

# news@UK

*The Newsletter of UKUUG, the UK's Unix and Open Systems Users Group*

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

---

Volume 14, Number 2

ISSN 0965-9412

June 2005

---

## Contents

<b>News from the Secretariat</b>	<b>3</b>
<b>Linux Technical Conference 2006</b>	<b>3</b>
<b>Announcement: 2005 UKUUG Award Winners</b>	<b>3</b>
<b>UKUUG Diary</b>	<b>4</b>
<b>Announcement: backstage.bbc.co.uk Open Tech 2005</b>	<b>5</b>
<b>Announcement: XML UK Member Presentation Day</b>	<b>6</b>
<b>Socitm Open Source Group report</b>	<b>6</b>
<b>BECTA Open Source Workshop</b>	<b>7</b>
<b>'Technical Contribution' Workshop: Bolton 17 March 2005</b>	<b>8</b>
<b>NAACE/SOCITM Teach-In on Building Schools for the Future</b>	<b>8</b>
<b>Book review: Linux Server Security (2nd edition)</b>	<b>9</b>
<b>Book review: Mind Hacks</b>	<b>11</b>
<b>Book review: Google Hacks, 2nd Edition</b>	<b>12</b>
<b>Book review: Linux Server Security (2nd Edition)</b>	<b>13</b>
<b>Book review: Internet Annoyances</b>	<b>13</b>
<b>Book review: Linux Cookbook</b>	<b>14</b>
<b>Book review: Revolution in The Valley</b>	<b>15</b>
<b>Book review: High Performance Linux Clusters with OSCAR, Rocks, OpenMosix, and MPI</b>	<b>16</b>
<b>Book review: Jakarta Commons Cookbook</b>	<b>17</b>
<b>Book review: Oracle SQL*Plus: The Definitive Guide, Second Edition</b>	<b>18</b>
<b>PePLink MANGA for University Embedded Linux</b>	<b>19</b>
<b>Contacts</b>	<b>21</b>



---

## News from the Secretariat

**Jane Morrison**

It has been another busy time for UKUUG. Since March we have been concentrating on bringing the details together for the Linux 2005 Conference (Swansea 4th – 7th August). The full information booklet and booking form has been sent to all members and you can also find all the up-to-date details at:

<http://www.ukuug.org/events/linux2005/>

Delegate bookings are arriving each day. Don't forget: if you wish to take advantage of the special early bird Tutorial and Conference fees you must book by 30th June. If you want to take advantage of the University Halls' Bed and Breakfast option please make your booking as soon as possible.

In May we organised another DNS tutorial. Jim Reid, the tutor, has provided tutorials for the UKUUG members on quite a few occasions and as usual we had a good number of attendees who found the day very interesting and worth while.

The next event for your diaries is the UKUUG Annual General Meeting which will be held this year on Thursday 22nd September at University College London at 6 pm. The agenda and further details will be sent to you automatically nearer to the date.

The next Winter Conference is planned for February 2006, we are currently looking at venues in York and Durham.

If you wish to have something published in the next issue (September) please note the copy date is: 2nd September.

---

## Linux Technical Conference 2006

**Jane Morrison**

We are still in the process of considering possible venues for the Summer 2006 Linux Technical Conference.

If you know of a venue which you think might be suitable, we would be interested to hear about it. We have found in the past that an important element in organising a good conference is being able to call upon some expert local knowledge from volunteers in the area: if you are a member working in a university or similar institution that has the right kind of facilities, and you would be able to help with local knowledge, please let us know.

---

## Announcement: 2005 UKUUG Award Winners

UKUUG (the UK's Unix and Open Systems User Group) has made its 2005 Award to Christopher Gutteridge of the School of Electronics and Computer Science, University of Southampton, for his work on the Open Archive Software: GNU EPrints.

UKUUG makes this annual award to give particular recognition to the development of free and open source software in the UK. Christopher wins a trip to the Open Source Convention in Portland, OR, USA, and support with costs thanks to O'Reilly, organizers of the convention and Gold Sponsors of the UKUUG Award, and a £500 cash prize from UKUUG.

'EPrints is both a practical tool and the crystallization of a philosophy', said Christopher. 'It enables research to be accessible to all, and provides the foundation for all academic institutions to create their own research repositories'.

The School of Electronics and Computer Science has been one of the prime movers in the global movement towards open access publishing. The University of Southampton is the first UK university to announce that it would be establishing its own institutional repository and requiring all its academic staff to self-archive their research.

The judges noted as 'Highly Commended' Dr Thomas Leonard's work on the free software projects ROX Desktop (a graphical desktop for Linux and Unix type systems, that he created to combine the elegance of the filer-centric RISC OS with the power of Linux) and Zero Install (which combines the simplicity of application bundles with the automatic dependency handling of package managers). Thomas was awarded his PhD in Computer Science last summer (2004) from the University of Southampton and now works there in its School of Electronics and Computer Science.

Further information about GNU EpPrints is available at:

<http://eprints.ecs.soton.ac.uk/>

<http://software.eprints.org/press>

<http://software.eprints.org/archives.php>

Further information about ROX Desktop and Zero Install is available at:

<http://rox.sourceforge.net/>

<http://zero-install.sourceforge.net/>

---

## UKUUG Diary

The UKUUG diary of events which are of possible interest to members is at

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

The following is just a small selection of summaries of some interesting forthcoming events (with starting dates) taken from the diary.

### **Linux Cluster Summit**

20th June 2005: Waldorf, Germany

<http://sourceware.org/cluster/events/summit2005/>

### **Defeat Virus and Spam the Open Source Way**

20th June 2005: Birmingham, UK

<http://www.openadvantage.org/events/oaevent.2005-04-06.2839567705/view>

### **LinuxTag 2005**

22nd June, Karlsruhe, Germany

<http://www.linuxtag.org/typo3site/8.0.html?L=1>

### **LUGRadio Live 2005**

25th June 2005, Wolverhampton, UK

<http://www.lugradio.org/live/2005/>

### **The low cost office**

4th July 2005, Birmingham, UK

<http://www.openadvantage.org/events/oaevent.2005-05-16.7744632940/view>

### **Institutional Web Management Workshop**

6th July, Manchester, UK

<http://www.ukoln.ac.uk/web-focus/events/workshops/webmaster-2005/>

### **Debconf5**

10th July 2005, Helsinki, Finland

<http://www.debconf.org/debconf5/>

**ApacheCon Europe 2005**

18th July 2005, Stuttgart, Germany  
<http://www.apachecon.com/2005/EU/index.html>

**What the Hack!**

28th July 2005, Den Bosch, The Netherlands  
<http://www.whatthehack.org/>

**UKUUG Linux Conference 2005**

4th August 2005, Swansea, UK  
<http://www.ukuug.org/events/linux2005/>

**LinuxBierWanderung 2005**

6th August 2005, Killin, UK  
<http://lbw2005.ziggur.at/>

**LinuxWorld Conference and Expo**

5th October 2005, London, UK  
<http://linuxworldexpo.co.uk/>

**EuroBSDCon 2005**

25th November 2005, Basel, Switzerland  
<http://2005.eurobsdcon.org/>

---

**Announcement: [backstage.bbc.co.uk](http://backstage.bbc.co.uk) Open Tech 2005****Saturday July 23rd - The Reynolds Building, Hammersmith, London W6 8RP**

<http://www.ukuug.org/events/opentech2005/>

Sponsored by [backstage.bbc.co.uk](http://backstage.bbc.co.uk), Open Tech 2005 is an informal one-day conference about technologies that anyone can have a go at, from “Open Source”-style ways of working to repurposing everyday electronics hardware.

So far, the line-up features:

- Ted Nelson, inventor of hypertext, on where the web went wrong
- The official launch of the [backstage.bbc.co.uk](http://backstage.bbc.co.uk) developer network, opening up BBC content for you to play with
- The ability to record an entire week of all Freeview TV and radio channels, probably the UK’s largest (fridge-sized) PVR

More speakers will be confirmed over the next few weeks – but, as the title implies, we’re very much “Open” to suggestions. If you’re reverse-engineering proprietary protocols, making useful information available in a way people couldn’t get at before, pioneering unexpected methods of knowledge sharing – or (equally likely) doing something so cool we haven’t even thought of it yet, then please get in touch via the submissions form at:

<http://www.ukuug.org/events/opentech2005/offer/>

The deadline for submissions is midnight UK time Saturday June 25th, and we’ll aim to notify everyone who’s submitted a proposal by July 1st.

We’ll be trying to fit in as many talks (and lightning presentations) as possible, so the shorter you can make yours, the better. Alternatively, if you have an idea for a panel discussion, or a

---

workshop, or anything else that's vaguely in keeping with the theme of the event, then we also can't wait to hear from you.

And there'll most likely be some sort of internet access at the event, but offline demonstrations are strongly encouraged, as bandwidth may not be guaranteed.

You don't have to suggest a session to take part; you can stay informed about the event by subscribing to our low-traffic announcement-only mailing list.

Send a blank email to: [notcon-subscribe@socialswirl.com](mailto:notcon-subscribe@socialswirl.com)

(your address will only be used to contact you about the event and will not be passed onto third parties).

You can email [opentech@ukuug.org](mailto:opentech@ukuug.org) if you've any other questions.

---

## **Announcement: XML UK Member Presentation Day**

We have received a call for papers and details of this event which will be held on 12th July 2005 at Reading Town Hall.

Full details of the event can be seen at:

<http://xmluk.org/memberpresentationday0705htm.htm>

---

## **Socitm Open Source Group report**

***Leslie Fletcher***

I attended a meeting of the Socitm Open Source Group (SOSS) at Bridgend on 13 May 2005.

Socitm, the Society of Information Technology Management, was founded in 1986 as the professional association for ICT managers working in and for the public sector. The Society is a significant provider of advice and guidance on ICT and e-government to local authorities. The Society's home page is at

<https://www.socitm.gov.uk/public/>

Socitm members receive a bi-monthly magazine and I hope to get an article about UKUUG published later this year.

The Socitm Open Source Software Group (SOSS) was established in mid 2003 with the objective of promoting the use of open source software in the public sector. The specific focus of SOSS is not the technical issues associated with OSS but with the business case for adopting OSS. Its home page is at

<https://www.socitm.gov.uk/Public/activity+groups/open+source+group/>

Given the remit SOSS has adopted, it could be just a talking shop but its meetings are much more useful than that. SOSS has good contacts within government and its members are key contributors to the Open Source Academy.

I feel that UKUUG's potential contribution is recognised and my presence at its meetings welcomed. The Open Source Consortium and Open Forum Europe are also usually represented. The other participants are mainly local authority officers with managerial responsibility for IT.

At the meeting on 13 May there was considerable satisfaction with the Socitm/Financial Times open source survey whose results were published in the FT on Monday 9 May. This SOSS initiative has had extensive coverage; for further details see

<http://www.ukuug.org/mediawatch/?p=214>

---

An officer from Bristol City Council described the steps they had taken to move to Star Office; a link to the report on this, again in the FT, is at

<http://www.ukuug.org/mediawatch/?p=209>

The story is, perhaps, a familiar one but still interesting. The underlying business case was the desire to provide uniform office productivity tools to all council users, replacing the ad hoc collection of packages which individual departments had acquired over the years. The licencing costs of Microsoft Office were prohibitive so the IT staff proposed the use of Star Office. Now, about three years later, the case has been made, concerns and suspicions overcome, a deal with Sun made, too-hasty decisions unpicked, staff training provided and the roll out of Star Office is underway. Disappointingly, the education department refused to be part of the project so Bristol schools won't see any benefit.

---

## **BECTA Open Source Workshop**

***Leslie Fletcher***

I attended this meeting in Coventry on Thursday 14 April, along with about 30 others, including decision makers from schools, DfES and LEA officials, commercial providers of systems and educational content for schools, all active in promoting Open Source in schools.

The objectives for the day were

- To identify what needs to be done to bring Open Source software into the wider education sector
- To identify what the Open Source community might do to support wider adoption of Open Source software
- To discuss the role of the commercial sector and market forces in relation to Open Source and the range of potential government interventions
- To discuss how the Open Source philosophy might be adopted to support improvements in educational content and learning platforms.

This was an extremely positive day for Open Source. Presentations from three BECTA staff each showed a high level of receptiveness to Open Source and recognition that it is a mature and serious alternative to existing proprietary offerings. Specific and useful guidance was given on the approach the Open Source community should adopt in dealing with the education sector. Issues such as interoperability and security, where Open Source is widely believed to have the edge, were highlighted. Most interesting and newsworthy of all was the announcement of the results of a TCO study which BECTA has been conducting in a small number of schools. This shows the annual total cost per PC in primary schools using Open Source is about 50% of that for proprietary systems; the corresponding figure in secondary schools is 80%. Across the schools in the study, reduction of support costs is the main contributor to the overall savings.

The proximity of the General Election created an unreal feel to some of the discussion because BECTA staff are bound by Cabinet Office rules

<http://www.cabinetoffice.gov.uk/guidance/two/08.htm>

during this period. In particular they may not "initiate new action of a continuing or long-term character" and it had clearly been decided that explicit support for Open Source is covered by this restriction. As a result, the slides of the presentations will only appear on the BECTA website after May 5 and participants were requested to be similarly discrete in their reporting

---

of the event. I believe it is very strongly in the interest of the Open Source community to respect this, but to be ready with high profile material to respond with when “normal business” is resumed.

---

## **‘Technical Contribution’ Workshop: Bolton 17 March 2005**

***Leslie Fletcher***

This meeting was one of a dozen or so being held around the UK in the next month following the DTI consultation on the EU Computer-Implemented Inventions Directive last December.

There were around 35 attendees at the event, mostly software developers together with a few lawyers with an interest in patents. The workshop was broken up into small groups who were asked to study five case studies of fictitious inventions and to judge whether they passed four possible wordings of a ‘technical contribution’ test. One of these was the formulation in the draft directive and one the test being proposed by FFII. The participants regarded these as serious contenders for a workable test but the two others, drafted by Patent Office staff on the basis of submissions they had received, were thought to be woolly and lacking in detail or substance.

Two of the case studies were software or business processes, both of which, on a show of hands, the overwhelming majority of participants thought should not be patentable. Moreover, the majority view was that these two cases did not satisfy the ‘technical contribution’ test in any of the four formulations. The other three case studies were “gadgets” and opinion regarding their patentability was divided.

Contrary to some of the posts on some of the anti-patenting lists, I thought the exercise was worthwhile. The main weakness was that non-lawyers were unable to analyse the case studies and the definitions “with intent” to achieve a certain outcome for a client.

The case studies and the possible wordings will be published on the Patent Office website on 8 April, just after the last of the workshops has taken place, and a results of the exercise will be published some time after that.

---

## **NAACE/SOCITM Teach-In on Building Schools for the Future**

***Leslie Fletcher***

I attended this day-long meeting in Leicester on Monday 21 March.

About 20 people were present, mostly LEA officers with a couple of quite senior people from DfES, a representative of BECTA and some private consultants. There was not much structure or a timetable for the day, and it veered around between:

- Philosophising about education, around some well-worn themes. LEAs are required to look to the future in this way as part of their bid for BSF (Building Schools for the Future) funding.
  - Case histories (AKA moans and anecdotes!) from LEAs who already have, or will shortly have, BSF funding. This was not very illuminating as DfES keeps changing the rules. However, it was clear that business – SERCO, Capita, ... – is expected to play a big part, as are construction companies. Normally they will procure the IT infrastructure and services as part of the building contract.
-



- Concern about the capacity of LEAs to manage BSF projects. This was the most interesting/alarming part of the discussion. The LEA officer mindset seems to be that they will decide what ICT schools will have and it is likely to be the local authority corporate solution, including an obligatory managed service.

There was not, as far as I could discern, even lip paid service to the concept of a level playing field for proprietary versus open source solutions. On the other hand, several participants volunteered to me that they used open source at home or were personally in favour of open source but the corporate culture was otherwise.

This was a depressing and irritating day. As a citizen and school governor I was dismayed by the prospect of LEAs high-jacking BSF and returning to the old ways of doing things to schools rather than for schools. From UKUUG's point of view, Open Source would not have got a mention if I had not been there. Long and hard thinking is required/needs to continue about how Open Source is to be promoted, starting with DfES. I intend to stay in contact with this group, to feed in some Open Source experience from a school perspective.

---

## **Linux Server Security (2nd edition)**

**Michael D Bauer**

**O'Reilly and Associates**

**ISBN 0-596-00670-5**

**544 pp.**

**£ 31.95**

**reviewed by Mick Farmer**

I haven't read the first edition of this book, entitled "Building Secure Servers with Linux", so I can't say how much the book has changed. The blurb on the back states that five or more new security topics have been added, including OpenLDAP for authentication and the popular Cyrus IMAP service.

My plan for this review was to understand each chapter and then to apply that understanding to the Linux machine that I use daily, where appropriate. Although my computer runs Apache for serving web pages, it is also the machine on my desktop – Red Hat 9, 2.4.20 kernel, kept up to date with apt-get and RPMs.

The book is based on the articles written for a monthly security column (Paranoid Penguin) in the Linux Journal Magazine, updated for the second edition with tips and techniques for Red Hat Enterprise Linux, Fedora, SUSE, and Debian. Each chapter covers a particular group of services.

Chapter 1 (Threat Modelling and Risk Management) sets the scene for the subsequent chapters by discussing what assets we're trying to protect and what levels of protection each asset requires. It starts with some general sections on data confidentiality and integrity, system integrity, and network availability. It follows with sections on who might attack our systems and some of their motives. It finishes with simple risk analysis using Annualised Loss Expectancies (ALEs) and Bruce Schneier's attack trees.

ALEs sound daunting until you see the calculation:

Single Loss Expected Annual

$ALE \text{ (cost/year)} = \text{Expectancy (cost)} \times \text{Rate of Occurrences}$

Working in a University has shielded me from itemising single-loss expectancies, but the author provides some nice examples based on an SMTP gateway – there's even a cost for sendmail bugs!

An attack tree starts with a root node, the attacker's objective (e.g. steal customer data), and a number of leaf nodes which are the subgoals (e.g. obtain backup media, intercept e-mail, or hack into file server). The process is repeated for each of the subgoals until you have a complete attack tree. Now comes the interesting bit – attach a “cost” to each node resulting in a cost for each possible attack path. You're now looking at risk analysis from the attacker's point of view.

Chapter 2 (Designing Perimeter Networks) looks at firewalls and different DMZ (demilitarised zone) architectures. Frankly, I skipped this chapter because I wasn't in the position to install a DMZ and because the topic of firewalls would be covered in the next chapter!

Chapter 3 (Hardening Linux and Using Iptables) started with the usual guidelines about what software should actually be running on a server. I was able to heed some of this advice and turned off a number of services, and learnt about Red Hat's chkconfig and Debian's apt-get at the same time.

I've been using iptables for some time, but finally bit the bullet and took the author's advice of dropping every packet by default and then allowing in and out only those packets I really wanted. This took some time as I slowly found services which weren't running correctly and deciding how this translated into IP address, protocol, and port combinations that would be accepted. There are excellent iptables scripts in this chapter and in the Appendix.

This chapter continues with a thorough discussion of port scanners, such as nmap, and security scanners, such as Nessus. I've used both of these to check both my work machine and my home ADSL router with built-in firewall.

The chapter concludes with Bastille Linux, a tool for hardening mainly Red Hat-centric systems. Again, I ran this on my work machine to really batten down the hatches. For those who want to know more about Bastille, there's a good review in the November 2004 issue of Sys Admin.

Chapter 4 (Secure Remote Administration) is perhaps mis-titled as it's all about the Secure Shell (SSH) which provides encrypted connections using public-key cryptography. It replaces telnet, rsh, rcp, and rlogin.

I must own up to only skimming over the next three chapters. These cover OpenSSL (Chapter 5), DNS (Chapter 6), and LDAP (Chapter 7). These are technologies that I don't need to consider in my workstation environment.

Chapter 8 (Database Security) is mainly concerned with configuring and administrating MySQL, which I use in association with my web pages (PHP being the glue). It was satisfying to see that my server and configuration matched the points made by the author.

Chapter 9 (Securing Internet Email) covers general e-mail security with details for securing two popular Message Transfer Agents (MTAs), sendmail and postfix. I use sendmail purely to send mail so most of the chapter was academic. I don't use postfix so find it difficult to comment on that material.

Chapter 10 (Securing Web Servers) is concerned with installing, configuring, and administrating Apache, with additional material on using CGI scripts in PHP and Perl. Very good advice it is too. I certainly picked up some techniques for hardening my web service.

Chapter 11 (Securing File Services) is really for people providing an FTP service, which I don't. It considers two replacements for traditional ftp, namely ProFTP and vsftpd, with details on how to install and configure both, especially for anonymous FTP.

This chapter finishes with a section on rsync, a file-transfer tool capable of performing differential uploads and downloads of files. It can work with SSH to provide encrypted connections.

Chapter 12 (System Log Management and Monitoring) starts with sections on configuring syslog and running syslogd. It follows with sections on syslog-ng, a flexible replacement for syslog that's already a standard package in Debian and SuSe distributions. Judging by the number of examples, syslog-ng is an extremely powerful package. It finishes with details about swatch (Simple WATCHer), which is a log monitoring utility written in Perl.

Chapter 13 (Simple Intrusion Detection Techniques) is concerned with host-based Intrusion Detection Systems (IDS) which alert you to unexpected changes in important system files. Most of the chapter is devoted to Tripwire, a popular IDS. It finishes with Snort, which also sniffs packets “on the wire” so it’s capable of reporting “attempted intrusion” detection while the attack is happening. I haven’t tried it yet, but it’s on my to-do list.

A most enjoyable and interesting read. The author has managed to condense down most of the information about a number of utilities which themselves are described in complete books (the Sendmail tome and the various MySQL books come to mind here).

Yes, my machine is now that bit more secure.

---

## Mind Hacks

**Tom Stafford and Matt Webb**

**O’Reilly and Associates**

**ISBN 0-596-00779-5**

**400 pp.**

**£ 17.50**

**reviewed by Mick Farmer**

First off, I should mention that this is not a computer book in the popular Hacks series; it’s a book about our brain and the psychological experiments devised that let us learn about how the brain functions.

Secondly, Josette invited me to an event at Foyles Bookshop (in London) where the authors attempted to demonstrate some of the more intriguing experiments, with mixed success. It was a bit like “The Odd Couple” with Tom playing the role of Oscar (Walter Matthau in the film version) and Matt playing the Felix role (Jack Lemmon). Matt is the serious one, explaining things carefully and methodically, while Tom is the psychologist who leaps around demonstrating this, then that, then something else. Anyway, the wine was nice.

Now down to business. This really is an amazing book! I thought initially that it would be a quick, easy read with a quick look round the numerous web sites housing many of the experiments. No way. I found myself reading each “hack” carefully, trying to comprehend the reasoning behind the experiment and the results, and then visiting nearly all the web sites to see what else was available.

It’s not possible to describe each of the one hundred hacks, but I’ll mention some that I really liked. The book consists of ten chapters (Inside the Brain, Seeing, Attention, Hearing and Language, Integrating, Moving, Reasoning, Togetherness, Remembering, and Other People). As you can see, most of the chapters are devoted to our various senses.

Hack #23 (Seeing) illustrates how our brain relies on previous experience to reconstruct three dimensions out of two. In the diagram on the left we use our knowledge of illumination and shadow to perceive that square B is the brighter of the two. However, the diagram on the right shows that both squares are actually identical in colour.

<http://web.mit.edu/persci/people/adelson/checkershadowillusion.html>

Controversially, hack #43 (Attention) maintains that we can improve our visual attention by playing video games.

<http://www.bcs.rochester.edu/people/daphne/visual.html#video>

Hack #51 (Hearing and Language) illustrates that we read a sentence phrase by phrase, rather than word by word. Sentences become cumbersome if they overrun the buffer required to parse them, and that depends on the lengths of the individual phrases. Consider the authors’ simple example:

*While Bob ate an apple was in the basket.*

It's grammatically correct, but a comma has been omitted. Note that you have to read the sentence again to break it into different phrases; you can't juggle the words in your head.

Hack #62 (Moving) is the classic broken escalator phenomenon. We know it's broken, but the brain's autopilot takes over regardless. Our brain cancels out the sensory consequences of its own actions.

I'll finish on hack #71 (Reasoning) which has generated an awful lot of hot air since 1990. It illustrates how bad we are at thinking about probabilities.

You're a participant on a game show, hoping to win a big prize. There are three doors, behind one of which is the prize, and behind the other two are goats (booby prizes). You choose a door. The game-show host throws open one of the other doors (not yours) to reveal a goat. He then gives you the choice of sticking with your choice or switching to the other unopened door. Do you stick or switch?

The answer is that you should switch. Read this book to see why or see vos Savant's book referred to below [1].

As you can tell, I really enjoyed this book. If you want to know more about how your brain works, then this is the book for you.

[1] *"The Power of Logical Thinking" by M vos Savant [St Martin's Press, New York (1997)]*

---

## **Google Hacks, 2nd Edition** **Tara Calishain and Rael Dornfest** **O'Reilly and Associates**

**ISBN 0-596-00857-0**

**480 pp.**

**£ 17.50**

**reviewed by Sam Smith**

This is the second edition of the best selling book. The first edition released about 3 years ago started the highly successful 'Hacks' series with a bang. Those years have been a busy time for Google. They have released many new services, and upgraded everything: this book compensates, covering new areas such as gmail, and adwords. Of course, with any book talking about new technology, it can be considered obsolete in some respects before it is even printed.

Looking back at the first book, while some hacks have been cut to make space for new ones, and some have been updated, there is relatively little in the first book that no longer applies. There are just other, new, flashier things that warrant discussion – hence version 2.

Much of this book is a very useful occasional reference, as well as the gems that come from dipping in to find something that you didn't know you could do. If all you ever do with Google is use the obvious, and don't need to dig deeper with their Web Services or other additional technologies then you will get limited benefit from this. However, if you have a website which would benefit from a customised search engine (e.g. [www.ukuug.org/search](http://www.ukuug.org/search)) then you will get value from the content and would benefit from a copy.

A full index of everything covered is at

<http://www.oreilly.com/catalog/googlehks2/>

---

## **Linux Server Security (2nd Edition)**

**Michael D Bauer**

**O'Reilly and Associates**

**ISBN 0-596-00670-5**

**544 pp.**

**£ 31.95**

**reviewed by John Collins**

This book is a second edition of a book previously entitled “Building Secure Servers with Linux”. Its new name is probably better as it assumes people have Linux servers to start with and now want to keep it away from the great unwashed.

It is alas, a cruel necessity that make you have to spend time checking and enhancing the security of your system. If you leave a single loophole, someone will jump through it soon. We get people trying to get in via SSH by brute force attacks with about 2000 user names about once a day. Once in a favourite trick is to exploit a loophole in some server process to get root access and then install “rootkits” which as well as creating server processes – probably to relay spam – insert doctored versions of **ls**, **ps**, etc which don't display the inserted modules or processes. And they **chattr** all sorts of things to make it awkward for you to clear up the mess.

The book mentions just about every topic I could have thought of, quite a lot I didn't think of, and not a few things I didn't know anything about before, for example alternatives to **syslogd**. Successive chapters deal with network and IP security, remote administration, OpenSSL, DNS, LDAP, database (particularly MySQL), email and various email servers, web servers, FTP, log management and intrusion detection techniques. Quite a lot of time is spent defining terms you will probably have already met to even get that far – for example “what is an algorithm”, and “what is public-key cryptography”.

As it covers so many topics, it cannot cover anything in very much detail. For example, I tried to look up PAM and MySQL security features “in anger” but found the former barely mentioned and the latter didn't really cover quite what I wanted. The style is to outline each particular problem and give you a collection of pointers as to how to tackle it, but without a lot of detail. It has an extensive bibliography and website list at the end of each chapter to tell you where to go and to download documentation and up-to-date replacement software.

You will need to get to grips separately with each aspect sooner or later and you certainly have to go for the more detailed documentation in each case but I think within the limitations of its scope this book is a good, well-written and fairly comprehensive overview and introduction to each topic.

---

## **Internet Annoyances**

**Preston Gralla**

**O'Reilly and Associates**

**ISBN 0-596-00735-3**

**256 pp.**

**£ 17.50**

**reviewed by John Collins**

This book is a collection of “FAQ” style questions and answers to various “annoyances” readers will have dealing the Internet.

Sections of the book deal with Email (including Spam of course), connections, wireless connections, web hosting and blogs, browser, dealing AOL, Instant Messenger, Search engines, Security (including viruses) and on-line shopping and auctions.

As will be obvious, it is very heavily Microsoft oriented, and over 50% of the book assumes you are using Windows. A different 30% of the book assumes you live in the USA, listing ISPs and helpline numbers in the USA and prices of things in dollars.

If you're not using Microsoft, and you're not using Internet Explorer and standard email clients on Microsoft, this book will not be a great deal of use to you. The author is some sort of Microsoft guru and the very scant references to other operating systems are a little disparaging. Some readers may disagree with some of the advice given in places. For example complaining that too many people are visiting your website and doing things about it is unlikely to be a high priority for most people who usually want the opposite.

It's hard to feel anyone reading the sections on eBay and Paypal will want to use their services with the references in the book to sites like [paypal.sucks.com](http://paypal.sucks.com) and class legal actions. Amazon doesn't get much more favourable a treatment.

I don't think readers of the UKUUG newsletter will want this book because of the "UK" and the "U" before the "UG" myself!

---

## **Linux Cookbook**

**Carla Schroder**

**O'Reilly and Associates**

**ISBN 0-596-00640-3**

**592 pp.**

**£ 31.95**

**reviewed by Roger Whittaker**

In keeping with its "cookbook" title, this book doesn't try or pretend to be a comprehensive Linux textbook, but offers a large number of solutions to problems (recipes if you like) grouped around various themes, each theme constituting a chapter. By doing so, in a way it gets more material covered in a useful way than many of the "textbook" style books out there.

I like this format, because you can pick up the book and usefully dip into it randomly. If you're looking for specific information, you can also use it as reference book, subject to the warning that it is not comprehensive.

The tasks covered are a mixture both in terms of difficulty and context: for instance, there are sections on package management (both for RPM-based and Debian systems) user management and Linux text editors (vi and joe) which are intended mainly as introductions to these topics for the new user. At the same time there are (among many others) chapters covering setting up NTP, CVS, postfix, Apache and DNS servers.

The emphasis is on the fundamentals of Linux and setting up server systems: this is not one of those books filled with desktop screenshots, and although much of its content will be useful to desktop users, it is not about running end-user applications on the desktop.

Although none of the chapters offers complete coverage of its topic, in most cases the typical tasks that you might want to carry out are covered, and the chapter gives you more than enough to get started with. For a person wanting to learn about a particular topic, the process of "getting started" is often the biggest hurdle. This book helps with getting over those hurdles.

The information is presented in a clear and readable way, and is far more palatable than most of the documentation that comes with the programs under discussion.

At 592 pages in 24 chapters, with four appendices, this is quite a big book, and it is well produced in the usual O'Reilly style. I found one or two tiny and insignificant errors, but the only one I'll mention here is the amusing typo "Unbuntu" in the preface.

The book could be useful to anyone, but I think it succeeds particularly well in its stated aim of trying to help someone with an intermediate level of knowledge to take things further. The preface describes the book as “light on theory and heavy on how-to-make-this-go”, which is a good description, and, in my view, quite a good aim. Once you know how to make things go, you can more easily assimilate the theory from other sources.

---

## Revolution in The Valley

Andy Hertzfeld

O'Reilly and Associates

ISBN 0-596-00719-1

320 pp.

£ 16.95

reviewed by Mike Smith

I usually elect to review O'Reilly (as opposed to other publishers') titles for the User Group newsletter because I like them (obviously). This one came through and I had to look twice, because its not in the usual O'Reilly Nutshell format. Of course its not the usual Nutshell technical guide either – but it is an O'Reilly book.

I remember now why I said I'd have a look – though not a complete Apple convert (yet) I enjoyed the classic “Hackers” by Steven Levy; Tiger and the new G5s are becoming more and more compelling all the time and I just know I'll succumb in the end. I think I've got Insanely Great too, by Levy, but can't remember a thing about it, so maybe never got around to reading it. Either that or it's rubbish.

Anyway, this book ... Its full of photos of people, computer parts, screen images, notes, designs, artwork. The first and last 10 pages or so are various note book entries. Its nice to see the early designs, scribblings, etc took place. There's also an picture of Bill, but enough of that.

Each chapter is an account, and part of the story that runs from 1979 to 1985: all about how the Macintosh was made. The book is split in to five parts. I'm not sure what the criteria were for these five parts – it may be there somewhere, but its not obvious to me at the moment. The chapters are mostly in chronological order and I guess they are phases in the development of the Mac. The last part is post Mac release (which I'm sure you'll know was 24th January 1984).

Whilst we're on the launch of the Mac itself, the chapter entitled “The Times They Are A-Changing” (there are no chapter numbers!) is an account of the launch event. Reading this, with the description of what happened, how Jobs built the tension and the reaction of the crowd, really got to me. In fact I was so excited I went straight to Google to find the videos (they're out there, [macbytes.com](http://macbytes.com) has a link). So Hertzfeld is pretty good at writing, as well as Mac design.

The rest of the chapters are about the interaction between the various designers and developers, including Hertzfeld, around the various Mac projects. You also learn insights into how Apple functioned and some of the problems they had on the project, but surprisingly (perhaps) not a lot about the design of the Mac itself (though there are one of two things on the software design side, but nothing technical that I spotted). The stories are brief – at just a few pages, small enough to pick up and read at any time.

In summary, it's a jolly good read, and absolutely worth your while picking up. I won't go as far as saying that it's Insanely Great – that would be too big a pun!

If you want to know more, head over to <http://folklore.org/> where I think many of the anecdotes also reside, but the book itself is well worth it.

## High Performance Linux Clusters with OSCAR, Rocks, Open-Mosix, and MPI

Joseph D Sloan

O'Reilly and Associates

ISBN 0-596-00570-9

368 pp.

£ 24.95

reviewed by Mike Smith

I've been trying to get my hands on an O'Reilly Linux Clusters book for years now. The original "Building Linux Clusters" by David Spector got a right trashing by the cluster community and was withdrawn – so I never got to read it. This means, unfortunately, that I can't compare this new title with the original one. But by all accounts I'm lucky not to have read the original in any case.

So I'll be as objective as I can with this review, and will probably just stick to the facts – I don't want to be accused of being biased either way in such a controversial area!

This is a getting started guide, so don't expect anything too advanced. The approach is to list common (and author recommended, following his own trials) software solutions in a number of areas – things like the cluster engine itself (if I can call it that), administration, management tools and ancillary stuff like PVFS.

Clusters still seem to be very much in the high-performance compute space, and this book is also geared towards this area. I'm personally very keen to see if and how clusters will evolve to meet the requirements of the commercial/enterprise market, but I haven't come across a great deal yet (other than Oracle 10g, which is perhaps making a move in the right direction). Get in touch if you know more.

In typical O'Reilly style, there are four sections, starting with the basics – an overview of clustering, and moving through a getting started guide, building clusters and then a section on programming (how you design workloads for cluster environments through parallel programming).

The first section is, as I mentioned, about the basics, so I want dwell on this.

Because there are a large number of elements to cluster computing, and indeed a large number of options for each element, cluster installation packages have sprung up. The last time I looked as clustering software I wasn't aware of them, and the two covered here are OSCAR and Rocks. Chapters 6 and 7 give pretty much step-by-step guides on the installation process using each of these tools respectively.

Chapter 8 (commencing the section on building a cluster, and the main meat on the cluster environment really) covers a few tools used for automated system building. You'll be familiar with Kickstart, I'm sure, for a scripted installation, but you could also use imaging tools – which ought to be quicker, and a number are covered. The author uses a machine called fanny for the examples in this chapter. Well, it amused me anyway.

Chapter 9 covers the development environment, and I skipped that (it talks about programming languages, libraries etc).

Chapter 10 covers a couple of complimentary management tools – C3 (Cluster Command and Control) for executing common tasks across the cluster, and Ganglia for monitoring.

The remaining two chapters deal with the scheduling software (OpenPBS here) and Parallel Filesystems (PVFS).

That's it for building, and the remaining 5 chapters cover the programming aspects using the MPI libraries. This covers some pretty obvious aspects about parallel programming but I am sure there are many more subtleties I don't appreciate.



I am sure there will be criticisms of this title, and I'd disagree that its a "comprehensive getting-started guide", but if all you want is an overview and some pointers in the right direction I do think it fits the bill.

It might just be the thing to persuade me to create a few VMWare virtual machines and have a go at setting a cluster up – not that I have any reason to, except for fun.

---

## **Jakarta Commons Cookbook**

**Timothy M O'Brien**

**O'Reilly and Associates**

**ISBN 0-596-00706-X**

**400 pp.**

**£ 31.95**

**reviewed by Mark Jones**

"Commons" is an open-source subproject of the Apache Software Foundation's Jakarta project. Commons contains a set of twenty-plus utility components which serve useful purposes and assist the development of Java-based applications. The components range from high-level, simple utilities such as text manipulation, to more complex issues surrounding XML data binding and SSL management, and also deal with lower-level issues on Networking and I/O. Although Commons covers a wide range of subjects, each component seeks to achieve the same goal of making life easier and quicker for the developer. In particular, when used together in the right combination to solve the right problem, Commons can prove very powerful.

The Commons Cookbook is aimed at experienced Java programmers and low-level designers, and leads the reader through the majority of the components currently within the Commons sub-project. Each utility is kindly summarised and introduced, before a set of problems, solutions, discussions and examples are used. The odd bit of history is also used, for example when discussing JavaBeans and XML, which will keep nostalgia fans content, but also give context and provides a brief respite from the repetitive nature of this book. The examples are generally very specific (no pun intended), however they do give a feel for the types of problems that Commons are meant to assist with solving. Other Jakarta projects, such as Struts and Velocity, are cited where applicable which gives a higher-level view on the Jakarta project family.

Having worked on several projects that use Commons, one of the main difficulties that developers and designers face is knowing what tools Commons actually makes available. For example, a developer doesn't often know that Commons provides Templating tools, doesn't know how to find the specific tool that they need, or doesn't know how to apply it. In this respect the Cookbook provides an invaluable, brief (about 360 pages) summary of the Commons utilities and, when coupled with the various examples, gives the reader access to the immediate, high-level information they need.

I find making criticism of this book difficult, in that it does exactly what it says on the tin, and very well. One could argue, however, that the progressive and evolving nature of open-source utility libraries such as Commons will render this book obsolete fairly quickly.

If you're currently using Commons or plan to in the near future, then this book will be of use as a high-level reference to the main Commons utilities. If you require a more comprehensive, thorough insight into one or two of the larger Commons components however, then I would look elsewhere.

## Oracle SQL\*Plus: The Definitive Guide, Second Edition

Jonathan Gennick

O'Reilly and Associates

ISBN 0-596-00746-9

592 pp.

£ 28.50

reviewed by Harry Newton

This is the second edition of O'Reilly's SQL\*Plus guide, originally published in March 1999.

### Description

The 554 pages of the book are divided into fourteen chapters and two appendices. Most of the chapters address specific types of task, the remainder essentially being related to configuration matters.

Introduction to SQL\*Plus is exactly what you would expect from its title, giving a little history of the product (surprisingly interesting!), and a description of the database that is used in the book's examples. Command-line SQL\*Plus shows you how to navigate in SQL\*Plus: entering commands, basic editing, and sending SQL and PL/SQL to the database. Browser-Based SQL\*Plus describes the use of iSQL\*Plus, but much of this is self-evident.

'A Lightning SQL Tutorial' purports to be a concise introduction to SQL (mainly DML). This was the weakest chapter in the book: I couldn't see the reason for it. The book is a guide to SQL\*Plus, and if the reader doesn't have an understanding of SQL, then I can't see why they would be reading it.

'Generating Reports with SQL\*Plus' is the heart of the book, getting to the heart of SQL\*Plus: reporting. The approach followed is very good: the journey from initial, plain query, to column formatting, to page formatting, and then summary-breakdown is logical and understandable.

'Creating HTML Reports' describes the generation of HTML using the SQL\*Plus HTML formatter. This is interesting, and sensibly separated from the previous chapter. 'Advanced Reports' follows on from Generating Reports with SQL\*Plus: totals, report headers and footers, etc but I can't see the rationale for separating it from the other chapter. 'Writing SQL\*Plus Scripts' is also about bread and butter SQL\*Plus use: scripting. I consider this to be the strongest of the chapters; it is concise and well constructed.

'Extracting and Loading Data' is another strong chapter. It describes the generation of CSV files with the standard reporting techniques, and the use of SQL\*Loader to import data to the database. There is also a discussion on the use of external tables to load data. None of this is SQL\*Plus specific, but it is useful, it does illustrate the use SQL\*Plus, and it fits in with the approach of using SQL\*Plus to do common tasks.

'Exploring Your Database' is an interesting chapter. It basically describes the approach to finding the structure of the database objects (DESCRIBE and the data dictionary). It doesn't of course really belong in a book on SQL\*Plus, but the author uses it to illustrate some of the reporting and scripting techniques he has already described. 'Advanced Scripting' continues from 'Writing SQL\*Plus Scripts' and like the advanced reporting chapter, should not really be split from its predecessor. This is another strong chapter: the discussion of bind variables, and their relation to substitution variables is very clear.

'Tuning and Timing' presents the features of SQL\*Plus that assist in SQL statement optimisation: the timers, and the auto trace facility. It also has very good material on the explain plan, and optimisation based on the explain plan results. This isn't SQL\*Plus specific, but it is useful, and it fits in, like the previous two chapters, with the ethos of the book, if not the exact title. The Product User Profile returns to genuine SQL\*Plus material. I suspect the product user profile functionality of SQL\*Plus (a mechanism for restricting SQL\*Plus commands by user) is not well known, and this is a brief description of the functionality. 'Customising Your SQL\*Plus

Environment' is the last chapter. The two appendices are on SQL\*Plus Command Reference and SQL\*Plus Format Elements.

### **Differences from the First Edition**

The book has been updated to cover the developments in Oracle over the five years since the publication of the first edition, specifically Oracle 9i and Oracle 10g, the first edition going to Oracle 8i. There are new chapters on iSQL\*Plus, generation of HTML reports, and primer on SQL. There is another new chapter, Advanced Reports, but this is a very slight expansion of part of the material in the Generating Reports with SQL\*Plus chapter of the first edition, the equivalent in the new edition having been slimmed down. One chapter has been removed: Administration with SQL\*Plus. My preference would have been to retain this at the expense of the Lightning SQL Tutorial.

### **Should I buy it? Should I buy it if I have the first edition?**

People tend to have scanty knowledge of SQL\*Plus, often using just a fraction of its functionality. Typically it is used to load PL/SQL code from file and accept elementary queries. This is a waste as one can do many things rapidly with the tool, and it deserves to be better utilised. If you use Oracle a lot, then I would recommend that you get more familiar with SQL\*Plus. Whilst I was reading the book in preparation for this review, I found many things that I had either forgotten or never known. This book is one of the few on the subject: amazon.co.uk has two others, both of which appear to be more elementary.

If you have the first edition then the decision rests on how useful you will find the chapter on HTML generation, and how inconvenient you will find the loss of the Administration with SQL\*Plus chapters. The Lightning SQL Tutorial will most likely be read once by most people, as will the iSQL\*Plus chapter. I thought the book did well when the author offered his experience in using the tool, and in the task-based layout. It did less well in some of the ponderous tutorial explanations.

Before you do buy it though, remember Oracle's documentation. The SQL\*Plus User's Guide and Reference for 9.2 runs to roughly the same number of pages as this book. It costs a little less than this book of course. And do not forget the SQL\*Plus help command!

---

## **PePLink MANGA for University Embedded Linux**

*Phillip Stevens*

PePLink is pleased to offer Universities/Colleges offering a course specialising in Linux for embedded applications an opportunity to purchase the MANGA Plus at over 60% discount off list price.

Unit price is £ 80 + VAT (as applicable) + Shipping (quantity dependent).

The MANGA Plus is ideally suited for course work, as the SDK source is fully available (GPL), and lecturers can select any level of complexity desired for students from entry to post-graduate.

- building the tool chain
- configuring and building kernel
- configuring and building standard applications
- creating and adding new applications

Interestingly the current MANGA SDK includes a full Python implementation. Thanks to Mickey Lauer:

<http://www.vanille.de/projects/python.spy>

The MANGA Plus is a ARM920T core with 32MByte of RAM and 4MByte of Flash, an integrated 5 port 10/100 Ethernet switch, and serial console port.

The MANGA Plus is configured with 1x mini-PCI, 1x VT6212 USB2.0. The mini-PCI slot allows VGA, WiFi, or other options to be added. USB2.0 is useful for robotic control, data acquisition, HID, or adding mass storage such as USB-IDE drives.

Each student/customer will get:

- MANGA Plus
- UK 240V Power Brick
- Ethernet cable
- White packing box
- MANGA SDK URL registration (to download as desired)

<http://forum.peplink.com/viewtopic.php?t=58>

As an example of the use of this device, see

<http://www.pamurray.com/manga/>

The article describes how Alex Perry installed Debian Woody on a Peplink MANGA Plus with the root filesystem on a USB storage device on the USB port, and a power consumption of 3 Watts.

---

## Contacts

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

Mike Banahan  
Ely  
mike.banahan@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

Newsletter  
newsletter@ukuug.org

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