

news@UK

The Newsletter of UKUUG, the UK's Unix and Open Systems Users Group
Published electronically at <http://www.ukuug.org/newsletter/>

Volume 13, Number 3

ISSN 0965-9412

September 2004

Contents

UKUUG Secretariat	3
Discounted Books from UKUUG	3
Announcement: UKUUG Apple Technology Briefing 2004	4
Announcement: SELinux Symposium	4
Request for historical information	4
UKUUG Media Visibility	5
Microsoft's Linux event: Manchester	6
UKUUG Open Source Advocate	7
Screen if you wanna go faster	8
Book review: Google: The Missing Manual	9
Book review: Hibernate: A Developer's Notebook	11
Book review: Digital Photography Expert Techniques	12
Book review: Web Database Applications with PHP and MySQL (2nd Edition)	13
Book review: CSS Pocket Reference, 2nd Edition	13
Book review: XML Hacks	14
Book review: Hackers and Painters	14
Book review: Learning Red Hat Enterprise Linux and Fedora	15
Book review: Network Security Hacks	16
Book review: High Performance MySQL	17
Book review: Cascading Style Sheets: the Definitive Guide	19
Book review: Eric Meyer on CSS, Mastering the Language of Web Design	19
Book review: Designing with Web Standards	19
Contacts	21

UKUUG Secretariat

Jane Morrison

The Linux 2004 event, held in Leeds from 5th - 8th August, was well attended and a great success.

Pictures and write-ups from the event can be found on the event web site. There is also a report in the October 2004 issue of Linux Magazine:

<http://www.linux-magazine.com/issue/47/>

For those of you who were unable to attend please find enclosed a copy of the conference CD. This is a membership benefit which entitles you to receive all UKUUG CDs.

Linux 2005 will be held in Swansea and we are almost ready to announce the dates!

The AGM this year was held on Thursday 23rd September at UCL.

At the meeting, Mike Banahan and Ray Miller were elected as full members of Council.

Charles Curran retired from Council having completed two terms of three years. Members at the AGM expressed their appreciation of the the work that Charles has done as Council member and Chairman over the years.

Full minutes of the AGM will be distributed shortly.

The AGM was followed by a talk by Jon Haslam entitled "Solaris Dynamic Tracing: the Observability Revolution".

We have received a letter from Julian Field, winner of the Open Source Award for his MailScanner software. He clearly enjoyed the trip to OSCON, describing it as probably the best conference he had been to, and thanked UKUUG and O'Reilly for giving him the opportunity to attend.

Looking further ahead to February 2005 we are working on details for the Winter conference which will be held in Birmingham. A hotel venue has been chosen as it provides good facilities and means that all speakers and delegates can stay under one roof.

Please put 23rd and 24th February in your diary now. The event call for papers is enclosed.

The next Newsletter will be the December issue and copy date is 23rd November.

Discounted Books from UKUUG

Ray Miller

Many of you already take advantage of UKUUG's deal with O'Reilly to provide books to members at discounted prices. We have some good news for you: our friends at O'Reilly are offering us a bigger discount, and we are now able to provide O'Reilly books to members at 27.5% off the list price. This includes books listed as "coming soon" on the O'Reilly site (they will be dispatched when available).

We have also reached agreement with UIT Cambridge, publishers of Philip Hazel's "The Exim SMTP Mail Server – Official Guide for Release 4", and are offering this book to members at 30% off the cover price. This is a complete guide that will let you take full advantage of Exim on your network, written by the author of Exim himself. You can order this from UKUUG for only £26.00.

UKUUG has recently become one of the UK's main importers of GNU Press books from the FSF. We have in stock a selection of titles covering the essential GNU utilities: GCC, GDB, Make, Emacs, RADIUS, as well as Richard Stallman's "Free Software, Free Society". These are available to UKUUG members at 10% off the UK-equivalent list price.

All of these prices include UK postage and packing and handling. Simply contact the UKUUG Secretariat to place your order.

Full, up-to-date information about the discounts available is available at:
<http://www.ukuug.org/books/>

Announcement: UKUUG Apple Technology Briefing 2004

Sam Smith

Following the very great success of last year's Briefing, we are pleased to announce a second event entitled "UKUUG Apple Technology Briefing – Mac OS X and Heterogenous Environments" to be held on the 1st November 2004 at the Institute of Physics, London.

The principal speaker from Apple will be Jordan Hubbard – engineering manager for Apple Computer, working on BSD technologies underlying the Mac OS X system. He was one of the founders of the FreeBSD project in 1993.

A full programme will be made available at:
<http://www.ukuug.org/events/apple04/>

Entry is free but preregistration is required and places are strictly limited. This event is aimed at all UNIX users, particularly those interested in its core technologies and the associated development environments.

Announcement: SELinux Symposium

We have received the following call for papers for this event.

First Security Enhanced Linux Symposium

The inaugural symposium on Security Enhanced Linux (SELinux) will be held March 2-4 2005 in the Silver Spring, Maryland, USA (near Washington, DC). SELinux brings the power of flexible mandatory access control such as type enforcement to Linux. The symposium is an opportunity to learn about SELinux and share technical and application experiences.

The call for presentation and tutorial proposals is now open until October 1, 2004. All proposals will be reviewed by the review committees for inclusion in the symposium agenda. To submit proposals visit

<http://www.selinux-symposium.org/>

Request for historical information

We have received the following request from Peter Salus.

I am embarking on a new (another?) history: this one will be on FOSS. Originally, it was going to be on Linux, but there is too much concerning FSF/GNU, the various BSDs, and their antecedents to limit the scope.

I'm interested in getting a wide range of input from contributors and early users. Please let me know if you have something to contribute – and also let others know of this.

peter@netpedant.com

UKUUG Media Visibility

Leslie Fletcher

Council is anxious to improve the visibility and credibility in the technical and general press of UKUUG and its mission to advocate open systems, promote free and open-source software, and advance open programming standards and networking protocols.

To help to achieve this, the UKUUG Office will maintain a database of individuals to whom press and media enquiries can be passed. Rather than publish the database, the Office will take details of any enquiry such as topic and whether a quotation or article or interview is being sought. These would be passed to the relevant expert who would be asked to contact the journalist or publication concerned. Council hopes that, in almost all cases, the expert will agree to follow up the enquiry as a matter of urgency. It would also be in order to suggest an alternative contact, or decline to be involved at all.

Individuals would speak to the press in their own right and would not be expected to adhere to a UKUUG official line, although they should be mindful of UKUUG's principles. It would be open to individuals to decide whether to give their contact details to enquiring journalists for future reference. While Council would encourage this, it will not be expected or required of those involved.

Council hopes that the following topics and areas will be covered:

- Unix, Linux and related operating systems
- Free and Open Source office and other user applications
- Security (technology and general issues)
- Internet and Web (technology and general issues)
- Licencing (GPL, BSD, ...) and intellectual property
- Software patents (limited specifically to "openness" issues, general queries being passed on to FFII)
- Credentials of major vendors: IBM, Microsoft, Red Hat, SUSE/Novell, ...

FLOSS in:

- UK central government and European Union institutions
- UK local government, including schools
- UK further and higher education
- Large corporations
- SMEs

UKUUG members are asked to put themselves forward, to a member of Council, or suggest others. Council will also approach individuals whom it believes have relevant expertise within the next few weeks. Council will decide on the initial list of expert contacts by the end of September. Council will only aim to respond to enquiries in areas where it is confident that an authoritative person is available.

The technical and general press in the UK will also be reviewed for news items to which UKUUG might respond. Members will be asked to volunteer to watch for such items in no

more than two publications and to contact UKUUG with the details of anything they find. Leslie Fletcher has agreed to manage this list of volunteers, to act as a clearing house for news items and to contact individual experts directly when a speedy response is required. Council hopes to launch this scheme in the autumn once the database of experts has been established.

Microsoft's Linux event: Manchester

Leslie Fletcher

“20:20 Seminar Series: Microsoft Windows and Linux – An open and honest technology discussion”

These comments on the Manchester event are intended to complement Alain Williams's write-up of the London event, which can be found at:

<http://slashdot.org/comments.pl?sid=110766&cid=9405641>

This (“in inverted commas”) is what Microsoft said was going to happen, together with my considered feelings in the days after the event about what actually happened.

“Reams have been written about Linux and Open Source Software. Some of it true, some of it hype. We think it's about time to open up this debate.”

Glitzy sales event it certainly was, and brilliantly choreographed; debate it was not. I found some of the tone patronising and there was much faint praise of Linux and other Open Source offerings. ‘Free’ was a word used many times by the platform juxtaposed with ‘open source’ in a manner which can only have been deliberately intended to confuse. Some of the so-called hype was manipulated in near subliminal fashion, such as several passing warnings of alleged licensing and indemnity uncertainties and leakage of IP which business users of Open Source should worry about. There was also an attempt to portray what is happening in Open Source as Linux v. Unix, of only academic interest to business users.

“Microsoft would like to invite you to an open and honest technology seminar about the pros and cons of Microsoft's technology platform in relation to Linux and Open Source. We want to share with you all the latest information we have, from all sides.”

Glossy sales event showing off some Microsoft goodies, not a technology seminar. It might have given warm feelings to existing Microsoft users that they had made the right decision and that nobody, including them, would get the sack for choosing Microsoft. There were classy videos of reassuringly big Microsoft users such as the Stock Exchange. In most substantial respects, the event was deeply uninformative. I was astonished to discover that there was no delegates pack with copies of the speakers slides, delegate list, ... so it is impossible to follow up any of the points made. There was no real sharing of information.

“In a half day event, which will include a lively panel discussion as well as views from an independent analyst, we will be sharing with you different perspectives surrounding this debate. We want to show you our side of the story, why we believe Microsoft's platform of technologies is still the safest, most cost effective route for many organisations and where we see the future.”

The panel discussion was dire and merely filled the time before lunch. The questions submitted beforehand were heavily filtered by the master of ceremonies (I think that is the only description for him); a couple from the floor were significantly more penetrating. The independent analyst was worth listening to, and the Linux expert sitting next to me felt that what he had to say was mostly fair comment.

“At the event you will also be able to listen to and question a number of IT decision makers who have already been through the process of looking at Linux and Open Source as a platform for delivering IT business solutions.”

There was no opportunity to question the speakers. The speakers were indeed IT decision makers, all of whom had decided that Microsoft provided the best platform for their business; the platform was not open to anyone who decided for Open Source. The main argument for deciding Microsoft seemed to be that other people/organisations in the business chain use Microsoft. There were many assertions of the superiority of Microsoft products but very little evidence.

“This will be a unique opportunity to question the business and technical experts at Microsoft about some of the issues you’re currently facing and some of the decisions you’ll need to make.”

As I have mentioned, there was little opportunity to ask questions. I would liked to have heard more from Nick Barley as he was the only one who got above ‘our goodies are better than their goodies’ point scoring, but he didn’t even stay for the lunch.

As you can see, I found the event very unsatisfactory and in almost every respect it failed to do ‘what it said on the tin’. However, it raised in my mind some important questions for the Open Source community, including:

- How can the business benefits of Open Source be better promoted?
- How can business decision makers be convinced that the systems they need can be built from or make effective use of Open Source components?
- Is the Open Source community really on top of the technology - visualisation, moving image, ubiquitous computing, ... and if it is, how would a user set about procuring an Open Source solution which included such things?
- Could/should there be an Open Source offering in every major area of IT and its applications?
- What are the aims of Open Source development over the next, say, five years, how convincing are they and how certain is it that they will be achieved?
- When and how will Linux on the desktop become available in a way that is credible to users without the need to evangelise?
- What is the critical mass for Open Source deployments in the UK, and when will it be reached? How can UKUUG and others in the field make as sure as possible that this happens and that it happens as quickly as possible?
- Could/should there be an Open Source road show from time to time?

I hope these questions are not naive or gratuitously contentious, particularly as this my first post to the newsletter! I believe that it is very important to have convincing answers to them and then to get the answers into the key arenas - business, politics, public service, education etc. - which is where I come in I guess.

Leslie Fletcher – UKUUG Open Source advocate

UKUUG Open Source Advocate

Leslie Fletcher

Following a lengthy discussion at the 2003 AGM, Council has asked me to take on the role of UKUUG Open Source Advocate. I will be working for Council on a part-time basis to improve the visibility and credibility of UKUUG and its mission in key arenas – business, politics, public

service, education etc. All concerned are aware that this is a huge task so I would welcome advice, offers of help and suggestions for tasks I could carry out.

My main first-hand involvement with Open Source is as chair of governors at Parrs Wood Technology College in South Manchester.

<http://www.parrswood.manchester.sch.uk>

This 2000-pupil community school, which is Manchester LEA's flagship, has invested very heavily in IT while I have been a governor, as part of a new building project. The school is a committed user of open source, in back-office applications so far, but moving towards an open source desktop as the technology improves and staff horizons are widened. The school's use of Linux thin-client technology to provide any child with home access to school IT resources is widely respected and has been commented upon (favourably!) by ministers. The school was sponsored by Red Hat and has been written about on the BBC web site and elsewhere so it has some visibility in the open source in education community. I have attended conferences in this area, giving me a wider view of the opportunities for Open Source in education.

See also:

<http://www.europe.redhat.com/news/article/26.html>

I also have some experience in IT management, having been head of the Department of Computer and Mathematical Sciences at Salford University for five years. Because of this, and my involvement with Parrs Wood, I have done some consultancy work for Manchester LEA on the educational and business process necessity of improved connectivity for schools; one of the reports I wrote can be read at:

http://www.night-shade.org.uk/Well_connected_city.pdf

Screen if you wanna go faster

Andrew Stribblehill

What is it?

Imagine the scene: you've got seventeen xterms running; some open, others iconised. A few of them are showing log-files, others are running mail, news and IRC client, and still others are logged into remote hosts that you're working on. We've all been there and it can get pretty boggling.

Enter `screen`: the virtual terminal manager. It does away with the 1-1 model of xterms to terminals. You can have, say, ten terminals running with no view of them, or conversely ten xterms all showing the same terminal.

An example would come in handy here. We start `screen` by typing `screen` followed by an optional command, say `ssh nntphost -l nntp`. The screen will clear and be replaced with this command running in a pristine virtual terminal. (If you had omitted the command, a shell would have been started.) Perhaps you start a job that turns out to run for longer than you expected. You can disconnect from your screen session, with `Ctrl-a Ctrl-d` which carries on as if you were watching it all along. To re-attach, use `screen -r`. Or you can start more virtual terminals from the same session and switch between them at will, rather like Linux's console. To make new terminals, type `screen [command]` or `Ctrl-a Ctrl-c`. Switching between them is just `Ctrl-a [digit]` to go to `screen [digit]`, or `Ctrl-a Ctrl-a` to go to the most recently used terminal.

All the commands within `screen` have the prefix `Ctrl-a` but if you dislike this, re-map it by putting one of the following lines in `~/ .screenrc`:

```
Make Ctrl- the control character: escape \034\
```

Or *. To get a real *, press *8. escape *8

Further features

As if that wasn't enough, you can log a terminal's output to a file a la script (Ctrl-a H), make a screen dump (Ctrl-a h), watch a terminal for activity (Ctrl-a M) or inactivity (Ctrl-a _), search back through a terminal's output for a string (Ctrl-a [?), along with a Byzantine troupe of other features.

Write-once, run anywhere

I occasionally need to run a command on each host in a cluster - usually a long-winded procedure. screen makes this a doddle. Using the at command, I can type on all or some of my screen-managed terminals simultaneously.

Example: Ctrl-a :at ssh stuff /bin/su - root [rootpw] will make me root on all my ssh sessions.

Multiple users

But the most interesting application of screen is as a multi-user multiplexing tool. You can start a screen session then make a window of it accessible to other users on your host. They can then attach to it on a read-only or read-write basis - rather like dual controls on a car. This can come in very handy for pair-programming or guiding a remote colleague through a task.

To do this, enter the screen command multiuser on by typing Ctrl-a :multiuser on. Now add the user with the acladd username command. They can connect to the window using screen -x your_username, then read and type at will.

Use it

screen is a power tool for console junkies. It will make you more productive and help you regain power over the dreaded proliferation of the xterms.

Andrew Stribblehill is a system administrator at the University of Durham.

Google: The Missing Manual

Sarah Milstein and Rael Dornfest

Pogue Press/O'Reilly

ISBN 0-596-00613-6

311 pp.

£ 13.95

reviewed by Damian Counsell

All prospective purchasers, expert or novice, are likely to sympathise with the back-cover "Missing Manual" logo. It shows an wide-eyed and bemused customer shaking a large, open software package with one hand and holding a single, lonely CD in the other. He waits vainly, like many of us have done, for anything else to drop out of the box. The "Missing Manuals" series looks like O'Reilly's attempt to abseil down from the cold, geeky peaks whence it sells tomes on programming embedded systems in C or administering Sendmail servers. Under its new "Pogue Press" imprint O'Reilly can release popular books and compete on publishing's lush, lay-reader lowlands, where beginners graze on "The Complete Idiots Guide to Windows XP" or "Rugby Union for Dummies", and it can do so without affecting the good and exclusive reputation that the parent O'Reilly brand has earned with the computing elite. Hungry Minds Inc., the publishers of the "For Dummies" series was recently acquired by Wiley, so O'Reilly's creation of the Pogue Press is probably the first strike in a proxy war between technical publishing giants. It might turn out to be ugly, but ordinary book-buyers should benefit from the clash of the Idiots, the Dummies, and the Missing.

You are reading this review in the UKUUG magazine so it's likely that you are comfortable climbing those geeky peaks of technical knowledge. You probably started using Google long before your friends – if you have any friends. You probably think you are pretty Google-savvy anyway. You probably think (as one of my coding colleagues exclaimed when he saw this book) that ninety percent of what you need to know about Google could be crammed onto two sides of A4 paper, never mind extended to 300 pages. And, if you picked up this volume and browsed the folksy language – describing interfaces as being “trippy”, for example – you'd probably toss it aside in lofty disdain. You'd be making a mistake.

Even if someone else writes another, longer book about how to use Google, it will have to be outstandingly good to beat this one. Though partly obscured by the “Ned Flanders” prose, there is so much information laid out here about the search engine and its offshoots that readers will wonder how they ever managed without some of the features it describes. If, when you read that previous sentence, you didn't know who Ned Flanders was, and were near a networked computer, you probably thought about googling his name. But this book will also teach you that Google can be used to search for phone numbers, survey ranges of digits, use wildcards for words you can't remember in the middle of phrases, filter your results on document type, and ask questions of real people. The Missing Manual not only covers this last (socially fascinating) feature – Google Answers – but it also describes every spin-off that I've heard of from what the company brochures probably call the “core functionality” of Google: Google Ads, Google Groups, the Google Toolbar, the Google Directory, and the book even claims that the “Missing Manuals” online companion site (<http://www.missingmanuals.com/>) will preview the relatively new Gmail (“unlimited” Web-based email) service, though, when I searched this resource, I couldn't find anything about Gmail myself. Further, the book covers third-party, non-Google, googling accessories and has the good sense to recommend alternatives for tasks for which Google would not be the best tool.

From my scanning of the bits at the beginning of the book and the bits at the end of the book, there was no mention of direct assistance being given to the authors by occupants of the famously secretive Googleplex (Google's global headquarters), but I suspect Sarah Milstein and Rael Dornfest had friends on the inside. When it comes to what might be considered the more technical side of Google: getting indexed, removing yourself from Google, making money from and paying to advertise via Google, this book has still more to offer. I was particularly impressed that www.missingmanuals.com even offers a spreadsheet for you to download to calculate how much you should be prepared to spend on advertising via Google. It would be a shame if übernerds missed out on the later, more specialised advice in the book because they were dismissive about the earlier more general stuff.

Despite all this, “Google: The Missing Manual” has, I must admit, the worst first line of any O'Reilly I've read: “By now there's no way to have missed the Internet.” The content that follows the inauspicious start probably contains mistakes. My approach to reading it was to skim through explanations of familiar things (like a description of what a 'Blog is, for example) and then pay close attention to the segments where new information lay, but I was so busy noting down new snippets of advice that, if there were any errors or typos, I didn't notice them. The structure is logical and the design tasteful. It's divided into sensibly chosen parts and chapters and illustrated with lots of worthwhile screenshots. The words are set in clear, unshowy fonts and arranged on each page in a pleasing way.

I was surprised to see this week that Google, the near-undisputed king of the search engines, has been around now for a gob-smacking six years – in Internet time that equates to something like two decades. What with that relative maturity, the emergence of Gmail, and Google's successful listing as a public company in the United States, this book is well timed. It deserves to sell by the shelfload. Ironically, the people who could most benefit from it probably won't buy it. I do scientific research, program computers, and run a 'Blog. Google is part of my life. I thought

I was quite an expert, but this book, which at first looks like it is intended for people who live on the plains, where non-technical users are at home, is extremely useful to those who think of themselves as residing high above the world of the “For Dummies” and “Missing Manuals” series. There are always greater heights of Googling to ascend (and probably there will continue to be). This book is a superb guide to anyone who wants to tour them.

If you are a heavy user of Google and would like to be a better user of Google, you should buy this Missing Manual. Google, after all, is still free.

Damian Counsell is a Bioinformatics Specialist at the Medical Research Council's Rosalind Franklin Institute of Genomic Research (formerly known as the MRC HGMP-RC). His academic homepage is at <http://www.hgmp.mrc.ac.uk/~dcounsel/> and his personal homepage is at <http://www.counsell.com/>. These Websites are not in any way related.

Hibernate: A Developer's Notebook

James Elliot

O'Reilly and Associates

ISBN 0-596-00696-9

190 pp.

£ 17.50

reviewed by Mats Henrikson

The O'Reilly book *Hibernate: A Developer's Notebook* is an introductory book about the Hibernate object persistence framework. It assumes that the reader has quite a good knowledge of the Java language and some knowledge of SQL and relational databases. It does not assume that the reader has any previous experience of Hibernate.

I have been looking at object persistence frameworks for a while now, to try and get away from having to manage large amounts of JDBC code in my projects. The only other framework I have looked at in detail is Java Data Objects (JDO), which unfortunately requires a byte code enhancer to be run on your compiled classes which makes it cumbersome to use. Hibernate on the other hand does not require a byte code enhancer, as it only requires that classes implement the `java.io.Serializable` interface. The rest of the magic is taken care of by putting Object/Relational (O/R) mappings in XML files. This makes Hibernate a much more attractive and simpler to use option than JDO. Hibernate also has a number of other nice features, such as Java source and database schema generation from the XML O/R mapping files.

The book reads a little like a Dummies book although it is not at all as basic. The text contains the occasional spelling mistake, but it is not any worse than any of the other first edition O'Reilly books I have read. The style of the writing is short and concise and not overly wordy, making reading and understanding it very easy. There are frequent explanations and code examples spread out in the text to illustrate how to use Hibernate by writing a small sample project.

All but the last chapters rely on the reader either completing the previous chapters exercises or downloading the source available from the book's web site. The first chapters are very easy if the reader is familiar with how to get and install Java libraries and the Ant build tool. The first chapter is all about installation and configuration, while the second two show how to actually use Hibernate for persistence. The next two chapters show how to store Java Collections using Hibernate. There is then a chapter on enumerated types, followed by a chapter on how to create custom types that can be used to wrap around objects not initially intended for persistence. The next chapter is the most interesting in the book, as it shows how to find and retrieve persistent objects in a programmatical way. This makes it possible to catch most errors at compile time, instead of at runtime as would usually happen with a query that originated as an embedded

string. The final chapter then goes into more detail on how to use the Hibernate Query Language (HQL), and also how it is possible to use SQL with Hibernate instead if the reader so prefers. After that follows three brief appendices and an index.

Overall, the book tells the reader what they need to know, and usually says when the online reference documentation have more useful detail on a particular subject if it is outside the scope of the book. It is not really a book that will be used very much for reference except for a few key parts, as the online documentation is quite good. It soon becomes apparent that the book will not remove the need to know SQL and relational database concepts in general. It would in fact be quite a difficult read for somebody with no knowledge about indexes and constraints on columns, how data is stored and retrieved from the database, and at least a little about good database design. What the book does show though is how to get by using a minimal amount of JDBC code, with the added benefits of being able to let Hibernate do most of the database schema for you and generate the Java source code required for your persistent classes.

The book could go into more detail in places, as well as discussing some more advanced topics, but the aim of the Notebook series is to be shorter, more of an introduction than a complete reference, and as an introduction it works well. It is quite short, and it is possible that there are other topics that could have been discussed without having made the book too advanced or too long. As it is it feels like quite an abrupt ending, although it does refer the reader to the Hibernate website and the reference documentation found there. Even so, I think that the author does what he set out to do, he manages to introduce the reader to Hibernate in a nice and easy way instead of being tedious, and I think this makes it quite a good read.

Digital Photography Expert Techniques

Ken Milburn

O'Reilly and Associates

ISBN 0-596-00547-4

467 pp.

£ 31.95

reviewed by Lindsay Marshall

There are hundreds of books on digital photography. Remainder bookshops are full of them, and mostly they are all the same: a bit like 'learn how to draw' books, where they seem to be telling you useful information but which in practice really don't help the novice at all. This book from O'Reilly is much better. It starts at the beginning and holds your hand pretty well all the way through, unless of course you run Linux in which case don't bother as all of the stuff on image processing is based on Photoshop. Not just Photoshop, but Photoshop CS, the latest version. The one that won't open scans of money. The one that nobody except professionals owns because it simply isn't worth paying the money for since the version you have still has more features than you know what to do with.

OK, so that's a niggle, most of the stuff that the author demonstrates uses filters that most photoshoppers already have, and the things he talks about are genuinely useful from dealing with problems in photographs all the way through arty manipulations to selling your digital images on the web and even printing and framing hints. The advice on not taking shots with problems in the first place is sound as well, though I must confess to being underwhelmed by the section on converting photos to paintings: yes, there are lots of filters to do that kind of thing, but most people have the good taste not to use them because the results are vile, and the examples bear this out admirably. The sections that I will spending most time with are the ones about making complex selections and sharpening and blurring : all hard to do well without good advice.

This is a big book with lots in it. There are high quality colour images on almost every single one of its shiny pages. It is expensive but it is the best book on getting going with digital photography that I have come across. All I need now is the time to spend fixing up all the crappy photographs I've taken over the years. And a fancier camera. And some talent...

Web Database Applications with PHP and MySQL (2nd Edition)

Hugh E. Williams and David Lane

O'Reilly and Associates

ISBN 0-596-00543-1

816 pp.

£ 31.95

reviewed by Lindsay Marshall

Lots of pages, big price. Hmmmm, I have some worries about this new edition. Yes, it covers PEAR and PHP 5 and MySQL 4.1 which is all to the good, but it also covers basic PHP programming, basic SQL and even some basic JavaScript. There are probably two if not three slimmer, more useful volumes struggling to get out of this chunky tome. And the new MySQL interface, `mysql_i`, is relegated to a few pages at the very end of the book; probably a victim of production scheduling clashing with software release timescales but a section for which it might have been worth waiting.

One of the biggest problems of this and indeed all books on PHP and MySQL is that the authors insist on writing about the features of quite old versions of the systems that probably most people don't use anymore. OK, both systems tend to be moving targets but that isn't any excuse for cluttering up the book with information that is of little use to most people. Flag the problem areas and point people to the web documentation for the older stuff. If you have any knowledge of PHP and/or MySQL finding anything new or useful is hard in this book. There is just too much here: too many long examples, too much syntax, too much advanced detail, too much other stuff: do we really need a 16 page manual for the `EZPDF` class? Scattered throughout the text are watchpoints about pitfalls – these are good, but there aren't enough of them to make up for the waves of cruft.

As it stands I cannot recommend this book to anyone. Take it apart and turn it into three nice thin (cheap!) books and the content will stand its own against any other but as it is it is just too hard to use both for novices and experts. I shudder to think what they could do with a 3rd edition!

CSS Pocket Reference, 2nd Edition

Eric A Meyer

O'Reilly and Associates

ISBN 0-596-00777-9

134 pp.

£ 6.95

reviewed by Sam Smith

The pocket reference comprehensively covers the syntax and options of CSS versions 1, 2 and 2.1. It also provides a brief overview of how CSS works, and the ordering that rules that are applied. Options are grouped by where they can be applied, so you can quickly and easily go

from what you have to find the style you need to use. Brief coverage is also given to non-screen display methods - both printed and aural.

If you need a large reference for CSS, you'll want both "CSS, The Definitive Guide" and this pocket reference (by the same author).

If you need a small reference for CSS, you'll want this pocket reference. It does exactly what you would expect of a pocket reference, and does it well.

XML Hacks

Michael Fitzgerald

O'Reilly and Associates

ISBN 0-596-00711-6

479 pp.

£ 17.50

reviewed by Sam Smith

The latest in the series of "Hacks" series is XML Hacks. In terms of form and style, it is very similar to any of the other books; 100 tasks and good ways of solving them. It's unlikely that everyone will use all of the Hacks, but, if you use XML regularly, it is probably worth a look for something new and useful.

The book is very Java based, with many Java examples. This makes the book very cross-platform (a significant number of the examples show Windows screenshots). It does however mean that the XML tools which are predominantly UNIX based get less of a consideration than they could have. There is less coverage of Perl or Python directly than I would have expected, but the methods and ideas described are relatively portable.

This also runs through the book, in that it doesn't feel very "UNIXy". I felt that the methods were significantly more closed in terms of extensibility, compared with the Google or Spidering books.

One of the benefits of the Windows content is that there is good coverage of XML and the Microsoft Office tools, and getting content from there to interoperate. Similar hacks are provided for OpenOffice, those being a lot simpler to do. One quibble I have, is out of the graphical XML editors it discusses, only one is available on non-OSX Unix, and then only for a fee. Whether this is a lack of a decent free XML editor (Vim/Emacs et al don't count), I'm not sure.

As with all the Hacks series, these are books you will generally have as a reference after quickly skimming once for background inspiration. The main use being to dip in to the Contents and finding a cute way of doing something specific when you have a problem you need to solve.

The book is good, and worthwhile for those who deal with XML; but I'd recommend having a quick skim read of the contents first:

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

Hackers and Painters

Paul Graham

O'Reilly and Associates

ISBN 0-596-00662-4

271 pp.

£ 15.95

reviewed by Sam Smith

This book, subtitled "Big ideas from the Computer Age", has an impressive billing to live up

to. The collection of essays from Paul Graham are on a set of topics which are of interest to hackers, in the purest terms of the word. The Planners, Architects, and Hod-carriers of the computer world.

To quote from the opening of the preface:

This book is an attempt to explain to the world at large what goes in in the world of computers. So it's not just for programmers. For example, Chapter 6 is about how to get rich. I believe this is a topic of general interest.

As a collection of essays, the chapters are very independent. With a glossary of technical terms mentioned, chapters are easily understandable, even if you are not aware of the area before. You will probably find something of value in the essays; even if you do not find something of value in them all. Equally, you do not have to agree with the point for it to be useful and valuable to read.

Whether you will buy this book depends on what you think about. All of the essays in the book, are also available for free from Paul Graham's website:

<http://www.paulgraham.com>.

So, if you are interested in reading, you can do so without buying the book. If you want to buy the book, you don't need me to recommend it.

Learning Red Hat Enterprise Linux and Fedora

Bill McCarty

O'Reilly and Associates

ISBN 0-596-00589-X

352 pp.

£ 28.50

reviewed by Roger Whittaker

This is actually the fourth edition of this book, which was previously called "Learning Red Hat Linux".

Aimed largely at the new Linux users and home/desktop users, it is a good general introduction to Linux of a certain type. It is a "how-to" book: the emphasis is on telling the reader what to click and how to use the GUI tools that come with Red Hat / Fedora. It doesn't neglect the command line altogether, but there is very little discussion of configuration files and what is going on behind the scenes.

When I was teaching Mathematics, I often became involved in discussions about whether it was enough (or right) to teach students "how to" get the right answer without teaching an understanding of the underlying concepts. It's easy to take a high-minded view on this in theory, but in practice there is always pressure to take the easy way out.

This type of introductory Linux book bothers me for the same sort of reasons. There are plenty of them around, and this is one of the best of them, but they leave me feeling disappointed because I feel that they take an easy way out which in some way doesn't help the reader as much as it appears to.

However, I know from recent bitter experience as junior joint author of a somewhat similar book (different publisher, different distribution...) how difficult it is to get the balance right between explanation and description, and how hard to decide what to include. But I do feel that most buyers of the book will find that it exhausts its usefulness fairly early on, even if they are total Linux beginners.

Clearly the publishers had to include the words “Red Hat Enterprise” in the title, but in practice the book is not really aimed at, or useful for, people who want to use Red Hat’s Enterprise versions for the first time for any serious purpose.

The book includes two CDs (containing Fedora Core 1). Although published in April this year, that’s almost two versions out of date. But that’s the way of the world.

The book is accurate in its information, well produced and has plenty of good screenshots and examples. I would recommend it to a Linux beginner, but on the understanding that the information it imparts needs to be supplemented from other sources and that it won’t answer any of those nagging questions about how things really work.

Network Security Hacks

Andrew Lockhart

O’Reilly and Associates

ISBN 0-596-00643-8

316 pp.

£ 17.50

reviewed by Mike Smith

Those of you who have read my previous reviews on the Hacks series will know that I like the format – it’s a quick reference guide with 100, usually relatively well documented, tips.

I’m not quite as happy with this particular book. A Hack, to me, implies a clever trick – something unusual, out of the ordinary or innovative. This interpretation is compatible with other books in the series I looked at, such as creating smart addons for eBay, and tips for Servers – but not really appropriate for security. Security is all about process and methodical control to reduce risk – so I think the scope for Hacks is limited. I know this is a general statement, so there may be exceptions. One example would be to establish a Port Knocking environment, for instance. This is still relatively new and innovative, with several different ways to implement it. However I didn’t spot port knocking being covered.

Some examples of “Hacks” here include scanning for suid and sgid programs, group and world writable files, and using sudo. ie General good practice. I could go on (use chroot, for instance)... So these things are not hacks as far as I am concerned. There are some interesting tips though – its not all bad. I’ll to get them later.

The book covers Unix (20 hacks), Windows (10 hacks), Network Security (23), Logging (7), Monitoring (6), Secure Tunnels (15), NIDS (14) and Recovery and Response (5). So the scope for each area is limited.

grsecurity is interesting – I hadn’t come across it. Its a kernel patch with various features – increasing entropy for the things requiring randomness, and locking down various areas better with ACLs (in both kernel and user space).

The obvious first Windows tip is to use HFNetChk. Also, a tool from sysinternals (who I know of) is recommended, and another from Foundstone (who I didn’t) that provides a feature I’ve always wanted – displaying network ports and the associated running processes (like lsof on Unix, but not quite as good I suspect). I didn’t know about the IP Security Policy Management snap-in for the MMC. This lets you set up firewall rules on the host – worth knowing, if you have to use Windows in your environment.

In the network section, there is some good stuff. One product is recommended for network scanning – Nessus, obviously. Also SFS to replace NFS, advice on securing MySQL and BIND and lots of other areas.

On to Logging. There's a tool for forwarding Windows Event Log events to a remote syslog. This helps with consolidation of events in a multi-OS environment. syslog-ng is covered too. Not a lot else: logwatch and swatch.

Monitoring is one of my favourite areas (I wrote a monitoring system years ago, before Big Brother and the commercial frameworks came on the scene – its still in use today, actually. That's stood the test of time!). The "Hacks" here include using Nagios, RRDTool, ntop and argus. These aren't really Hacks, as I discussed at the beginning of this article, just product recommendations and a few screen shots really.

There's a good chapter on tunnelling. One of the other Hacks books (the server one, I think) had some tips on setting up tunnels and similar areas are covered. Large commercial environments tend not to tunnel across the Internet much. I think they should, but they don't – it's not a question of security (which tunnelling addresses), but of guaranteed availability. Although the Internet infrastructure is resilient, when things do go wrong with something you're relying on, there is rarely compensation. Having "dedicated" (hardly anything is dedicated these days, of course) communications infrastructure permits service level agreements to be enforced.

FreeS/WAN is covered – the final release of FreeS/WAN came out in April and is no longer being developed. So I'm not sure what'll happen in future, or of the wisdom of using it for any long term solution. There are also tips on setting up IPSec on FreeBSD and OpenBSD, PPTP on Windows (but not on other OSs), SSH, stunnel, httptunnel, VTun, OpenVPN and using PPP with SSH.

The Nids section can be summarised by the following list: Snort, ACID, sguil, SnortCenter, Snort_inline, SnortSam, Oinkmaster, Banyard, honeyd, sebek and writing Snort rules. i.e. (Mostly) Snort, Snort, Snort, Snort, Snort, Snort, Snort! That's a little one-sided perhaps.

There is a short chapter on "recovery and response" – i.e. forensics. Use tripwire and rpm to check for changed files and chrootkit to scan for root kits. A good tip (the very last one) is to use geekttools for whois lookups. This was from Rob Flickenger – he gets all over the place! I have has a script `awhois.sh` that I must have picked up somewhere that does a similar thing – selects the right nic for a whois query based on the domain or netblock.

Despite my initial reservations, I do again quite like the book. The format isn't quite as appropriate for security, but it does cover a lot of ground (very briefly, obviously) and it provides a good starting point to build on. Alas, nearly as good would be to get the list of contents and simply Google for them (as long as Google is up, unlike what we saw at the end of July!)

High Performance MySQL

Jeremy D Zawodny and Derek J Balling
O'Reilly and Associates

ISBN 0-596-00306-4

294 pp.

£ 28.50

reviewed by Mike Smith

High Performance. That makes you think of speed, but there's more than just raw speed to think about. This book is well thought out (however wait until you read the whole of this review before making your own minds up), covering some of the basics but reminding us that there are wider issues to consider than just tuning the app server and storage to produce a high performance service. So the book also looks at replication schemes and load balancing to support higher throughput requirements. However, surprisingly, it also covers Backup and Recovery strategy and Security.

It may be worth checking out the associated website, mentioned only in the introduction, so I missed it the first time I skimmed through the book - <http://www.highperformancemysql.com/>. At the time I looked, which was the first week of August, there was just a blog with four or five entries talking about the book being published - though the book itself promises scripts and future tips will be put there. (The last entry was back in April, so I have my doubts!)

There's a short chapter on basics - nothing of note here really. Then a chapter on storage engines - the main options being MyISAM, InnoDB, Heap and NDB. The engine you use for a particular table (and you can use a combination of all engines in a single database) depends on the characteristics you need supported. Only some (InnoDB and BDB) support transactions, for instance. You probably knew that already.

The next chapter covers Benchmarking - why its important, strategies and three examples (including a tool written by one of the authors). I use Winrunner and Loadrunner in commercial environments, but they're ever so expensive. The tools here are MySQL specific, and just a brief example is given for each..

Now we are getting to the main attraction. There are a few chapters on indexing, analysing and optimising queries and tuning at the server level. The indexing chapter provides a fairly basic overview - what they are and the different types - B-Trees, hashing etc. I hadn't realised MySQL now supports spatial indexes (with R-Trees), but it does. I remember doing interesting stuff with convex hulls years ago, and wonder whether these features would have helped! So in making sure you have optimal performance you need to chose the right table and right indexing schemes (obvious, really).

The server tuning chapter is a little odd. It talks about CPU, Memory, disk types, selecting a filesystem, swap etc. There were a lot of descriptions, but not a lot of recommendations with specific regard to MySQL. I found it a bit disappointing. There was also discussion about the economics of buying servers. Its all relevant I suppose, but didn't quite hit the spot for me.

Replication and load balancing. I quite liked this area - lots of diagrams and discussions about the relative merits of various schemes. I generally avoid multi-master replication because it can be a nightmare to manage and recover data consistency if things go wrong - I wouldn't recommend it unless its well controlled and well understood. The point of setting up replication in this book is not for disaster recovery (which is where I mainly use it) but to permit multiple query servers to be set up to meet a large query demand. There are explanations on how to setup replication, with example config files. Load balancing just builds on this - but although there's a lot of discussion on load balancing schemes and hardware, it doesn't go into specifics - such as using Alteon, Cisco CSS, Foundry etc. LVS is only mentioned on the last page - so rather lacking in specifics.

A strange inclusion is the chapter on backup and recovery. I can see potential in this particular book for more detailed discussion on the best backup schemes for high performance MySQL services. For instance exploiting replication to keep the backup load away from a production service. Although this point is covered in three small paragraphs, the whole chapter is far more general and quite disappointing. It covers online and offline backup schemes, dump (like Oracle exports) and raw copy backups, and taking snapshots at the filesystem or storage level.

The final chapter is on security. I just can't see why they bothered including it - they covered nothing specific to high performance. Just general MySQL accounts again, firewalls, tcp wrappers, chroot and the like.

MySQL 5.0 has Stored Procedures and 5.1 will add triggers. These features should also facilitate application performance improvements, though there was no discussion on how one might capitalise on them.

I have a mixed view on this book, but mainly disappointment. There was a lot of potential in the subject matter, but the book didn't cover it sufficiently. It was also peppered with "Jeremy did

this, and Jeremy did that...". Quite annoying ... what did Derek do apart from write about what Jeremy did at Yahoo! ? (... actually, I think the answer is in the introduction - help him finish the book by writing the last 20% for him. Not good.)

Cascading Style Sheets: the Definitive Guide

Eric A Meyer

O'Reilly and Associates

ISBN 0-596-00525-3

528 pp.

£ 28.50

reviewed by Andrew Macpherson

See the combined review below.

Eric Meyer on CSS, Mastering the Language of Web Design

Eric Meyer

New Riders

ISBN 0-735-71245-X

350 pp.

£ 34.99

reviewed by Andrew Macpherson

See the combined review below.

Designing with Web Standards

Jeffrey Zeldman

New Riders

ISBN 0-735-71201-8

456 pp.

£ 27.50

reviewed by Andrew Macpherson

About 3 months ago I was putting up Mambo OS (an excellent PHP content management system) for a client, when the client had the cheek to submit the site to the W3C conformance test.

It looked bad at first, but in fact there were about 20 distinct errors on the template and everything else was down to the somewhat interesting markup in the content that the client was responsible for. However the exercise showed me how things had moved on since I had started writing HTML. It triggered a week of reading round the W3C site's links – the site itself is fairly impenetrable.

Eric Meyer and Jeffrey Zeldman (A List Apart) were common references, and the books came highly recommended. Certainly the content in the New Riders books was thought provoking.

Let's deal with the O'Reilly technical book first. It's a different intended audience, much the same as this newsletter. Buy it. Now, before you write another bit of HTML. Enough said? It really leads you into the heart of using CSS to separate content from presentation – the alternate

CSSes for print, presentation or display, and some of the flaws in current browsers. The style is fairly light, the examples, though black and white are clear, with the contrast needed to illustrate the points being made.

All 3 books mention that most current browsers are in a poor standards compliance state. Summary: nothing with a version less than 5 even heard about standards, Internet Explorer doesn't (full stop), Mozilla is the current leader. But Zeldman later points out that IE has a non-compliant, but sane box model whereas the standards have a thoroughly insane box model which is not fixed until CSS 3 (with a mode switch). This is a big issue and is covered to some extent in all, but most thoroughly in Zeldman.

Meyer on CSS is a nice book with the colour illustrations one would expect from the subject matter, A bit wide for reading on the sofa – it is better on a desk – to accommodate a mass of marginal notes, The text fortunately is in a sufficiently narrow column to be readable, though a blacker text font would be more accessible, and this failing was highlighted by good strong headings.

This is an ideas book that takes one off onto the trail of “How could I?” and deeply stimulating ideas. Not a must have, but despite being 2 years old it has not aged, and I am glad to give it space on my bookshelf.

As Hamlet said “Methinks the lady doth protest too much.” If the only book in this trio you buy is Jeffrey Zeldman you will conclude that web standards have been deliberately screwed up by factional interests. He tries very hard to make the case for dropping tables as layout elements, using the box model which prevents proportioned layout, and putting in the CSS workarounds that enable one to work round the IE incompatibilities.

Actually he does a good job of the last point, but reading The Definitive Guide first will let you understand how some of the tricks work. Much of Zeldman's hard content is good, though spoiled greatly by lack of colour in the illustrations, the shades of grey were insufficiently distinct to communicate some of the differences he was trying to communicate.

The soft content was infuriating. Perhaps if one had met him it would be fine; perhaps. I kept going because I did want to get the technical content but I found it hard, and to some extent it has dated by trying to document the differences in 2002's browsers, and though IE does not seem to have moved on Mozilla and Opera certainly have. I can't recommend this one.

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