news@UK

The Newsletter of the UK UNIX Users Group
Published electronically at `http://www.ukuug.org/newsletter/`

Volume 10, Number 3 October 2001

Table of Contents

Cover CD
CD Number 16 3

UKUUG News
News from Council 4
News from the Secretariat 4
UKUUG AGM 5
Linux Developers' Conference 2001: review of *some* talks 5
Perl news 6
Shell Programming Puzzle 7
UKUUG Open Source Award 7

Reviews
The Perl CD Bookshelf 8
Effective awk programming 8
Learning Web Design: A beginner's Guide to HTML, Graphics and Beyond 9
A Beginner's Guide to HTML, Graphics and Beyond 10
LPI Linux Certification in a Nutshell 10
SSH, The Secure Shell: The Definitive Guide 11
SSH, The Secure Shell: The Definitive Guide 11
Server Load Balancing 12
Linux Device Drivers 13
Professional XSL 14

Soap Box
From O'Reilly and Beyond 15
More News from O'Reilly 15

Letters to the Editor
From Stephen Forster 17
User Friendly 17

From the Net
Why I Have Disconnected from Email 18
Mozilla update 19

Forthcoming Events 21

Contacts 22

ISSN 0965-9412

CD Number 16

(David Hallowell)

Members' Requests

Once again here is a CD featuring members requests. All the suggestions I've received were either on the *Members' Requests* CD sent out at the beginning of the year or are on this CD. As the CD is produced for the benefit of our members then we'd like to know what you'd like to see on future CDs. Whether it is a request for a single piece of software for the next request CD or a suggestion for a theme for future CDs then I'd like to know. My email address is listed at the back of this newsletter and on the CD. Once I get some ideas I'll be asking for comments on the UKUUG chat mailing list.

This CD contains a variety of software as well as extensive online documentation, most of the documentation is from the Linux Documentation Project but we also have a copy of the O'Reilly Book "Linux Device Drivers" which has been released under the GNU Free Documentation Licence.

This CD features improved navigation as all the HTML navigation pages now make use of the LINK element which the latest releases of Mozilla support. To see it in action, enable the `Site Navigation Toolbar` in Mozilla from the `View > Show/Hide` menu. Unlike many magazine cover disks which are a few months behind with their Mozilla releases, the builds on our cover disk are nightly builds from the 10th October. This is because LINK element support is only a recent addition and was not in the 0.9.4 release. The 0.9.5 release of Mozilla will be out soon and will feature support for the LINK element, but this was not released in time for the CD. The HTML for this CD was written by hand but in future CDs this process will be automated and will also feature additional information on the software packages. Any suggestions let me know.

X11 Desktop software

- Ximian Evolution
- Mozilla
- OpenOffice
- X-Chat
- Opera
- The GIMP

Project Management Tools

- Bugzilla
- CVS
- FAQ-O-Matic
- IRM
- Mr Project
- XPlan
- PyGantt

Distributed Computing

- MOSIX
- Linux Terminal Server Project
- PVM

Documentation

- Linux Device Drivers, 2nd Ed
- Linux HOWTOs
- Securing and Optimising Red Hat Linux
- Linux Kernel 2.4 Internals
- Advanced Bash Scripting
- Linux Network Administrators Guide
- Linux Administrators Security Guide
- Linux Admin Made Easy
- The Linux System Administrators Guide
- LDP Author Guide

ERRATUM: the file `ldd/index.html` is missing from the CD, this means the link to the Linux Device Drivers book from the main index page does not work. However, to view the book just go into the `ldd` directory and view the files contained in there, each chapter is contained in an individual PDF file.

News from Council

(Charles Curran)

Editorial

This is the first newsletter put together by the newsletter working-group and there are certainly rough edges, however, I trust they will not put you off reading this issue.

Unfortunately, it did not make it into the last issue, but that was the last issue to be edited by Sue Small, who has now stepped down as Editor of news@UK. Sue has edited this publication for many years and has done admirable job. We thank her for the work she has done and for her various contributions to UKUUG.

Newsletter Contributions

We welcome material for the newsletter, and are particularly interested to hear from anyone either willing to write a regular column or act as the editor for such. If that sounds like you, or if you have concrete suggestions about the content of the newsletter, please contact us: newsletter@ukuug.org

December's newsletter will be the first of the slimline issues. The deadline for copy is mid-November, so please send us what you have as soon as possible.

Recent Events

Events since the last newsletter have included: UKUUG's Fourth annual Linux developers' conference held in Manchester, 29 June – 1 July. This was a most enjoyable event. There is a review of some of the talks elsewhere, but the papers, etc are available at http://www.ukuug.org/events/linux2001/papers/. We had arranged a talk in London by Tim O'Reilly at the beginning of August, but this had to be cancelled. Fortunately, Ted Ts'o, who was in London for the 51st IETF, stepped in and to talk on *A Decade of Linux*. Andrew Findlay, of Skills 1st Ltd, gave a talk after the AGM on *Building an Open Source PKI for E-Mail* .

Future Events

Elsewhere, you will see a repeat CFP for our winter conference which is being held in London 13 — 14 February. The venue is still being decided for next year's Linux conference; it is likely to be somewhere between Bristol and Birmingham at the end of June; this will be determined soon. We are also considering various tutorials/workshops. Topics suggested include IPv6, Zope, XML, XSL, Apache, Information Privacy, Samba, Embedded Systems, Perl for Systems Administrators, and Security. If have particular interests, please inform the Office.

Membership Classes

Council are about to consider what might be done about the different classes of membership. Anything will be considered. One suggestion is that, at the top end, we introduce sponsoring memberships — gold, silver — aimed at companies who would be willing to make an annual contribution to the running of the group. Please let the Council know your suggestions: ukuug@ukuug.org .

UKUUG's ISP

The new Council has decided to move from OA5 to Manchester University's service. The changeover should take place sometime later this year.

Working Groups

We welcome all members to participate in the running of the group; this you can do in a very positive way through our working-groups. If you wish to help in any of these — the newsletter, events, schools, liaison with industry and education, the web, publicity — please contact us: ukuug@ukuug.org .

News from the Secretariat

(Jane Morrison)

In the last issue Newsletter this piece was referred to as 'News from Owles Hall', we did in fact move from that address in January this year. I think because the office was located at Owles Hall for so many years we have had problems shrugging off the name. Any way, from now on I suppose this column has to be known as 'News from the Secretariat'. I do not think 'News from PO Box 37' sounds quite right!

I was kept busy over the summer months with the Linux 2001 Developers Conference. This year was even bigger than ever before. We had a large number of delegates and after securing sponsorship were able to provide an excellent Conference Dinner at a well-know Chinese restaurant. I have just finished the evaluations from the conference questionnaires and the event appears very well received.

At the time of writing we are very near the AGM date and are planning to hold up this issue to include an update for members on this meeting.

I would like to use this column to thank our book review team. This issue includes some ten or so reviews and I hope that you find these reviews interesting and helpful. The book reviews have been an important part of this Newsletter for many years and it would be helpful to know if the members do enjoy reading them and find them helpful — if you do could you let me know?. I suppose if you do not you should also let me know! Please email office@ukuug.org .

UKUUG AGM

(Jane Morrison)

Fifteen members attended the tenth Annual General Meeting of UKUUG, held on Thursday 27th September at the Institute of Education, Bedford Way, London. Members will receive their copy of the minutes in the post with this Newsletter.

It is such a shame that more members do not attend and we can only think that by not attending you are quite happy with all aspects of the current running of the group!

The only nominations for the two places on Council were for Charles Curran and David Hallowell who were both re-elected for a second term of three years.

The accounts for the company are very healthy and our reserves do put us in a safe position. However, these reserves were built up over many years (in the days when we had 400 or so delegates at events!). We continually monitor our expenditure and look very seriously when setting the costs for events so that we maintain a break-even status. Remember we are a non-profit making company and we aim to use any profits to fund member benefits, such as the Newsletter and CDs.

We hope you will be able to attend next year.

Linux Developers' Conference 2001: review of *some* talks

(Tom Wilkinson)

Linux Clusters in Enterprise (Dave McAllister, Egenera Inc)

This talk, the first of the conference, discussed the applications of Linux as an accelerant for enterprise clustering applications. Dave began by discussing the traditional responses to enterprise level applications, such as SMP and NUMA, and defining what he meant by a cluster — a collection of processes organised to store resources, requiring easy addition and removal of resources. Various considerations for clustering were then discussed, including load balancing, failover capacity, management software and shared storage solutions. Both open and closed source solutions were considered, including Mosix, Kimberlite, LVSP and Failsafe. This incorporated a discussion of the storage solutions, including shared SCSI and NFS. Finally the use of Linux as an accelerant in Processing Area Networks was discussed — having the source open enabled Egenera to work much more quickly to adapt it to their needs.

Design and Management of a Brute Force Cluster(Bob Gautier)

Contrary to the preceding talk, Bob's requirements for a cluster were very different. The University of Wales in Aberystwyth's computational Biology group had much simpler requirements for a cluster — each processor needed merely to continue to process a job until it had finished. There was no need for process accounting — the jobs are, in the main, database searches and prolog programs. Another requirement is that the programs should not need rewriting for use in the cluster as the database software in use is closed source. The hardware was discussed — the cluster consists of 40 650MHz Athlon machines with 256MB of memory and a 20G hard disk. The technology used to preserve the aims (1 job per node, long jobs, resources shared evenly and preserving work in progress from failure) included PostgreSQL, TCL as the job control language used, and Apache with CGI scripts for user control. Finally a comparison between the cluster as built and the university's Sun E6500 server (20 400MHz Ultra Sparc II processors, 10G memory) was given — the cluster came out of this comparison very well, giving 12-60x (depending on the jobs being run) cost benefit.

Value of Linux (Nick Davis, IBM)

Nick discussed the growing marketplace of acceptance of Linux ad the open source movement in terms of its growing marketplace acceptance, as an industry-wide initiative, mulitplatform OS which has formed the basis for innovation across the industry (with the notable exceptions of Sun Microsystems and Microsoft). He then went on to discuss how IBM intend are reacting to this change in the industry. This includes the fact that all IBM hardware currently in production (and for the foreseeable future) will run Linux. Also mentioned was IBM's Ready, Set, Linux programme which aims to train new users in Linux, including RHCE status and free development use of IBM software.

Porting Linux to x86-64 AMD Sledgehammer (Bo Thorsen, SuSE Gmbh)

The next generation of processors are approaching; Intel's Itanium and AMD's Sledgehammer processors are set to ultimately replace the current generation of CPUs. While Intel's new chip is completely incompatible with the current x86 range, AMD aim to keep their processor compatible with existing software while at the same time exploiting the power available in a 64-bit architecture. Bo explained clearly the various operating modes of sledgehammer, which include a native 64-bit operating mode and a compatibility mode for running 32-bit (and lower) applications. He then went into greater detail, talking about the differences between an x86-64 and an Athlon, highlighting the new registers and features. Finally a demonstration of the processor in emulation was made — there being no actual processors in existence. This was running a version of GNU/Linux ported by SuSE. One point of note here was that SuSE were ahead of Microsoft in that they would have software ready for release on the day Sledgehammer is launched.

x86-64 Vsyscalls (Andrea Arcangeli, SuSE)

A syscall is one of the most expensive operations which need to be executed in an operating system, due to the large overhead of changing from user to kernel mode and back — on a P166 gettimeofday() can take 125 cycles, and on an Athlon, 266. gettimeofday() is fast, so would be better executed in user space, however cannot be moved due to its requirement for interrupt data. This is a typical use for a vsyscall, which is implemented by using a kernel page in userspace. Vsyscalls use locking numbers in order to ensure a correct response to whichever function has been called. On benchmarking the vsyscalls, when compared to syscalls, come out up 250% faster.

The Journalling Flash Filesystem (JFFS2) (David Woodhouse, Red Hat Inc)

The talk began with an explanation of the different types of flash memory available — these include NOR and NAND flash. The original JFFS is used in the Compaq Ipaq's Linux implementation, which uses NOR flash. NAND flash is unsuitable for the implementation of a filesystem due to the rapid wear time — each bit of memory can be unusable after just ten write/ reset requests. While NOR flash is less wearable, it suffers from similar problems so part of the filesystem's requirements are that wear should be balanced across the entire device. In addition some form of journalling is required, due to the likelihood of power failure on a portable device. The garbage collection algorithm was explained in detail — this code also performs the wear balancing on the flash device.

Perl news

(Simon Cozens)

At last year's Open Source Conference, Larry Wall and the other Perl maintainers got together and decided that the best way to drive Perl forward was to make some radical changes to the language and its implementation.

Over the past year, the Perl community has been consulted on how to make Perl better — the new features that people want to see, the things that people want to get rid of, and new directions to explore. As the appointed Perl 6 Language Designer, Larry has been weighing up the suggestions, and, together with his assistant Damian Conway, has been steadily putting together the design of the Perl 6 language. www.perl.com and dev.perl.org publish the results of Larry's deliberations, the Apocalypses (http://www.perl.com/pub/a/2001/10/02/ apocalypse3.html) as well as Damian's more practical tutorials, the Exegeses (http://www.perl.com/pub/a/2001/ 10/03/exegesis3.html)

While the language will still look pretty much like Perl 5, anyone who has ever hacked on Perl will be glad to hear that the internals are being completely redesigned from scratch. Dan Sugalski has been designing the Perl 6 interpreter, code-named Parrot, with the goal of being a portable bytecode interpreter which can deal with other dynamic languages as well as Perl, such as Python, Ruby and Tcl. I have been helping Dan out by holding the release "pumpkin" —

the responsibility for committing patches, gearing up for releases and generally herding developers.

Currently, we have seen 3 Apocalypses from Larry, and expect one roughly every month or so; Parrot is currently at version 0.0.2 and its assembly language can deal with basic string, integer and number operations. Parrot 0.0.3 is expected within two weeks, and will see the addition of simple "scalar" types which can morph between strings, integers and numbers at will. We expect a full design of Perl 6 and a first cut at the implementation to be no more than 12 months away, and certainly to have some very interesting prototypes within 6 months.

But you do not have to wait! One of the big benefits of Perl 6 is that it is inspiring plenty of improvements in Perl 5 right now; check out the work that is being done in the `Attribute::*`, `Class::Delegation`, and other modules by Damian Conway on CPAN, and also have a look at the `Perl6::Parameters`, `Perl6::Variables` and `Perl6::Interpolators` modules for a glimpse of how to program Perl 6 code in Perl 5.

Simon Cozens is an open source programmer and author; he is a columnist for the Perl Journal, the author of Beginning Perl by Wrox Press, and co-author of Using Perl and C. He has recently joined the UNIX team at Oxford University (OUCS).

Shell Programming Puzzle

(James Youngman)

Write a script which takes as input (on stdin) a list of files (one per line) and which emits a list of files with identical contents, each group on one line, and does not emit "groups" which contain only one file.

Example input:
a
b
c
d
e
f
g

Example output:
a c d
e b

The winning entry will be the bug-free entry which is most portable (i.e. runs on the widest selection of vanilla-installed UNIX systems available to the judges). In the event of a draw, the entry occupying the fewest bytes will be declared the winner. There are no prizes, but the winner will be announced in the next issue, and runners-up may get honourable mentions.

Entries by e-mail to james.youngman@ukuug.org.

UKUUG Open Source Award

(Charles Curran)

As mentioned at the AGM, UKUUG's Council is initiating an Open Source Award which is open to current students in UK Higher Education.

The prize — £ 500 — is to be awarded annually (provided submissions of sufficient merit are received) for a significant contribution to open source; this might be in the form of an article or paper, software product, or other contribution. The winner will be expected to deliver a talk at the annual UKUUG Linux conference, which in 2002 will be held around the end of June.

The judging panel will include representatives from UKUUG, the Open Source community, and UK Computing Science departments.

The closing date for submissions is Friday, 5 April 2002. Email any enquiries to office@ukuug.org.

The Perl CD Bookshelf

O'Reilly & Associates, Inc., May 2001
672 pages, £56.95
ISBN 0-596-00164-9

(Reviewed by Joel Smith)

O'Reilly have updated their *Perl CD Bookshelf* . In the process, it has shrunk from six books to five. So what is going on? The common books are *Perl in a Nutshell*, *Advanced Perl Programming* and *The Perl Cookbook* (all 1st edition). *Programming Perl* gets updated to the current 3rd Edition, and *Perl for System Administration* replaces *Learning Perl* and *Learning Perl on Win32*. As before, the bookshelf comes with a paperback copy of *Perl in a Nutshell* as a bonus.

Although some might complain that the CD is less value than before, as it only covers five books, rather than the six of the first edition, I think that the balance is far better now than it used to be. Was there really a need to have *Perl in Nutshell*, *Learning Perl*, **and** *Learning Perl on Win32* on the CD? My feeling is that *Perl for System Administration* adds more depth to the CD than the two omitted books.

I also suspect that it might have been partly to do with logistics. *The Perl CD Bookshelf* was published in May, whilst the 3rd edition of *Learning Perl* did not come along until July. The 2nd edition dates back to 1997, and is more than a little long in the tooth. Putting the older version on the CD would lead to it rapidly becoming out of date.

Enough of the books, and on to the CD itself, and in particular the usability of the search engine. Firstly a word of caution. There is a bug on the CD which prevents the search engine from working with anything other than Windows (and versions of Mac OS prior to X). It fails under Linux, Solaris, MacOS X (native and classic mode).

The problem is that the Java search page refers to an archive `qagent.zip`, which is a file which does not exist on the CD (DOH!), even though the expanded directory does exist. The Windows version works because there is an alternative reference to a `qagent.cab` file which does exist.

For O'Reilly to produce a Perl CD which does not work under UNIX is a serious faux pas, and they appreciate this. I received this email back from them when I reported the problem.

You're right about expecting better QA testing on the disk, especially for Linux from us. This was definitely a blunder, but we have made changes in our QA process for those products, and it should not happen again!

They have provided a fix for the problem which can be downloaded from http://examples.oreilly.com/perlcdbs2/ . The file is called `perlcd_updates.zip`. This supplies the missing file as well as providing several fixes for broken links and errata from the books. Of course, to be able to apply the fixes, you need to install the files on your hard disk, rather than running from the CD, but if this bothers you, you can always cut the fixed files to a new CD.

But how good is the search engine once you have it installed? On one hand, it is streets ahead of the rather dismal engine that came with the older edition. It runs directly from the browser, without requiring any server processes running. It does this by using *a Java applet interface to query a pre-generated index of terms derived from the full text of all book pages on the CD-ROM* . It runs on any browser supporting Java 1.1.2 or later, and it does indeed seem to be cross-platform.

Unfortunately one other rather serious blunder is with the indexing (once I got the search applet working). Bearing in mind that this is a bookshelf on Perl, I would have thought that I should have been able to search on functions such as `$arg`, or `ExtUtils::MakeMaker`, or `$_`, etc. It appears that anything with a special character (not that rare in Perl!) is excluded from the search engine. So to do a search on `ExtUtils::MakeMaker`, you would have to search on ExtUtils and MakeMaker, rather than the full module. Thankfully, there is also the master index on the CD which allows you to browse through a linked alphabetical index (exactly as you would find in the back of the books).

Since you have had to install the files on your hard disk, there is nothing to stop you running `ht://dig` or your favourite indexing software on it yourself, and creating your own search engine. However, it should not be necessary to have to resort to this. O'Reilly are looking into ways to fix things in future editions.

To sum up, it is very useful to be able to have an electronic version of these books. In spite of the reduction in the number of books on the CD, the breadth of information is greater, and the CD hence more useful. Oh yes, it also works out much cheaper than buying the books individually, but is probably less amenable to bedtime reading!

Joel Smith is the Senior UNIX Administrator for IFonline.

Effective awk programming

Arnold Robbins
O'Reilly & Associates, Inc.
421 pages, £28.50
ISBN 0-596-00070-7

(Reviewed by Jan Wysocki)

This book claims just one version of a freely adaptable manual. It is licensed with the Free Software Foundation's Free Documentation Licence and is said to be available in electronic form for you to copy and modify. However, I (well ok — Google and Teoma) was unable to locate an online copy of the whole book.

I realise that awk may be so old and obscure to today's Unix newcomers that I should explain what awk is. It is an interpreted language designed to allow you to search and manipulate large text files. The text might be in a terse database format, for example, a few years ago, you could download the whole Internet Movie Database and the associated awk programs to allow you to do much the same searches that you can conduct via the website today. You might use it today to filter news items on a press agency feed, or to tidy your email archives.

Long ago I wrote and maintained fairly large awk programs. Nowadays I rarely use more than single line awk programs in the shell scripts that I write. I know that I could often write the whole shell script more effectively as an awk program, but habit and memory loss get in the way. When I saw that this new edition was coming out I decided that reviewing it would be just the fillip I needed to get back into awk programming. As it turns out I have not actually written much awk since receiving the book, but I have got excited enough to actually grab and compile version 3.1 of the gawk source code, so that I could try out some new features. When work requires it, this will be the book I choose to consult when writing awk programs in future.

The book is arranged in two parts and works conveniently as either an introduction to awk programming, or as a reference manual for the whole language. I am a little surprised that there is only one reference to executable awk programs, right at the start. It is such a convenient feature, of modern Unix variants, that I feel that beginners need a later reminder. However, I cannot fault the comprehensive nature of this book. It is clearly laid out with an adequate guide to the contents, frequent cross references to other parts of the book and a comprehensive index. A book like this is never going to be unputdownable, but if you need to write or understand awk programs on your system, then you might want this book on your desk. I say might because, if you do not need the newer or more complex features of awk, then O'Reilly's "Sed and Awk" 2nd edition, 1997, may be sufficient.

Oh, I had better tell you that what got me excited enough to move from gawk 3.0 to 3.1 is the introduction of two way pipes that allow awk to communicate with subprograms either on the same host or across a network using TCP/IP. I have not actually written anything useful yet, but writing trivial awk programs to communicate with remote hosts is quick, simple and unaccountably pleasing.

Jan is currently administrating a Solaris / Apache / Weblogic / Oracle / Java development environment for Virgin Money.

Learning Web Design: A beginner's Guide to HTML, Graphics and Beyond

Jennifer Niederst
O'Reilly & Associates, Inc.
388 pages, £24.95
ISBN 0-596-00036-7

(Reviewed by Lindsay Marshall)

The blow-ins say "The Internet Runs on O'Reilly Books", and let's face it, it's true, at least as far as the web is concerned. How many times a week do you reach for a reference book on some web issue and pick up something that is not an O'Reilly book? I just glanced at my shelf of frequently used references and they are all O'Reilly apart from one book (Russell Winder's Java book since you ask). O'Reilly have always had the "just the facts, ma'am" market pretty well sown up for a few years now, but their less terse books are (in most cases) just as good.

This book on Web Design is a case in point. Written y the author of Web Design in Nutshell, it takes a slower more developmental approach to the whole process of web design, from basics right through to sound advice on Flash (only where appropriate). Whatever your level of skill and experience you will certainly learn something by dipping in somewhere in the book, and for people starting out, it will bring them up to speed very quickly.

Of course, it won't turn you into a graphic designer overnight and the author makes pains to point this out, but it will point out the major pitfalls to avoid, and the best ways to approach a particular task. The book is bursting with examples, tips, pointers to use-

ful software and references books that will provide more information on particular topics. The illustrations are printed in grey-scale, but are no less clear because of this (and a small colour insert shows many of them as they would be seen in a browser). There are some nice "how not to" examples as well — the animations on a black background homepage is wonderful.

If you want to get to grips with web design then I can definitely recommend this book. For more experienced practitioners — though they will certainly learn something — it is a less useful purchase. They will know a lot of the day to day stuff already and may feel the need for lower level detail.

A Beginner's Guide to HTML, Graphics and Beyond

Jennifer Niederst
O'Reilly & Associates, Inc., March 2001
388 pages, £24.95
ISBN 0-596-00036-7

(Reviewed by Paul Webb)

Learning Web Design is intended to introduce the novice web designer to a broad range of web related technologies. Niederst therefore begins with an Internet-related definition of terms and moves, via discussions of HTML and Graphics, towards lucid descriptions of technologies like Flash, DHTML, Audio and Video.

The book is therefore explicitly intended for the newcomer who wants both an introduction to the subject and a reference to more advanced resources. Niederst writes in a chatty and accessible style and demonstrates techniques and concepts with reference to sample web pages and images which are intermittently reproduced as coloured prints in a `Gallery' section. She also defines any terms as they occur and provides a balanced appraisal of the pros and cons of any particular technology. (Chapter 18 on `Usability' is essential reading for those designers who choose special effects before assessing audience needs). The book is also liberally sprinkled with a range of sidebars which contain invaluable hints and tips, useful URLs and pointers to books for further study.

In short, if you would like to get into web design, buy this book. After all, it seems that we now have a companion volume for the same author's Web Design in a Nutshell.

LPI Linux Certification in a Nutshell

Jeffrey Dean
O'Reilly & Associates, Inc., May 2001
570 pages, £28.50
ISBN 1-56592-748-6

(Reviewed by Neil Levine)

Certification has unfortunately become more of an accepted industry standard over the past few years with MSCE or CCIE qualifications becoming more important than the number of hours you spent at home hacking your bash startup scripts or XF86Config file at 3:00 in the morning.

With the popularity of GNU/Linux, it was inevitable that exams such as the RHCE (Red Hat) or the more general LPI (Linux Professional Institute) would crop up. LPI Linux Certification in a Nutshell documents the first two exams in the LPI course and as such provides a good introduction to the basics of UNIX whilst giving enough insight into the specifics of maintaining GNU/Linux box to hopefully make the Certification worth something.

The first exam, as detailed in the first half of the book, covers general UNIX concepts such as file and stream manipulation, file system administration and file system hierarchy. The brief 20-odd page given to each section consists of a mixture of manpage-"lite" descriptions accompanied with more verbose three to four page explanations of the topics. As you would expect each topic has an accompanying set of examples and review questions. What is less practical is the setting out of "Objectives" detailing why you would need to perform certain tasks. Listing "User Support" as an objective under the Documentation section or "Maintain An Effective Data Backup Strategy" appear more biblical in their hopes than useful for exam purposes.

The second half of the book, and exam, cover more Linux specific areas such as package installation and kernel compilation as well as networking topics. Again the level of information is sufficient to push people in the right direction armed with the relevant tools and concepts if not the absolute answer.

The exam aside, this book is a very competent introduction to using UNIX for the uninitiated and goes some way to illuminating the atomic nature of the many individual UNIX features which can be combined to produce powerful results. The latter half of the book would also help as a compact FAQ for

those who have got their GNU/Linux box up and running already and want to start customising and installing applications.

SSH, The Secure Shell: The Definitive Guide

Daniel J. Barrett & Richard Silverman
O'Reilly & Associates, Inc., January 2001
558 pages, £28.50
ISBN 0-596-00011-1

(Reviewed by Bob Vickers)

SSH is a very useful collection of tools, providing secure replacements for programs such as `telnet`, `rlogin`, `rsh`, `rcp`, `ftp`, `xon`, and `rxterm`. Any administrator concerned about passwords and other sensitive information being stolen by packet sniffers needs to be familiar with it.

Once set up SSH is generally easier to use than the less secure equivalents, but before you get to this happy state there are a number of choices to be made and hurdles to be jumped. This book is aimed primarily at the UNIX system administrator or advanced UNIX user and provides a huge amount of useful information.

The book starts with an introduction to the concepts of SSH, then moves on to the meat: detailed information about SSH for UNIX users and administrators. Finally, it discusses some implementations for Windows and Macintosh users.

There are currently three main implementations of SSH for UNIX: SSH1, SSH2 and the new but rapidly developing OpenSSH. They have much in common and the authors have elected to describe all three, pointing out differences as they go along. This is an ambitious approach, and inevitably adds some clutter to the text, but the authors have risen to the challenge and written an extremely clear and helpful book. They cover server administration, client configuration, and the clever port forwarding features that allow you to tunnel other protocols such as X11 and IMAP across an insecure network.

I would recommend this book to people who simply want to get SSH working as quickly as possible as well as those with an interest in cryptography.

Bob manages the Computer Science computers at Royal Holloway, University of London.

SSH, The Secure Shell: The Definitive Guide

Daniel J. Barrett & Richard Silverman
O'Reilly & Associates, Inc., January 2001
558 pages, £28.50
ISBN 0-596-00011-1

(Reviewed by Virantha Mendis)

Until the Internet became a fashionable commodity in the mid-90s, there was not much talk about its security aspects. The emergence of ISPs (Internet Service Provider) brought many household connections to the Internet. A new breed of hackers started to emerge with these connections. These young people had access to powerful computers, lots of time, as well as a wealth of information on how to become a hacker. All these hacking activities have resulted in many high profile security breaches in last few years. *SSH, the Secure Shell: The Definitive Guide* by Daniel Barrett and Richard Silverman tries to educate the user on how to conduct their business in a secure manner when performing interactive connections, i.e. telnet. When the Internet was part of the original ARPA network, one of main objective was the ease of use. One example of this is the "r" utilities such as `rlogin`, `rsh` and `rcp`. These are written purely to provide the necessary functionality with minimum effort. For example, login onto remote machines without a password. Unfortunately, these are the areas hackers started concentrating their efforts most. Also, most of the network traffic generated by computers travels across the network in plain text.

Once this became common knowledge, the main target for hacking shifted towards capturing user account details and passwords. Two common utilities, `telnet` and `ftp`, that are used to connect to remote servers, became the main source for revealing the user account details. One way to stop snooping of the user account details is to encrypt the network traffic. In the mid-nineties, Tatu Ylonen used public encryption mechanisms to encrypt the network traffic between the server and the client, and released the first version of Secure Shell. The Internet community adopted this as a protection against snooping attacks. There are now two versions of the secure shell protocol and various versions of the client software, which run on various UNIX flavours, Windows and even on epoch based machines.

The Secure Shell suite is comprised of a daemon program (`sshd`) and a collection of client programs (`ssh`, `scp`, etc.). The daemon and the client pro-

grams use public-key authentication to identify and verify the client. Once a connection is established, a secure tunnel is created and the network traffic is encrypted. This prevents any would be attackers from snooping the user account details such as usernames and passwords.

The authors have organised the book into several "tracks". This allows readers with varying skills to identify the relevant sections for reading. The tracks are grouped according to the functionality of the software: installation and configuration for system administrators, how to use the secure shell client programs for UNIX users and several sections for Windows and Macintosh users.

As there are now two versions of the protocols available, the authors have done their best to explain the differences. This is augmented by detailed explanation on how to configure the two versions of the `sshd`. Given the importance of using software to encrypt the network traffic and the use of the public-key encryption to authenticate clients, Secure Shell has gained wide popularity. For any user or system administrator who is concerned about the vulnerability in the network and decides to use the software, this book provides valuable wealth of information. The book is written in easy to understand English and there are numerous diagrams that are used to convey complex technical details behind the secure shell protocol. This book is a must for any person who uses the Internet for communicating with remote servers.

Virantha is a Senior UNIX System Support thingy with BT Internet Data Centres.

Server Load Balancing

Tony Bourke
O'Reilly & Associates, Inc., August 2001
192 pages, £24.95
ISBN 0596000502

(Reviewed by Joel Smith)

This new book assumes no more knowledge of IP and Ethernet than a familiarity with their basic workings, and the workings of the Internet in general. Although it briefly touches on DNS round robin and clustering solutions, this book is firmly based around hardware Server Load Balancing (SLB) solutions.

The book is divided into three sections. Part I concentrates on the theories and concepts of Server Load Balancing. Part II concentrates on the implementation and network topology of load balancers. Part III is a configuration guide to four major load balancing products on the market: Alteon WebSystems, Cisco's CSS (formerly Arrow Point), F5's BIG-IP and the Foundry ServerIron Series.

The first four chapters are a clearly laid out treatment of the theory behind SLB and the essential terminology and concepts used. There is the inevitable explanation of the 7 layer OSI model and networking basics. Perhaps some of the ground covered is a little too basic to be really necessary — I am not sure how necessary it is to cover concepts like "read-only", "super-user" and "server". A quick definition of terms in the introduction should be sufficient, but to be still devoting pages to such concepts in chapter 2 smacks of bulking out the text. That said, the information is communicated well, and the illustrations aid the text well.

Part II gets down to the nitty gritty of designing a network to work with SLB. There are two basic approaches using either a flat-based SLB architecture, or a Network Address Translation-based architecture. A NAT-based architecture is where the load balancer routes network packets between different subnets, and translates them appropriately. A flat-based architecture is where the load balancer is on the same subnet as the servers and does not route packets between subnets. This section deals with both of these two architectures in some detail again using clear illustrations to communicate the differences being discussed.

The final section gives a walk-through to configure SLB products from four of the major vendors in this area. The walk-through takes you from a new out of the box load balancer through setting it up in the each of the two basic architectures. The walk-throughs are conducted via the CLI wherever possible, rather than any web interfaces. This section is useful if you have these particular load balancers, but I would think is of limited value if you do not.

A very nice touch is the Quick Command Guide in Appendix A. This appendix provides a quick reference to commonly performed administration tasks involving the load balancers featured in the book. It is designed to save time and help in a crisis situation, when reading through a chapter would take too long. Since load balancers are by their very nature fairly critical pieces of equipment — you would not buy an expensive load balancer for an unimportant service — I thought that this could prove to be a very useful feature of the book.

Appendix B covers how to set up Direct Server Return (where the server sends the return packet back without returning it via the SLB) for various different server systems. It covers the configuration required

for Solaris, Linus and Windows 2000 operating systems, and Apache as the web server.

Although the final Appendix C is entitled Sample Configurations, this is in fact vendor neutral. The "Configurations" referred to are network configurations rather than hardware configurations. Consequently it is filled with network diagrams showing several different possible network topologies and implementations.

This book would be useful for those deciding to implement SLB on their networks. It is a help for those new to their load balancer, but does not delve into the more complex features on offer, for which you are referred to the documentation. At present the book is up to date with the various load balancer OS revisions. As the book ages, the information it presents will become less relevant. As half the book is devoted to the OS specific configuration, I suspect that it will date quite badly, unless frequently revised. But for now, this does not present a problem. I suspect I will find this book an easy source of information.

Joel Smith is the Senior UNIX Administrator for IFonline.

Linux Device Drivers

Allesandro Rubini and Jonathan Corbet
O'Reilly & Associates, Inc.
570 pages, £28.50
ISBN 0-59600-008-1

(Reviewed by Matt Palmer / Richard Ibbotson)

Contents

Linux like all operating systems uses device drivers to get the software to drive a piece of hardware. This book is all about those drivers. How to understand them and what to do with them.

The open source revolution has given us access to several things which were previously inaccessible. One of the reasons I got into Linux in the beginning was I thought "hey, I can see how an operating system works". One look at the kernel code scared me off, and I went into other things. The second edition of Rubini's book now with help from Jonathan Corbet goes a long way to dispelling my fears to do with technology that scares me because I do not understand it. While it is primarily written for those who wish to support new hardware and program within kernel space, it is impossible to do these things without knowing a fair bit of information about core kernel activities. All of these things are explained well in the book.

The primary focus, of course, is device driver functionality. At the outset, Rubini and friend make sure the reader is familiar with the concept of a loadable module, and the things which are required to understand the book. These are:

- A linux machine, running a 2.4.x kernel
- A C compiler — preferably the GNU one
- Working knowledge of C programming or C++

It is probably a good idea to upgrade that last point to a firm grasp of C programming or something similar. If you are capable of running C code in your head while reading it, your comprehension of the code snippets provided will be greatly improved. The coding style is open and readable, and it makes for a greatly improved understanding.

There is only a very small part of the book which actually requires you to work with hardware — almost all of the basic concepts of device driver writing and use can be demonstrated purely in software — wait queues, task queues, ioctls, I/O, and a host of other things. Debugging, the bane of every programmer, is dealt with well, especially when one considers just how hard it is to debug something which you cannot start and stop at will, and which will bring down the system if it crashes. You will also pick up a working knowledge of real live hardware issues — it is obvious that the author has `been there, done that' with real devices. He has also been gifted with an ability to explain concepts very clearly, with the aid of diagrams, where appropriate. This is, in fact, one of the strongest points in the book — all of the hard technical details are spelt out, in detail, to the level where you can immediately say "I get that now".

After a thorough grounding in the basics, there are some advanced topics as well — handling versioning within modules, DMA access, network drivers, and peripheral buses, to name a few. The extensive advances made in newer kernels — up to and including 2.1.43, are also covered, allowing some measure of advanced design. If you follow the advice given throughout the book, however, your own modules should work, with minimal modifications, with the newer 2.4.x source tree. I would be interested to see a book by Rubini concentrating on the function of the Linux kernel in detail. Chapter 16, Physical Layout of the Kernel Source, gives a taste of what the authors are capable of.

In short, if you want to support hardware, or want to

know what goes on under the hood of your kernel, and have a good working knowledge of C, I would recommend buying this book. You will not regret it, even if you never actually write a device driver that ends up in the Linux kernel.

Matt lives at Wollongong in Australia. Richard is from Sheffield. Sheffield LUG is the link.

Professional XSL

Kurt Cagle, *et al*
Wrox Press Ltd., April 2001
727 pages, £38.99
ISBN 1-861003-57-9

(Reviewed by Alan McLintock)

I was quite pleased when I got this book for review: I thought it had a whole chapter on FOP. It doesn't!

What it does have is a chapter on CSS and XSL:FO. I guess I do need to read all of the CSS section to know what I can and cannot do with font styling and the like. So I get to the XSL-FO part of the chapter. It starts off with a good description of where FO fits into the world, and how the page is broken up into regions. It lists the 56 XSL-FO elements, and points out that many of them have HTML analogies (e.g. the "list" ones). And then I flipped back and forth trying to find an explanation of each one. I flipped in vain.

The explanation of the structure of an XSL-FO Document was good. I think it is better than the documentation with FOP — but that documentation has improved since I first read it.

It talks about creating XSL-FO files from XSLT — something I have not yet needed to do since I've generated the FO directly.

And then for some reason the author uses Antenna House XSL Formatter. This seems to be an add-on for Microsoft Internet Explorer and so is useless for generating PDF. Hurrah. A few pages later we find FOP mentioned. The exact file they use in their examples is `fop-0_15_0-forBeginners.zip` at http://xml.apache.org/dist/fop/ which does not seem to exist any more. Perhaps we can re-reinstate it — or redirect users to a the current release.

Sadly it does not really talk about FOP in any more detail than enough to get it going.

What else might you like to know...

It has a whole chapter on SVG! Unfortunately this chapter seemed a bit disappointing too. It did not give me any examples I wanted, for instance, on dotted lines. (I have some colour line graphs which I want to produce with different styles of dotted lines for rendering on a black and white printer).

The next thing I want to learn about is Cocoon. Cocoon is mentioned — but once.

Oh well. Never mind.

Alex works for the Open Source Consultancy OpenWeb Analysts Ltd.

From O'Reilly and Beyond

(Josette Garcia)

General interest: O'Reilly and Pearson launch joint venture

O'Reilly and Associates, Inc. and the Pearson Technology Group announced joint venture at the beginning of October, launching a new online information retrieval service for the enterprise IT market--Safari Tech Books Online. It's O'Reilly, and the best of the rest; take a look at Safari at http://safari.oreilly.com/mainhom.asp?home.

No UG discounts are being offered on this service due to already low pricing. To kick off the Safari launch, we [ORA] are offering a chance to win a $1000 tech book shopping spree. No purchase necessary. Enter the draw at: http://oreilly.com/news/safari_drawing_0901.html For the complete press release see http://press.oreilly.com/safari_jvlaunch.html.

According to SeyboldReports.com, Safari Books Online is Inexpensive, easy to use, and loaded with high-quality content -- the Safari online reference service for IT professionals will provide stiff competition for Books24x7.com. http://www.seyboldreports.com/ebooks/features/011002-oreilly.html [subscription needed :-(, ed].

An interview with the User Friendly gang

We visited Columbia Internet and spoke with members of the popular comic strip. http://oreilly.com/news/rootofevil_1001.html .

Mac: Nine things I wish I had known about Cocoa

Practical advice to programmers getting started with Apple's Cocoa, especially programmers used to command-line environments http://mac.oreilly.com/news/cocoatips_1001.html .

Shell scripts, command lines, and classpaths

For a look at Java programming on the Mac, read: http://www.oreillynet.com/pub/a/mac/2001/10/02/osx_java.html .

Open Source: keep publicly funded research public

Just as scientific research must undergo peer review, software developed by publicly funded research should be released under an open source or free software license, according to OpenInformatics.org. If you agree, sign their petition at: http://www.openinformatics.org/dsp_petition.php .

Larry Wall's Apocalypse 3

As the design of Perl 6 unfolds, Larry continues his analysis with a discussion of operators: http://www.perl.com/pub/a/2001/10/02/ apocalypse3.html .

python.oreilly.com

If you have not seen O'Reilly's latest Python page, you may want to have a gander at: http://python.oreilly.com/ for the latest Python articles and books.

Oracle: Oracle9i's improved RMAN

Coauthor of O'Reilly's upcoming "Oracle RMAN Pocket Reference," Darl Kuhn, writes about the significant improvements made to Oracle's Recovery Manager: http://oracle.oreilly.com/news/oraclerman_1001.html .

More News from O'Reilly

(Josette Garcia)

Java

O'Reilly author Eric Burke provides ten valuable tips for using XSLT with Java. To learn how to apply XSL transformations in real-world situations with Java, do not miss Eric's upcoming *Java and XSLT*. Learn more: http://java.oreilly.com/news/javaxslt_0801.html.

Learn about servlet filters from this ONJava.com article: http://www.onjava.com/pub/a/onjava/2001/08/28/filters.html .

P2P

For several decades, the military has been using large-scale client-server systems to build networked environments where soldiers can train in simulated

battle conditions. Now the military is looking at peer-to-peer technology as a way to build these simulations without a vulnerable central server. Michael Macedonia, the Chief Scientist and Technical Director for the U.S. Army Simulation, Training, and Instrumentation Command (STRICOM), talks with O'Reilly editor Richard Koman about how the military simulates battle, how peer-to-peer technology could change that, and the advances that have made a $69 flight simulator program as valuable as the multi-million dollar systems of a few years ago. P2P Goes to War: http://www.openp2p.com/pub/a/p2p/2001/08/28/ p2pwar.html.

Security

Server load balancers can perform double duty as cost-efficient firewalls. Learn why and pick up some useful techniques, including configuration examples, in this article by Tony Bourke, author of O'Reilly's recently released *Server Load Balancing*. Load balancers as firewalls: http://sysadmin.oreilly.com/news/ loadbalancing_0801.html .

The PC world continues to use methods that simply do not protect users against viruses, worms, trojans, and other infectious code. Here are the preventive modifications that should be made to every Windows PC, according to Roger A. Grimes, author of O'Reilly's recently released *Malicious Mobile Code*. Fighting Malicious Mobile Code in a Windows Environment: http://security.oreilly.com/news/ maliciouscode_0801.html .

Apple/Mac

Teaching a few programming basics to kids will show them what is really going on behind the monitor, and REALbasic is a great tool for the task. Here are some fun programming projects that children can grasp, by Matt Neuburg, the author of O'Reilly's upcoming *REALbasic: The Definitive Guide, 2nd Edition*. Learn more: http://mac.oreilly.com/news/ realbasic_0901.html .

Open Source

Recognising the natural fit between an open documentation standard and open source development, O'Reilly & Associates has released *DocBook: The Definitive Guide* under the Free Software Foundation's GNU Free Documentation License version 1.1. Learn more: http://web.oreilly.com/news/ docbooklic_0901.html .

CrossOver Brings QuickTime Movies to Linux

Writing for O'Reilly's Linux DevCenter, Derrick Story rejoices that QuickTime multimedia capabilities are finally available for Linux users, and he examines the new CrossOver plug-in that makes this possible: http://linux.oreillynet.com/pub/a/linux/2001/09/06/ crossover_partone.htm .

Josette is Marketing Manager for O'Reilly's UK operations.

Dear Ed,

I do not understand why we have a review of *DHCP for Windows 2000* in the UKUUG newsletter.

Surely, if we want a review of a book covering DHCP, then we could have a review of a book which concentrated on a UNIX based implementation?

This Windows centric book and review seem quite out of place and it is certainly not why I subscribe to the UKUUG.

There are plenty of books on UNIX/Linux related topics still to be covered.

Steve

(Stephen Forster)

Dear Steve *et al* ,

Council did notice a spate of books being reviewed that were not apparently relevant to our members. However, some of the texts are borderline, and if someone thinks `from the cover' that such a book, or their review of it, may be of interest to members, we will at least offer them a copy for review.

Charles Curran

User Friendly

(Illiad)

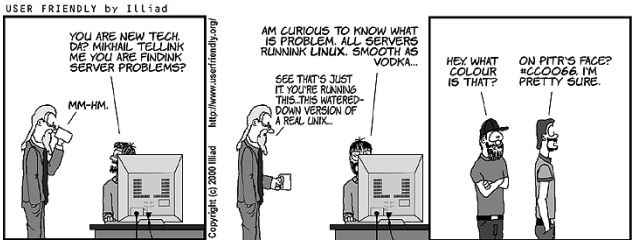

Why I Have Disconnected from Email

(Steve Talbott)

I suppose it's a strange thing for the editor of an on-line newsletter to disconnect from email. Well, not quite "disconnect"; I will still have my email account and I, or an assistant, will have to conduct obvious NetFuture business such as passing manuscripts back and forth, receiving letters to the editor, and processing financial contributions — maybe 5% of my overall email burden. (This is to leave aside the 80% or more of my mail that is spam — is there no end to this worsening plague?). But correspondents will receive an automated message saying that if they need substantial interaction with me, we will have to conduct it via phone or postal service.

Long-time readers will know that I've agonised and vacillated over this issue for years. Shutting off this avenue of contact feels too much like turning a cold face toward all those readers and supporters of the newsletter whose only "sin" is that they would like to get in touch — readers who have energised my activity over these years and contributed so much to my understanding. How could I possibly leave messages sitting unanswered in my mailbox, some of which are, and nearly all of which should be, the attempt of one human being to speak to another?

But now, through extremity, my spine has finally stiffened, and you will find me almost combative on the issue. Actually, the stiffening is more than metaphorical. Nerve pain in my neck, worsened by work at the computer, brought me face to face with the prospect of disability. It is true that the neck problem also derives from my inadequate handling of the stress and pressures of my overall work life, a challenge that needs to be addressed "from the inside out". But even there, it became clear to me that one of the best ways to reduce the pressures and make them more manageable was to remove the main burden of email from my life. I say "burden" only in the sense that some of the activities we may want to pursue most can become burdensome if we do not maintain proper balance in our lives. The burden, in other words, is a feature of the way I manage my own life; I have no quarrel at all with those who have approached me through the perfectly acceptable medium of email.

Balance is the decisive thing. Even if I had no neck problem at all, I still would disconnect from email. As I've mentioned before, every healthy community needs to breathe both inward and outward, between the cultivation of its own life on the one hand, and intercourse with the wider society on the other. Similarly, every individual must hold a balance between focused pursuit of his own work and openness to whatever may come from without.

But there is no denying it: in our society today, the centrifugal, interrupting, distracting tendencies have gotten out of control, jerking us around with almost demonic violence. The way television's intrusion has re-shaped household schedules around programs, fragments of programs, and advertisements; the way we are tempted to click through the endless and chaotic doorways presented to us on the web; the vastly greater ease with which one person can approach another through email — an approach we can initiate with less weight of significance or personal presence than before — in these and many other symptoms you will recognize the contemporary forces that would throw us off our own foundations. What chance do we have to breathe harmoniously between a centred, reflective devotion to our own ever-deepening work (upon which society depends), and thoughtful response to the needs approaching us from the outside (without which we as individuals wither and become selfish)?

I admit to some remaining twinges of discomfort over my decision. I have, after all, made a few very good friends over the Net, and there are, additionally, many casual contacts that nevertheless seem freighted with the kind of significance one cannot easily dismiss. I am unilaterally forcing some of these people into a different, perhaps less "convenient" means of communication, and this feels a little arbitrary.

But if our frenetic and de-centred culture invites us to abandon ourselves, we need to recognize that the most critical and fateful abandonment of self typically arises from the abandonment of those nearest us, beginning with our own families. (I am not proud of my record in this regard.) Conversely, to abandon oneself is to abandon society, except insofar as society is a collection of sleepwalkers and mechanisms that run by themselves, without need of self-recollected individuals who can take its concerns conscientiously in hand.

So I guess my decision can be seen as a vote for self-recollection. I am happy to make it a public vote as long you do not take me to be suggesting what *your* appropriate means of self-recollection might be. Coming to ourselves so that we are useful to others — if there is any process that is a matter of individual insight and choice, surely this is it! It is a matter of rousing yourself to discover who *you* are. A lot of the pathology of the modern society has to do with displacing this responsibility of the wakeful self onto social machinery, which we allow to carry us where it will.

For me, it just happens that I came to the point where I needed to stand firm within myself and say, in one particular regard, "Stop!" I am sure you will understand.

(Go ahead, make my day: Send me an email! Actually, humour aside, I'll always be delighted for you to do that, and will try to see that any business you have gets taken care of. As a bonus, you'll get to see my spiffy new automatic-response message.)

*[You may redistribute this newsletter for non-commercial purposes. You may also redistribute individual articles in their entirety, provided the NetFuture url (*http://www.netfuture.org/ *) and this paragraph are attached.]*

Steve Talbott is the author of The Future Does Not Compute: Transcending the Machines in Our Midst and Editor of NetFuture, where this article originally appeared.

Mozilla update

(David Hallowell)

It has been a while since I have written about the Mozilla project. At the time of writing, the next release will be 0.9.5. A lot has happened since my last article: a series of rewrites have taken place to improve performance with image rendering, caching as well as various performance improvements in the mail application.

Netscape released version 6.1 [1a] of their browser which was based on the Mozilla 0.9.2 branch, this represented a vast improvement over Netscape 6.0 which most people considered was no where near ready for release, Mozilla is becoming a very stable browser for most people with most crasher bugs getting fixed quickly. Netscape 6.2 is scheduled for release in the near future. This is going to be a fairly minor release compared to 6.1 with most of the improvements being in the mail client. In the longer term Netscape have a release code-named MachV [1b] planned (which will probably be called Netscape 6.5) this release will probably be based on Mozilla 1.0 and many new features are planned including a built in download manager (so all the downloads are contained within one window), better printing support (e.g. fit to page), print preview, full support for the Platform for Privacy Preferences (P3P), full screen support and improvements to bookmark management.

The Mozilla project is taking longer than many of us expected back in 1998 to produce a finished product, most of the reasons I mentioned in my previous article, however the definition of 1.0 is a subjective manner, many people use Mozilla regularly as their daily browser already and consider it a release quality product, other people will still not consider Mozilla ready for prime time even when the 1.0 release is out. The only way to know for certain is download a build and try it out.

The main problem I see with Mozilla is that it requires a lot of memory, it is often the application on my system that takes up the majority of memory. This makes Mozilla unsuitable for low end machines, but with memory prices at an all time low this is not as severe a problem as it could have been. One of the main reasons that memory footprint is so high is because Mozilla uses its own toolkit (XPtoolkit — XP meaning "cross platform" no relation to Microsoft's latest OS) the main reason XPtoolkit was created was to make the codebase across platforms to be as similar as possible so applications could be developed using the Mozilla framework without having to write platform specific user interface code.

Mozilla.org are currently working on ways to lower the barrier to entry for contributions. One of the latest ideas is a tool called 'patch maker' [2] this is useful for people who would like to work on the user interface (which is written in XML) but do not want to download the entire source code just to submit patches. Also planned is a 'bug week' which is due to be held late October/early November — it is an IRC based event where people who are experienced with contributing to the Mozilla project will spend more time than usual on IRC helping people who are interested in getting involved make their first steps. More information on this event will be published on the MozillaZine [3] site when the details have been finalised. If you are interested in contributing then it is worth joining IRC at sometime during bug week to ask any questions you may have, even if you miss bug week there is likely to be someone able to help on irc.mozilla.org. The default Mozilla startpage (http://www.mozilla.org/start/) also provides information to people interested in getting started with the project.

Mozilla is still considered a Netscape project by many people. In some senses this is fair — the Mozilla project was created by Netscape, originally based on the source code of what was to be Netscape 5 but the source code was scrapped and Mozilla is an almost total rewrite bearing very little code from previous Netscape releases. Netscape are also the largest single commercial contributor to the project so in many senses they do have a lot of control in the direction the project heads. However, as this is an open source project if anyone is unhappy with the direction they consider the project to be heading they

are free to fork the project, most of the time I do not see any problems with Netscape's involvement with Mozilla, there is the occasional time when disagreements occur, but this happens in many open source projects, the number of independent contributors is increasing and there is also some commercial interest in Mozilla, people who say that Netscape is the only contributor to the Mozilla project are basing that comment on old facts.

The Mozilla project also has many lessons to corporations who want to adapt an open source method of development. The Mozilla project made many mistakes, but also got a lot of things right, there is no time here to look back at what lessons can be learnt; maybe that is the subject for another article.

The future looks bright for the Mozilla project, it is now at the stage that although it would be a massive blow if AOL decided to cease Netscape's involvement with the project there would still be enough volunteer effort to keep the project alive. Mozilla is not for everyone, the excessive memory footprint is enough to put many people off (if anyone has time to improve memory footprint then I am sure your efforts will be appreciated), however it is my browser of choice and the fact that I can also use it when I have to use Windows for some reason is also a bonus. Linux finally has some decent choice in the browser market, at one time it was just Netscape 4, Amaya (the W3C testbed browser) or Lynx. Now in addition to Mozilla we have Konqueror [4] which is the browser that comes with KDE and is definitely a serious consideration if you use KDE as your desktop environment, Opera [5] is a small, fast browser but personally I think it has the worst user interface of any mainstream browser I have ever tried, but other people like the interface so it is a matter of preference — Opera is a commercial product but you can download a free ad-ware version. Then you have various browsers that are based on the Mozilla rendering engine (Gecko)[6] but use a native user-interface such as Galeon [7], there is even competition to Lynx on the text browser front with Links [8] which is meant to offer many improvements over Lynx I have included two text based web browsers on this months UKUUG CD.

[1a] http://home.netscape.com/browsers/6/
[1b] http://www.mozilla.org/xpapps/MachVPlan/
[2] http://www.mozilla.org/hacking/patch-maker/
[3] http://www.mozillazine.org/
[4] http://www.kde.org/
[5] http://www.opera.com/
[6] http://www.mozilla.org/projects/embedding/
[7] http://galeon.sourceforge.net/
[8] http://artax.karlin.mff.cuni.cz/~mikulas/links/

Forthcoming Events

UKUUG LISA/WINTER CONFERENCE 2002

13 – 14 February 2002

UKUUG is holding its next Winter Conference in London on 13 and 14 February 2002. The UKUUG Winter Conference is historically an event where not only prominent topics are discussed within the Conference presentations but also where members and friends meet, learn, and enjoy lively debate on a host of subjects.

We wish to encourage discussion on all aspects of systems and their administration and are especially interested in papers covering theory and practice, high-availability, performance, network management, novel solutions to practical problems, integration, interoperability, and security (including the business, legal and moral issues).

You can see details of the last winter conference at http://www.ukuug.org/events/newcastle_2001/.

The event will take the form of a series of presentations, each followed by a discussion on issues raised. To this end, papers are invited from interested parties on the general themes, and related topics.

As always, we wish to encourage work-in-progress presentations, and student project posters; proposals for these should be submitted in the same way as for full papers.

You do not have to be a member of UKUUG to submit a paper. Submissions from speakers from outside of the UK are welcome.

You can see details of the last winter conference at http://www.ukuug.org/events/newcastle_2001/.

Significant Dates: **Closing date for abstracts**: 10 November 2001, **Authors notified**: 17 November 2001, **Programme published**: end November 2001, **Final papers due**: 12 January 2001.

Initial abstracts should be sent as e-mail to winterconf@ukuug.org. Abstracts should be accompanied by a short biography, and, ideally, should be about 250-500 words long. Final papers should normally last 30-40 minutes. If you need more time for your presentation, please tell us when you submit your abstract. Potential authors may request further information via winterconf@ukuug.org

SANE Conference

May 27 – 31

Maastricht, The Netherlands

Technology is advancing, the systems administration profession is changing rapidly, and you have to master new skills to keep apace. At the **Third International SANE** (System Administration and Networking) conference you'll find a wealth of opportunities to meet other system administrators and network (security) professionals with similar interests, while attending a program that brings you the latest in tools, techniques, security and networking. Complete program and registration information will be available in December 2001. If you need assistance, sane2002-info@nluug.nl. Announcement and Call for Papers (http://www.nluug.nl/events/sane2002/CfP.html) .

2001

- October 30-November 1 Linux Expo Canada (http://www.linuxexpocanada.com/) Toronto
- November 6-10 Annual Linux Showcase (http://www.linuxshowcase.org/) Oakland
- November 9-11 BSDCon Europe 2001 (http://www.bsdconeurope.org) , Brighton
- December 2-7 15th Systems Administration Conference (LISA 2001) (http://www.usenix.org/events/lisa2001/) San Diego, CA

2002

- January 28-29 File Storage and Technologies Conference (http://www.usenix.org/events/fast/) Monterey, California
- January 30-February 1 Linux Expo Paris (http://www.linuxexpoparis.com/) Paris, France
- February 13-14 UKUUG Winter Conference (http://www.ukuug.org/events/winter2002) London
- February 18-22 NordU2002 (http://www.nordu.org/NordU2002/) Helsinki
- February 19-20 Linux Expo Amsterdam (http://www.linuxexpoamsterdam.com/)
- May 27-31 3rd International SANE Conference (http://www.nluug.nl/events/sane2002/) Maastricht
- June 9-14 2002 USENIX Technical Conference (http://www.usenix.org/events/usenix02/) Monterey, CA, USA

Contacts

Council

Charles Curran (Council Chairman)
Phone: +44 1865 273 269
E-mail: charles.curran@ukuug.org
News from Council · UKUUG Open Source Award

James Youngman (UKUUG Treasurer)
E-mail: james.youngman@ukuug.org
Shell Programming Puzzle

Alasdair Kergon (Events)
E-mail: alasdair.kergon@ukuug.org

David Hallowell (Website)
E-mail: david.hallowell@ukuug.org
CD Number 16 · Mozilla update

Roger Whittaker (Schools)
E-mail: roger.whittaker@ukuug.org

Dr A V LeBlanc
E-mail: owen.leblanc@ukuug.org

Contributors

Alan McLintock
E-mail: alex@OWAL.co.uk
Professional XSL

Bob Vickers
E-mail: bobv@cs.rhul.ac.uk
SSH, The Secure Shell: The Definitive Guide

Josette Garcia
E-mail: josette@oreilly.co.uk
From O'Reilly and Beyond · More News from O'Reilly

Joel Smith
E-mail: joel@firstsoftware.co.uk
The Perl CD Bookshelf · Server Load Balancing

Jan Wysocki
E-mail: jan.wysocki@ukuug.org
Effective awk programming

Lindsay Marshall
E-mail: lindsay.marshall@newcastle.ac.uk
Learning Web Design: A beginner's Guide to HTML, Graphics and Beyond

Matt Palmer

Neil Levine
E-mail: levine@yoyo.org
LPI Linux Certification in a Nutshell

Paul Webb
A Beginner's Guide to HTML, Graphics and Beyond

Richard Ibbotson
E-mail: richard@sheflug.co.uk

Simon Cozens
E-mail: simon@netthink.co.uk
Perl news

Steve Talbott
E-mail: stevet@ora.com
Why I Have Disconnected from Email

Tom Wilkinson
E-mail: tom@dentrassi.net
Linux Developers' Conference 2001: review of some talks

Virantha Mendis
E-mail: virantha.2.mendis@bt.com
SSH, The Secure Shell: The Definitive Guide

Local User Groups

Piete Brooks (Cambridge LUG)
E-mail: postmaster@cl.cam.ac.uk

Andrew Findlay (London LUG)
Phone: +44 1628 782 565
E-mail: andrew.findlay@skills-1st.co.uk

Secretariat

Jane Morrison
UKUUG Secretariat
PO Box 37
Buntingford
Herts
SG9 9UQ
Phone: +44 1763 273 475
Fax: +44 1763 273 255
E-mail: office@ukuug.org
News from the Secretariat · UKUUG AGM