004 Malaga	5
UKUUG Winter Conference	7
FLOSSIE Conference 2004	9
OSS Watch launch conference: December 11th 2003	9
Opening the JVM	11
Book review: Google Pocket Guide	13
Book review: The Art of UNIX Programming	13
Cookbooks Introduction	14
Book review: sendmail Cookbook	15
Book review: Active Directory Cookbook	15
Book review: Learning Python, 2nd Edition	16
Book review: Postfix: The Definitive Guide	17
Book review: Computer Science and Perl Programming	18
Book review: Cisco Cookbook	19
Book review: Apache Cookbook	19

<b>Book review: Learning XSLT</b>	<b>20</b>
<b>Book review: Teach Yourself Linux</b>	<b>21</b>
<b>Book review: The Best of the Joy of Tech</b>	<b>22</b>
<b>Book review: Mac OS X The Missing Manual Panther Edition</b>	<b>22</b>
<b>Book review: Running Mac OS X Panther</b>	<b>22</b>
<b>Contacts</b>	<b>24</b>

## News from the Secretariat

**Jane Morrison**

Firstly I want to thank everyone who kindly sent in their subscription payments so promptly. We have received a good number of early payments. Those remaining outstanding will be reminded this month and any still outstanding at the end of March will not receive the next issue (June) of the Newsletter.

The Winter conference and tutorial held on 25th and 26th February was very successful. We achieved the highest attendance numbers at a Winter event for many years and this was due to the strong programme put together by Ray Miller. I wonder if the sea-side venue of Bournemouth was also a factor in this? [The snow on Thursday afternoon was an added attraction ...RW]

There is an account of the conference elsewhere in the pages that follow, and the enclosed CD contains both the tutorial notes and the papers from the conference.

We are currently looking at venues in Birmingham for the 2005 event.

You will have received in a prior mailing details of our next event - Advanced Topics in DNS Administration, which will be held on 20th April in London. The tutor, Jim Reid has provided tutorials for us before and as these have always been very popular, you are encouraged to book your place early to avoid disappointment.

The Linux 2004 Technical Conference will be held in Leeds from 29th July to 1st August. Details, including a web form for potential speakers' submissions is available at <http://www.ukuug.org/events/linux04/>

The next Newsletter will be the June issue and the copy date is Friday 21st May.

We welcome any comments about past or future events, and in particular any contributions from members to this Newsletter.

UKUUG Secretariat  
PO Box 37  
Buntingford  
Herts  
SG9 9UQ  
Tel: 01763 273475  
Fax: 01763 273255  
[office@ukuug.org](mailto:office@ukuug.org)  
<http://www.ukuug.org>

---

## **Tutorial: Advanced Topics in DNS Administration - by Jim Reid**

**20th April 2004**

**Imperial Hotel, London**

Who should attend: DNS administrators who wish to extend their understanding of how to configure and manage name servers running BIND9. Attendees should have some experience of running a BIND8 or BIND9 name server and be familiar with DNS jargon for resource records, as well as the syntax of zone files and `named.conf`. This tutorial will answer the question: "I've set up master (primary) and slave (secondary) name servers. What else can I do with the name server?"

Topics include: The BIND9 logging subsystem  
Getting the most from the name server's logs

Managing the name server with rndc  
Configuring split DNS: internal and external versions of a domain  
Using the views mechanism of BIND9 to implement split DNS  
Setting up an internal root server  
Securing the name server  
Dynamic DNS (DDNS)  
The Lightweight Resolver Daemon, lwresd  
Secure DNS (DNSSEC)

Jim Reid started using a PDP11/45 running V7 UNIX over 20 years ago and has been working with UNIX systems ever since. He worked for three years at Origin on behalf of Philips Electronics, where he wrote a DNS management system and designed, built, and ran the DNS infrastructure for the corporate network, one of the biggest in the world. He has over a decade's experience in writing and teaching training courses ranging from kernel internals, through system administration and network security, to DNS administration. He's a frequent speaker and tutorial presenter at conferences and workshops in Europe and the USA. He chairs the DNS Working Group at RIPE and is active in ENUM, serving as Chair and Technical Manager of the UK ENUM Trial Group.

Places are limited - see booking form and further details:  
<http://www.ukuug.org/events/dnstutorial04/>

---

## Emergency General Meeting

### *UKUUG Administration*

The Emergency General Meeting was held as announced in Bournemouth on Wednesday 25th February after the first day of the Winter conference.

Those present unanimously agreed to increase the number of Council members to eight.

The minutes of the EGM are enclosed with this Newsletter mailing.

---

## Open Source Award 2004

Win £500 and a trip to the 2004 Open Source Convention, USA.

UKUUG is again seeking significant contributions to free and open source software; this could be in the form of an article or paper, software product, or other contribution.

Do you have a contribution that is or might be significant? Student projects are eligible. If your work is part of a joint project, please state clearly your own contribution.

Please send a one-page summary (accompanied by an abstract and a short biography of about 250-500 words) of your project, software, book, or idea to UKUUG's Secretariat.

The prize for the best contribution is £500 and, thanks to O'Reilly (UK), a pass to the 2004 Open Source Convention (USA in July) with monies to cover travel and accommodation.

The deadline is 2 April 2004. Please enter. Further details are at:

<http://www.ukuug.org/osa>

---

## **Announcement: AUUG 2004 Call for Papers**

We have received notification from David Purdue of the AUUG 2004 event which will have the theme "Who Are You? Identification and Authorisation Issues in Computing".

Proposals for papers and tutorials are invited on matters related to identification, authentication and authorisation, as well as cryptography, internet security and privacy and other security issues.

Other more general submissions in the area of open standards, open source software and networking are also welcome.

Abstracts and proposals are due by 7th May 2004.

Further details are available here:

<http://www.auug.org.au/events/2004/auug2004/>

---

## **Announcement: Linux User and Developer Expo 2004**

The LinuxUser and Developer Expo 2004 will be held at Olympia on the 20th and 21st April 2004. The event will incorporate a conference.

Full details are here:

<http://www.linuxuser.co.uk/expo/>

UKUUG will be represented in the .org area at this event.

---

## **Announcement: LISA '04 - call for papers**

We have received notification from Lee Damon of the USENIX LISA '04 event which will be held between the 14th and 19th November 2004 in Atlanta, Georgia, USA.

The organisers invite abstracts and outlines for refereed papers, and are also looking for suggestions for talks by invited speakers and ideas for possible tutorials.

Details of the event and the full call for papers can be found here:

<http://www.usenix.org/lisa04/>

Submissions are due by April 20th.

---

## **OSWC-2004 Malaga**

***Charles Curran***

### **Introduction**

Where is the ideal place in Europe to hold a conference? OK, after Dublin, where? – Malaga. What is a hot topic (still) that you could get thousands of representatives from all areas (not just techies) to come along to? – Open source/free software/software libre. I was fortunate to already be in Andalusia when the (first) Open Source World Conference was held in Malaga's Palacio de Ferias & Congressos (that sounds better than the Trade fair and Congress Centre, and the setting was indeed very pleasant – very different from the seedy exhibition halls generally used in the UK) from 18-20 February.

The stated purpose of this event was 'to go beyond the purely technical or instrumental aspects and tackle free software as an element capable of extending the Information and Knowledge

---

Society, not only as the basic component of access and mobility, but also as an enabler in making a reality the use of services and other factors that could generate benefits at personal and social levels'. So, they tried to keep it simple! Regarding terminology, since the (Spanish) title used Software Libre, I shall use that or rather the abbreviation SL. Very few wasted their time arguing the difference between open source, free software, FLOSS, etc.

To speak or otherwise participate, they attracted those from all levels of development and implementation of solutions based on SL: government representatives, those from public administration, directors of corporations and SMEs, those from academia, hackers, representatives of users groups and trade unions, et al, and from all around the world. If you had your eye on the programme weeks before the event you would have noticed some big gaps; they even listed names of key players in SL who had not even been invited as well as those who had not responded to their invitation. On the UK front, it was a great pity that the invitations and other alerts to government, companies proposing SL solutions, and other commercial groups were not taken up. The organizers pulled out a couple of trump cards at the beginning of February to get the crowds in: Prince Felipe (not Philip), and possibly a bigger draw, the former prime minister, Felipe González. With them on board, the registrations increased from around 500 to more than 2000.

What they hoped to cover were key areas of SL and its adoption, such as the future trends of SL; the role of public administrations and their expansion and development; the viability of current business models; sustainable development and SL; globalization and SL; large successful projects; new social spaces and their implication with SL; public policies; consequences of adoption (on, e.g. environment and social space). And to find answers for the issues raised.

Who put on the event? Well, the main sponsor was the Andalusian Junta (i.e. regional government) who have a vested interest in SL and they appear to have been the first anywhere to legislate about SL (decree 72/2003). They are also distributing in a big way GuadaLinEx, which is based on the LinEx distribution from Extremadura; this was aimed initially at local government, libraries, and open-access centres and now at domestic use.

<http://www.guadalinux.org>

<http://www.linex.org>

But the people who seem to have put it together were Juan Reig and Manolo Martín Matas, and JuanTomás García and their teams.

### **The Event**

There were a few plenary sessions but mainly there were about five parallel tracks at once. The conference web site contains the programme and proceedings with some of the papers (a lot are only in Spanish) (available only as a single 14MB PDF).

<http://www.opensourceworldconference.com>

Most of the talks were in Spanish but simultaneous translation was available for talks in the main auditorium. Alongside the conference, key Debian developers held their first derivative meeting to sort out the architectural needs, etc. of the various Debian-based distributions. There were also some short workshops. There was also a sizeable exhibition area.

First up was the local mayor who said that as well as being at the vanguard of artistic innovation, Málaga was also there with SL. He welcomed the generosity of programmers in making SL available and said we owed a special gratitude to Richard Stallman and Linus Torvalds. He said that SL means that we can reassess our needs with a greater degree of competitiveness, reliability, and compatibility. So he or his speech writer was on the ball.

Sketches of some of the talks I heard(\*) are included in my fuller report. There were about 125 speakers grouped in sessions on: \*Inaugural talks from presidents and princes; \*Keynote on Convivial Software; \*Public Administration and SL; SL Development and Tools; Benchmarking and Case Studies; \*Business Models, Services, and Legal Framework, SL Applications;

Office Automation, Integration, and Desktop; \*Plenary on Globalization, Information Society, and SL; \*Corporate Strategies and SL; \*SL in the World; Open Technology; Universities and E-Learning; \*SL in Europe; Industrial and Sector Applications; Corporations–Mobility and Telecomms; \*Plenary on New Social Spaces and SL; \*Plenary on the Gurus Vision on the Future of SL; \*Security and Innovation models.

### Summary

One of the main objectives of the event was ‘to lift the spirit of free software for those working with SL or considering doing so and to encourage the sharing of knowledge and universal collaboration without barriers’; and that it achieved. The positive feeling from this event was certainly very cheering however it would be easy to become complacent. For a long time, SL remained largely in the domain of universities and R&D departments of enlightened or unknowing companies. The code sharing mechanisms and communities were somewhat different from what exists and is possible nowadays. So, what’s new? Well, there has been a decade of building quality, portable and interoperable software and systems, built on the foundations of what went before. It is the resulting quality products that are attractive, it is this that is telling those responsible for procuring IT that SL is for them. Things are only just beginning to change in favour of SL and it would be very easy for companies big or small, other organizations or even just individuals to undermine the work of SL. Of course, some, although attracted, are distracted by vendors of proprietary systems and their free lunches or scare stories; it to such areas that we need to pay particular attention and offer what support we can.

A fuller report will be available at:

[www.ukuug.org/reports/OSWC\\_2004-CC.xhtml](http://www.ukuug.org/reports/OSWC_2004-CC.xhtml).

---

## UKUUG Winter Conference

*Thomas Bengtson*

### 25th and 26th of February 2004

This years winter conference was held at Bournemouth University while the delegates stayed at the comfortable Queens hotel where the preceding tutorial was also held. Ray Miller who organised the event made sure that everything ran to schedule with no hiccups, appropriate given that the topic of this years conference was high availability. As usual the full proceedings of the conference will be on the CD so here I will only cover a few main points.

Lars Marowsky-Brée from SUSE labs kicked off the tutorial going through the main features of Heartbeat and discussing high availability in general. The second half of the tutorial was held by Alan Robertson from IBM Linux Technology centre who started the open source Heartbeat program and is also project leader of the Linux-HA project. Alan was an energetic speaker who was clearly passionate about his work – often interrupting his associate.

Alan clearly has a special place reserved in his heart for people who don’t read the documentation to his project, as he mentioned this for the third time people looked around to see whom he could be referring to, as the good people of UKUUG actually enjoy reading documentation. All in all the tutorial was an excellent introduction to high availability and it was hard not to be affected by Alan Robertson’s enthusiasm for Heartbeat.

As a newcomer to UKUUG I had very little idea what to expect of the conference. While there is little doubt that companies like IBM are increasingly swinging behind Linux, the conference was largely an academic affair with very few suits. The atmosphere was relaxed and focused on new technology rather than on slick sales talk. A quick show of hands, during the tutorial, showed that 90-95% of the delegates were system administrators with a few developers making up the rest.

The conference itself started with a talk from Matt Asay from Novell about the business opportunities that Novell see in open source. This was in fact a good lecture but perhaps the most salesman-like of the conference. I got the feeling that the audience was slightly sceptical but denouncing Bill Gates still went down very well. The lecture ended with a pledge that Novell will make key pieces of software open source. Stephen Mayo from Hewlett-Packard made no excuses for referring mostly to HP equipment as he was describing the recent developments in making hardware 'high availability'.

The status of IPv6 deployment was covered by Tim Chown in a brief but interesting talk. He admitted that the progress has been slow but pointed out that it has recently been given a boost by the US DoD deciding to procure only IPv6-capable systems. Also promising for IPv6 advocates is the fact that many ISPs have now got IPv6 prefixes as well as 30 UK universities.

A slightly nervous Matthew Bloch held a very interesting talk on User-Mode Linux. User-Mode Linux is a virtualiser that allows you to run several fully operational Linux kernels inside Linux. He also talked about his experience of setting up Bytemark Hosting, a server hosting company utilising User-Mode Linux rather than individual servers for each client making them very cheap indeed.

Next was Julian Field who was in combative mood as he described the advantages of his open source Mail Scanner over the commercial alternatives. With a long list of prestigious users and combined with free open source virus scanner this does indeed seem to be an extremely attractive proposition.

The next day a talk on NetRaid by Peter T. Breuer from Universidad Carlos III de Madrid convincingly showed the limitations of traditional RAID when it comes to extremely large data stores since doing a complete re-sync can simply take too much time. He described how the FR1 driver for the Linux Kernel was capable of doing an intelligent re-sync thus only restoring the corrupted data.

Philipp Reisner continued in a similar vein describing DRBD which has an activity log to avoid doing a full re-sync should a node rejoin a cluster.

Alan Robertson was back to explain about the evolution and goals of the Linux-HA project. He stressed how HA would be more widely available as cost were falling and HA would spread into new un-thought-of areas. Lars Followed up by describing how the new cluster resource manager will allow Heartbeat to support more than two nodes.

Richard J. Moore described the work that has been going on in getting 'post-mortem' analysis built into the kernel by getting a snapshot of the memory should the system crash.

David Axmark co-founder of MySQL then described in the spirit of the conference the high availability features of MySQL while at the same time arguing what an excellent product MySQL is. Among other things we learned the history of MySQL and who My is.

Tony Finch from the University of Cambridge gave an excellent talk on something I personally find very interesting, namely e-mail systems. He explained how they had improved the e-mail system at Cambridge university by switching from the widely used UW-IMAP to Cyrus – he also rather gave the impression that the students at Cambridge are spoiled when it comes to e-mail.

Matthew Grove from the University of Portsmouth had the unfortunate honour of being the last speaker of the conference as most people by that time were feeling the effects of a long day of talks or the previous evening in the pub. Matthew's talk was on a jGMA a Java based tool for grid monitoring and he managed to end the conference on a high note.

The conference showed just how strong the software is that is coming out of the open source community – it is a shame that neither the UK private nor public sector has realised this.

## **FLOSSIE Conference 2004**

***Roger Whittaker***

The FLOSSIE (Free Libre and Open Source Software in Education) 2004 conference was held at the Institute of Education in London on the 18th February this year.

The conference was jointly organised by the Association for Free Software, Schoolforge UK and The Learning Machine. Although it was the first event under this title and these auspices, it was essentially part of a series of events which have taken place in the past few years bringing together those interested in the use of free and open source software in schools, and followed on from last April's successful event at Anglia Polytechnic University.

It is possibly a sign of the times and of the progress which has been made that the keynote speaker was Professor David Hargreaves, Chairman of BECTA, and that there was an MP (Richard Allan, Liberal Democrat, Sheffield Hallam) on the discussion panel at the end of the day.

The conference was well attended and there was an accompanying exhibition around the edges of the hall. Richard Rothwell (Chairman of Schoolforge UK) opened the meeting, and the event was very capably chaired throughout the day by Ian Lynch of The Learning Machine.

Professor Hargreaves' talk was largely an attempt to propose the utility of what might be termed the 'open source spirit' and the free software mentality as a way of revivify parts of the educational system. So he was only marginally referring to software as such, but drew some very strong analogies, referring to the poor levels of collaboration between teachers within institutions and between institutions and education authorities more generally, and seeing these problems as something which could be overcome by adopting an attitude akin to that of the open source and free software movements.

Ian Lynch then described the St Monica's Language Lab Project, in which a single school was able to develop its own language lab system based entirely on free software for approximately the same cost as that of licensing one of the proprietary software solutions on the market for this purpose. This theme was taken up later in the day when there was considerable discussion of how pooling resources could enable the development of software capable of replacing certain specialised tools on which schools currently pay very considerable licence costs.

Martin Jowett of Wakefield City High School spoke about how he had been able to integrate the large scale use on FLOSS on desktops with existing systems in the school.

Mark Ray and Tony Whitmore then spoke on FLOSS in further education and used the example of the excellent Koha library system to make some of their points.

The last talk of the day was by Steve Alexander, a lead developer on SchoolTool, a project funded by the Shuttleworth Foundation which aims to create an open source high class school administration system. This exciting project, if successful, could offer schools in the UK and elsewhere a radical alternative to the currently available administration software and save the education system considerable sums of money. The day ended with a lively panel discussion.

---

## **OSS Watch launch conference: December 11th 2003**

***Jan Grant***

This was a mostly enjoyable trip. Some new faces, some old familiars.

It's not clear exactly who the conference was aimed at, or the motivations of the attendees. I'd presume many are looking at providing services to support education and research, often on a shoestring. A number of times I heard the old HR-inspired chestnut that people were effectively

“free” because they lived under different budget heading and they’d be there to support commercial or open software, regardless; and that they therefore should (or should not, conclusions drawn from this hypothesis varied) come into figuring out TCWPYP (Total Costs of Whatever Proves Your Point). Lots of “let’s be open-minded and pragmatic” talk. At least, for most of the day. Seems to me that the important issue is that if you expect to run a professional service then you must be prepared to invest in support, whether that’s locally-provided or a combination of local and external. Spending money to get your staff expertise level up (whether it’s on commercial software or open-source) is going to be a part of this, and the devil is in the detail of exactly how large a part this might be.

The day opened with a quick run-down of the findings of the initial OSSW scoping survey, which raised many questions. The prime factor for institutional managers appeared to be “where do we go for support?”

Jim Farmer gave a lively talk about JA-SIG and uPortal in particular, which is a bit of a darling in the Academia and OSS world. He highlighted that support was available through several channels (commercial and otherwise). The future of the software was “interoperability, stupid” – i.e. the importance of adopting open standards.

The second session I attended was Jon Maber on “how do you make an open source project”, talking for far longer than he had expected about his work on the Bodington VLE (the other speaker could not attend due to illness and Jon manfully managed a rather interesting session).

He talked about reasons why Bodington had been open-sourced: amongst other things, there had been internal pressure to “commercialise” the project which, for a number of reasons, had not come to fruition. There was a round of applause when he said that since he was a public servant the taxpayer had already paid for the software once already!

Developer details were sketchy since it’s early days of the actual open-sourcing of the project. They use SourceForge for CVS and mailing lists, and (theoretically) bug tracking, but currently bug reports tend to be one-line patches to the mailing list. They’ve not practice yet at QAing and merging external contributions (but that’s likely to come soon) so I look forward to a future talk from Jon on that subject.

One other thing that came out of the discussion (that struck me, anyway, because I worked for a stint doing support for FE) was that there’s plenty of scope yet for supporting FE colleges, who tend to be somewhat alone in the wilderness.

The third session was entitled “Practical approaches and support”. Paul Browning spoke on hard-won experience with Open Source deployment – in this case, Zope. Much of the message would have applied regardless of the source of the software, but his “lessons learned” included the need to actively engage the community involved with the software you’re using, since it’s often the first source of your support. (Also true with respect to commercial software, e.g. sunmanagers, but more so since there may not yet be commercial support available).

The second speaker was Henrik Omma, talking about TheOpenCD. This is a grassroots project aimed at (I guess) user reeducation: supplying best-of-breed open-source alternatives on a Windows platform to common desktop apps. I guess his was the only real “advocacy” talk I attended.

The final session was “Does open source matter?” - a question studiously unanswered by either speaker, into which one might, I suppose, read an affirmative.

Jeremy Wray from IBM presented himself as a pragmatic man’o’the world, giving IBM’s position as basically, “we’ll support pretty much anything if you pay us to”. His repetition of the “well, somebody pays for everything, eventually” line grated on a few members of the audience, notably a developer from Apache UK/Cocoon. He presented a fairly down-to-earth “who cares about [any specific piece of software], what you need to know is well always be here and always be supporting you” which is a message aimed at IT managers. It’s pragmatic and probably a fair representation of the truth (although it still felt like a careful pitch), although his dismissal

of professional-level support available though other channels didn't go down well with David. He didn't seem well-briefed on prominent OS projects (which form the basis of some of IBM's offerings(!)) but that may have been part of the "who cares about the specifics? Trust us with your money and we won't let you down" line.

The other speaker was Nick McGrath from Microsoft UK. He gave a generic and ill-received talk (subtext: Linux sucks, watch out or SCO will own you) but again at least pointed out / acknowledged that interoperability was important. There was certainly an undercurrent of "we don't care how cheap it is, you're Evil with a big E" in the audience :-)

The question time was painful, rude, and unproductive in the extreme. I found none of my technical opinions challenged by it and the whole level of politically-motivated slanging to be a total turn-off. Grabbed a "free beer" and departed post-haste. All in all, a (mostly) worthwhile day.

---

## Opening the JVM

*James Youngman*

I read on the web today that IBM has again asked Sun to move to an open source development methodology for Java. This isn't the first time they've done this, but perhaps the renewal of the request represents a change in IBM's approach.

For a long time now, IBM's developerWorks programme has produced some very useful Java technology, much of which has ended up as the basis for very popular and widely-adopted technologies (Java SOAP implementations, for example). Java is a crucial technology for IBM in several ways; its Websphere product suite includes a J2EE application server as its central component, and "write once, run anywhere" could be a boost for IBM's hardware portfolio.

Java is, of course, strategic to Sun, too. The portability of Java helps Sun to stave off the onslaught of Microsoft, and they've fought to retain the portable quality of Java (which led of course to the well-known 1997 lawsuit). The high performance of Java on Windows is, of course, a double-edged sword. It's necessary for Sun to make sure that Java performs well on the Windows platform, because poor Java support for Windows would in practice have severely limited the adoption of Java. However, good support for Java on Windows makes it hard for Sun to demonstrate a clear performance lead on their SPARC systems.

The competition that we've been used to seeing between Sun and IBM has centred around Solaris on SPARC versus AIX on POWER (with, of course, HP occupying a third corner). These days of course there is another major front for their competition - the SunONE product suite versus the Websphere product suite. Both suites are centred around a Java application server and were initially integrated by a change in marketing strategy. Actual technical integration came later.

Sun's control of the Java platform has always been an advantage in this competition, but historically the J2EE market leaders have been Websphere and BEA's Weblogic. It's not surprising, then, that IBM has been calling on Sun to open the Java platform for some time. Somewhat more surprising is the fact that Sun doesn't seem to have tried to make use of its control of the Java platform to disadvantage IBM. Sun's advantage of control has led IBM to make the occasional unwise move, for example the dubbing of their excellent Java development environment "Eclipse". It's not surprising that Sun declined IBM's offer to join the Eclipse project.

It's clear that a more open development process for Java would benefit IBM. However, the interesting questions are, would it benefit other users of Java, and would it benefit Sun themselves?

Java has its legs comfortably under the enterprise table these days. The largest businesses are very happy to use Java to build vast and enormously expensive systems and don't see the Java

---

platform as being a risk. Any relaxation in Sun's degree of control over Java might raise the possibility that the future direction of the technology will not be so beneficial to such large customers. It's probably fair to say however that the existing level of Java adoption together with the prospect of Sun and IBM working in a more aligned way on Java will carry Java through any initial confidence wobble.

Smaller customers are more likely to occupy areas of the market that are well targeted by Microsoft's products. While it is possible to use Microsoft components all the way from operating system to application server, there are no Java equivalents for many of these things. Using Java means adopting a more heterogeneous environment, which can seem to smaller shops like a more risky, or at least harder, approach. While Java probably brings a wider range of possible components and add-ons, selecting them is itself an additional effort and each choice can be looked at as being a risk. The ability to choose among alternatives is a benefit, but making that very choice requires additional effort and consideration.

My personal view of the impact on smaller players is that it will be positive. Opening the development process for Java will allow smaller companies to exploit technology niches in the Java arena more effectively, and smaller users of Java will be able to find suppliers who address their needs more closely.

Linux would also stand to benefit. While even the latest beta of the JVM, version 1.5, is available for Linux, it can be fiddly to install it on a random Linux box. Package maintainers go to some lengths to make it possible to install the official Sun distribution at all via the package management system, but there's no question that an Open Source Java platform would be easier to install on any Linux system and would be less susceptible to the minor headaches of installation. It would be a great leap forward if the licensing terms of Java changed such that it became possible to make Java part of the default install of Linux systems, which isn't possible now.

Linux is not alone – many platforms would benefit from a more open Java. It would be possible to make Java run on any platform - the old "write once, run anywhere" would actually become true. Given the right licensing terms, Java could ship as standard with every computing device.

What about Sun, though? It could be giving up a large fraction of its control over the Java platform. If it gives up too much control, that might make it possible for Microsoft to make a second attempt to embrace-and-extend Java (though in any case its .NET platform is doing well these days). If Sun holds back from giving up that much control, they stand to miss out on the full benefits of community development. Sun would have to have a compelling reason to give up the advantage it has over IBM - and of course everybody else in the Java marketplace - by virtue of its control of Java. Although it could be argued that weakening Microsoft's position could actually be that compelling advantage, it doesn't seem terribly obvious that the prize is worth the cost for Sun.

I would be surprised, but very pleased, if Sun did in fact release Java as an open-source system, or even came close. IBM would likewise be very pleased. However, I think that IBM's Rod Smith will have to be very persuasive indeed. His main challenge will be to convince Sun that they can open-source Java without making it vulnerable to fragmentation and embrace-and-extend.

A great leap forward is still a leap, and the larger the leap the more intimidating it seems. Does Sun have the courage to take this opportunity? Would it be leaping across the chasm or into the abyss?

*(Java is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries. The UK Unix Users' Group is independent of Sun Microsystems, Inc. All other trademarks mentioned are the property of their respective owners.)*

## **Google Pocket Guide**

**Tara Calishain, Rael Dornfest and DJ Adams**

**O'Reilly and Associates**

**ISBN 0-596-00550-4**

**140 pp.**

**£ 6.95**

**reviewed by Gavin Inglis**

The joy of a web search engine is that it is intuitively easy to use and immensely powerful: type a couple of words, click the Search button and you are offered results from anywhere in the world. Of course effective searching of a 4.2 billion page database is a different matter.

Google Pocket Guide is a short instructive guide to the Google search engine. It is sensibly broken down into four sections: what Google can do, how to ask for what you want, how to interpret the results you get, and additional services.

Primarily the book is concerned with Google syntax and features. This begins from basic Boolean concepts and works through negation, wildcards and grouping. Experienced users can skip a few pages but then will likely come across something they are unaware of, such as using the date range: syntax to find only fresh content. The examples are useful, and occasionally entertaining - such as the beautiful story of John Doe who sets a world record for eating hotdogs and immediately afterwards rescues a baby from a burning building.

The text in fact promotes intelligent thinking about searches, and encourages the narrowing down of search results to a manageable number of more relevant links. One interesting idea is to use regional slang as an informal filter. For a newcomer these concepts alone are probably worth the low cover price. The more advanced user will be interested in the various forms of specialised syntax, and how to combine them.

With a nod to this book's big sister, 'Google Hacks', the third section briefly describes how to modify the URL directly. This can make changes which normally require a trip to the Advanced Search or Preferences pages. The remaining section considers Google's other services, such as Groups, Images, News and language tools. Finally, the five page appendix summarises syntax for the various filters. This would make a decent quick reference card.

Less a pocket reference and more a short introduction, this inexpensive book is a good choice for anyone who wants to improve their web searching with a light investment of money and time.

---

## **The Art of UNIX Programming**

**Eric S Raymond**

**Addison Wesley**

**ISBN 0-13-142901-9**

**528 pp.**

**£ 30.99**

**reviewed by Owen LeBlanc**

I had a bit of difficulty coming to terms with this book at first, partly because I was not sure what it was trying to do, and partly because I found the very beginning irritating. Eric has very definite views, and he states them bluntly. The book begins with a statement of seventeen rules of Unix philosophy, and I did not appreciate the tone in which they were presented.

The following sections, including a history of Unix and a comparison with other operating systems, was more interesting; I knew most of the data, but it looks a little different from Eric's point of view. These introductory chapters take up about 80 pages.

After this the book begins in earnest. Chapters on Design (over 200 pages) and Implementation (80 pages) are followed by a section entitled ‘Community’, which covers the evolution of C, standards, documentation, ‘open source’ (from the characteristic Eric S Raymond point of view), and the future. Four appendices and a concluding remark are followed by the long index.

Despite my misgivings, and the irritation which some sections of the book caused, I find much more to like than to dislike in this book. Eric’s ideas are interesting, and if you ignore the occasional pontification, there is quite a lot that’s worth serious consideration. I’ve been working on an application while reading the book, and I think I made four or five small changes in the overall design as a result of considering some of Eric’s advice. He analyses strong points and problems in a wide variety of software, using many examples to make his points, and they are usually very well taken.

The comments from twelve other ‘contributors’ are occasionally amusing or informative; I was certainly glad there were not many of these comments, and most of them are mercifully short.

There are some difficulties in the printing, which I feel should have been fixed. It is worth saying that the editing appears to have been fairly thorough and able, but a few errors have remained uncorrected: there’s a couple of odd page breaks (for example, on p 473), and some sentences like this one: ‘Less remembered is that the Gabriel’s central argument was about...’ (p 298), where the extra ‘the’ should have been removed. But these are minor blots on a fairly satisfactory product.

---

## **Cookbooks Introduction**

***Mike Smith***

In previous newsletters I’ve had the opportunity to reivew some of the titles in the O’Reilly Hacks series. There’s a mixed acceptance of these books – many don’t like them because they just list a few (a hundred, actually) brief facts that are probably just as easy to find on the web.

This month, I looked at a couple of titles from the Cookbook series. This is a half-way house between the Hacks and the full-brown reference manuals. For instance in the case of the Active Directory Cookbook, there are around 350 “recipes” each with quite a bit of information.

As you will have just gathered, the Cookbook theme continues inside the text, with recipes. Each recipe is structured into “Problem”, “Solution” and “Discussion”. They aren’t really “Problems”, just lists of tasks that you have hypothetically been asked to perform – but I suppose its better than having “Ingredients” and “Method” – that would be far too annoying.

The details in the cookbooks are more extensive than in the Hacks series, so these are likely to have wider appeal. Of course they will be quite good if your specific “Problem” is covered, and even if not, as long as you’ve had time to read the book it will give you some hooks into the subject matter. Then when you’re asked to do that tricky configuration thing, hopefully you’ll be in a better position to do it. Its all a bit artificial though, I make my own decision on what needs to be done - I don’t wait for someone to ask me!

Hello Joel. Following our discussion in the Midlands, are you reading this review?!

---

## **sendmail Cookbook**

**Craig Hunt**

**O'Reilly and Associates**

**ISBN 0-596-00471-0**

**408 pp.**

**£ 31.95**

**reviewed by Mike Smith**

Note the small 's' in sendmail. Very important.

I like Exim, but that doesn't temper my affection for sendmail. I know its old and quirky, has that super-duper configuration file, and our favourite security challenges, but its simply a part of interweb history and is to be treasured. Okay, that might be an old view, it has been around for 24 years now - its come a long way with IDA and KJS versions about 10 years ago, and now the commercial stream.

This book uses sendmail version 8.12.9 on Redhat Linux. [Hello also to Alasdair - new Redhat employee!] So its moved on a little, and will do so again by the time you read this, no doubt.

Recipe number 1.1 amused me - "Problem": You must keep your sendmail software up-to-date, and covers how you download and install it. There are some other trivial ones, but the recipe on setting up sendmail to use LDAP is an example of precisely the type of activities I wanted to see. I couldn't find the recipe to "build the `sendmail.cf` file from scratch" though. Nor the one about setting it up as an open relay - oh wait a minute, that might be recipe 3.1.

There are some good sections on Authentication and Encryption - and a chapter on features of the Queuing system.

Things are a lot easier these days, you know. There's a chapter on Masquerading techniques, and you've got various m4 macros which hide all the detail. The same is true of many of the features covered in the book that relate to the config file. I remember the days when ... [removed by editor: We're in the 21st century now!]

Other important areas with their own chapters are Spam, Routing and Security. We use Exim and SpamAssassin with some additional MessageLabs rules to fight our spam, and it works pretty successfully. Of course most users in corporate environments (and home for that matter) end up using Outlook or Outlook express with the preview pane on. Have you noticed the crafty spammers have been using html tags that include your email address in urls? So you don't even need to reply for them to confirm that you've received the mail, you just need to preview it. [Added by editor: many expletives.]

All in all, I like the style of the book. A bit like the hacks series, but with more meat. You can dip into a section and get a quick overview of what you can do in that area - and its a good building block for exploring further. I'm going to read it again and do a bit of playing, just for fun.

---

## **Active Directory Cookbook**

**Robbie Allen**

**O'Reilly and Associates**

**ISBN 0-596-00464-8**

**622 pp.**

**£ 31.95**

**reviewed by Mike Smith**

Active Directory is becoming increasingly important. I run a team of technical consultants in one

of the outsourcing companies and we're always getting asked for Directory skills, as it underpins many areas, not least ... no, perhaps least of all, large scale Windows 2003 implementations!

AD is based on LDAPv3, so it does have its uses (!) We have a web based front-end to our directory at work (that's not AD based, but the principles are the same). Its really great - everyone maintains their own details, and you can include your calendar, movements, skills, personal details and best of all - your photograph. Because its in LDAP, you can also link to your manager - which of course means that you can display an Org chart automatically. Really, really good.

So the Cookbook, with "Over 325 Recipes", covers a lot of areas. I hadn't realised that there were so many things to consider when designing an AD solution. But its not going to stop me from criticising those technical consultants - they need to keep on their toes.

Although a Cookbook, its organised well and starts with some basics - a bit of background and design. I'll list the main chapters: Forests, Domains and Trust. We know most of this from NT4 world. Domain Controllers, Global Catalogs (sic) and FSMOs. "What's an FSMO?" I hear you ask - AD is a multi-master system that replicates information across domain controllers. However some (5 to be precise) critical operations can only be done on a single controller. You can distribute these 5 operations - they don't all need to be on the same controller, but you can only have one master at any one time. Some FSMO (Flexible Single Master Operation) roles have cool names too - like the "RID Master".

There's a chapter on Searching. In this book, the "Solution"s are presented in several ways - using a graphic interface for the experienced windows users, and also command line tools and VBScripts scripts. There are some perl scripts too, to widen the appeal.

A bit on OUs and GPOs. Some good acronyms along with FSMO, which I used successfully in a presentation earlier this week. They were impressed anyway.

There are some more interesting chapters on Schema and Site Topology - its where the theory gets a bit closer to practice when you have to consider the realities of customer networks etc. Then we have replication, DNS, Security and Logging chapters. One on Backup, Recovery and maintenance of the DIT (use Google for that one.) Finally some application and integration stuff.

As you see, its a real action-packed book. If you need to work with AD, I'd recommend it, along with other references (from O'Reilly, of course!) I think the combination of different styles will help you rapidly cover what you need to know for a quick start, and then into more advanced topics.

So I've picked a few things up and I'll be posing a few tests for those consultants next week; to see how much they really know ;-)

---

## **Learning Python, 2nd Edition**

**Mark Lutz and David Ascher**

**O'Reilly and Associates**

**ISBN 0-596-00281-5**

**620 pp.**

**£ 24.95**

**reviewed by Sam Smith**

The second Edition of "Learning Python" has been updated to cover version 2.3 of Python, and weighs in at 591 pages. Quite helpfully, it arrived on the morning that I needed to start working with some Python code. As a perl programmer whose knowledge of Python could be summarised as "it thinks whitespace is important", the book proved extremely helpful.

It opens with a nice overview of where Python has come from, who uses it, and what for (the answers including Monty Python, lots of people, for lots of things). Chapters 2 and 3 cover running python code. This book is not focused on running solely on UNIX, and covers Win32 versions of Python as well. As a result, it does not assume that you are at a shell prompt, but all the examples work within the Python interpreter prompt. This allows for the book to mainly be platform independent.

What this book does provide is detail of Python's richness, depth, features and functionality in an informative environment rather than in the form of a reference guide. That is not to say that Learning Python can not be used as a reference guide, as its index is comprehensive and the chapters cover topics such as documenting python code, a (seemingly) comprehensive coverage of the advanced data types available in Python, advanced functions, modules, classes and Object Orientated Programming.

New programmers may find themselves thrown in at the deep end a little; but the exercises and questions at the end of each chapter should help somewhat. An overview chapter would help significantly here; although I suspect it is not a major omission.

As a book for programmers who want to learn Python, it does a very good job. The coverage is informative and well ordered; making it easy to find what you're looking for.

Overall, if you do some work with Python, you will benefit from owning this book. If you work with Python constantly, the chances are you already do.

---

## **Postfix: The Definitive Guide**

**Kyle D Dent**

**O'Reilly and Associates**

**ISBN 0-596-00212-2**

**264 pp.**

**£ 24.95**

**reviewed by James Youngman**

Postfix is a mail transport agent written by Wietse Venema (also the author of TCP Wrappers). In Wietse's words, "Postfix attempts to be fast, easy to administer, and secure, while at the same time being sendmail compatible enough to not upset existing users. Thus, the outside has a sendmail-ish flavour, but the inside is completely different."

The "Postfix Book" has been much anticipated; certainly since at least 1999. When it actually appeared in print, I was eager to review it.

The appeal of the Postfix MTA is that it is easy to set up, fast and secure. This combination is hard to come by. Some years ago I used to have a copy of the O'Reilly "bat book" on Sendmail by Costales and Allman. It's a good indication of how easy Postfix is to work with that this book is only 260 pages while the "bat book" is 1232 pages long.

This of course brings into question the necessity for having a Postfix book at all. As usual with Unix, the answer is that Postfix comes with a fine set of reference material, but the book provides an overview of the program's features, explaining how they fit together and how you might configure Postfix for various specialist purposes (there's no need to cover how one uses it for everyday usage, since Postfix will do that out of the box).

The book is fairly well-rounded; in its 15 chapters the range of Postfix's capabilities is addressed, including basic configuration, management, delivery to various mail stores, domain hosting, spam, database and LDAP integration, and authentication. It's certainly readable – I read it in less than a day. However, I did find that I skipped bits, but that may be because I've used Postfix for some years and so I'm very familiar with some of the more basic points about Postfix.

The book is a good introduction to and guide for Postfix. It covers everything you might need to know to configure Postfix to do some fairly complex things. However, if I have a gripe about the book it's that while two whole chapters (31 pages) are devoted to UCE and mail filtering, I think that more treatment of this topic would be worthwhile. Postfix provides basic content filtering "out of the box": it can reject email having contents which match criteria in various back-end lookup mechanisms (lists of regexps, for example). However, to actually do a reasonable job of filtering spam, you need to plug in one of the many anti-spam tools. Because I've been using Postfix for a long time, one of my priorities is to improve the fraction of spam that I never have to see. This is therefore the main topic I was hoping to find a treatment of in the book, and so the extent of the book's coverage was less than I had hoped. On the other hand, filtering spam is largely a job for a spam-filtering tool, and so perhaps I should have been reading a book on SpamAssassin.

If you are installing Postfix – or considering installing it – in anything but a home environment, I think you should read this book. At the very least, you will be able to ensure that you're familiar with the range of capabilities that Postfix has. If you are considering Postfix for home use, then I suspect that the "out-of-the-box" configuration will be sufficient for your needs. However, I'm off now to install SpamAssassin...

*James once spent a couple of years administering firewall systems which had Sendmail installed on them, and hand-hacked the `sendmail.cf` files (with, of course, the aid of the "bat book"). Understandably therefore he installed Wietse's "VMailer" (later Postfix) when it first came out as a public beta. Since then he's not installed any other variety of MTA on any kind of system. James lives in Manchester with no dogs, no rabbits, and no cats.*

---

## Computer Science and Perl Programming

Jon Orwant (editor)

O'Reilly and Associates

ISBN 0-596-00310-2

758 pp.

£ 28.50

reviewed by James Youngman

This book is the first of three volumes which contain a collection of articles from The Perl Journal.

I wanted to review this book because I'm a fan of Perl, having started using Perl 4 to do things that just weren't possible in any other scripting language back then (for example fix the relational integrity errors in a proprietary database by finding and deleting orphaned records). The title was particularly attractive because I think there isn't enough computer science in professional software development.

However, this review is very, very, very late. The reason is that I didn't get round to finishing the book. If you think that sounds bad, it is. I read a large number of books and this one wasn't interesting enough for me to get around to reading the last third or so.

The fundamental idea of the book is to collect together articles from TPJ which have something in common and present them in themed sections. This works reasonably well, but although the collection is fairly varied, the end result is mixed without being appealingly eclectic.

I'm glad that nobody tried to introduce me to Perl with a book like this. I would never have become an enthusiast. There are many good articles in the book, but I think it just doesn't work as a collection.

## **Cisco Cookbook**

**Kevin Dooley and Ian J Brown**  
O'Reilly and Associates

ISBN 0-596-00367-6

908 pp.

£ 38.95

**reviewed by Raza Rizvi**

Now I like books, and I like Cisco routers. So this book had me nervous with anticipation. I didn't know whether it would be regurgitation of the excellent Cisco website (like some of the Cisco Press books) or whether it would add something to all the Cisco IOS and router technology books I already had.

I was stunned. This is a great work, full of config snippets, router output and sensible explanations (note this is for router IOS and is not applicable to Catalyst Switches or PIX firewalls). Like any cookbook, you can dive in anywhere, pick off what you want, and then sit back and see the results.

The authors have picked 23 popular topics and drilled down to the sweet essence of each with sensible subsections set out in the same problem / solution / discussion / see also format.

Each topic – and you can see this online at:

<http://www.oreilly.com/catalog/ciscockbk/toc.html>

starts with a concise and accurate summary to preface the scenarios that follow.

I spent a great deal of time reviewing this book. I used it as it was intended - as a first stop in the answer to a question, I also used it as an aid memoir, and I also read whole topics from start to finish. In each case I found what I needed and only what I needed.

Try as I might, I couldn't find anything that was wrong or explained badly, and I am a critical person (okay if you push I would say DSLw will fall on deaf ears in the UK). I thoroughly recommend this book to people who want to do more with their Cisco routers, and especially to those who support networks professionally. It goes without saying that this book will be of great use to those studying for Cisco certifications to CCNP level with topics on the four main routing protocols, queuing, ACLs, and HSRP.

It even explains step by step tunnels and VPNs without blinding the reader with science – fantastic!

*Raza Rizvi is Technical Manager at REDNET. He deals with networks every day and worries that he is now only months away from his CCIE Lab exam.*

---

## **Apache Cookbook**

**Ken Coar and Rich Bowen**  
O'Reilly and Associates

ISBN 0-596-00191-6

234 pp.

£ 20.95

**reviewed by John Collins**

This book is a set of “recipes” for doing various tasks with Apache, which is by far the most popular software for driving web servers.

Successive chapters cover Installation on both Windows and Unix or Linux systems, common modules, logging, setting up virtual hosts, aliases, security issues, SSL, dynamic content, error

handling, proxies, performance issues, miscellaneous topics. Appendices tell the user about regular expressions and some troubleshooting.

I've never much cared for the "cookbook" style of presenting things myself. I do like to understand what I'm doing. It may take longer to start with, but it saves time in the end unless the explanations are detailed and give you scope for experimentation and variation. I don't really feel this book does this. It refers you to other documentation all the time, sometimes without really telling you where to find it or telling you to run commands which may or may not work on your system (if you can figure out what `"/path/to/serverroot"` might be for example).

In places, such as in two consecutive "recipes" in the security section, it presents a problem and then tells you you can't really solve it in Apache. Some snippets of Perl code are offered by way of a workaround for one of those. There are some large lumps of Perl scattered throughout the book, to a degree at which it seem patronising to me to have the appendix on regular expressions, virtually the bread and butter of Perl programmers.

I think this book gallops through the topics without properly introducing and explaining them to a degree that it's not going to meet many needs. You might be lucky and find exactly what you want to do in here, but probably you won't. If you do, you won't properly understand what you're doing just from this book. If you're really at the level of needing to understand regular expressions, I suspect this book will leave you totally confused. (It's hard to imagine that you'd be let loose on a web server if you were).

I don't think this is one of O'Reilly's better efforts myself.

*John Collins – Xi Software Ltd – [www.xisl.com](http://www.xisl.com)*

---

## **Learning XSLT**

### **Michael Fitzgerald**

#### **O'Reilly and Associates**

**ISBN 0-596-00327-7**

**368 pp.**

**£ 24.95**

**reviewed by John Collins**

This book introduces the reader to XSLT and also XPath.

XPath is a language for finding parts of an XML document you are interested in. XSLT is a language for doing things with XML documents or the bits you've located with XPath by transforming them into other things such as more XML documents, HTML or text.

I'm a bit of a reluctant user of XML myself. It is horribly verbose at the best of times and it is quite ugly-looking all the time. Still it's standard, there are standard tools to manipulate it, it gives people a nice warm fuzzy feeling that you're using it because everyone else is. Above all it isn't ASN.1. For more comments see Aaron Crane's

<http://xmlsucks.org>

XSLT is actually written in XML so you end up with one piece of XML transforming another possibly to yet another. XPath isn't XML as such, but you end up burying it in XML as elements and attributes of what your "program". You'll have to get used quite quickly to the "functional" approach as opposed to the so-called procedural approach of programming languages such as C, Perl etc.

To get anywhere with this book you will have to download the examples from O'Reilly's website – a ZIP file of over 6 MB. Most of the examples given involve use of Xalan which you may have to download from the Apache website and probably build from the sources on your computer, a

process which may require a little effort. Some graphical XSLT processors, for Windows only, are mentioned in the first chapter and at the end.

The book takes the reader through all the features with some care and with patient explanations, I thought. Many of the later examples will be useful as a basis for adaptation by the reader. After covering the features of XSLT and XPath as they stand, various people's extensions are discussed in one chapter and XSLT 2.0 and XPath 2.0 in another. There is a chapter about writing XSLT interfaces in Java and C#. A final chapter refers the reader to the definitive works on the subject and an appendix lists various XSLT processors and how to install them.

I think this book is very well put together and succeeds in making what could be a very tedious subject relatively interesting. I am sure that anyone starting to work with XSLT will find it most helpful but you will need to get Xalan going as well as download the examples.

*John Collins – Xi Software Ltd – [www.xisl.com](http://www.xisl.com)*

---

## **Teach Yourself Linux**

**Robert Billing**

**Hodder and Stoughton**

**ISBN 0-340-81204-4**

**226 pp.**

**£ 8.99**

**reviewed by Roger Whittaker**

I remember the "Teach Yourself" series from long ago – books with yellow covers which tended to be simply introductory textbooks which were sometimes a little disappointing in their scope, but cheap to buy and reasonable cheerful.

The current list has certainly changed since the days when I had a few of these on my shelves, but the aim of the series seems to be the same as ever – to provide cheap introductory books on popular subjects.

The fact that the series now includes "Teach Yourself Linux" is therefore an interesting sign of the times – it's yet another sign that Linux has "arrived" and has established itself in the consciousness of the wider public.

There are some introductory Linux books for beginners which don't really tell the reader anything at all – I've seen books which are mostly filled with screenshots of KDE showing you how to launch a program from the KDE menus and how to create a spreadsheet with OpenOffice.org. This is definitely not such a book: it packs a very large amount into its 226 pages and might well be a little intimidating to a reader whose only exposure to computers had been through using Windows.

The book is well written and has been carefully proof-read: because of its size, depth has had to be sacrificed to breadth, but I was impressed with the way in which quite difficult ideas are at least introduced: the reader can then go off and find out more elsewhere.

The book recommends and discusses RedHat, SUSE, Debian, Gentoo and Knoppix: most of the discussion is distribution-neutral, and the specific setup tools peculiar to particular distributions are not discussed in any detail. In the few cases where it makes a difference, there is a leaning towards RedHat.

The emphasis throughout is on understanding as well as doing: after a chapter of background information and a general overview of the installation process, there is an introduction to using the bash shell (which comes before any discussion of GUIs).

There is a good chapter on basic networking, which goes on to give enough information to get a new user started with NFS, samba and mounting shares from Windows machines. There is even a brief introduction to firewalling using iptables with a sample iptables script.

As well as a very brief survey of the programming and scripting languages available on Linux (the author seems to have a particular liking for TCL/Tk), there is a section on “Advanced topics” which covers building custom kernels, as well as an “Inside Linux” section which explains the init scripts and discusses kernel modules and hardware drivers.

I would recommend this book to an intelligent new user of Linux with an interest in understanding the system – it would make a good supplement to the two books I usually recommend in such cases: “Running Linux” and “Linux in a Nutshell” from O’Reilly.

---

## **The Best of the Joy of Tech**

**Nitrozac and Snaggy**

O’Reilly and Associates

ISBN 0-596-00578-4

192 pp.

£ 9.95

reviewed by Lindsay Marshall

See the combined review below.

---

## **Mac OS X The Missing Manual Panther Edition**

**David Pogue**

O’Reilly and Associates

ISBN 0-596-00615-2

782 pp.

£ 20.95

reviewed by Lindsay Marshall

See the combined review below.

---

## **Running Mac OS X Panther**

**James Duncan Davidson**

O’Reilly and Associates

ISBN 0-596-00500-8

328 pp.

£ 28.50

reviewed by Lindsay Marshall

If you’re not a Mac person then you probably don’t read the Joy of Tech. You may not even have heard of the Joy of Tech – after all there are so many cartoons and strips available on the good old Interweb that sometimes I wonder if I am the only person not actually drawing one. The JoT is a fairly regular series of cartoons and short strips usually with a Mac theme, though frequently they are Linuxy and sometimes other stuff entirely. Quite a lot of them rely on you knowing some of the latest Mac gossip, but there are also lots which just poke fun at Windows and Bill

Gates, so there is something for everybody (apart from Windows aficionados of course). The really important thing about the cartoons is that they are well drawn (and rendered) and on many occasions are genuinely, laugh out loud funny (unlike, say, User Friendly which has been dreadfully dull for a long time now.)

The book is nicely produced in colour and a lot of my favourites are in there. I could have done without the the listings of the user poll results that went with each cartoon. - tedious filler. I'm not sure who's going to buy this though: perhaps the person who doesn't know what to get a geek. (A real geek will of course read it online as soon as it appears and download the image and save it) I suppose it might be a geek coffee table book, if geeks had coffee tables.

On the other hand, "Mac OS X the Missing Manual (Panther Edition)" certainly isn't a coffee table book. Unless your coffee table is reinforced to take the weight of the 761 pages. And another thing, it definitely isn't missing anymore. Every time I go in my office I fall over a pile of previous editions. What to say about this that I haven't said about previous editions? err, nothing - it's just the same as it always was but has some stuff about the new things that have been added in OS X 10.3. If you're new to OS X and want a comprehensive book on nearly everything you can possibly think of then this you might want to buy this. I can't see any point though, if you have an earlier versions - you can find out about all the new features of Panther on the net!

Of course, once you have got used to OS X and have read through OSXTMM(PE) and are fairly on your way to being a guru (if you weren't already through your Unix knowledge), you'll find the stuff that isn't there. Then you'll need "Running Mac OS X Panther" (though there is overlap). However I would guess that the vast majority of Mac users will never need to know anything about some of the stuff in this book: turning on journaling, LDAP and such like. And a good thing too, there is dangerous stuff in here and there is a liberal sprinkling of mantrap icons highlighting those areas (usually root related) where things can get hairy.

As with all OS X books, if you know Unix (is that like knowing kung fu?) so much is familiar that much of the text seems redundant, and then there is a little special Mac wrinkle that catches you out. You have to be careful of things like the fact that OS X Applications are directories and the GUI hides this from the user, but not at the command line.

There is useful stuff in here if you have to do this kind of thing for a living. Panther introduces enough newness that you probably do need to keep up with the editions this time too.

---

## Contacts

Charles Curran  
Council Chairman; Events; Newsletter  
Oxford  
Tel: 07973 231 870  
Tel: +34 954 371 381  
charles.curran@ukuug.org

James Youngman  
UKUUG Treasurer  
Manchester  
james.youngman@ukuug.org

Sam Smith  
Website  
Manchester  
sam.smith@ukuug.org

Alasdair Kergon  
Events  
Reading  
alasdair.kergon@ukuug.org

Alain Williams  
Watford  
alain.williams@ukuug.org

Roger Whittaker  
Schools; Newsletter  
London  
roger.whittaker@ukuug.org

Ray Miller  
Events; Newsletter  
Oxford  
Tel: 01865 273 200  
ray.miller@ukuug.org

Jane Morrison  
UKUUG Secretariat  
PO Box 37  
Buntingford  
Herts  
SG9 9UQ  
Tel: 01763 273 475  
Fax: 01763 273 255  
office@ukuug.org