

news@UK

The Newsletter of the UK UNIX Users Group

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

Volume 12, Number 2

ISSN 0965-9412

June 2003

Contents

UKUUG Secretariat	3
Announcement: lecture by Tim O'Reilly	3
FreeBSD's 10th Birthday – UK Meeting	4
Linux for Business Conference - 10th June	4
Asian Enterprise Open Source Conference 2003	5
Australian Open Source Symposium	6
SCO's Big Lie	6
Letter to the Editor	7
Book review: Perl Graphics Programming	7
Book review: C Pocket Reference	8
Book review: Building Secure Servers with Linux	9
Book review: Practical TCP/IP	10
Book review: Sequence Analysis in a Nutshell	11
Book review: The Exim SMTP Mail Server: Official Guide for Release 4	13
Book review: Macintosh Troubleshooting Pocket Guide	14
Book review: sendmail Performance Tuning	14
Book review: sendmail 3rd edition	15
Book review: Programming Web Services with Perl	16
Book review: Apache: The Definitive Guide 3rd Edition	17
Book review: Webmaster in a Nutshell 3rd Edition	17
Book review: Practical C++ Programming, 2nd Edition	18
Book review: Practical UNIX and Internet Security	18
Book review: Google Hacks	19
Book review: Linux Server Hacks	20
Contacts	22

UKUUG Secretariat

Jane Morrison

Firstly, a thank you to all our book reviewers. We have some 16 book reviews in this issue and I know that many members find these very informative.

It has been another busy time for UKUUG. Since March we have been concentrating on bringing the details together for the Linux 2003 Conference (Edinburgh 31st July - 3rd August). Please see the flyer enclosed with this issue.

The speakers are now in place and 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.

On Monday 23rd June at 7 pm, Tim O'Reilly will be speaking at an evening meeting in London to be held at City University. See the announcement below.

The next event for your diaries is the UKUUG Annual General Meeting which will be held this year on Thursday 25th September at the Institute of Education, London at 6.00 p.m. Further details, agendas etc. will be sent to you automatically.

The next Winter Conference is planned for February 2004. We would welcome any suggestions for suitable venues.

If you wish to have something published in the next issue (September) please note the copy date is: 22nd August.

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

Announcement: lecture by Tim O'Reilly

UKUUG Secretariat

The Open Source Paradigm Shift: LAMP (Linux/Apache/MySQL/PHP|Perl|Python) as the 'Intel Inside' of the Next Generation of computer applications

On June 23rd, at 7pm, the UK Unix Users' Group and O'Reilly UK present Tim O'Reilly, speaking at City University in London. Full details are available at:
<http://www.ukuug.org/events/TimOReilly/>

The computer industry has gone through a sea change in the past few years. The killer applications of the web era turned out not to be PC-based software packages like the web browser, but web hosted applications like google, mapquest and amazon.com. These applications are built on top of Linux and Apache, yet they are themselves fiercely proprietary. But what would most developers do with their source code? These massive systems are valuable for their data as much as for their programs. And by opening up XML web services APIs to that data, the most innovative of these sites are creating new opportunities for hackers to "scratch their own itch." One of the greatest challenges for open source in the next few years is to understand and adapt to the paradigm shift implicit in network computing, and to shed the legacy thinking of the desktop era.

Online registration is free, but required for entry
<http://www.ukuug.org/events/TimOReilly/registration>

UKUUG is grateful to O'Reilly UK for organising Tim's visit, to the Department of Computing, at City University for providing accomodation and to David Dodson of the ACM British Chapter for local organisation.

FreeBSD's 10th Birthday – UK Meeting

Sam Smith

On 19th June 1993, FreeBSD was released on the world.

10 years (and two days) later, there will be a meeting of the FreeBSD UK User's group in Cardiff to commemorate this event.

The meeting will take place at Dempseys, opposite Cardiff Castle in the city centre on Saturday 21st June, 2003 starting at 7pm.

All FreeBSD users are welcome to attend, although please let Paul Richards know in advance (paul@freebsd.org) so there is a rough idea of numbers. However you can just show up on the day.

The first usage of the word FreeBSD was in an email to David Greenman and is now available on the FreeBSD website at
<http://www.freebsd.org/news/1993/freebsd-coined.html>

Membership of the FreeBSD UK User's group is free – see:
<http://ukuug.uk.freebsd.org>

Linux for Business Conference - 10th June

Netproject

Unlike other Linux events, 'Linux for Business' has been specifically put together for those who need to investigate the business implications of moving toward a Linux and Open Source / Free Software solution rather than offering just a review of the technical alternatives available.

Packed with case studies, given by IT Directors who are in the process of transforming their businesses, speakers include:

Dr Martin Armitage, Senior Vice President & Head of Global Infrastructure Organisation (GIO), Unilever.

Paul Martin, IT Director, Nottingham City Council.

Gareth Lloyd, IT Director, Hill House Hammond.

Paul Friday, Head of IS, West Yorkshire Police.

Also speaking will be leading experts in the field of Linux and Open Source.

Registering for the conference will provide you with:

- unlimited access to all conference sessions;
- unlimited access to the 'Linux for Business' exhibition;
- full conference documentation including copies of all speaker notes;
- refreshments, lunch and evening reception.

Limited exhibition space is also still available.

For further information or to register please visit:

<http://www.netproject.com/conference/lfb>

or email: tracey@netproject.com by return and register for only £100 + VAT per delegate (places subject to availability)

Asian Enterprise Open Source Conference 2003

First call for papers

October 29th-31st, 2003 Singapore

This annual conference and expo aims to draw world-renowned experts in the exciting field of Open Source and showcase the state-of-the-art of Open Source Computing in the IT hub of South-East Asia, Singapore. AEOSC started out as the Singapore Linux Conference in 1999 and was run as that in 2000 and 2001. This year, with the global recognition of Linux and Open Source in general, the Singapore Linux Conference has been reinvented as the Asian Enterprise Open Source Conference.

In this fourth edition, the conference will focus on the growth of Open Source Software as a broadly available, trusted, secure and value for money platform.

Conference

Papers are invited for presentation during the Conference, which will be held on Wednesday 29th, Thursday, 30th, and Friday, 31st October 2003. Following the event success in the past, there will be two tracks, "Business" and "Technical". The topics below are meant to indicate the possible variety of areas and should not be regarded as exhaustive.

Agenda

Linux, *BSD, Mac OS X

Apache, Perl, PHP, Python, MySQL, JBoss

Policy issues from a adoption as well as a governmental/regional perspective

Educational issues: SchoolForge and the OpenSchools effort

Hosting of the Asia Open Source Symposium follow-on event

Linux Users' Group (Singapore) Awards

Exposition/Exhibition

Prospective authors are requested to submit full papers for review. Electronic submissions are strongly encouraged through email. Hard-copy submissions are also allowed, and three copies of the manuscript should be sent to the address below.

The closing date for submissions of full papers is August 1st, 2003

Papers will be rigorously reviewed by the Technical Program Committee, and reviewers' comments will be relayed to the authors on request in the interest of transparency.

Tutorials

Tutorial sessions will be held on Wednesday 29th October. Tutorials have a strong technical bias, and are intended for a more mature audience, in terms of Computing and Programming exposure. We will also consider tutorials that are in the introductory in nature as well.

Prospective tutorial speakers are invited to submit proposals to the Conference Secretariat. Each proposal should include a summary, a course outline and a brief biography of the speaker.

Proposals complete with the required documents should be sent to the address below:

Mr. Mohsenruddin Moonshi

TechWorx Solutions Pte Ltd

9 Temasek Boulevard
31-02 Suntec Tower Two
Singapore 038989
Tel: +65 9745 2310
Fax: +65 6356 7045
mmoonshi@techworx.net

Australian Open Source Symposium

UKUUG Secretariat

We have also received notification of the fifth Australian Open Source Symposium (AOSS 5) which will be held in Brisbane on Wednesday July 16th, 2003.

Abstracts of proposed submissions are due by June 11th: more details are available here:
<http://www.auug.org.au>

SCO's Big Lie

Bruce Perens

This article contains Bruce Perens' reaction to Novell's statement made on 28th May on the SCO case and is reprinted from:

<http://www.perens.com/Articles/SCO/BigLie.html>

We knew that SCO's attack on Linux was a lie. But we never dreamed of the big lie behind it.

This morning, Novell announced some of the terms of the company's 1995 agreement to sell its Unix business to SCO. The shocking news is that Novell did not sell the Unix intellectual property to SCO. Instead, they sold SCO a license to develop, sell, and sub-license Unix. The title to Unix copyrights and patents remains with Novell. To back up this assertion, Novell refers to public records at the Library of Congress Copyright Office and the U.S. Patent Office.

In their announcement, Novell refers to recent letters from SCO asking Novell to assign the Unix copyrights to SCO. So, apparently SCO's management team knew that they did not own Unix while pursuing their sham campaign against Linux.

Along with this revelation, Novell is reiterating its support of the Linux and Open Source developer community, and its status as a partner in that community. Novell rejects SCO's accusations of plagiarism. Novell management says they do not intend to stand in the way of the development of the Linux kernel, its companion GNU system, and other Free Software.

It would be an understatement to say that this leaves SCO in a bad position. The company has loudly and repeatedly asserted that they were the owner of the Unix intellectual property, all of the way back to AT&T's original development of the system 30 years ago. They've lied to their stockholders, their customers and partners, the 1500 companies that they threatened, the press, and the public. Their untruthful campaign caused the loss of sales and jobs, and damaged Linux companies and developers in a myriad of ways. And now, SCO will be the lawsuit target.

Microsoft executives also have egg on their faces. The company self-servingly rushed to buy an SCO license one business day after the threat letter, bringing a senior attorney to the office on a Sunday to tell the press how much Microsoft "values intellectual property." Microsoft's management could have taken the time to analyze SCO's claims, if the company had wanted this license for practical and technical reasons. Their decision to buy when they did must have been

motivated by a desire to add to SCO's fear campaign. Of course they'll grab any opportunity to spread fear about Linux, but this time Microsoft bought a pig in a poke.

SCO management, if they insist on standing in the way of a train, could still claim that software they developed in the years since 1995 is being infringed by the Open Source developers. That claim, always a dubious one, will be difficult to take seriously now that their prevarication throughout this campaign has come to light. SCO would be well advised to drop their suit against IBM in exchange for IBM's agreement not to counter-sue. But IBM might not feel that charitable toward SCO.

In contrast to SCO, Novell's made a friend among the Free Software developers. We're always happy to see people using our software. But a real partnership between an IT vendor and our community is an equal partnership, with the company donating services and new software in exchange for the value it receives. Novell has already placed important software under Open Source licenses. Today, the company has done us a tremendous service, by stomping upon an obnoxious parasite.

Bruce Perens is a director of Software in the Public Interest, Inc., an Open Source development organization. He operates an independent consultancy and is a senior research scientist for Open Source at George Washington University's Cyber Security Policy Research Institute.

Letter to the Editor

Ramanan Selvaratnam

Thanks again for the GNU Win II CD in the latest newsletter. As with the Knoppix CD distributed at Bristol this has already proved to be very helpful.

I had overlooked the usefulness of GNUWin as super advocacy material.

The two uses today was when provided mail client solutions two of my business clients through Mozilla. They are yet to discover the power of free software fully :-)

Just a suggestion – if you are going to develop this CD distribution... why not highlight that Mozilla mail is a good mail client. Most people seem stuck with the Outlook express “mumbo jumbo” and never get around to downloading alternatives and suffer for eternity.

Maybe Mutt on CygWin might be a good idea too to be included especially for those who want a uniform mail clients in all the different environments one is sometimes forced to work on. I will suggest this to the [EPFL.CH] people.

Best regards

Ram

Perl Graphics Programming

Shawn Wallace

O'Reilly and Associates

ISBN 0-596-00219-X

462 pp.

£ 28.50

reviewed by John Collins

This book covers all the graphic and animation formats used today and explains how to generate, manipulate and convert between them using Perl modules.

After defining and explaining the various types of graphic file formats, the book spends a chapter on the Perl modules available for working with the various formats: `GD`, `Image::Magick` and `GD::Graph`. The Gimp is dealt with in Part 1 on Raster-type images, SVG, SWF (Flash) and Ming in Part 2 on Vector-type images and finally in Part 3, Documents and Printing are discussed with discussion of PostScript and PDF Modules. An extensive set of appendices cover how to provide graphics on web pages, how to use the Gimp, `Image::Magick` file formats and ActionScript (for Flash).

I have used `GD` and `GD::Graph` quite a lot and hit a few problems early on in the area of background colours and transparency which I would have avoided had I had this book. It gives a good number of worked examples which illustrate concepts well. I noticed this particularly in `GD::Graph` where it isn't at all clear from the "man pages" where the various titles and legends go and how you can adjust their style and positions. The section on `Image::Magick` with which I'm much less familiar, seemed equally well-written as did the chapter on Perl plug-ins to the Gimp.

The chapters on SVG are equally helpful with copious interesting examples in areas I haven't worked with to date. I tried a few and they all worked fine. I just wished there was enough time to really explore all the interesting features of some of these packages as with SWF and Ming in the next sections.

The PostScript and PDF chapters likewise have a wealth of useful information and examples of techniques for solving real-life problems.

At various places there are "boxes" explaining some of the more obscure features of the format or the Perl pack and unpack operators, for example, in the SWF chapter (as this is required), and tables of functions and operators. Most chapters end with references for further reading, web links and alternative Perl packages to look at in the same area as that which the chapter covers.

In summary, I think this is a very well-written and helpful book. I learned quite a lot from it quite rapidly. Someone with a graphics problem to be solved (and quite a few other problems as well) will soon be put onto the right track with the aid of this book as to what packages to try and what lines to follow. I am sure I will refer to it frequently in the future.

John Collins – Xi Software Ltd – <http://www.xisl.com/>

C Pocket Reference

Peter Prinz and Ulla Kirch-Prinz

O'Reilly and Associates

ISBN 0-596-00436-2

134 pp.

£ 9.95

reviewed by John Collins

This book is intended as a pocket reference for C programmers.

The first 72 pages of the book go through every aspect of the C language and the remainder covers library functions.

I have to say that I was very disappointed with this book. It seems to me a mistake to discuss the finer points of C syntax in a pocket reference. No one is going to learn C from a book like this and the old hands aren't going to want to thumb through 72 pages to refresh their memories on the layout of IEEE floating-point (which you are not guaranteed to have in all implementations of C of course) or the layout of bit-fields. I should have thought a much snappier summary of syntax would have been appropriate leaving it for the larger works to go into all the details. This is supposed to be a "pocket reference" not a "suitcase reference".

I was not very happy about the treatment of library functions. I would like a pocket reference to present a clear table, with clear listing of arguments, return values and error indications. If you need more detail, you should go elsewhere. I should also like to see comments and warning messages about compatibility between various implementations. Instead each function is presented, some in completely unexpected places, with a about 10 lines of verbiage to describe the action and arguments. In places system calls get muddled with library routines. Process management and signals get put in a peremptory page or two at the end. Most of the definitions are Microsoft based. I noticed that in another place `mktime` is shown in a diagram but it isn't defined anywhere.

Particularly unforgivable, in my view, is that no mention is made of security problems with functions, for example `tmpnam` is presented without mentioning that it is not generally thread-safe and better alternatives are often available. Most alarming of all is that `gets`, the scourge of security everywhere is presented without a health warning. You can wreak havoc with `scanf` and `sprintf` too, but no mention is made of those dangers. `snprintf`, the safer alternative to `sprintf`, is mentioned only briefly in the blurb about `sprintf` and `sscanf` which appear in the section on string handling. When mentioning the format arguments, the reference is just to "earlier in this book" no page number is given.

Producing a good pocket reference is a difficult task to get right. However these people have not done it in my view. It is illogically set out, incomplete and hard to find your way around. It is most disappointing as so many of the pocket references O'Reilly produce are excellent.

John Collins – Xi Software Ltd – <http://www.xisl.com>

Building Secure Servers with Linux

Michael D Bauer

O'Reilly and Associates

ISBN 0-596-00217-3

448 pp.

£ 31.95

reviewed by Mike Smith

The books I got to review for this newsletter have turned out to be rather a good bunch. Interestingly, this one has some similarities with *Linux Server Hacks* book which you should find reviewed nearby. Subjects like SSH, securing mail and DNS services, tunneling, and Tripwire are covered in both books. As you would expect, the other book has some brief tips on these subjects whilst this one covers them in more detail - a sort of half-way house. Of course much more detail on some of these things is covered in their own O'Reilly titles (*DNS*, *Sendmail* etc) - but the emphasis here is on security, naturally. Hang on, I'm getting a little ahead of myself here - lets get back to the beginning.

We start off with a discussion on modelling threats and performing risk assessments. This is quite interesting - we're not jumping straight into configuration parameters, but starting from the beginning - and that's a very good place to start too. Thankfully, its not a long chapter, but it is important to understand the context of security threats when building not only servers, but designing solution architectures.

And that's precisely where we go next - network design. I (like many readers I suspect) have been used to multi-layer, multi-security-zone Internet designs for years now, but when you're new to them it may not be immediately apparent why we do things like we do. I remember my management saying that you'd never design such a complex mess (as they usually turn out to be) ... but actually there are reasons for it. It is unfortunate, as we do end up with more complex

routing designs, firewall rulesets and higher management overheads - but that's just the way it has to be.

Right, so we've covered firewalls as part of that - quite a bit on iptables (we're on Linux, of course) but there is some detail on CheckPoint and other commercial solutions too. Then there's some info on hardening Linux servers. Good advice all round.

A little test: what do you do once you've set up a secure environment then?

Test it of course. In this text we look at nmap and Nessus. (I remember the days when I was playing with SATAN - we haven't really come on very far, have we.) I don't think its mentioned, but when you actually have to prove you've setup a secure environment, you do generally have to go to an external approved organisation. For Government, you need CLAS certified consultants to do this from the likes of Portcullis or Insight (see the CESG website for details). Security in this space is very regimented - has to be - so writing official documentation such as an "ADS" is essential too.

Next up - remote management. SSH. 'nough said. (Well, there is more, but you know it already.)

Finally, the book covers securing services like DNS and Mail, as I started out by saying. For DNS, it looks at running it in a chroot environment. It also looks at the main alternative - djbdns, but nothing on nsd.

For mail, there's quite a discussion on whether to Sendmail or not! This guy is a Postfix fan, but does also mention qmail and Exim. He obviously doesn't want to upset anyone, and recognises that sendmail is the most prevalent MTA, so this is covered in most detail.

Next up is Apache, ftp (ProFTP, of course) and we're done on services. Finally some aspects of monitoring and system management are covered. In terms of monitoring, we're just looking at logs really. Then there's a bit on IDS - Tripwire and Snort. This is perhaps a bit light. I've been looking at tools in this area and there's some good stuff about (not necessarily Linux based though, and commercial too.) If you're interested, have a look at the products from Mazu and Netforensics, for example.

Stuff missing? Hmm. There's no mention of threats in protocols such as BGP, OSPF etc ... though to be fair you don't usually run BGP on Linux boxes in production environments (I think). Didn't notice anything about UML - a big step up from chroot. Or LIDS. Or the "secure" distributions (SELinux etc).

Anyway, its another thumbs up from me. Perhaps I just like reading this stuff. Sad man.

Practical TCP/IP

Niall Mansfield

Addison-Wesley

ISBN 0-201-75078-3

864 pp.

£ 38.99

reviewed by Mike Smith

Well, I suppose I must have asked to review this book, but it isn't an O'Reilly title - its an Addison-Wesley one. The first similarity is that there is a picture of an animal on the cover - albeit very pixilated. I think its a racoon.

Its a bit of a tome - 27 chapters in 4 parts, and nearly 900 pages. I got a little concerned when I read the introduction and it started talking about NT4. Yes NT4, not Win2K or 2003 or XP. An explanation was given - that NT4 is the common denominator, but I don't buy it. Anyway the reason its so big, I suspect, is that it covers both Windows and Unix. I'll retract that comment later if I discover another cause after reading a few chapters.

It has a chapter 0, and that gets a big thumbs up from me - I like counting from zero. This covers some TCP/IP basics - though introducing the novice to tcpdump on page one is an interesting approach. Chapter one goes more into tcpdump, and chapter two covers arp, ping, netmasks and suchlike.

Jumping on a bit, there's routing, DNS, the application layer - with telnet, mail etc and then back to diagnostics with ethereal and ngrep.

As we're covering windows as well as Linux, Netbios and WINS are covered. And SAMBA actually.

Back to that point about the book being so large, there are other reasons:

0: There are a lot of examples

1: There are many, reasonably good, diagrams too

2: There are labs at the end of many chapters, in which you are asked to do certain things (like setup a small network with Windows and Linux servers on)

Other things covered: Firewalls, dial-up, PPP, NAT - even Cisco's VRRP protocol. There's a whole hotch-potch. IPsec, VPNs, the list goes on. IPv6 in the appendices, for example. (And RJ45 pinouts and Null modem wiring diagrams!) So its another one of those books which tries to cover everything - and I'm not sure on the wisdom of covering everything Linux and everything Windows together - its too much these days. Still, a lot of good information in there.

I hadn't realised there had been a practical implementation of rfc1149! See <http://www.blug.linux.no/rfc1149/>

Fantastic.

Yesterday my daughter asked me if I had an old thick book - so that she could cut the middle of the centre pages out and hide a camera in it (must have been something on TV). I had half a mind to give her this one - its not an O'Reilly after all, but actually I think I will hang on to it. Its okay: covers lots of area, has the odd good bit in it, and has lots and lots of references to other materials. I had a look on the web to see what other stuff he's written - 10 years ago he wrote "The Joy of X". Oh dear.

Oh yeah, did I mention X, LDAP, NFS, FTP, netcat, NTP, Shares, PDC, Elections, MIBs, HTTP, UCE, Virus Scanning, ICMP ... You get the picture.

Sequence Analysis in a Nutshell

Scott Markel and Darryl Léon

O'Reilly and Associates

ISBN 0-596-00494-X

£ 20.95

reviewed by Damian Counsell

I spent a brief (and unpleasant) time working for an online scientific publishing company. Back then I liked that unoriginal rule of thumb: "Any book entered via its index is better implemented electronically". That "rule" often comes to mind when I pick up volumes in O'Reilly's "Nutshell" series. I try to avoid printing out documents out whenever I can, and, most of the time, I solve my computer problems more quickly by Googling than by reading, but I still find the dead tree format pleasing. (Of course, most O'Reilly publications, including this one, are available online via the Safari service

<http://proquest.safaribooksonline.com/> so this point is becoming academic.)

Of all the currently fashionable topics in scientific, technical and medical publishing - both online and on dead tree - bioinformatics is among the hippest. Roughly, the term "bioinformatics"

describes any use of computers to handle biological information. In practice, most people use the term to mean “computational molecular biology” – the use of computers to characterize the molecular components of living things. Biological molecules are generally polymers; ordered chains of simpler molecular modules called monomers. Think of the monomers as beads or building blocks which, despite having different colours and shapes, all have the same thickness and the same way of connecting to one another.

These beads may make pretty necklaces (as any number of 3D graphics of molecular helices suggest), but some of the popular shine has begun to come off the new cell and molecular biological revolution: PPL Therapeutics, the company that cloned Dolly the Sheep, has decided not to build a £42m factory to manufacture drugs based on this technology and many pharma and biotech firms are shrinking their bioinformatics departments. Despite this, O’Reilly has, in the past few years, gone from rejecting bioinformatics book proposals for the want of a buyers, to finding bioinformatics books amongst its bestsellers. Its most recent bioinformatics-related title is “Sequence Analysis in a Nutshell”.

Chains of DNA or protein monomers can be treated computationally as letters of an alphabet, put together in pre-programmed arrangements to carry messages or do work in a living cell. Since the “completion” of the Human Genome Project, the mission to read the order of all the monomers in human DNA, we have a great many more of sequences to analyze. There are, unsurprisingly, a whole range of techniques for interpreting these “biological stories”. Most of these techniques come under the heading of sequence analysis. Most of the analysis is done on UNIX boxes.

Although there is still a search for standards in the relatively young science (engineering discipline?) of bioinformatics, certain data formats and collections of analysis tools have become more widely used than others in the area. The authors of “Sequence Analysis in a Nutshell”, Scott Markel and Darryl Léon, have made some shrewd choices about which of these to cover in the 300 pages or so of this handbook. This is not surprising, given that they are both experienced PhD bioinformaticians. Of the data formats, they describe FASTA, GenBank/EMBL/DDBJ, SwissProt, Pfam and PROSITE. Of the software tools, they deal with more specialized packages like Readseq, BLAST, BLAT, ClustalW, HMMER, and MEME/MAST, plus the ultimate Swiss Army Knife (or perhaps that should be “Swiss Army”) that is EMBOSS. The book also includes appendices containing various tables of information useful to practising bioinformaticians.

So what lifts this book above the level of Google searches? Firstly, the authors have done the hard work of gathering surprisingly scattered chunks of information together in one mass – a neat, glossy mass which should fit easily on a shelf near your desk. Secondly, their work is packaged and produced to the usual high O’Reilly standard of typesetting and layout: the text is clear, consistent and tasteful (with a striking cover image of a liger). Thirdly, by the simple act of making an informed selection, Markel and Léon, have served the field by more clearly defining the de facto standard bioinformatics standards and systems.

This reference is sensibly aimed at the generalist, possibly in a commercial, administrative or service bioinformatics role who just needs to get things done. “The liger book” would also be especially useful to relatively inexperienced bioinformaticians or ones only superficially familiar with the tools it covers, for example, students tackling a research project. Both groups in particular would find it a handy “meta tool” to help themselves and help others.

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 and personal pages are at <http://www.hgmp.mrc.ac.uk/> and <http://www.counsell.com/> respectively. These Websites are not in any way related.

The Exim SMTP Mail Server: Official Guide for Release 4

Philip Hazel

UIT Cambridge

ISBN 0-954-45290-9

621 pp.

£ 37.00

reviewed by Oliver Gorwits

Exim is a mail transfer agent that can be run as an alternative to Sendmail on most Unix and Unix-like systems. At my organisation we use it to relay around half a million messages per day, although it's suitable for many other types of installation including those with local delivery, and far larger (or smaller) ISPs.

A bit of history, first. Exim is currently in its fourth version, and is developed by Philip Hazel at the University of Cambridge Computing Service. The third release was accompanied by an O'Reilly book, also written by Philip, but there were enough fundamental differences that this release warranted its own volume. And what a book: more than 600 pages straight from the horse's mouth (as it were); you can't go wrong.

Philip begins with five chapters that introduce the reader to Internet mail, Exim, and some rudimentary runtime configurations. There's nothing to fear here, as the text is beautifully self-contained, covering topics from the DNS to routing lookups. As Exim's runtime configuration is both flexible and easy to read, the quite technical examples given early on can be understood without flicking to and from other chapters in the book. The next four chapters cover in a rather succinct manner the parts of Exim that route and transport your messages. By this point you should have a grasp of the philosophy and design of Exim, which allows Philip just to give you the details. This section does feel most like a reference manual but I'm not sure there's another way he could present the information without confusing the reader. The remainder of the book covers each of the Big Features of Exim, one per chapter. I'm guessing that Philip just kept on writing until he ran out of features, rather than time or space! These chapters feel far more like the heart of the book, and the author treads a fine line between thorough process description and distracting technicalities. The two appendices cover regular expression syntax and special variables (both being available to Exim's configuration).

The book would be ideal if, for example, you manage a mail system on your own and don't have a great deal more admin experience close at hand. Its great strength is the vast number of scenarios that Philip has thought up; it seems that if you can think of something that you want the application to do, it'll be in there somewhere. At my site however we do have a good number of people who are familiar with Exim, so armed with a copy of the (equally well written) reference manual we can usually get along just fine.

Those expecting the chatty, irreverent style of an O'Reilly text may be in for a disappointment. Philip writes in a clear, precise manner, and obviously knows the subject matter (literally) inside-out; but there's no messing around and you have to be committed to learning about the subject in question. Having said that, I don't want these last two paragraphs to put you off. If there's even a whiff of a chance of you having to come into contact with Exim or its runtime configuration, then I can do nothing else but strongly recommend this book. The detail's there in spades, it reads very well, and is a fine complement to the reference manual.

<http://www.uit.co.uk/exim-book>

<http://www.exim.org>

Oliver Gorwits is a member of the Network Software Group at the University of Oxford's Computing Services.

Macintosh Troubleshooting Pocket Guide

David Lerner and Aaron Freimark

O'Reilly and Associates

ISBN 0-596-00443-5

80 pp.

£ 8.95

reviewed by Raza Rizvi

I have never been a fan of IT pocket books. These small guidebooks have their place when you are somewhere far and distant and need the advice of someone on what to see and what to do, but in a business, your needs are not set by something you read in a book.

Having said that, I did find this one a rich mix of the weird ('How do I clean my screen') and the obvious. Some of the facts have been technical support favourites since the early days of the Macintosh SE (if not before). Overall, perhaps the greatest use of the book is as a crash course to MacOS 9 and OS X for diehard MacOS 7 and 8 users.

Ultimately, like all FAQ based books, it suffers from the fact that if you carry it round with you, the questions are never quite what you see or you fail to interpret them as they are laid out in the book, and of course once you have read it there is nothing else you can do with it, as it is not a reference book.

If you are new to the Macintosh and travel on the train, then the small format and the small price is worth it for that nugget of information that gets you out of a jam. If you have been using Macs for a while, then there is no need, especially as the FAQ used as the basis for the book is available online at the Tekserve website.

Raza Rizvi is Technical Manager at REDNET, a business ISP and Cisco Premier Partner, based in High Wycombe. He has always liked the Macintosh.

sendmail Performance Tuning

Nick Christenson

Addison Wesley

ISBN 0-321-11570-8

272 pp.

£ 21.72

reviewed by Raza Rizvi

'This book is great.' is what the quote from Eric Allman says on the front cover.

This quote from the creator of sendmail is truly and absolutely spot on. The book is a superb guide to the entire area of tuning mail platforms based on sendmail for increased performance. And when I say the entire area, I mean it. Not only are the options within sendmail covered in a fully descriptive manner, but there is proper documented evidence of the performance achievements in benchmarked scenarios.

Every aspect of the hardware is evaluated and the role it plays in mail throughput clearly examined – from network latency to filesystems to RAID to memory.

After a general introduction to what is meant by tuning of mail systems and to the background of sendmail implementations, Nick goes straight into email relays, those boxes that act as the message handlers passing mail to and from the initial and the final mail systems. There is a very readable and logical flow through what is done and how one might act upon different parts of the system to ensure maximal throughput.

The same methodology is then applied in turn to email receipt and email sending (though whether we should thank him for giving the spammers more throughput is a topic in it's own right!).

After discussing security and architectural aspects of sendmail, the final chapters show how system tools/logging and testbeds can be used by the readers to improve their own scenarios, and to provide baselines on which performance changes can be realistically measured.

Although the major focus of the book is sendmail itself (primarily on Solaris or LINUX), information is given on the user interfaces to mail - POP3 and IMAP4.

There are few good books on performance tuning (and Nick lists all of them in his bibliography) and although this is specific to mail, it includes information that is useful to all senior systems administrators. For those of you with the responsibility of running sendmail systems, this book is as much part of your essential library as the 'Bat' sendmail reference book - good advice falls from every page of this book.

To quote from part of the section dealing with queues:

'At first glance, even more complex strategies seem appealing. However, further consideration will reveal that moving too far down this path leads to madness.'

This book stops such madness from eating away at sysadmins. I fully recommend it as preventative medicine.

sendmail 3rd edition

Bryan Costales with Eric Allman
O'Reilly and Associates

ISBN 1-56592-839-3

1232 pp.

£ 42.50

reviewed by Raza Rizvi

The 'Bat' book. Probably the largest, thickest, most feared O'Reilly book and growing bigger each edition (this one is a full 400+ pages bigger than the 1st edition).

Such a large book paradoxically requires only a short review since what it does is bring the story up to date. This book covers the building, installing and administration of sendmail up to the current 8.12 release (even though the sendmail.org site shows the 2nd edition book as the latest!).

All the details one could want about sendmail are here in reference format: every option, every parameter, every error message, every minutiae of detail, broken down in the now familiar four parts.

Part 1 deals with Build and Install. This is the where and what of sendmail. The breakdown of the constituent parts that come with the distribution.

Part 2 deals with Administration. This is the how of sendmail. The day to day operations of security, DNS, SPAM, queues, and aliases.

Part 3 deals with The Configuration File. This is the why of sendmail. Why it does the things it does, and how to manipulate the voodoo to suit your whims.

Part 4 is the appendices, including a very useful summary of the changes since 8.8 (as documented in sendmail 2nd edition).

The index is as good as ever, yet still not quite good enough, and with a book of this sort it is a make or break factor. You won't read this for bedtime, you turn to it when there is real trouble

and you need to find information fast. I would like to see the author put himself in the shoes of the panicked mail administrator who needs to lookup things like 'running the queue'?

This is not a book for the faint hearted. You know if you need it and if you need it, you have to have it, there is no substitute, only this will do.

Programming Web Services with Perl

Randy Ray and Pavel Kulchenko

O'Reilly and Associates

ISBN 0-596-00206-8

486 pp.

£ 28.50

reviewed by Raza Rizvi

This book covers that area of integration brought about by the increasing use of web servers as the visible front entrance of knowledge repositories. In order to build upon them, or indeed to steal from them, methods are needed to enable access to that information by other servers. Those methods use XML (Extensible Markup Language) over HTTP using SOAP (Simple Object Access Protocol).

The introductory chapters provide a kickstart to XML and to XML-RPC, which is an offshoot from the early SOAP specifications, and is considered by many to be a less complex means to reaching the same goals as SOAP.

A short but example rich chapter on programming with XML-RPC is included to show practical examples, even though in at least one of the cases the authors do make a side reference to RSS as an alternative means to collect live 'news' information from a target web site.

Now in the middle of the book we go into four chapters on SOAP, starting with an introduction that also serves to lead on to the chapters that deal with the transport mechanisms that are open to SOAP programmers - namely, HTTP, and non-HTTP (such as SMTP and the open instant messaging platform Jabber, although others are included). Programming SOAP concentrates on the two available PERL modules (the eponymously named `soap` and `soap::lite`)

The last third of the book covers related XML use within Web Services through WSDL (Web Service Definition Language) and the discovery of suitable services that make themselves available for use through UDDI (Universal Description, Discovery and Integration). There is also a chapter on REST, a practical design philosophy for web services that in some ways is counter to the approach taken by SOAP. Here the service is separated from the implementation and the end object is processed as desired by the author rather than the publisher.

You need to be a good Perl programmer in order to make best use of this text. Additionally although the authors go to great lengths to provide an introduction to XML, it certainly helps to have had a reasonable exposure to it before starting, if only so that you don't become bogged down in the middle chapters, since SOAP is basically a mechanism for using complex XML documents.

There are plenty of examples and code snippets given, both within the text and in the appendices, though as the authors expect you to be Perl literate, the comments focus on the XML rather than the Perl.

If this all spurs you to do greater things, one of the authors has also published a title called 'Programming Web Services with Soap' (also from O'Reilly).

Apache: The Definitive Guide 3rd Edition

Ben Laurie and Peter Laurie

O'Reilly and Associates

ISBN 0-956-00203-3

588 pp.

£ 28.50

reviewed by Lindsay Marshall

Please take as read all the usual good stuff about O'Reilly books that I have written so often before: this book has the same production values etc. etc. and I'm sure you don't want to read that all again. So what is there to say? Well not a lot really which is why you are getting all this filler. The name of the book tells it all - this really is the definitive guide to using Apache. It isn't entirely up to date of course as the current version was 1.3.26 when it came out and it is now 1.3.27. Not a big deal. Version 2 has of course gone up a few numbers more, but, as the authors point out, hardly anybody needs or uses version 2 yet.

So what's covered? Everything from building the source to using the APIs to write modules. And, unlike a lot of "definitive" guides, it is not just a rehash of the Apache documentation, in most cases there is genuinely useful additional content. This may of course be because the Apache documentation is terrible, but whatever the reason, you get do get some added value here for a change. I went looking for some information about authentication and found what might be the answer (I need to do some experiments naturally) almost at once and I have previously spent ages with the Apache provided documents without finding the slightest clue.

The material on Apache 2 is also pretty good, though it does have a slightly rudimentary feel too it - the authors' hearts are really not in it. They can see the good stuff coming down the line but realise that it isn't worth getting all excited about yet. However, having had a go at setting up a version 2 system I can appreciate the material that is here - the stuff about setting up SSL is extremely useful and reassuring, as the process is entirely different from setting it up for a 1.3 server. (No, I am not still running 2, I went back to 1.3.27, mainly because PHP doesn't play nicely with 2 yet, and I bet that that is how most Apache webmasters feel. When it does, there will be a lot more installations done.)

The only downside of the book is that for most people it is too definitive. There are whole chunks that you will never look at - I for one will never (except under severe duress) look at the chapters on using Perl or Java. There is also no coverage of some of the third party modules that exist, mod_kerberos for example or mod_dav, only a minor quibble though.

You run apache? You need to understand it better? (The chapter on running a big web server is sound.) Get the book. The person who pays your wages can certainly afford it.

Webmaster in a Nutshell 3rd Edition

Stephen Spainhour and Robert Eckstein

O'Reilly and Associates

ISBN 0-596-00357-9

576 pp.

£ 24.95

reviewed by Lindsay Marshall

Hmm, this book feels mean. The print is small and set tightly, the book is the typical nutshell size - smaller than the other O'Reilly books. The content is mean too - CSS is 11 pages? Come off it. There is material on http, HTML, CSS, JavaScript, XML, CGI, Perl (boo), PHP (hurrah) and Apache! But not really enough about any of them to be useful to someone who has to use

them seriously. Oh, and there is nothing at all about Web Services - the authors think that that is not something webmasters need to worry about.

Most books in the nutshell series offer condensed information for people in a hurry that will get them up to speed quickly, or just act as a quick reference guide. The trouble with this book is that the information has been condensed so much that all the nutrition has been boiled away. O'Reilly have excellent books on all the topics listed above, for many of them in fact, the best books available. Much better to spend a little more money and get the full story on the topics you need than to get led a little way and then be abandoned just as you see the eyes of the wild animals glinting amongst the trees. Go first class with proper guides, not tourist.

Practical C++ Programming, 2nd Edition

Steve Oualline

O'Reilly and Associates

ISBN 0-596-00419-2

574 pp.

£ 28.50

reviewed by Peter Waller

When I picked up Practical C++ I was very sceptical, I've read numerous books aimed at beginner/intermediate readers and very few of them left any lasting impression on me. This book however, did! It's an excellent book, and it feels like it's written by a programmer and not an academic as most tend to. Not only are you going to learn about C++ in the easiest way possible, you're going to learn a lot of tips from someone who's been developing in C++ a long time.

This is not a Windows-oriented book, although there is a mention of Visual C++ in chapter 1, it focuses a lot more on the gcc tools, and I found the chapter on debugging in gdb invaluable. There is also some material on optimisation, when and how to apply it. There are questions at the end of the chapters (with answers) which are always fun to test your knowledge. The chapter on Linked Lists was also well written and demonstrated, a lot of people often find Linked Lists difficult, and it's well backed up with an excellent chapter on pointers. However, I thought the class on fixed-point was overkill, on FPUless machines I tend to use macros for simplicity and speed. There is also a lot of mention of C and sometimes you get the feeling that some knowledge of C is required. But I can't say enough good things about this book, it covers a lot more than C++, and it has a whole chapter on program design. This book is all you need to get started with C++. The title says Practical C++ and I'd say it's very practical. Highly recommended.

Practical UNIX and Internet Security

Simon Garfinkel, Gene Spafford and Alan Schwartz

O'Reilly and Associates

ISBN 0-596-00323-4

984 pp.

£ 38.95

reviewed by Sarah Loyd

There are some books which deserve to be in the libraries of everyone who works with UNIX or computer networks, the first and second edition of "Practical UNIX and Internet Security" by Simon Garfinkel and Gene Spafford are such books.

February saw the release of the third edition of this book and the addition of a third author Alan Schwartz. Maintaining the same size as the previous edition (just under 1000 pages) it's an imposing book to dive into.

The size of the book shouldn't put you off though; the book is packed with well written and accessible information for everyone from the total beginner to the expert.

It contains sections on fundamental security questions, "UNIX history and lineage", "Policies and Guidelines", "User Passwords and Authentication", "Users Groups and the Superuser", "Filesystems and Security", "Cryptography Basics", "Physical Security for Servers", "Personnel Security", "Modems and Dialup Security", "TCP/IP Networks", "Securing TCP and UDP Services", "Sun RPC", "Network-Based Authentication Systems", "Network Filesystems", "Secure Programming Techniques", "Keeping up to date", "Backups", "Defending Accounts", "Integrity Management", "Auditing, Logging and Forensics", "Discovering a Break-in", "Protecting against Programmed Threats", "Denial of Service Attacks and Solutions", "Computer Crime" and "Who do you trust".

Each section provides comprehensive guidance and solutions in the subject area. Plenty of links and additional reading suggestions are provided for the person who wants to delve deeper or expand their knowledge beyond what is provided.

The section on logging is excellent and nicely explains the occasional syslog entry sysadmins see on the lines of "Captain there are Klingons on the starboard bow".

Is the book perfect? No, I would have liked to see a chapter on Intrusion Detection Systems (IDS) such as "Snort" and lighter weight options such as "Portentry". A chapter on VPNs such as "FreeSwan" and "PPTPD" would also be a valuable addition.

If you don't have a copy of the book I would strongly recommend getting one. If you have one of the earlier editions there is sufficient new and revised material to justify getting the new version for your "better half". A well read copy should stand proudly on the bookshelf of every Systems Admin, UNIX Geek and network manager.

If my house caught fire and I had only time to save one computer security book, this would be the one I would carry out of the flames.

Sarah Loyd is a Senior Technical Consultant (Security Practice) at LogicaCMG.

Google Hacks

Tara Calishain and Rael Dornfest

O'Reilly and Associates

ISBN 0-596-00447-8

352 pp.

£ 17.50

reviewed by Mike Smith

Maybe I've been working too hard recently and ignoring the rest of the world, but the O'Reilly Hacks series is a new one on me. The format is simply to list 100 tips on the subject matter at hand - some are simple one-liners and others more complex (perhaps a bit of code, or an explanation or something). I really like this style. The tips are categorised into sections and its easy to dip into when you have a spare moment. Both of these books are therefore excellent. You can probably pick all of this information up from the web (if you knew what to look for) but its nice to have it altogether in one place.

There is a list of all the hacks in these books (but not the details) at:
<http://hacks.oreilly.com/>

This gives tips all the way from specifying search terms in the Google search form, to writing apps to use the Google API. Even the early ones I found interesting. In the bad old days when using AltaVista I used to use the search modifiers (such as putting a minus sign to remove a word). Google has similar and more powerful functions that I hadn't realised (though I have done the odd `link:www.long-clawson.com` in the past to see who's linking to the village website I constructed). So its pretty good value all the way from the start.

Some of the more interesting hacks are in the chapters about the API. There are two types of activities documented here - those hacks which you can do yourself and those which make use of other peoples' work. In the latter case performing searches by email and using the XooMLE interface are examples. So I now have my own developer key! Some of hacks here are more perl coding exercises than google related, but still useful for ideas.

There is a fun chapter near the end on the obligatory Google games - such as Whacking, Art, the Google Mirror site etc. One thing I will try is the Recipe finder, which locates appropriate meals based on what you've got the fridge. Though I'm pretty sure I know the things I can make with a Tesco Vindaloo ready meal tonight - not a lot of scope in that.

The last chapter is for the web master - interesting stuff on Adwords and getting removed from the Google Phonebook, for instance.

I didn't spot my own little hack - a one click google search. Add the following to your Favourites (I put it on the toolbar.) Then you can highlight a word in a webpage and click on the link to Google it. Probably only works in IE though.

```
javascript:Q=document.selection.createRange().text;
if(!Q){void(Q=prompt('Search...',''))};
if(Q)location.href='http://www.google.com/search?num=100&hl=en&q='+escape(Q);
else location.href='http://www.google.com/';
```

Ah, having just said that, check out hack number 28 - it looks like it covers this area.

Linux Server Hacks

Rob Flickenger

O'Reilly and Associates

ISBN 0-596-00461-3

240 pp.

£ 17.50

reviewed by Mike Smith

This book is written by Rob Flickenger, of NoCatAuth fame (and now NoCatSplash), and if anything, I think its better than the Google one. Again there are several (8, in fact) sections - this time covering things from basic Linux, through networking, scripting and some specific sections for backups and monitoring.

I have only one claim to fame, as some people who know me will be aware - I was a technical reviewer for the O'Reilly SSH book (along with Dug Song and others). There is a section on SSH so I headed straight for it. There are some fairly elementary hints - for instance how to setup ssh-agent, tunnel X11 and do port forwarding. Good to get you started anyway.

There are tips on a number of tools - stuff like ntop, lsof, ngrep, rsync etc. You know, all the good ones. I hadn't come across "watch" before - that's quite good, and simple.

There's a chapter on three applications in common use on Linux: Bind, MySQL and Apache. There's some good information in here, like using Views in Bind 9 - but I'm not sure I'd trust this in high security environments - to deliver one domain view to the Internet, and another internally (in the past I've done this completely separately). For the paranoid, maybe keep away from Bind anyway. The MySQL tips are good too - with some excellent hacks to glibc, kernel and MySQL for performance - not for faint-hearted !

What else? Well there is information on mod_proxy and mod_rewrite which does the same sort of thing as layer 7 switching on Alteon loadbalancers (and similar). Various ways to tunnel (IPIP, vtun etc), a bit on chroot, and the list goes on.

There some light-hearted comments, and I did notice one typographical error which amused me - an instruction to type "99<Enter>". Lets hope the newsletter preserves this [I think we succeeded (Ed)], or the moment will be lost for me when I re-read it.

There's also a chapter on RCS and CVS. It's a good idea for the SysAdmin to make use of these tools (CVS mainly, these days) for maintaing versions of config files, though I must admit I never have.

Summary

Thankfully I didn't get the third book in the series (as it currently stands) - it might just have given me an excuse to buy a mac (not that I don't want one, its just a money issue!) I do like these books - short, to the point and revelant - just like my book reviews ;) I've just realised (reading the back cover) that they are also trying to reclaim the word Hack for its original meaning (obvious, really, when you think about it). Excellent.

Contacts

Charles Curran
Council Chairman; Events; Newsletter
Oxford
Tel: 07973 231 870
charles.curran@ukuug.org

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

Sam Smith
Website
Manchester
sam.smith@ukuug.org

Alasdair Kergon
Events
Reading
alasdair.kergon@ukuug.org

Alain Williams
Watford
alain.williams@ukuug.org

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

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