

news@UK

The Newsletter of the UK UNIX Users Group

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

ISSN 0965-9412: Volume 10, Number 4

December 2001

Contents

December CD: Open documentation and UNIX tools for Win32	3
News from the Secretariat	3
Linux 2002 – UKUUG Linux Developers' Conference	4
Shell Programming Puzzle: December Puzzle	4
Shell Programming Puzzle: October Results	5
News from O'Reilly and Beyond	6
Book review: Web caching	9
Book review: Programming C#	10
Book review: C# Essentials	11
Book review: DNS and BIND, 4th Edition	11
Book review: DNS and BIND, 4th Edition	12
Book review: Cisco IOS Access Lists	13
Book review: Oracle PL/SQL Best Practices	14
Book review: Learning XML	16
Report on UKUUG's 2001 Linux Developers Conference	16
Microsoft – Where do we want you to go today?	21
Events in 2002	22
Contacts	23

December CD: Open documentation and UNIX tools for Win32

David Hallowell

This CD has been compiled based on feedback on the UKUUG 'chat' mailing list. The main items on the CD include UNIX tools for the Win32 environment such as Cygwin, and some documentation that's available under a free documentation licence, this includes two O'Reilly Books, Using Samba and Open Sources, the Zope book and various GNU manuals. VNC has been included: it's a useful remote access tool which can be used to access a Windows desktop from any other Windows, UNIX, Mac, or even WinCE device that's running the VNC client. It can also be used for remotely accessing a UNIX desktop from any machine with a VNC client.

One popular suggestion that didn't make it onto this CD was a small bootable Linux distribution that could easily be used as a rescue disk. There was no time to produce such as disk so this suggestion will be considered.

The CD features HTML navigation which makes great use of the LINK element (see October newsletter), therefore the navigation is enhanced if you view these pages in a capable browser. The latest Mozilla release (0.9.6) is included on this CD. Link element support can be enabled by going to View... Show/Hide... Site Navigation Bar. To navigate the CD just open the index.html file in your browser.

We also include Xsnow, a Christmas background for X, and Netscape 4.79 (requested because Netscape 4.x supports a wider range of UNIX platforms than Netscape 6.x or Mozilla). However, Netscape 4.x has poor support for the W3C standards.

Please email me with feedback with what you think about the CD and any suggestions for the future. If you know of any software that you think other UKUUG members may be interested in or if you've written any free software that you'd like to see included on a future UKUUG CD, let me know.

If you'd like to discuss ideas for the CD then it's worth joining our 'chat' mailing list. Details of how to join are on our website; if you have any problems, email office@ukuug.org. The chat mailing list is unmoderated, but it is very low volume. Most suggestions we get about CD content come through this list. Remember this CD is for members, so let me know what you want and we'll do our best to provide it.

News from the Secretariat

Jane Morrison

We have almost finished the details for the Winter Conference which will be held on 13 and 14 February 2002 in London. Full information and booking form should be enclosed with this mailing.

This will be the first event of 2002 and the programme suggests that this will be popular with members.

The venue for the LINUX 2002 Developers conference has now been confirmed and full details appear together with the call for papers below.

Do you need to buy some books? Don't forget that your UKUUG membership subscription allows you to receive a discount of 21.5% on O'Reilly publications – see the UKUUG web site for a list of titles available and prices.

Talking of subscriptions, the annual invoices will be sent out in January, please look out for your invoice and as always prompt payment will be gratefully received!

Unbelievably it is that time of year again and I would like to wish you all a very Happy Christmas and a Peaceful New Year.

The Secretariat will be closed from 20 December to 2 January 2002.

Linux 2002 – UKUUG Linux Developers' Conference

Bristol, 5-7 July 2002 (PROVISIONAL)

CALL FOR PAPERS

Planning for UKUUG's 2002 Linux Developers' Conference is already underway; we are seeking ideas, speakers, and sponsors.

We invite speakers on all types of Linux development to contribute. The programme will cover a variety of subjects, including kernel and desktop development, tools, applications, and networking. Any topic likely to be of interest to Linux developers and enthusiasts will be considered.

The topics presented in recent years have included: ARM Linux, Clustering, CORBA, Package Management, Enterprise Filesystems, Exim, GNOME, I20, Mail Servers, Memory Management, Performance Programming, Samba, Security, SMP, and VMware.

Further details of earlier conferences can be seen at:

<http://www.ukuug.org/events/linux2000/>,

<http://www.ukuug.org/events/linux1999/>, and of course at

<http://www.ukuug.org/events/linux2001/>

Abstracts should be emailed to the conference organisers at linux2002@ukuug.org. They should be accompanied by a short biography, and, ideally, should be about 250-500 words long. Final papers should normally last about 40 minutes, including 10 minutes for questions and answers. If you need more time for your presentation, please tell us when you submit your abstract. We shall acknowledge all submissions.

Significant Dates:

Closing date for abstracts: 1 March 2002

Accepted authors notified by: 22 March 2002

Final papers due by: 21 May 2002

Particular queries should be sent either to the UKUUG office, or to the Linux 2002 mailing list linux2002@ukuug.org.

To keep the conference fees low, we are seeking sponsors and exhibitors. For further information about sponsoring, exhibiting, or attending the event please contact the office – details below.

Information about the event will be updated on a regular basis on the website:

<http://www.ukuug.org/events/linux2002/>.

Shell Programming Puzzle: December Puzzle

James Youngman

In an attempt to be seasonal, this issue's challenge is to write a script which draws a picture of a Christmas tree; there may be up to two winners:

1. The shortest entry (by byte count) which produces a believable image of a Christmas tree
2. The most aesthetically pleasing rendition of a Christmas tree in the opinion of the judges.

No restriction on output format is made, the output may be ASCII, curses, PNG, PostScript, etc. Also, no restriction is made in respect of the scripting language used; the only restriction is

that the program must work on a system which is not connected to the Internet. The judges will make an effort to get all entries to work (apparently-not-working entries however cannot win, so if in doubt supply instructions!). To avoid file format problems, please supply your script uuencoded.

Entries should be emailed to puzzle@ukuug.org. The award of an actual prize will occur if a sponsor is found to provide one.

Shell Programming Puzzle: October Results

James Youngman

The winning entry is that from Bob Vickers; although it was not the shortest entry, at 216 bytes it was the shortest correct entry. Bob's entry is also unique in the sense that it was generated by a program (also supplied) from a version of the program which included comments.

Thanks to a generous offer from O'Reilly, Bob gets a prize of either an O'Reilly book or a T-shirt.

Special mention goes to Stuart Mitchell, whose impressive 89-byte entry was the shortest; the single problem with Stuart's entry was that it was unfortunately fooled by two test files which had the same checksum.

The task was to produce a program which reads a list of filenames from stdin and emits a list of groups, one group per line, where all the members of each group are identical to each other. An example input and output was provided in the previous newsletter.

The input data was largely pedestrian (files containing 'hello', empty files, files containing white space and so on) apart from the pair of files which have the same checksum, which were as follows:

```
$ od -c 514 ; od -c 260
0000000 002 002  \0  \0
0000004
0000000 004 001  \0  \0
0000004
```

Winning Program – Bob Vickers

```
#!/bin/sh
while read x
do
r="$r $x"
done
while [ -n "$r" ]
do
set - $r
r=""
a=$1
shift
if [ -n "$1" ]
then
s=""
for f in $*

```

```
do
if cmp -s $a $f
then
s="$s $f"
else
r="$r $f"
fi
done
[ -n "$s" ] && echo $a $s
fi
done
```

Runner-Up – Stuart Mitchell

```
sum `cat`|sort|awk '{x[$1]=x[$1]" "$3}END{for(i in x){print x[i]}}' \
|awk 'NF>1{print $0}'
```

My attempt

Since I couldn't resist the problem myself, I had a go and came up with the 120-byte program below. Its main flaw is that it 'poisons the well' rather, creating in the current directory a file called ' ' (because this is shorter than '/dev/null').

```
:
z(){
d=
a=$1;shift;g=$a\ ;for x
do
diff $a $x>\ &&g=$g\ $x||d=$d\ $x
done
echo $g|grep \ ;[ -z "$d" ]||z $d
}
z `cat`
```

Input of cunning from Olly Betts provided the second entry which used only 104 bytes but produces a spurious diagnostic from diff.

```
(sum `cat`|sort|while read f g h;do diff $o $h>\ ||echo
echo -n \ $h;o=$h;done;echo)|sed 's/./;/ /pid'
```

News from O'Reilly and Beyond

Josette Garcia

W3C patent policy opposed by O'Reilly

The Web's success depends on fully open standards, implemented without restrictions, according to O'Reilly & Associates' official comment on the W3C's proposed Patent Policy Framework:

<http://www.oreillynet.com/cs/weblog/view/wlg/761>

Writing for O'Reilly

We're always on the look-out for new authors. To find out what our ideal requirements are, topic areas of interest to us, and style info go to:

<http://www.oreilly.com/oreilly/author/intro.html>

Bioinformatics: why biologists want to program computers

Bioinformatics is fast-rising, and scientists are finding new ways to explore their data through programming. James Tisdall, author of O'Reilly's upcoming release 'Beginning Perl for Bioinformatics' explains the benefits of biologists learning to program, in this article:

<http://www.oreilly.com/news/perlbio1001.html>

Bioinformatics is an avenue seasoned programmers can take to expand on their skills. O'Reilly is offering a Bioinformatics Conference this January—more to come on this. We've also released 'Developing Bioinformatics Computer Skills', a book that Perl and C# programmers in particular may want to take a look at.

Java: learning command objects and RMI

William Grosso introduces the basic ideas behind command objects and remote method invocations (RMI) in this ONJava.com article.

<http://www.onjava.com/pub/a/onjava/2001/10/17/rmi.html>

Will you see open source J2EE implementations?

O'Reilly & Associates' Executive Java Editor Mike Loukides gives us his take on the current state of open source enterprise Java and the hurdles open source projects like JBoss experience. He also wonders about Sun's current open source licensing policies regarding J2EE.

<http://www.onjava.com/pub/a/onjava/2001/10/10/osjava.html>

An interview with Richard Monson-Haefel

Richard Monson-Haefel is the author of just released 'Enterprise Javabeans, 3rd Edition'

http://java.oreilly.com/news/ejb_interview_1001.html

Mac: with REALbasic, object-oriented programming is easy

Object-oriented programming is useful, powerful, and fun; and REALbasic is a great way to learn it, says Matt Neuburg, author of 'REALbasic: The Definitive Guide, 2nd Edition'.

http://mac.oreilly.com/news/realbasic2_1001.html

.NET: Hailstorm in hand

Mark Lucovsky, chief architect of Microsoft's .NET My Services initiative (formerly Hailstorm), explains what's available in the software developer kit just released at Microsoft's Professional Developers Conference.

<http://www.oreillynet.com/pub/a/dotnet/2001/10/23/myservices.html>

Open Source: Perl debugging for beginners

Debugging is like car repair: a novice might pull out the transmission to fix a broken light. Here are ten Perl debugging tips that will save you hours of frustration, by John Callender, author of O'Reilly's upcoming 'Perl for Web Site Management'.

http://web.oreilly.com/news/perlmanagement2_1001.html

An interview with Philip Hazel, the creator of Exim

Exim is an open source mail transport agent that has been steadily growing in popularity, and Philip talks about why he developed it and what to expect in the next version. Philip is also the author of O'Reilly's 'Exim: The Mail Transfer Agent'.

http://sysadmin.oreilly.com/news/eximinterview_1001.html

apachemonth.com

A new service called ApacheMonth <http://www.apachemonth.com/>, is being started by the team that runs LinuxMonth, <http://www.linuxmonth.com/>.

This is a heads-up if you enjoy the style of LinuxMonth—Q&A tips, articles, etc.

O'Reilly wins award for Linux books

'Most Indispensable Linux Book' awarded to O'Reilly's 'Linux in a Nutshell, 3rd Edition' and 'Running Linux, 3rd Edition', by LinuxJournal Reader's Choice Awards:

<http://www2.linuxjournal.com/cgi-bin/frames.pl/ljissues/issue91/5441.html>

If you voted... thanks!

Palm: Ten common Palm OS programming pitfalls

the title of this article speaks for itself...

http://palmas.oreilly.com/news/palamosprog_1001.html

Monitoring applications with Palm OS and Tomcat

Learn how to build a Palm Query Application that communicates through a Tomcat server to monitor and manage mission-critical systems.

<http://www.onjava.com/pub/a/onjava/2001/10/05/monitor.html>

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

Wireless: Oddball wireless devices

How can you resist at least trying a text pager that wants to be a cell phone, or a PDA that doubles as a 900-MHz cordless phone? We couldn't! Here's what we discovered.

<http://oreillynet.com/pub/a/wireless/2001/10/19/oddball.html>

Five steps to adding physics-based realism to your games

If you want to spice up your computer games, check out this overview of the major steps you must take to develop physics-based simulations, by David M. Bourg, author of O'Reilly's upcoming 'Physics for Game Developers'.

<http://linux.oreillynet.com/pub/a/linux/2001/11/01/physics.html>

An interview with Tim O'Reilly

In this interview with La Fnac, France's largest book and music retailer, Tim O'Reilly talks about next-generation operating systems, Linux in embedded systems, how open source can play a kingmaker role in the battle between Java and .NET, and more.

<http://www.o>

Getting Started with LDAP

The second article on LDAP (Lightweight Directory Access Protocol) has sparked a lot of interest... In his first article, Luke Kanies explained LDAP. Now Luke shows you how to set up a basic LDAP directory to store Unix user accounts, along with a script to pull those accounts to a Unix system:

<http://linux.oreillynet.com/pub/a/linux/2001/11/08/ldap.html>

David Pogue on the real Mac OS X

Mac OS X is unlike anything Mac users have ever seen before, and the latest Missing Manual, which covers version 10.1, will help you understand Mac OS X's Unix roots, says David Pogue, author of O'Reilly's upcoming 'Mac OS X: The Missing Manual'.

<http://www.oreillynet.com/pub/a/mac/2001/11/09/missingmanual.html>

Josette can be contacted at josette@oreilly.co.uk

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

Web caching

Duane Wessels

O'Reilly and Associates

ISBN 1-56592-536-X

318 pp.

£ 39.95

reviewed by Andrew Cormack

As one of the developers of the harvest and squid cache programs, Duane Wessels would be ideally qualified to write a book on the details of running a web cache. However this is not that book. Instead he has taken several steps back from the detail to review the overall design of caching systems and how they fit into a typical network. The content is addressed as much to network managers, web server administrators, content providers, and users as to those who will administer the cache itself.

Caches are not the most familiar of topics, so there is some basic information that all readers will need to know. The first chapter therefore explains the components of the World Wide Web and how caches fit in. There is an introduction to the protocols used by browsers, caches, and servers with a discussion of the advantages and disadvantages of caches in different situations. Issues of privacy, censorship, and copyright often arise in any discussion of caching. Although the specific legislation discussed in the book applies only to the USA the general principles are universal; caching is one area where common sense seems to be bringing laws closer together in any case.

Caches need not work alone and are often combined for performance or reliability. Cache hierarchies tend to be used to improve hit rates or speed of access but these raise both technical and political issues. There is a discussion of the various protocols that caches can use to co-operate and, in appendices, full details of these including the best description I have found of how Bloom filters work. Within an organisation caches can be clustered to increase capacity or reliability but here there is only a discussion of possibilities with no technical detail to help in making the cluster work.

These and other design considerations are brought together in a chapter discussing how to go about buying or building a cache; the list of products in the first chapter may also be of use to those who have reached this stage.

A cache needs to have traffic directed to it and there are a number of ways of achieving this. Conceptually the simplest is to configure browsers to use the cache rather than a direct connection to the web server and there are good instructions for doing this either manually or using an auto-configuration file. Alternatively, network devices can simply intercept web requests and direct them to caches. This has a number of problems, both politically and at the protocol level, but for those who wish to follow this route there are descriptions of what needs to be done and instructions for constructing a suitable device using Linux, FreeBSD or other operating systems.

At the other end of the web connection, servers can do some things to help caches and a lot to hinder them. The caching headers in HTTP/1.1 are all described, though it might be less intimidating to describe the function of the headers before listing all of them rather than after. The reasons for returning correct headers are stated clearly along with instructions for ensuring that the Apache web server does so.

Cache performance is often a concern so the instructions for monitoring, using ucd-snmp and rrdtool, and for benchmarking are very valuable. Benchmarking is complicated and it is all too easy to measure the performance of the test rig rather than the cache itself. An appendix includes useful data on the behaviour of a cache in the real world, but these are taken from a parent cache so show rather different behaviour from a typical first level caches.

Many cache products have good operating manuals; this book's wider perspective should help

to avoid expensive or inconvenient mistakes at other stages in the process. It provides a good general introduction to caching and will also provide helpful guidance through the process of designing and obtaining a caching solution. If all web sites followed the suggestions on cache-friendliness then the web could be a much faster and more efficient place.

Andrew Cormack was responsible for the first large web cache in Wales (so far as he knows) and now worries about how caches can be abused to steal bandwidth. He is Andrew@Cormacks.org

Programming C#

Jesse Liberty

O'Reilly and Associates

ISBN 0-596-00117-7

680 pp.

£ 39.95

reviewed by Drew Durkin

The latest programming language to be released from Redmond has been with us for just over a year now. Primarily C# was released to be used in the development of .NET applications, which provides for a new application programming interface to the services of Windows and support for COM+ and XML.

C# has the potential to be used for development on UNIX and there are already projects underway to bring the .NET framework across. However, with this in mind, .NET is mainly going to be used on Microsoft operating systems for the short-term future at least.

The book places the language plainly in the context of the Microsoft .NET platform in the development of desktop and web applications. If you're looking for a book on how to get started developing .NET for UNIX then unfortunately this isn't it. If you're someone who currently develops for Windows, you'll find the book a good tutorial with lots of example C# code to illustrate points. Content is structured into three main areas. The first allowing the reader to get to grips with the fundamental aspects of the language, including a traditional Hello World example to start you off. The section then leads on to discuss areas ranging from classes and objects to events and delegates.

If you've already had programming experience of C++ or even Java, you'll notice the similarities to C# almost immediately, and it will help a great deal while learning the language. The section read quite well and left me with a good understanding of the basics and the organisation of content allowed me to quite easily flick to a specific chapter when I needed to refer back in the later sections.

The second section discusses application development, both on the desktop and for the Internet. Again, it reads quite well, starting you off designing simple user interfaces for the desktop before leading to into ADO.NET including explaining common database terminology. The web development chapters also touch on the use of XML, SOAP, and passport.com, but on a whole seem to lack the depth of content provided in the earlier desktop development chapters.

The more advanced concepts are explained with some good example material in the third section. This section certainly is comprehensive in its explanation of some quite complex features including data streams, ActiveX controls, and threading.

The book definitely takes a no-nonsense approach to explaining C# and where appropriate, the .NET framework. Although it hasn't convinced me to drop everything and start developing Windows applications using .NET yet, it certainly makes good reading for someone who perhaps develops using ASP, COM, and ADO and who wishes to learn more about what Microsoft is pushing as the "Next Big Thing".

Overall, I'd say the book is a valuable resource to refer to for anyone learning, or wanting to learn more about C#.

C# Essentials

Ben Albahari, Peter Drayton and Brad Merrill

O'Reilly and Associates

ISBN 0-596-00079-0

216 pp.

£ 13.95

reviewed by Peter Waller

O'Reilly has attempted to produce a concise and compact text that teaches anyone wishing to learn Microsoft's new C# language. This book is by no means a tutorial to teach C# as a first language, so I suggest you look elsewhere. But will this book teach the competent programmer the necessities to help them get to grips with C#? Let's find out shall we? At first glance, C# looks just like C++ with a hint of Java; I found all the samples very familiar and quickly wrote several programs of my own using Microsoft's C# SDK (which you may download from www.microsoft.com for free, although I believe it only supports WINNT based systems.) You are right to assume that some knowledge of Object Oriented Programming is required for you to understand the language. The book itself does a good job at feeding you the syntax; I found all the explanations very clear and concise. I don't think there was anything left unanswered. The structure of the book is simple, after a quick introduction of the language; the book quickly dives into the language reference, then teaches you about the .NET Framework and then an overview of the Base Class Library. After the appendices you'll find a complete index which is usually a let down in books of this kind. I'm certainly not slow to admit that this book is a worthy buy, C# looks like an exciting new language that builds on other O-O language's, and with the .NET Framework you can quickly build and deploy web based applications as well as traditional client/server programs. Any programmer wishing to get into this new language may want to pick up this book for definite. Priced at 13.95 it's certainly cheap and doesn't take up half your shelf space as other massive volumes tend to. If you're at all interested in learning a great new Object-Oriented language, then I recommend you give this book a go. O'Reilly also have another book that gives the .NET Framework the same treatment, this may be a worthy edition so you can use Microsoft's new .NET technology to its full potential.

Peter can be contacted at sulphur@ukscene.org.

DNS and BIND, 4th Edition

Paul Albitz and Cricket Liu

O'Reilly and Associates

ISBN 0-596-00158-4

622 pp.

£ 31.95

reviewed by Joel Smith

This book has been the bible for DNS administration since 1992. It goes without saying that this is the book to help you set up and administer DNS servers. This fourth edition brings it up to date and it now covers BIND 9.1.0 and 8.2.3, occasionally mentioning the older 4.9 versions. The book as a whole has been updated with examples changed to cover the more recent versions of BIND. It also has a fair amount of new information, primarily in the areas of Advanced Fea-

tures and Security, but also covers accommodating Windows 2000 clients, servers and Domain Controllers within a BIND DNS domain.

There is more extensive coverage of dynamic update and NOTIFY (zone change notification), including signed dynamic updates and BIND 9's new update-policy mechanism. Incremental zone transfer, and forward zones (which support conditional forwarding) are introduced, as well as support for IPv6 forward and reverse mapping using the new A6 and DNAME records.

The security section now covers TSIG or transaction signatures which is a new mechanism for authenticating transactions. There is an expanded section on securing name servers and expanded coverage of dealing with firewalls. There is also a section dealing with the DNS Security Extensions (DNSSEC), which allows DNS servers verify their authority using public key cryptography. I am sure that this will become far more important as attackers on the internet find ever more imaginative ways of interrupting service.

The section on dealing with Windows 2000 is useful because although Microsoft has now decided to drop WINS in favour of DNS, Win2000 clients and DHCP servers have a nasty habit of deleting address records owned by the same domain name as the clients or servers. For example, a Win2000 client booting up with the same name as one of the pre-existing DNS records would attempt to delete the conflicting record and replace it with its own. If the name in question happened to be one of your servers, this could cause quite a disruption. (For some reason, I can't help thinking about Invasion of the Body Snatchers!) Thankfully there are several suggested workarounds.

I can't fault this new edition of the book. The first edition served me well when I was setting up my first DNS server. The book still achieves what it sets out to do, and explains DNS and BIND. This has got more complicated (sorry, feature rich!) over the years, but this book still explains it in clear terms. O'Reilly rightly made their name through publishing titles like this.

As ever, Administrators starting to set up their own DNS servers would be well advised to have a copy of this book. Similarly if you are considering the migration to BIND 9, it would be useful to have a copy of this new edition to explore the new features. Whether it is worth replacing your older copy will be down to your budget and your needs.

Joel Smith is the Senior Unix Administrator for IFOonline.com

DNS and BIND, 4th Edition

Paul Albitz and Cricket Liu

O'Reilly and Associates

ISBN 0-596-00158-4

576 pp.

£ 31.95

reviewed by Raza Rizvi

Coming in at two hundred pages more than the first edition, this worthy tome includes information on versions of BIND up to v9.1, going as far back as 4.8.3 which is still included in some vendor implementations. It has been a long time coming, the previous edition was issued in 1998, and many a DNS administrator has attempted to struggle with BIND 9 and cursed the lack of available documentation.

I will concentrate on the newer parts of the book as it carries the same chapters (albeit with updated content) as the previous editions - why, what, and where DNS fits in, setting up DNS name servers, Mail services using MX records, host configurations, day to day operations, sub-domains, DNS tools, debugging, and troubleshooting.

Chapter 10 covers 'Advanced Features' starting with the use of access lists. These detail the permitted (and occasionally list the restricted) sources of access. Their first application is in DNS

Dynamic Update, a process which permits the authorised modification of a set of records on an authoritative name server, in accordance with RFC 2136. Although this is of most importance to DHCP implementations striving to get DNS to reflect the DHCP lease pool, it can also be used to, say, modify a single record for a load balancing application.

In order that a server with changed data can reflect the new information on the secondary servers, there is now a notification mechanism which also supports incremental zone transfer. Known bogus name servers can be ignored (although there is still no BIND mechanism to prevent zone spoofing).

A good explanation of system tuning leads to the final topic in this mish-mash of a chapter, IPv6 – which is supported directly by BIND v9.

Chapter 11 details the security changes, primarily TSIG, which (from BIND v8.2) allows the use of Transaction Signatures to authenticate DNS messages. Security Extensions (DNSSEC) additionally permit the secure exchange of keys using public key cryptography.

Usefully, the rest of chapter 11 covers how to minimise the risks caused by unauthorised zone transfers, or indeed queries, together with specific firewall issues. The splitting of DNS function for serving and resolving is covered in detail, including a section on particular configurations for internal DNS roots.

I mentioned dynamic update earlier, and you will no doubt be pleased that Windows 2000 clients, servers, and domain controllers use this feature heavily. Well, the authors have included information towards the back part of the book on how to live in such an environment. There in fact an other O'Reilly text that deals exclusively with DNS and Windows issues.

This book is as useful now as it was back in the mid 90's. Buy it if you have to do any more than be a simple user of DNS. As a measure of how times change, the appendices no longer show you how to compile and install BIND on a Sun operating system, it is now shown with Linux.

Raza Rizvi is Technical Manager at REDNET. He only edits zone files from time to time these days, but he does remember to change serial number.

Cisco IOS Access Lists

Jeff Sedayao

O'Reilly and Associates

ISBN 1-56592-385-5

260 pp.

£ 39.95

reviewed by Raza Rizvi

Access lists are by their nature both important and somewhat dull to read about. This book certainly does nothing to improve the latter, with reader concentration being required right from the start. There are a number of occasions when concepts are suddenly introduced without explanation, or with explanation occurring later on in the text, e.g. Metrics, and the application of Access Lists to an interface. Even global configuration mode doesn't feature until page 41, and given that this book will surely be used by novice IOS engineers, it would not have been indulgent to explain that a '?' will invoke the IOS help feature!

Certain very useful concepts seem to feature in a throwaway fashion, with very little time spent on them – the most notable early one being named access lists, which scrapes by with only 1 page.

There are some good tips included within the book, often highlighted or picked out from the main text, for example, the 'IOS feature' which corrects an entered IP address supplied with a

mask and stores instead the start address of the network block. This can be hard to understand, and dangerous in a configuration, unless you understand that it occurs.

Within chapter three, there is a good walkthrough of a setup of access lists for a public-facing web server, followed by a description of inbound access lists, which of course have already been described – a situation which is only saved by contrasting their use with outbound access lists, together with a well crafted piece on reflexive access lists. This access list type opens inbound holes for valid outbound traffic, akin to a poor man's stateful inspection firewall.

Chapter 4 is devoted entirely to the use of access lists in route filtering, a tricky subject, and one prone to error in the real world. The examples are accurate but this is very unusual mechanism for manipulating routes, and to his credit the author does end the chapter with alternatives mechanisms for route selection, although without the emphasis they would have received in a general purpose routing volume.

A good chapter on access list debugging is followed by a useful chapter on route maps, with internal routing protocols but in particular with those associated with BGP. Examples are based on weights, Multi-Exit Discriminators (MEDs), and communities, with the debugging theme returning towards the end of this chapter 6, again dealing with BGP.

The final chapter is a set of three case studies, based on a multi-site WAN, a firewall, and an Internet routing. All come with detailed topology maps and IOS config dumps.

The latter chapters save this book, although it does strike me that it may perhaps have been a long time in development and then was rushed out to appear alongside the other recent access list titles. There is too little time spent on IOS features of recent years (named access lists), and complete missing areas (remark options, time based access lists, context based access lists, IPSEC, and queuing). Even the 'burro' on the front cover looks reminiscent of an earlier Cisco based O'Reilly publication!

So in conclusion, one for IOS buffs only, as an adjunct to other texts, and for the case study material.

Raza Rizvi is Technical Manager at REDNET, a business ISP, and Cisco Premier Partner, based in High Wycombe. He is not recruiting.

Oracle PL/SQL Best Practices

Steven Feuerstein

O'Reilly and Associates

ISBN 0-596-00121-5

202 pp.

£ 19.95

reviewed by James Youngman

Oracle PL/SQL Best Practices is a good length for a computer book - 182 pages. Too many computer books these days have many pages and little value.

This book aims to provide a concentrated set of 'best practices' for PL/SQL programming. The book includes about 120 separate rules. There is a section at the back which just lists the headlines (i.e. the statements of the recommendations themselves). There is also – buried in the Preface where only reviewers will see it – a short-list of the eight most important rules.

The main body of the book is divided into sections: The Development Process, Coding Style and Conventions, Variables and Data Structures, Control Structures, Exception Handling, Writing SQL in PL/SQL, Program Construction, and Built-in Packages.

Each section discusses the best practice in each area and provides a small number of short rules for best practice for that area. Each rule is accompanied by an Example section showing

the use of the rule, a Benefits section explaining why this is a good idea, a ‘Challenges’ section which outlines the challenges involved in successfully applying the rule, and a Resources section pointing the reader to related resources (not every rule has one of each section). The structure is well laid out and the text is clear.

Having said this, the book has one unavoidable but significant problem. The people who most need to read it are those who are least inclined. That is, the people who are likely to buy the book are those who have a strong interest in improving the quality of the software they write, and will certainly already have adopted many of these best practices. I’m not suggesting that these people should be seen as contrasted with the rest of the PL/SQL-using world, but the fact is that best-practice recommendations only hit home to people who have been bitten. People don’t rave over this sort of stuff unless they are already enlightened about the importance of the material. A similar effect exists with similar books relating to other languages.

The other potential problem with a work like this is that the techniques only work well when everybody on your team follows them. Many of the recommendations are things like ‘adopt a common approach to ...’ which is something that one simply can’t do unilaterally. So, don’t buy a copy of this book; instead, buy 12 and make everybody on your team read it – only that way will you get the full benefit of the best practice.

Once you begin to put the recommendations into use, you’ll rapidly wish you had a concise list of all the recommendations. One is provided in the back of the book, but the inevitable fate of that section is to be photocopied and pinned to a vertical surface.

I trust that I will not be spoiling the surprise if I whet your appetite by reproducing here – in order – the top 8 recommendations (but naturally, not their discussion, examples and associated discussion) :-

SQL-00: Establish and follow clear rules for how to write SQL in your application (see chapter 6).

MOD-01: Encapsulate and name business rules and formulas behind function headers (see chapter 7).

EXC-00: Set guidelines for application-wide error handling before you start coding (see chapter 5).

PKG-02: Provide well-defined interfaces to business data and functional manipulation using packages (see Chapter 8).

MOD-03: Limit execution section sizes to a single page using modularisation (see Chapter 7).

DAT-15: Expose package globals using ‘get and set’ modules (see Chapter 3).

DEV-03: Walk through each other’s code (see Chapter 1).

STYL-09: Comment tersely with value-added information (see chapter 2).

These are the headline items, and so will inevitably seem a little facile; many of the remaining 112 rules are detail-packed and thought-provoking. The code examples from the book are available at <http://www.oreilly.com/catalog/orbestprac/>.

James has spent the last 18 months designing and implementing a large system based mostly on PL/SQL. Many of the detailed practices here are no longer news to him but some of them were discovered the hard way.

Learning XML

Erik T Ray

O'Reilly

ISBN 0-596-00046-4

368 pp.

£ 34.95

reviewed by Mark Ashley Jones

Learning XML provides a good gentle introduction to the subject matter, is full of good worked examples using a clear referencing scheme to breakdown complex XML documents. There is an excellent Taxonomy of standards section which allows the book to be used as a pointer to other relevant text/URLs.

The Introduction gives you a flavour of just how powerful XML can be and makes you want to read on to find out more. It also explains the history of XML, who is responsible for standards, etc, and the types of applications the author would recommend to view, test, and parse documents. The markup and core concepts chapter provide a good solid understanding which allows you to read an XML document within a relatively short time. It explains the syntax of a well formed document and how DTDs can be used to validate content.

The chapter on connecting resources with links is less interesting than the previous two chapters which are quite humourous, this may be down to the subject matter. However, the chapter does seem to cram a lot of information into a few pages. The next chapter on CSS explains how Cascading Style Sheets are used to provide a range of different display options. It has a good explanation of how to manipulate display attributes.

The DTD chapter makes effective use of the Barebones DocBook DTD standard to illustrate how a good DTD design can make a document easier for an XML developer to work with. It also gives examples of how a bad DTD design can reduce productivity.

The Transformation section is full of good examples, but is a little long and hard to follow. The short chapter on character subsets provides a fascinating aside to the development of Unicode and character subsets such as UTF-8. This was very relevant to the project work I have recently undertaken in XML-RPC.

The last chapter was probably worthy of a complete book in its own right and there are many texts devoted the subject area. The author made a good job of condensing the subject material into a single chapter and the author was not afraid to provide his own opinions on the relative merits of XML protocols such as SAX and SOAP.

Overall the book was an excellent introduction to XML and I am sure that I will explore some of the recommend further reading, since the author clearly knows his stuff.

Dr Mark Ashley Jones is an experienced software engineer, who has a PhD in Nuclear Physics and is currently working for a top software consultancy.

Report on UKUUG's 2001 Linux Developers Conference

Richard Ibbotson

The UKUUG developers conference was held at UMIST in Manchester this year. A great deal of thought was put into this year's venue and it paid off. Many people attended. Everyone said that they enjoyed themselves. The UKUUG say that about 170 people attended. The main sponsors were IBM and Sistina. Charles Curran started the fourth developers conference session with his usual welcome to the delegates. He kept it short and we quickly got on with the business

of discussing some of the more recent developments to do with GNU/Linux. There was much interest in the kind of things that were to take place over the next three days. This writer sat down in his favourite spot next to the mains socket with his notebook and proceeded to make notes about anything that moved or didn't move. Or . . . should I have painted it if it did move ?

David McAllister who is a director of Strategic Technologies gave us the intro to the rest of it with his talk on Linux Clusters in the XSP environment. That's big iron to you and me. His main problem was the jet lag he was suffering from. He had just arrived from San Francisco and it was 5 a.m. in the morning where he came from. In spite of that he gave us a good talk about the commercial side of this topic. His talk was based around the fact that with the increased commercial use of the internet more and more companies are coming to rely on the net. A new approach is required to deal with this in the shape of integrated hardware and software and networking. Rob Gautier came along next to chat on brute force clusters . His talk was about Grendel which is a Beowulf cluster that is used by the University of Wales at Aberystwyth by the Bioinformatics Research Group. Most of the work consists of large database searches that can take days to complete. The design and management of the cluster is fairly labour intensive and it took Rob a long time to describe it. Next up was Nick Davis. He gave us the commercial stuff that IBM are very much interested in just now. IBM are doing their utmost to support GNU/Linux and to educate their own employees to work with free and open source software. One of the gizzits that came with the conference was an IBM information pack with the title Partner World for Developers. Nick gave us recent and future developments at IBM. He explained that the open source movement is no longer driven by just one company. It is now an industry wide phenomenon. There are some very big companies ordering large Linux networks from IBM. Linux is now a well established and well understood part of the world of e-commerce. IBM are presently investing what can be described as a complete government budget in Linux.

Bo Thorsen's talk about the porting of Linux to 64 bit systems might well be thought to be the highlight of the whole three days. Bo works at the SuSE labs and so he should know what he's on about. He did say 'don't underestimate what's going to happen with 64 bit Linux in the next five years'. So, keep your eyes on www.x86-64.org if you want to know more. His excellent presentation was well received by the crowd and we even heard laughter as well. Hopefully we'll see much more of him at some time in the future. Andrea Arcangeli came along next and gave us his stuff on vsyscalls. Andrea also works at the SuSE labs and his main thing is kernel developing. Having made a conscious decision to port Linux over to the X86 64 bit architecture the task in hand just now is to get it to work. Andrea discussed this at some length in front of us. His talk greatly complimented Bo's talk in many ways and also gave practical insight into something that is a closed world for most of us.

Professor Frank Sumner got up next and gave us the birth of computing in Manchester. This was completely incredible from beginning to end. Frank explained in graphic detail how computers got to be invented and built and how the dramatic British attempts to join in with the technology race were so hard to sustain in pre-internet times. In spite of that a small team of highly academic people managed to do some amazing things back in the 1950's and 60's. Following his talk we were taken over to the Manchester museum of Science and Industry where the Baby which was one of the original computers was demonstrated. Friday finished with a trip to the Yang Sing restaurant where pre-dinner drinks were followed up with a Chinese banquet. This was the highlight of the day. It was well organised and everyone enjoyed it.

On Saturday morning when most of us were still suffering from the previous nights over indulgence Professor David Aspinall gave us the second part about the birth of computing in Manchester. This was an incredibly fascinating presentation and the kind of thing that would easily make a compete TV program or series on it's own without help. He explained the rest of the Manchester computing project to us. It was quite an amazing talk and it finished with a

deafening applause. Continuing with the shock'em dead approach to Saturday morning Werner Heuser gave us his first presentation of the day which was about the use of Linux on a laptop. This did make quite a favourable impression on the author who wrote this report on a laptop using SuSE 7.2 and Star Office. There are quite a few people out there who haven't yet discovered that Linux can be used on a desktop and so Werner's presentation did draw quite a lot of interest from the crowd. Almost all of us could hear an imagined and open-mouthed deafening silence from the direction of Redmond.

Luke Leighton got up next and gave us Samba the next generation. Luke says he'd like to get people interested in open source software and the distributed computing environment and remote procedure calls. He'd like to get them really mad so that they go to their computer and start coding right away. He thinks that open source software is well out of date when it comes to DCE and RPC. The crowd in the lecture theatre looked as though they actually agreed with him every bit of the way. Perhaps a bit less rage and a bit more use of the keyboard would be helpful. Chris Benson from Treepax got stuck into DCE and why Linux developers should know about it. He gave us a brief history of DCE the service that it provides and the current availability on Linux. The Open Software Foundation was founded in 1998 as a balance to the UNIX International Consortium which was set up by Sun and AT&T. One of the strong points of DCE is integration.

At this point the conference split between two different buildings. Some of us went over to the KDE talk by David Faure at Stafford House. Others stayed behind in the Renold Building and listened to Chris Benson from Treepax talk about Distributed Computing Environment. By now John Southern and myself were pretty much confused. What to do ?

David Faures' presentation on KDE architecture and development was well worth taking in. Even some of the KDE developers we have spoken to like to watch him in action. His explanation of how the whole things works and how the developers work together was easy to understand for anyone and even those of us who are not developers can understand what he said. At the same time Martin Hamilton gave us Heartbeat and the JANET virtual server. He works as a developer on the JANET web cache service. Martins' talk was all about how to build scalable and reliable internet services. As we all know our Universities don't have much money and any attempt at producing a low cost and reliable service internet service for staff and students is to be welcomed. The Linux Virtual Server project provides for these conditions and combined with heartbeat all of the usual problems to do with traffic handling and load balancing can be quickly sorted out in a way that helps everyone. After Martin David Woodhouse came along and explained the journaling flash file system. David works for Red Hat UK where he is responsible for kernel drivers for memory technology devices. The journaling flash file system is all about the work that Axis Communications which is a Swiss based company have come up with in the shape of JFSS. The talk was about the limitations of flash technology and how JFSS can be implemented to improve on those limitations.

A wander across the road again revealed that . . . Michael Meeks was giving a presentation on the GNOME architecture and development. Something that most of us think he should do more of. Christian Egli was also given a chance to explain rapid application development with GNOME and Python. This quite honestly made Borland Kylix look a bit childish or silly. For those of you who haven't yet seen Michael give a public explanation of the way that GNOME and Bonobo components work should try to be there for the next one. He's very good at taking an extremely complex piece of technology and explaining in a few simple words how all of it works. Even people who are not anything to do with the world of information technology can understand what he's on about. Christian Eglis' chat about how to use GNOME came after Michaels. He created an application called gPizza for the overly large crowd that was in the room. If you have Glade or Anjuta and Kdevelop it's very difficult to find other tools that are just as flexible and with the same range of possibilities. gPizza worked first time without a hitch. Back over at

the Renolds building Andrea Arcangeli of the SuSE labs gave us an excellent talk on NUMA. What is NUMA, I hear you say? Well, it's a technology that opens a new class of optimisations so that the full power of a machine can be utilised. Andrea showed us the optimisations that are available on wildfire alpha boxes. He also showed us Zerocopy O_DIRECT. This is something that does massive I/O stuff to storage devices in the filesystem. There are many advantages in doing things in this which are too many to mention in this small space.

Steve Whitehouse came along from Sistina software and gave us Zerocopy. This further enhanced our understanding of Andreas work. Steve is the DECNet kernel code maintainer and recently completed a PhD on error resilient image compression at the University of Cambridge engineering department. He runs his own consulting company which is ChyGwyn Limited and runs the European Sistina office full time. Steve explained the recent improvements made in the GNU/Linux kernel under the 2.4.x kernel releases. There are many sides to this explanation too many to list here.

Going back to the Stafford building once again we found Tom Rathbone giving an explanation about fractal generation using GIMP MathMap. This is quite an eye opener for many highly trained photographers who deal in the daily use of digital photography. If only more of them were able to use GIMP rather than the well known and highly expensive Adobe Photoshop. MathMap is the scientific end of image manipulation. The bit that happens before photographers and commercial artists get to use GIMP. It's a GIMP plug in that allows a user to describe image transforms and combinations. It's very fast and only requires a reasonably good understanding of maths. It's also work in progress. Tom describes himself as a recreational mathematician. He also holds a B.Math in computing science. His understanding of GIMP is quite incredible. Luke Leighton turned up in what looked like a happy smiling mood. This is good to see considering the kind of things that he's been through recently. He gave us his Pymmetry – Python Trust Metrics presentation. Pymmetry has a maximum flow model at the centre of it. This is used to evaluate trust metrics which are cascading or hierarchical access control lists. This provides for an easy way to integrate trust metrics into any project. Examples are DNS protection or web site access rights.

The close of the day on Saturday was seen after Werner Heuser showed us some GNU/Linux on PDAs. This pulled in quite a crowd at the Renold building. The other presentations in the Stafford building had finished and people preferred to stay and watch Werner rather than go home. He reviewed the present situation as an introduction and then went on to show us some examples. The most notable was the use of GNU/Linux on the Compaq iPaq. Werner pulled out his own iPaq only to find that several other people in the auditorium had then as well. Even a friend of mine from Bolton had one in his hands next to me. The Debian distribution seems to be a favourite with these machines and their users. This was followed by a friendly question and answer session where everyone discussed the events of the day and possible future events.

Sunday turned out to be a typical sleepy British Sunday. The rail lines over the Pennines had been removed for maintenance and this made things even more sleepy. A few of us arrived by bus.

Randy Appleton started the day in the Renold building with benchmarking the GNU/Linux kernel. Three students and a professor spent a complete term benchmarking nineteen different versions of the Linux kernel from 2.01 to 2.4. The results show that many functions did improve in speed and others didn't. Randy showed that some improvements were in fact a step backwards and also that the total number of lines of code has now tripled on an exponential growth curve. Next Dave Jones from the SuSE Labs attempted to break us out of our slumber with tuning the Linux kernel with Powertweak. David has recently completed a degree in computing science at the University of Glamorgan. He spends all his time hacking kernel code and anything else that might be of interest to him. Powertweak is the first performance-enhancing tool for GNU/Linux. David's talk covered what there is in the present implementation.

Meanwhile back over in the Stafford building Stephan Richter got started on his talk about Zope which is the open source web application server. There is a very steep learning curve before a user can understand Zope. The end results justify the initial pain in the head. Stephan discussed Zopes architecture and it's ability to be scalable and reliable at all times. He was originally from Dresden and is now studying physics and chemistry in Memphis Tennessee.

Tom Rathbone made a welcome return at this point to explain how to simplify dynamic web sites with mod_perl and the GIMP. This was one of the highlights of the three days for those of us who like graphics rather than darkened rooms and sober people. He described a template system for mod_perl which is part of the Apache application which is used on most of the servers around the world.

Over in the Renold building Richard Moore arrived in a great rush explaining that he'd only just got back from Boston where he had been for about an hour. In spite of his sleepiness Richard was able to give his talk on dynamic probes for user and kernel space and showed us that he was in fact very much awake. Dprobes as he calls it is a generic and pervasive debugging system that can operate under the most extreme conditions with minimal system disruption. He followed on from this one with a talk about GKHI which is a generalised kernel hooks interface. This allows for many enhancements that otherwise would not be possible. Richard's presentation did drag on for a long time but it did contain many aspects of kernel design and use that most people don't get to see at all. Richard originally graduated from the University of Kingston-upon-Hull where he took a degree in pure mathematics. As a final session Alex Perry gave us the Flight Gear flight simulator. This was a kind of a cross between a top class joystick game and a flight simulator for 747s or commercial aircraft. He showed us some Californian airports and also some in built air corridors. One of the features are the virtually real instrument displays which closely mimic the real instruments that are to be found on commercial aircraft. The question and answer session at the end of this one was fairly intense and even witty and humorous in places. Alex very kindly gave us a demo of the flight simulator at the end of his talk. Alex originally took his second degree at the University of Cambridge in England. In the present day he works at Quantum Magnetics in San Diego.

After a general discussion about past and future events we trooped off downstairs to lunch where a large pile of sandwiches were consumed in just a few minutes.

We now look forward to next year which as all of us will be aware will not be anything like as boring and uninformed as all of the other 'technical seminars' always are.

Useful URLs

Linux Virtual Server

<http://www.lvsp.org/>

Robert Gautier – Design and Management of a Brute-Force Cluster

<http://www.ateb.co.uk/packages/farmer/>

Bo Thorsen – Porting Linux to 64-bit Architecture

<http://www.x86-64.org/>

Andrea Arcangeli – SuSE Labs

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

David Faure

<http://www.kde.org/>

Tom Rathbone

<http://www.aceldama.com/~tomr/papers/2001/web-gimp/>

Richard Ibbotson is the Chairman and Organiser for Sheffield Linux User's Group. You can view their web site at <http://www.sheflug.co.uk/>.

Microsoft – Where do we want you to go today?

Anonymous

There has been a lot of mentions in the news over the past few years about various viruses that have affected e-mail as well as things such as the Code Red worm which affected web servers. One thing that was not mentioned in many news broadcasts is that these in general just affected Microsoft products. For some reason many people just accept these viruses and worms as part of everyday life, even though they are costing companies millions. This rant is not about what I think about the quality of Microsoft products it's more about what rights I think users of Microsoft products should have.

Unlike Linux, Microsoft operating systems are not free and in a big corporation the cost of buying licences for each machine can become very expensive. However when you buy such licences you agree not to hold Microsoft liable for the quality of software that's why Microsoft has never been successfully sued over the quality of their software. However, in this day and age it seems ridiculous that people can get thousands of pounds compensation if they trip over and get a minor bruise but can't get a penny if the software that they paid good money for ends up costing them or their company a fortune.

You may be saying that it's the buyers fault for choosing to accept the Microsoft licence but what choice do most people have? Most applications are written for Windows, most documents are sent in MS Word format and most people are trained in a Windows environment so switching to an alternative and more reliable solution is often not an option. Because of anti-competitive business practices Microsoft was found guilty of being an illegal monopoly and therefore is proof that the vast majority of the computing world relies on Windows (although the courts later relaxed Microsoft's punishment they were still found guilty). Therefore Microsoft's licence agreements should be made illegal meaning that if anyone suffers financial loss due to negligence in Microsoft products they'd have the right to sue Microsoft for this loss. This would mean Microsoft would have to exercise higher quality control in their products.

The law would have to be fair, Microsoft can't be blamed for something if they've already released a freely available patch to fix the bug. It is then the system administrators responsibility to make sure all patches are applied. These patches would have to be somewhere easily accessible such as the Windows Update site. Microsoft's liability would then be restricted to the period before the patch was available plus a small length of time in order for administrators to deploy the patch through the organisation.

There are some of us who don't run Microsoft software unless we have to and even these people can still get affected by Microsoft's negligence in an indirect manner. For example, although Linux and the Apache web server are immune to Code Red many Linux users were still getting many attempted attacks from Code Red infected Windows machines, although they couldn't infect the Linux server it was wasting bandwidth which on a commercial web hosting server may result in the company losing money as not as many people would be able to connect to the server simultaneously. It would also increase your bandwidth bills if your hosting company charged for the amount of bandwidth used. Personally I can see a case where a Microsoft-free company who has never agreed to Microsoft's unfair licences could sue Microsoft for their negligence affecting their business. However, I'm not a legal expert so I don't know if that is feasible. I think most admins in Microsoft-free organisations were just happy that they suffered far less problems than those who used Microsoft solutions.

Events in 2002

January 28-29

File Storage and Technologies Conference, Monterey, California

<http://www.usenix.org/events/fast/>

January 30-February 1

Linux Expo Paris

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

February 13-14

UKUUG Winter Conference, London

<http://www.ukuug.org/events/winter2002/>

February 11-14

Usenix BSDCon 2002, San Francisco

<http://www.usenix.org/events/bsdcon02/>

February 16-17

Fosdem, Brussels

<http://www.fosdem.org/>

February 18-22

NordU2002, Helsinki

<http://www.nordu.org/NordU2002/>

February 19-20

Linux Expo Amsterdam

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

May 27-31

3rd International SANE Conference, Maastricht

<http://www.nluug.nl/events/sane2002/>

June 9-14

USENIX 2002 Technical Conference, Monterey, CA, USA

<http://www.usenix.org/events/usenix02/>

July 5-7 (Provisional)

UKUUG Linux Developers Conference, Bristol

<http://www.ukuug.org/events/linux2002>

November 3-8

Usenix LISA'02 Systems Administration Conference, Philadelphia

<http://www.usenix.org/events/lisa02/>

December 9-11

Usenix Symposium on Operating Systems Design

<http://www.usenix.org/events/osdi02/>

Contacts

Council

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

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

David Hallowell
Website
Tyne
david.hallowell@ukuug.org

Alasdair Kergon
Events
Reading
alasdair.kergon@ukuug.org

Dr A V LeBlanc
Newsletter
Manchester
owen.leblanc@ukuug.org

Roger Whittaker
Schools; Newsletter
Borehamwood
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