

# news@UK

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

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

---

**Volume 15, Number 1**

**ISSN 0965-9412**

**March 2006**

---

## **Contents**

<b>News from the Secretariat</b>	<b>3</b>
<b>UKUUG/Apple Technical Briefing</b>	<b>3</b>
<b>AUUG 2006 call for papers</b>	<b>4</b>
<b>EuroOSCON 2006</b>	<b>4</b>
<b>FSFE Press release: Microsoft compliance</b>	<b>4</b>
<b>FSFE Press Release: "Microsoft offers poisoned apple"</b>	<b>5</b>
<b>Free Knowledge Foundation announcement</b>	<b>6</b>
<b>From the UKUUG Diary</b>	<b>7</b>
<b>Proactive wireless networks with hostapd</b>	<b>9</b>
<b>Introducing the Template Toolkit Part 2</b>	<b>11</b>
<b>All-Party Parliamentary Open Source Group – what is happening now</b>	<b>16</b>
<b>Report on CISL II, Malaga, Feb. 2006</b>	<b>19</b>
<b>Book review: UML 2.0 in a Nutshell</b>	<b>21</b>
<b>Book review: JBoss: A Developer's Notebook</b>	<b>22</b>
<b>Book review: Essential PHP Security</b>	<b>22</b>
<b>Book review: Digital Identity</b>	<b>23</b>
<b>Book review: Learning SQL</b>	<b>25</b>
<b>Book review: XML Pocket Reference</b>	<b>25</b>
<b>Book review: XSLT Pocket Reference</b>	<b>26</b>
<b>Book review: Classic Shell Scripting</b>	<b>26</b>
<b>Contributors</b>	<b>27</b>
<b>Credits</b>	<b>28</b>
<b>Contacts</b>	<b>29</b>



---

## News from the Secretariat

**Jane Morrison**

Thank you to everyone who has kindly sent in their subscription payments so promptly. We have received a good number of early payments. Those remaining outstanding will be chased this month and any not paid at the end of March will not receive the next issue (June) Newsletter.

The UKUUG Spring conference is being held at Durham University on 21st, 22nd and 23rd March and advance bookings are coming in quickly. Have you booked your place yet? Bookings are open until Thursday 16th March.

At this time we are also working on a joint event with Apple. Please check the web site for up-to-date information.

The Linux 2006 Conference will be held in Brighton this year. We have confirmed the venue, University of Sussex, and the dates: Friday 30th June to Sunday 2nd July. Please put these dates in your diary. A call for papers is enclosed in the same mailing as this newsletter.

As in previous years we are hoping to be able to achieve sponsorship for the event which will allow us to keep delegate fees at a minimum. If you know of any company who may be interested in sponsorship please let me know.

The next Newsletter will be the June issue and the copy date is Friday 19th May.

For any comments about past or future events, or if you have something to say about the UKUUG or this Newsletter please contact

[newsletter@ukuug.org](mailto:newsletter@ukuug.org)

---

## UKUUG/Apple Technical Briefing

**Sam Smith**

We are happy to announce another UKUUG/Apple Technical Briefing, which will take place in London on the 20th April 2006. We are pleased to welcome Eric Albert to speak at this event. At time of going to press, additional speakers and other details are unconfirmed, but will be announced on the web site:

<http://www.ukuug.org/events/apple06>

For over five years, Mac OS X has combined the power of UNIX and the ease of the Macintosh for users around the world. This January, Apple shipped Mac OS X for Intel processors for the first time. The entire Macintosh product line will transition to Intel processors by the end of 2006.

Eric Albert will discuss:

- The challenges of migrating an entire operating system to a new architecture
- How the Intel transition affects scientific computing on Mac OS X
- How to make your own software run natively on both PowerPC and Intel-based Macs
- Rosetta, the translation engine which enables existing PowerPC software to run on Intel-based Macs
- Intel's published processor roadmap and how it may affect the future of personal and scientific computing

### **About the main speaker**

Eric Albert is a Senior Software Engineer in Apple's Core Technologies group. During Mac OS X's transition to Intel processors he has worked on almost every part of the operating system, from the OS kernel and compiler tools to applications like iPhoto and Keynote. Prior to working on bringing Mac OS X to Intel, he worked for Microsoft on the shared-source .NET project ("Rotor") at Microsoft and for Apple's Java team.

This event is aimed at all UNIX users, particularly those interested in its core technologies and the associated development environments. The UKUUG is grateful to Eric Albert for agreeing to speak, Massimo Marino and Alan Bennett at Apple for providing the main speaker, venue and publicity; along with a raft of other support.

---

### **AUUG 2006 call for papers**

We have received notification from Lawrie Brown of the AUUG 2006 conference which will be held at the Rendezvous (formerly Duxton) Hotel in Melbourne during October 2006.

The call for papers for this event is available at:

[http://new.auug.org.au/events\\_store/auug2006/cfp](http://new.auug.org.au/events_store/auug2006/cfp)

---

### **EuroOSCON 2006**

We have received the following announcement regarding the EuroOSCON event which will take place in September.

The Call for Participation for EuroOSCON 2006 is now open:

[http://conferences.oreillynet.com/cs/euos2006/create/e\\_sess/](http://conferences.oreillynet.com/cs/euos2006/create/e_sess/)

The O'Reilly European Open Source Convention is where coders, sys admins, entrepreneurs, and business people working in free and open source software gather to share ideas, discover code, and find solutions. At EuroOSCON 2005, nearly 500 delegates took part in sessions and tutorials across eleven technology tracks, learning about the newest features and versions from creators and experts. We anticipate that EuroOSCON 2006 will be even more successful – the place for the open source community to meet up, debate, make deals, and connect face to face with other open source enthusiasts from across the continent and around the world. EuroOSCON 2006 will take place at the Hotel Le Plaza in Brussels, Belgium, 18-21 September 2006.

For more details, see:

<http://conferences.oreillynet.com/euos2006/>

---

### **FSFE Press release: Microsoft compliance**

#### **FSFE to Microsoft: stop complaining, and start complying!**

"After several years of investigation, the original ruling in 2004, and a European Court case lasting close to two years, we now have to conclude that Microsoft never had any intention to comply with the antitrust ruling," comments Georg Greve, president of the Free Software Foundation Europe (FSFE). "We were forced to witness years of delays, stalling and playing for more

---

time during which Microsoft has made no attempt to allow interoperability and competition with its competitors, including Free Software such as Samba.”

“It makes a very sad and worrisome statement about the quality and reliability of software engineering at Microsoft if they indeed do not have proper interface specifications and documentation for their considerably complex and organic systems, as Microsoft have repeatedly maintained,” explains Jonas Öberg, vice-president of FSFE. “That ‘hundreds of Microsoft employees and contractors’ [1] were apparently not able to produce this documentation in almost two years further undermines the confidence in Microsoft’s technological prowess.”

“That Microsoft now questions the competency and integrity of a Trustee they themselves helped appoint is outrageous. All the parties involved in the case found the Trustee showed an intimate competence and understanding,” says Carlo Piana, the lawyer representing FSFE on the case: “How do they now dare maintain that he is biased? Because he is not willing to lie against all evidence? We support the Commission all the way, and possibly beyond. And: If Microsoft wanted to know how it could easily comply, they just had to ring Volker Lendecke, member of the Samba Team, one of our technical experts, as we offered them”.

“Microsoft has behaved as if they consider themselves above the law and any decision by the European Commission. The aggressive stance they now take towards a Commission that was unbelievably patient with Microsoft further confirms that view,” concludes Georg Greve. “Microsoft strove very hard to be the first company to leave the European Commission no choice but to impose daily fines for the first time in European antitrust history. Microsoft deserves to be granted what they worked for so vehemently and be brought to compliance the hard way.”

Reference:

<http://www.microsoft.com/presspass/press/2006/feb06/02-15EUStatement.msp>

### **About the Free Software Foundation Europe**

The Free Software Foundation Europe (FSF Europe) is a charitable non-governmental organisation dedicated to all aspects of Free Software in Europe. Access to software determines who may participate in a digital society. Therefore the freedoms to use, copy, modify and redistribute software – as described in the Free Software definition – allow equal participation in the information age. Creating awareness of these issues, securing Free Software politically and legally, and giving people freedom by supporting development of Free Software are central issues of the FSF Europe, which was founded in 2001 as the European sister organisation of the Free Software Foundation in the United States.

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

---

## **FSFE Press Release: “Microsoft offers poisoned apple”**

### **FSFE: “Microsoft offers poisoned apple” – “Praises premature in hindsight”**

The Free Software Foundation Europe (FSFE) critically reviews Microsoft’s offer to publish source code instead of protocols.

Georg Greve, President of FSFE explains: “The European Commission asked Microsoft to publish their protocols to allow others to write interoperable software and to re-establish competition. That does not appear to have happened, at least the terms under which the protocols are supposedly available are not clear. What they have done instead is publish source code, which they weren’t asked to do.”

In a press release [1] Microsoft announced yesterday “to license all the Windows Server source code”. This would underscore Microsoft’s “commitment to resolve compliance issues”.

According to FSFE this would not be of any help, but would cause additional troubles. Carlo

Piana, FSFE's lawyer explains: "This 'commitment' is poisoned – they do it in a way that worsens the situation for Free Software: Because that source code is under Microsoft copyright, developers who have seen the source code cannot reimplement it in Free Software for fear of copyright violation."

Even worse, "Microsoft might try to shut down competition by making claims of copyright infringement based on the argument that the developer 'could' have looked at the source code", Piana forecasts.

Georg Greve adds: "Indeed. How exactly are developers supposed to write interoperable software without looking at the source code if the specifications are not available? If you eat from that apple, you find yourself poisoned by their copyright."

After their welcoming of Microsoft's new 'Shared Source' Licences [2] which were substantially improved over their previous licensing proposals, FSFE was willing to give Microsoft the benefit of the doubt.

"It now seems clear this was just another marketing ploy: Yesterday Microsoft acted like a robber that, when asked to please put away the gun, tosses you a grenade. Our praise seemed premature in hindsight", Greve finishes.

References:

<http://www.microsoft.com/presspass/press/2006/jan06/01-25EUSourceCodePR.msp>

<http://mail.fsfeurope.org/pipermail/press-release/2005q4/000120.html>

---

## Free Knowledge Foundation announcement

### **FKF and FSFE teaming up: FKF official associate organisation of FSF Europe**

The Free Knowledge Foundation / Fundación Conocimiento Libre (FKF) and Free Software Foundation Europe (FSFE) are proud to announce their new official associate status, working together for the promotion and protection of Free Software in Spain.

Following a decision at the last foundation board meeting of FKF to co-operate more closely with FSF Europe, as well as a decision by the members of the FSFE Team, both organisations are delighted to announce their official co-operation. The FKF is the first associate organisation of the FSFE in Spain, a country with a strong Free Software community.

"Our adversaries are usually multi-national groups, funded heavily by multi-national companies. If we are to defend Free Software and all that is related to it against their encroachments, then we need to be equally present," explains Georg Greve, President of Free Software Foundation Europe. "That is why FSFE started out as a multi-national organisation from the outset, building teams across cultural and language barriers to work together to further Free Software. We are very happy to welcome another group to our large family."

"While we are focused on Free Knowledge, Free Software, Free Standards and Intellectual Wealth dissemination, Free Software is at the centre of our work, as the initiating and inspiring phenomenon of this new way of understanding knowledge.", says Pablo Machón, President of FKF. "The FSFE has always been an inspiration for our work. Becoming an associated organisation of FSFE is, thus, a natural consequence of sharing ideals, a long-term vocation, and working together towards the same goal."

**About the Free Knowledge Foundation (FKF)**

The Free Knowledge Foundation is a non-profit organisation based in Spain and focused on Free Knowledge, Free Software, Free Standards and Intellectual Wealth dissemination. The FKF promotes the concept of knowledge as being publicly and freely accessible for both usage and contribution.

<http://www.libre.org>

---

**From the UKUUG Diary**

The UKUUG maintains a web diary of future events of interest at

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

The following events are a small selection of those currently listed.

**FOSS in Business****16th March 2006: Spires, Belfast, Northern Ireland**

Richard Stallman and Bruce Perens will keynote at in Belfast on March 16th. This whole-day event is tailored towards IT decision makers in medium sized businesses and public administration – but the content will have a wide appeal. Among other topics, Richard Stallman will discuss GPL version 3, and will take questions on the current draft and the drafting process. Bruce Perens will deliver a CEO / CIO targetted opening keynote. Other speakers will discuss the practical issues of transitioning to Free Software and there will be opportunities for networking.

<http://belfoss.org/Home>

**Libre Graphics Meeting 2006****17th March 2006: Lyon, France**

The first Libre Graphics Meeting will be held in Lyon, France in the Ecole d'Ingénieurs CPE on the university campus at La Doua, Villeurbanne, from 17th to 19th March 2006. LGM will be a melting pot, with free software graphics developers and artists meeting each other, exchanging ideas and tips, and planning the future of free graphics. Graphics professionals interested in learning about the state of the art in free software are also welcome. LGM will bring together developers and users of the best of free software graphics applications – the GIMP, Inkscape, Scribus, Blender and more. It promises to be a fun ride through what is on offer, with tutorials and presentations of applications, and lots of time and space to chat, meet up with old friends, make new ones. Plans will be made, hacking and drawing will be done, fun will be had.

<http://wiki.gimp.org/gimp/LibreGraphicsMeeting>

**UKUUG Spring Conference 2006****22nd and 23rd March 2006: Durham**

UKUUG's annual Large Installation Systems Administration (LISA) conference will take place in the historic city of Durham on Wednesday 22nd and Thursday 23rd March 2006. This is the UK's only conference aimed specifically at systems and network administrators. It attracts a large number of professionals from sites of all shapes and sizes. As well as the technical talks, the conference provides a friendly environment for delegates to meet, learn, and enjoy lively debate on a host of subjects. This year, we are also planning a series of talks on the BSD family of operating systems to run alongside the main systems administration stream.

<http://www.ukuug.org/events/spring2006/>

**OSS Watch Conference: Open Source and Sustainability****10th April 2006: Said Business School, University of Oxford**

Open source has proved itself as a development and distribution model that can deliver soft-

---

ware which is functional, efficient, innovative, and cost-effective. What is the long-term future? What is the best model for sustainability? Do different models succeed equally well? This conference will look at open source sustainability from the point of view of: strategic funding bodies seeking to maximise their investment through open source development; universities and colleges entangling their infrastructure and pedagogical environments with open source software; software development projects seeking to perpetuate themselves through building communities; businesses basing their sustainability on an open source business model.

<http://www.oss-watch.ac.uk/events/2006-04-10-12/>

### **MySQL Users Conference 2006**

**24th April 2006: Santa Clara, California, USA**

The MySQL Users Conference brings over 1,000 open source and database enthusiasts together to celebrate the huge MySQL ecosystem. The theme for the 2006 conference is "Discover. Connect. Succeed. Scale Your Business with MySQL." The conference mirrors the product's wide-ranging appeal and capabilities, encouraging community members to help create the best database software, gain knowledge, develop tools, and build applications that help grow businesses and nurture individual skills.

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

### **SambaXP**

**24th April 2006: Goettingen, Germany**

From April 24th to 26th 2006, developers and users will meet again in Goettingen, Germany at the fifth international SAMBA conference, the "samba eXPerience 2006". The first four conferences have been very successful events, attracting more than 160 participants from all over the world. The sambaXP is the leading event with focus on the most important free alternative to proprietary SMB/CIFS servers.

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

### **Linux Audio Conference 2006**

**27th April 2006: Institute for Music and Acoustics, Karlsruhe, Germany**

LAC2006 aims at bringing together developers and users of Linux and open source audio software with the target of information sharing, project discussion and music. Highlights: Public Talks and Workshops; Concerts; Internet Cafe; Proceedings; Live Audio/Video Stream. Free admission to all events except the concerts.

<http://lac.zkm.de/2006/>

### **LinuxTag 2006**

**3rd May 2006: Rhein-Main-Hallen, Wiesbaden, Germany**

LinuxTag 2006 is presented in cooperation with the German Unix Users Group (GUUG). As every year since 1996, LinuxTag invites you to take advantage of its unique synthesis: LinuxTag is a complete conference program and a professional exhibition; a unique information platform where .COM meets .ORG. Come to LinuxTag and find out all about the latest developments and trends, direct from leading IT companies and free development projects!

<http://www.linuxtag.org/2006/en/home/aktuelles.html>

### **Debconf6**

**14th May 2006: Oaxtepec, Mexico**

Debian Conference is the annual Debian developers meeting, an event filled with coding parties, discussions and workshops – all of them highly technical in nature. It will be held in Oaxtepec, Mexico, from May 14th to May 22th 2006. Previous Debian Conferences have featured speakers from around the world. They have also been extremely beneficial for developing key Debian software components, including the new Debian Installer, and for improving Debian's

internationalization.

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

## **SANE 2006**

### **15th May 2006: Aula Congresscentre, Delft, The Netherlands**

The SANE 2006 conference offers 3 days of training, followed by a 2-day conference program filled with the latest developments in system administration, network engineering, security and open source software, and practical approaches to the puzzles and problems you wrestle with. You'll also have many opportunities to meet other system administrators and network (security) professionals and chat with peers who share your concerns and interests. The venue for SANE 2006 will be the Aula Congresscentre, located on the campus of the University of Technology in Delft: the city of Delft Blue, the world-famous painter Johannes Vermeer and its historical ties to the Royal House of Orange. But also a lively modern city for funshopping, going out for a great dinner or wandering around to experience its special atmosphere.

<http://www.sane.nl/sane2006/>

---

## **Proactive wireless networks with hostapd**

***Reyk Floeter***

A new daemon, the new OpenBSD hostapd(8), will be included in the upcoming 3.8 release. The daemon helps to improve roaming and monitoring of OpenBSD-based wireless networks. A powerful feature, event rules, will introduce a new kind of proactive wireless monitoring. WLAN is a battle field, and hostapd helps to keep track of your systems, unwanted wireless visitors, and rogue accesspoints.

Read on for some amazing (and possibly controversial) examples of how this new tool can be used.

The OpenBSD hostapd(8) was initially implemented to support the Inter Access Point Protocol (IAPP), the IEEE 802.11f roaming standard for Wireless LANs. I initially developed it during the 21c3, a large annual hacker event in Germany, to improve roaming and station tracking in a network with about 40 OpenBSD-based accesspoints. Because its initial purpose was IAPP only, I intended to just call it "iappd". Admittedly, this was a stupid name. During the 3.7-current release cycle right before the c2k5 hackathon this daemon was imported into the OpenBSD tree. After a discussion with Theo, the name was changed to "hostapd" because we realised that it shouldn't be focussed on a single IEEE protocol; its real purpose being a "HostAP assist daemon" indeed.

hostapd can be used with a single accesspoint, e.g. using a box running ath(4) in hostap mode, or even an entire wireless network with several accesspoints. An accesspoint, typically equipped with a wired and a wireless interface, tracks the state of stations associated to it. Of course, the kernel-level hostap recognizes when a station joins the network but if the station moves away it just disappears for the AP. As a solution, hostapd, which is listening on the wireless interface and the kernel-level station states, will send out notifications to a multicast or broadcast address on the wired interface for every newly associated station. It also listens on the wired interface for messages from other hostapds, and an "ADD.notify" message for a station previously associated to itself instructs hostapd to clean up any resources for this station in the kernel.

Tracking and monitoring of station movements is possible as well. The hostapd messages will be sent immediatly after recent station associations and every listening hostapd will be able to log the movements in the entire wireless network. hostapd will turn into a passive monitoring daemon if the wireless interface is omitted and if it's listening to the wired "IAPP" interface only. This is a typical case for a wireless controller and monitoring server, connected to the same multicast group or in the same broadcast domain. hostapd is designed for highly decentralized

and scalable wireless networks, regardless of how many accesspoints or monitors are actually interconnected. In contrast, most vendors of wireless network solutions are using CAPWAP (Control And Provisioning of Wireless Access Points) based designs with a single, fragile, server controlling all the IEEE 802.11 wireless states for a bunch of dumb WTPs (Wireless Termination Points) instead of real accesspoints.

The most interesting feature of hostapd is the implementation of so-called “Event Rules”. During the c2k5, I started to work on a way to handle specific IEEE 802.11 traffic received on the wireless interface. Human-readable event rules in hostapd.conf(5) similar to rules in pf.conf(5), which emerged to be an OpenBSD-typical configuration style, can be used to trigger an action on specific events in the wireless network. Currently, there are five different actions available: frame, iapp radiotap, log, node, and resend. An example of such a rule is

```
hostapd handle type management subtype beacon bssid !<my_hostaps> \
with log limit 2 sec
```

which locally logs beacon frames from unknown accesspoints, limited to every 2 seconds. Like pf, hostapd provides support for powerful and fast tables to match against multiple addresses in a single rule (based on hashed tables in OpenBSD 3.8, changed to red-black trees in 3.8-current). The previous example uses a table, based on the same syntax as in pf.conf. It has to be declared on top of hostapd.conf, like the following table which defines a list of MAC addresses (BSSIDs) of my accesspoints:

```
table <my_hostaps> { 00:02:6f:21:e0:e9, 00:90:4b:0d:fd:c8 }
```

It’s even possible to use address masks to match against the MAC address vendor part, e.g.

```
table <cizzco-eee> { 00:40:96:00:00:00 & ff:ff:ff:00:00:00 }
```

and to use this table to trigger a specific action

```
hostapd handle type management subtype probe request \
from <cizzco-eee> with log
```

Furthermore, the frame action turns hostapd into a WIPS (Wireless Intrusion Prevention System), because I extended the bpf(4) interface with the ability to send raw IEEE 802.11 frames to the wireless network. This action can be used to affect the wireless management states of any stations or accesspoints in range. By exploiting the weakness of the IEEE 802.11 protocol, a former DoS attack turns into an useful countermeasure against rogue accesspoints in your wireless territory, e.g.

```
hostapd handle type data bssid !<my_hostaps> \
with frame type management subtype deauth reason auth expire \
from &bssid to &from bssid &bssid
```

which sends deauthentication frames to stations sending data frames in unofficial and “rogue” wireless networks.

Besides the log and frame actions, “iapp radiotap” sends the received and successfully matched IEEE 802.11 frame attached to a hostapd-specific message to the IAPP domain; “node” can be used to alter the kernel-level station table (i.e. remove specific associated stations); and “resend” simply re-injects the wireless frame. All the possible scenarios for using these and future actions are hard to determine, for example with “iapp radiotap” hostapd could act as a sensor for a WIDS (Wireless Intrusion Detection System), including radiotap headers with signal levels and all the relevant information for wireless location tracking.

Future work for hostapd is in progress, so expect some new features and improvements. A new “rate” keyword in 3.8-current extends the event rules with the ability to detect flooding against wireless networks - attacks like the void11 wireless DoS. The log action has been extended to print the actual rate if it has been specified. Thus, for the curious, the simple event rule to detect any flooding of non-beacon management frames will be

```
hostapd handle skip type management subtype ! beacon \  
with log rate 100 / 5 sec
```

(The “log” action will be triggered if there are at least 100 management frames, except beacons, within at most 5 seconds.)

All in all, the term “Proactive Wireless Monitoring” has been chosen to describe the advanced abilities of hostapd. It is fully functional and ready-to-use in OpenBSD 3.8; future work will concentrate on extending the IEEE 802.11f IAPP roaming support and event rule matching and action capabilities. Finally, hostapd is small, simple, and designed to be secure with privilege separation, OpenBSD principles, and a clean and reviewed code base.

A dedicated tutorial about wireless support in OpenBSD, including lectures and details about hostapd, was held on the EuroBSDCon 2005 in Basel (Switzerland) on November, 25th.

---

## Introducing the Template Toolkit Part 2

**Dave Cross**

### Using `ttree`

Last time we were processing one template at a time using `tpage`. Often you will want to process a set of associated templates at the same time. The Template Toolkit comes with a utility program called `ttree` that allows you to do just that. It processes a set of files from an input directory and puts the processed versions into an output directory.

There are many other ways that `ttree` is more powerful than `tpage`. It supports a huge number of options that control exactly how your templates are processed. You can get a preview of what all of those options are by typing `ttree -h` or a more detailed description by typing `man ttree`. We’ll cover some of them in more detail later.

### Networking Configuration Files

In this article we’ll create some configuration files that define a simple network. Specifically, we’ll generate an `/etc/hosts` file and some of the files that are required to configure BIND. We’ll be looking at the simple network that which is described below: “The Sunnydale Network”.

### Getting Started with `ttree`

As I mentioned before, `ttree` takes templates from an input directory, processes them, and writes the processed versions to an output directory. Therefore when starting with a new `ttree` project, I like to create a project directory and subdirectories called `in` and `out`.

You can configure `ttree` in a number of ways. The easiest is probably to use a `.ttreerc` file. By default, `ttree` looks for `~/ .ttreerc`, but you can override this by either using the `-f` option or by setting the `$TTREERC` environment variable. As I often like to have slightly different `ttree` configurations for different projects, I set `$TTREERC` to be `./ .ttreerc` and put a new `.ttreerc` file in each project directory.

One nice little touch is that if you run `ttree` and you don’t have a `.ttreerc` file (either in your home directory or in the location defined by `$TTREERC`) then `ttree` will offer to create a sample file for you. This file has most of the common `ttree` options in it together with copious comments that make it easy to edit. We won’t use that this time, we’ll edit our own `.ttreerc` from scratch. Here are the first three lines

```
verbose  
src = ./in  
dest = ./out
```

The second and third lines are pretty self-explanatory. They define the source and destination template directories. The first line puts `ttree` into verbose mode where it tells you everything that it is doing.

With this in place we can test our first template. We'll create a template that expands our network files into a hosts file. Here's the template.

```
#
# /etc/hosts
#
[% USE networks = datafile('data/networks.txt') -%]
[% FOREACH network = networks -%]
# Network: [% network.netname %]
# IP:      [% network.number %]

[% USE hosts = datafile("data/net_${network.netname}.txt") -%]
[% FOREACH host = hosts -%]
[% host.IP %]    [% host.hostname %]    [% host.alias %]
[% END -%]
# End of network [% network.netname %]

[% END -%]

# End of networks
```

If you put this in a file called `hosts` in the `./in` directory, you can then run `ttree` like this

```
$ ttree
```

If all is well, you will see output that looks something like this `ttree 2.75` (Template Toolkit version 2.13)

```
Source: ./in
Destination: ./out
Include Path: [ ]
Ignore: [ ]
Copy: [ ]
Accept: [ ]
Suffix: [ ]

+ hosts
```

You'll see that `ttree` has reported on the source and destination directories that it has processed. It also reports on a number of other options which we will explain soon. At the end of the output you'll see the name of our template file. The '+' sign next to it indicates that the template was processed. `ttree` has laziness built in. It will only process the templates that are necessary. It works out which templates to process by comparing the contents of the source and destination directories. A template is only processed if it doesn't exist in the destination directory or if the source version is more recent than the destination version. In fact it works very much like `make`. If you try to run `ttree` again immediately then the output will be identical to the first run except that the last line will be replaced with

```
- hosts                                (not modified)
```

This indicates that as we haven't updated `hosts` there is no need to process it. If you want to process all of the templates without checking their timestamps, then just give `ttree` the `-a` option.

The template itself is pretty simple. There's not much in it that we didn't cover last time. The only interesting thing is that we are using the `datafile` plugin twice – once to open the mail `networks.txt` file and then again within the loop to open each individual network's data file. In this second usage, we have to use the more explicit syntax `#{network.netname}` to reference the network's name. Without it the parser would have had little chance of realising

what we meant. The Template Toolkit parser can usually work out what you want it to do, but in rare cases (like this one) it needs a bit of a hint.

If you look in the `./out` directory, you will see the results of processing our template which should look like this.

```
#
# /etc/hosts
#
# Network: sunnydale
# IP:      192.168.1/24
192.168.1.1    buffy    slayer
192.168.1.2    willow   witch
192.168.1.3    xander
192.168.1.4    spike
# End of network sunnydale

# Network: los_angeles
# IP:      192.168.2/24
192.168.2.1    angel
192.168.2.2    cordelia    cheerleader
192.168.2.3    wesley
192.168.2.4    spike
# End of network los_angeles

# End of networks
```

### Some more `ttree` options

Let's fill in a few more options in our `.ttreerc` file, so it looks like this.

```
verbose
recurse
ignore = \b(CVS|RCS)\b
ignore = ^#
copy = \.png$
copy = \.gif$
src = ./in
dest = ./out
lib = ./lib
```

The `recurse` option tells `ttree` to look in any subdirectories below your source directory and to recreate the same directory structure under the destination directory. The `ignore` option lists files that should never be processed. The arguments to this option are interpreted as Perl regular expressions and filenames that match the regular expression are ignored. In this example the first `ignore` line matches CVS or RCS thereby removing any sourcecode control files from consideration. The second line matches files that start with `#`, thereby ignoring emacs backup files. If you're a vi user you might like to replace that with `ignore = ~$`.

The next option is `copy`. This lists files that are simply copied from source to destination without processing. Here we are copying `.png` and `.gif` files. The final option defines a library directory. This is an additional directory where `ttree` will look for templates. This is often used to store templates that are included in other templates and that aren't intended to be processed on their own.

Two other useful options that we won't be using in this example are `accept` and `suffix`. `accept` is the opposite of `ignore` as it defines the set of files that will be processed. You normally only need to use one of `ignore` or `accept`. `suffix` gives you a way to change the extension of files as they are being processed. For example you might want to have a standard extension of `.tt` for templates, but convert that to `.txt` for the output files. In that case you could have a `suffix` option which looked like `suffix tt=txt`.

### Creating More Files

The main advantage that `ttree` has over the simpler `tpage` is that it processes a complete

directory of templates in one go. So far our example only processes one template. So let's add another.

Another file that could potentially be derived from our network definition data is a BIND configuration file, so here is a template that could be used to create such a file.

```
[% PROCESS config;
file = 'db.' _ main_domain;
FILTER redirect(file);
PROCESS soa domain=main_domain -%]

[% main_domain %] IN NS [% dns %].[% main_domain %]

; Hosts
[% USE networks = datafile('data/networks.txt');
FOREACH network = networks;
USE hosts = datafile("data/net_${network.netname}.txt");
FOREACH host = hosts -%]
[% host.hostname _ '.' _ main_domain _ '.' | format('%-32s') %] IN A \
[% host.IP %]
[% IF host.alias -%]
[% host.alias _ '.' _ main_domain _ '.' | format('%-32s') %] IN CNAME \
[% host.hostname %].[% main_domain %].
[% END;
END;
END;
END -%]
```

One thing that you'll notice immediately is that because this template uses a lot of directives we have started to combine multiple directives within one tag set. The Template Toolkit parser allows you to do this as long as you separate the directives with semicolons.

This template is far more complex than anything that we've seen before so it's worth going through it in some detail. It starts by processing another template called `config` which is shown below.

```
[% main_domain = 'whedon.example.com'
ttl           = '3h'
dns           = 'xander'
hostmaster   = 'hostmaster.whedon.example.com'
-%]
```

This is a good example of a library template. All it does is defines some variables that we will need elsewhere. We don't want to put it in the source directory as then it will be processed by `ttree` and we will end up with an extra unnecessary output file. Therefore you should put the config template in the `./lib` directory.

The next thing the template does is to create a new variable `file` which contains the name of the required output file. In this example, `file` will get the value `db.whedon.example.com`. We do this because we will eventually want to create a number of BIND configuration files and it will be nice to create them all using the same input template. To actually write the output to the correct file, we use the `redirect` filter. This takes one parameter which is a filename and writes the output from the filter to that file. Everything from the opening `FILTER` directive to the matching `END` will end up in the new file. In this example, the `END` that matches our `FILTER` right at the end of the template, so everything is written to the given file.

The next directive processes another external template called `soa`. This template provides the "start of authority" block for the BIND file. Again, we've created a separate template as we would like to use it from several different templates. The template is shown below.

```
$TTL [% ttl %]
[% domain %]. IN SOA [% dns %].[% domain %]. [% hostmaster %]. (
[% serial %] ; serial
```

```
3h      ; refresh
1h      ; retry
1w      ; expire
1h )    ; caching TTL
```

This is another file which we don't want `ttree` to process, so once again we put it in the `./lib` directory. Notice that the `soa` template uses a variable called `domain` and that this is passed in as a named parameter in the `PROCESS` directive.

The next directive in the template simply adds the NS record to the output file. It uses simple variable expansion that we've seen many times before.

Then we come to the part of the template which creates the A and CNAME records for the hosts on the network. This uses the same kind of logic that we used for the hosts file to loop through the data contained in the various data files and display the correct records. One nice touch is that we use the "format" filter to ensure that the domain name part of the record is always padded to the same length. Here we use the short syntax for the `FILTER` directive where the `FILTER` keyword is replaced by the pipe character (`'|'`). This makes it read a bit like a Unix filter command like `ls -l | sort`.

The template generates an A record for each host and a CNAME record for any aliases.

There's one thing missing from this description of the template. Sharp-eyed readers (and BIND experts) will have noticed that the `soa` template uses a variable called `serial` and that hasn't been defined anywhere. As the serial number needs to be incremented for each version of the configuration file I thought that it was pointless to include it in any of the template files. You could, of course, include it in the "config" template, but you would need to remember to update it each time you processed the template. In my opinion, it's much easier to pass this value on the command line to `ttree` and `ttree` supports the same `--define var=value` option as `tpage` does. You can therefore process both of our templates with a command like this

```
$ ttree --define serial=1
```

And you'll see that both the `hosts` template and the `db` template are processed.

### Other BIND files

Of course one `db` file doesn't make a complete BIND configuration. You'll need to define reverse lookup files for the `1.168.192.in-addr.arpa` and `2.168.192.in-addr.arpa` domains as well as for the loopback domain. You'll also need the actual `named.conf` file that pulls all of these files together. I don't have space to demonstrate creating all of these in this tutorial, but I hope I've given you some ideas on how you might go about it.

### Template Complexity

This month's templates have been a lot more complex than the ones that we saw last month. It would be easy to argue that they were too complex. Part of this complexity comes from trying to do too much in a template. Templates should really only be concerned with presentation logic.

We can simplify our templates significantly if we use Perl to gather the data that we want to display and only use the Template Toolkit language to control how we display that data. We'll look at how you do that in next month's tutorial.

### The Sunnydale Network

Throughout this tutorial we'll be using examples based on a very simple network. The network has two subnets and a single bridge between them.

The first subnet is `192.168.1/24`. It has three computers called `buffy` (`192.168.1.1`), `willow` (`192.168.1.2`) and `xander` (`192.168.1.3`).

The second subnet is `192.168.2/24`. Its three computers are `angel` (`192.168.2.1`), `cordy` (`192.168.2.2`) and `wesley` (`192.168.2.3`). The bridge between the two networks is called `spike` and it has the two IP addresses `192.168.1.4` and `192.168.2.4`.

Data about this network is held in a number of files. `networks.dat` contains details of the subnets and `net_sunnydale.txt` and `net_los_angeles.txt` contain details of the computers on each of the subnets.

The contents of the files are as shown below.

```
# networks.txt
netname : number
sunnydale : 192.168.1/24
los_angeles : 192.168.2/24

# net_sunnydale.txt
IP : hostname : alias
192.168.1.1 : buffy : slayer
192.168.1.2 : willow : witch
192.168.1.3 : xander
192.168.1.4 : spike

# net_los_angeles.txt
IP : hostname : alias
192.168.2.1 : angel
192.168.2.2 : cordelia : cheerleader
192.168.2.3 : wesley
192.168.2.4 : spike
```

As we saw last time, these files are deliberately designed to be in the default format used by the Template Toolkit ‘datafile’ plugin, but it would be equally simple to get the data from other file formats, XML documents or even a database.

### Specifying ttree Options

In this article we have mainly been controlling `ttree` by putting option definitions in the `.ttreerc` file. It’s also possible to pass options on the command line. This can be useful if you want to override a value from your `.ttreerc` for one or two processing runs and it’s not worth the effort to edit `.ttreerc`. You can get a complete list of these options from `ttree -h` but here are a list of the values that we have used in this tutorial.

General options:

```
-a      (--all)           Process all files, regardless of modification
-r      (--recurse)      Recurse into sub-directories
-v      (--verbose)      Verbose mode
-s DIR  (--src=DIR)      Source directory
-d DIR  (--dest=DIR)     Destination directory
-l DIR  (--lib=DIR)      Library directory
-f FILE (--file=FILE)    Read named configuration file
```

File search specifications:

```
--ignore=REGEX          Ignore files matching REGEX
--copy=REGEX            Copy files matching REGEX
--accept=REGEX          Process only files matching REGEX
```

---

## All-Party Parliamentary Open Source Group – what is happening now

*Leslie Fletcher*

An article in December’s Newsletter explained some of the background to the proposed All-Party Parliamentary Open Source Software Group (APOSG); this article brings members up-to-date with what has happened so far and what I hope will be achieved in the next few months.

**What is being aimed at?**

The objective of this initiative is to establish an effective means of communication between the Free/Libre Open Source Software (FLOSS) community and UK legislators. The community has important messages for a parliamentary audience about the benefits to UK business, citizens and the public sector of FLOSS and open standards. Although many MPs appreciate the negative economic and social impact of monopoly profits, there is not much awareness of software lock-in and its disadvantages nor of the practicality of avoiding these by implementing FLOSS. The need to enforce conformance to verifiable open standards in order to protect citizens' rights is another subtle but important message.

However, the communication may not be all one way! At the end of January1 "Government plans for an ambitious £6.2bn IT project designed to revolutionise the NHS were plunged into confusion when iSoft, the software firm charged with supplying systems for half of England, admitted the rollout had been delayed". Some opposition MPs believe this IT project is close to melt down and have challenged the open source community to say how it could help prevent this.

**Approaches to avoid**

At a critical stage in the EU software patents brouhaha, The Register offered this advice to those lobbying against the infamous directive: While [politicians] are a peculiar and varied breed, there are three things you can be fairly certain will not hold much sway with them:

- Ideological argument. Politicians are nothing if not pragmatic. Their very survival is based on seeing which way the wind is blowing and adjusting accordingly.
- Little-man defence. Politicians will not risk upsetting rich and powerful people and companies unless there is a principle at stake: that principle being that the government ultimately decides. Therefore arguing a point on the basis that it will restrict or impair a powerful body is counterproductive.
- Criticism. Politicians do not respond well to criticism. In fact, the more they get, the more stubborn they become. Flattery is the surest route to their heart, and this means making them feel important. Wining and dining, listening, applauding their insight and then putting your point across [sic].

**Working with Open Forum Europe**

Although UKUUG first raised the possibility of APOSG, it is important to the success of the project that it engages as wide a representation as possible of the FLOSS community. A significant step is the relationship which has been established with Open Forum Europe (OFE) with its mission to accelerate, broaden and strengthen the use of FLOSS within business. OFE and its parliamentary lobbyists arranged a meeting in Westminster last December of the Conservative Technology Forum attended by parliamentarians and parliamentary advisors. The meeting heard a talk entitled "Open Source in Government"; although this covered familiar ground, the value of APOSG was endorsed during the discussion which followed.

A follow-up meeting in January was disappointing in that no MPs came. As Basil Cousins of OFE commented afterwards, "It is going to be a long, hard road raising the Open Source and Open Standards issues to become ones of immediate political interest in the UK."

**Meeting with John Pugh MP**

By contrast, John Pugh, Liberal Democrat MP for Southport is very keen for APOSG to come into being, and be a campaigning group. At a long meeting I had with him in January he was very receptive to the concerns of the open source community. He has offered administrative support from his Westminster office, a welcome complement to the support which UKUUG is

garnering from the open source community. He is checking with the Clerks Office what needs to be done to bring the group formally into being.

Many UKUUG members will be aware that Richard Allan, former Liberal Democrat MP for Sheffield Hallam, was a vocal advocate of FLOSS in the previous parliament. He has given a great deal of support and advice on the project so far – and says that the coming of APOSG is one of the few things which could cause him to regret no longer being an MP!

### **Provisional timetable and programme**

An important outcome of my meeting with John Pugh, and of earlier discussions with Richard Allan, is a provisional timetable for setting up APOSG. This sees a launch of the group at a reception-style event in Westminster in late May, followed by a business or investigative meeting in early June.

John Pugh emphasised the importance of careful preparation, particularly for the May launch event. All parliamentarians will be invited, but it is important to identify in advance some individuals who are committed to attending. A couple of speakers should be asked to give an overview of the potential of FLOSS in the public sector. My suggestions are somebody from FLOSSpols3 to give the Europe-wide view and a senior officer from a UK local authority which makes significant use of FLOSS. Demonstrations of working systems would provide a talking point at the launch.

A programme of activities for APOSG needs to be prepared. John Pugh thinks MPs will attach the greatest importance to the savings which could be gained if public sector procurement policies were more "FLOSS-friendly". There are also important implications for UK public policy of interoperability and adherence to common and open standards.

An all-party group works by gathering evidence on a specific topic within its terms of reference and writing a report. Possible topics which APOSG could take up include

- FLOSS in schools – it is through this that I became involved so I feel confident about advising on the issues and identifying suitable expert witnesses.
- FLOSS in the NHS – opposition MPs are keen to know if and how FLOSS might help.
- FLOSS in UK local government – according to authoritative reports, this lags behind other EU countries.
- Government policy on the use of FLOSS in central government – again the UK lags behind.
- Open document formats – some MPs are aware of the intense interest in this from public bodies and campaign groups in the UK and around the world.
- Other topics, including FLOSS licencing and UK law, the economic benefit to UK plc of FLOSS, FLOSS and computer security, ...

### **What can UKUUG members do?**

Feedback on the provisional programme outlined in the previous section would be very welcome. Apart from that, it is much the same as in the December Newsletter. Further financial support would be very helpful; please send details of organisations who could be approached to me.

Next, it will be necessary to convince MPs that APOSG is likely to have a useful purpose. Around 20 MPs and peers who are already sympathetic have been identified and will be contacted individually over the next couple of months and asked to become involved in APOSG. I would like to have more names, particularly of Labour MPs, who might be receptive; if you think your MP is one of these, please let me know. At some time in the next couple of months members will be asked to press their MP to give their support; please watch for a further announcement about this.

---

## Report on CISL II, Malaga, Feb. 2006

**Charles Curran**

This was the second *Conferencia Internacional de Software Libre* with the English title of *Open Source World Conference*, a discrepancy commented on again. Judging by the registration of a new domain name, this might change for next year. Very few spent time arguing the difference between terminology: open source, free software, FLOSS, etc., and many used the term interchangeably. Herein, I shall use SL (software libre) unless a speaker appeared to stress otherwise. It was also interesting to note how the language – both ES and EN – has been tuned in the last two years and how related and supporting concepts such as developers' community/ies, co-operation, open standards, interoperability, and transparency were mentioned frequently now.

The event was held in Malaga's Trade Fair and Conference Centre 15-17 February 2006 and included an extensive exhibition, tutorials+, and conference presentations. There were about 150 speakers in about 30 sessions, usually five running concurrently, each with a separate theme. There were about 30 workshops/seminars/tutorials before and during the conference. Altogether the organizers said more than 6500 attended (they gave out 6000 shoulder bags); I suspect that many of these were students bused in from around Andalusia+. From my observations, the conference proper did seem to attract a solid body of 1500-2000, with the larger auditorium packing in over 1200. On fees, there were none except from sponsors and exhibitors. The sit-down, four-course meals were free too.

The conference had the theme *Freedom and Innovation* building on what essentially was a pre large-scale roll-out stance of the previous conference two years ago. In general, the conference reviewed recent developments and tendencies of the various sectors using SL and examined their strategies, and discussed business models incorporating SL, potential legal problems, social issues, social inclusion aspects, and, of course, various more technical aspects of IT.

Apart from the hour-long Keynote and the Inauguration, the format for nearly all presentations was for the chair to give a short or very extended introduction to the topic, then for each of five or six presenters on the panel to give a 10-15 minute talk, all followed by 'discussion'. In all, there were almost 30 conference sessions/ sets of talks; They covered:(Plenary sessions) the keynote on the \$100 laptop – One Laptop Per Child. Round-table sessions: Patents/Licences/IP Rights; Long-Term Industry Vision; Technology & Gender Issues–the Case of SL; Social Dilemmas, Are We Building New Society Models?; Foresight on SL: Gurus' Visions. There were also another ~100 talks arranged as four sets of parallel sessions, each with five sets of speaker panels: 1) Experiences in Educational Environment; SL As an E-Inclusion Tool; Co-operation and Collaboration Development Methods; Spanish (Linux) Distributions; SL Projects in Commercial Applications. 2) SL in Europe; SL: Development Experiences; Applications by Sector: E-Health, E-Administration, etc.; SL Business Models; Experiences in Public Administration Projects 3) Regional/Spanish SL Policies; Systems & Architecture; E-Inclusion & Accessibility–Citizen-Oriented Projects Education, Research, and E-Learning Development Models in Public and Private Environments 4) SL Worldwide; Applications, Integration, & Distributions; SL'S Contribution Sustainable Development; Applications and Services for Business(es); Platforms, Lan-

guages, & Development Tools. I managed to get to eleven sets of these sessions, almost 20 hours of talks (most of which were less than 15 minutes) in the two and a half days. Most presentations were in Spanish but there was simultaneous translation in the large auditoria; I think most of the international presenters didn't have fluent Spanish.

The event was again organized by the Juntas (i.e. regional government) of Andalusia and Extremadura with support from various interested parties. Their purpose in holding and supporting this event remains 'to go beyond the purely technical or instrumental aspects of IT and tackle SL as an element capable of extending the Information and Knowledge Society, not only as the basic component of access and mobility, but also as an enabler in making a reality the use of services and other factors that could generate benefits at personal and social levels'.

Keynote: This was given by Michail Bletsas, the Chief Connectivity Officer of the *One Laptop per Child* (OLPC) project. This initiative was first announced by Nicholas Negroponte at the Davos World Economic Forum in January 2005 and it made its debut last November at the UN Summit on the Information Society in Tunis.

MB started by explaining the social-interest background of Nicholas Negroponte and MIT's Media Laboratory and projects they had been involved in since the early 1980s. Their aim is to ship 5-7M laptops around 2007. Overall, they are aiming at third-world children aged 5-16. The target is to reduce production costs below US\$100; initially they will aim at \$120, eventually towards \$85. He examined the typical current laptop costs and showed what needed to be trimmed in order to get near a \$100 cost. For instance the display normally accounts for 20% of the costs; they have got this down to \$35. The laptop, to be made by Taiwan's Quanta, will be Linux based, modified to have a smaller and more stable kernel, WiFi, 3@USB, AMD 500MHz GX533 processor, 128MB of DRAM, no hard-disk-500MB flash memory using a compressed file system, and a 7" dual-mode display: B+W 1110x830 sunlight readable / colour 640x480 backlit, >150 nits, and low (1W) power consumption. The IPv6 based network will be of a mesh form, each laptop able to talk to its nearest neighbours, creating an ad hoc LAN. The laptops will use innovative power (including wind-up). There are more, and presumably up-to-date details on the web site, URL below. The development prototypes should be available 2006-Q3. He mentioned that the choice of Linux was technical and strategic rather than economic; he mentioned that Microsoft and Apple had offered their OSes for free.

Their backers/sponsors/collaborators are Google, AMD, News Corp, Red Hat, and Nortel. The initial countries they will help are Egypt, Thailand, Brazil, Nigeria, and Argentina, although he mentioned that they are talking with Mexico too. The governments must commit to paying \$100 plus the distribution and other costs.

For the 'why' of the project, he suggested reading an article that appeared in *The Economist* last autumn; I assume this is the one that appeared 2005-09-29 but which is only available on *The Economist's* web to subscribers.

I am writing/sketching this just hours after the end of the event and just in time to make this newsletter. I have a long report in preparation and will make this available in the next weeks on UKUUG's web.

Organization: the pre-event organization and access to information was very poor. I registered, heard no more, and they didn't reply to e-mails. It was only that I knew the organization pattern of the previous conference that I thought I would get in. The event itself was well-organized, no mean feat when you are handling such crowds. Registration was almost instant; they saved time by printing badges-on-demand, rather than searching through endless trays.

If you are thinking of attending next year's event – to be in Badajoz, Extremadura 7-9 February 2007, URL as below – don't be put off by language problems, you won't (shouldn't)

have a problem.

**URLs:**

Event and other years'

<http://www.FreeSoftwareWorldConference.com>

Andalusia's Ubuntu-based distribution:

<http://www.guadalinex.org>

<http://www.juntadeandalucia.es/repositorio/>

Extremadura's distribution:

<http://www.linex.org>

OLPC

<http://www.openhardware.net/>

A fuller report is available at:

<http://www.ukuug.org/reports/CISL-2006.xhtml>

---

## **UML 2.0 in a Nutshell**

### **Dan Pilone and Neil Pitman**

**O'Reilly and Associates**

**ISBN 0-596-00795-7**

**234pp.**

**£ 24.95**

**reviewed by Mark Jones**

“UML 2.0 In a Nutshell” is a very complete and precise reference to the latest UML standard. The book is also refreshingly lightweight, weighing in at just over 200 pages. That's not to say it's simply a yellow-pages style reference, because it does contain a wealth of information and explanations of everything it discusses.

An introduction to UML is first covered, after which the book moves quickly through the different types of UML diagram and related components. In particular, the book summarises class diagrams, package diagrams, composite structures, deployment diagrams, use-case diagrams, statechart diagrams, activity diagrams, interaction diagrams, tagged values, stereotypes and profiles. The book concludes with a brief discussion of how to produce effective UML diagrams. Two Appendices are also present which discuss Model-Driven Architecture and the Object-Constraint Language.

Although very concise, the explanations and examples given are of sufficient detail to enable the reader to dip into this book and come away with much knowledge. I've often found reference books to be a very dry read and only of use as a guide when you're in the field. However, by reading or skimming this book from cover to cover, the reader is given a good overview of the new UML standard. Having said that, to get the most out of this book, you do need to have some previous experience of UML and systems-modelling.

If you're looking to learn UML from scratch, I would recommend you explore an introduction- or tutorial-style book first. However, if you're a modeller or designer already using UML, I strongly recommend that you take a look at this book. Skimming it from cover to cover will give you a good overview of the new standard, and it will provide an invaluable desktop reference.

---

## **JBoss: A Developer's Notebook**

**Norman Richards and Sam Griffith**

**O'Reilly and Associates**

**ISBN 0-596-10007-8**

**172pp.**

**£ 20.95**

**reviewed by Mark Jones**

The first thing that struck me about this O'Reilly book was the "non-standard" cover art. Instead of the usual animal adorning the front page, the cover is somewhat plain with a couple of coffee stains being the sole decoration. There is a reason behind this, namely that "JBoss A Developer's Notebook" is not a theoretical book full of abstract diagrams and design patterns. Nor is it a book about programming full of Java-application code. What it presents is a very concise, example-driven guide to JBoss, one of the leading open-source J2EE application servers.

I can't emphasise how concise this book is. At just over 140 pages, it runs through downloading and installing JBoss, deploying simple and more complex applications, database connectivity and data persistence, security, logging, monitoring, and production roll out. All major aspects of JBoss are covered in a short, snappy manner with plenty of example configuration and screenshots.

The book is also very easy and quick to read, helped by a risky but effective format which uses a "grid lined" page background complete with coffee and water stains. Scribbled notes are occasionally placed in the margins to add to the "laboratory" and "hands-on" feel of a notebook. The book contains a lot of "In this chapter ..", "What Just happened?" and "How Do I Do That?" summaries, which are useful and effective.

You'll get the most out of this book by working through the various labs and examples whilst sat in front of a computer running JBoss. However, skim-reading the book whilst not in class or getting your hands dirty also gives a good, high-level overview of the capabilities of JBoss and is thus suitable for designers and architects as well as programmers.

Some people may find this book too short. If you're looking for a detailed discussion of a particular area of JBoss or want examples of Java applications, then I would advise you look elsewhere. However, if you want a light-hearted, example-driven tour of the main features of JBoss whilst sat in front a PC or server, then I can most definitely recommend this book. Designers and architects looking for a quick overview of the capabilities of JBoss should also take a look.

So how to summarise? A daring format which works well, a style that lends to quick reading and progress, and an all-round enjoyable read. I look forward to other O'Reilly books in this series.

---

## **Essential PHP Security**

**Chris Shiflett**

**O'Reilly and Associates**

**ISBN 0-596-00656-X**

**124pp.**

**£ 20.95**

**reviewed by Alain Williams**

I have been writing PHP scripts for several years. Although I have not had a problem I always have a nagging feeling that I ought to be more security aware, are there big holes that I am not aware of ?

Chapter 1 starts with warning about having `register_globals` (yes there is still code about that depends on it) and some `php.ini` settings (these are expanded in appendix a). He continues with an overview of security principles: defence in depth; least privilege; simple is beautiful; minimise exposure. Then onto: checking input (obvious, but many programs that I come across just assume that input is OK); escaping output (vital because 'plain' text could contain HTML markup).

Chapter 2 is about forms and URLs. Should you use GET or POST ? Spoofing form submission, cross-site scripting. Chapter 3 is databases, how unescaped data put straight into SQL could do interesting things. Chapter 4 is about cookies and sessions – how sessions might be stolen. Chapter 5: do people really do an `include` with part of the file name straight from a form ? Chapter 6: is much the same when it comes to opening files in general.

Chapter 7 interested me: authentication and authorisation (how do you log someone in and then decide what they can do). Some good advice on choosing something appropriate with sample code. Chapter 8: how you can mitigate security problems if you use a shared server, this is an aspect of security that most of us probably ignore – the ISP knows what they are doing, everyone else on the machine is benign – right ?

Appendix a: talks about the `php.ini` settings that you should check. Appendix b: warns you of potential problems with functions, like not validating everything that you give as an argument to `exec()`.

It was nice to see that being secure is not that hard - as long as you write the code with security in mind. While this book is not the last word on PHP security what it contains should be known by all those who use PHP. If you only do one thing it must be to validate all input.

The addition of a checklist would be good as it is hard to remember all the things that one should do – although a checklist may give false confidence since security has to be in mind at all stages, not just checked before delivery to the customer.

Summary: If you write PHP scripts, get a copy.

---

## Digital Identity

Phil Windley

O'Reilly and Associates

ISBN 0-596-00878-3

254pp.

£ 24.95

reviewed by Greg Matthews

Identity is a hot topic right now. Sociologists are debating the meaning of individual identity, politicians are arguing over ID cards. What is perhaps less well known are the big issues being tackled by the technical departments of corporations and organisations, surrounding the concept of digital identity.

It seems fairly obvious that when I log into my computer in the morning, my identity is checked against a directory, authentication takes place and then I am authorised to perform most of my day to day activities on this system. On closer inspection, it turns out I have multiple identities existing on various directories of many different kinds and that's just the organisation where I work. I have an identity stored in OpenLDAP, another in eDirectory, several more for various parts of our internal and external web space, a few more for various corporate systems, I could go on. These various identities are used for different purposes and contain greater or lesser

amounts of overlap in the information that they contain and optionally make available. Most of the examples above are for authentication or authorisation or both (although they needn't be, I also have a "white pages" identity) and each can have completely separate authentication tokens. This can be a nightmare for the user to remember and manage all those passwords. It can also become very difficult to maintain and synchronise all these various repositories of information.

This is only the tip of the proverbial iceberg. Identity verification is required for just about every digital transaction. For instance, I can go to any newsagent in the country and buy my newspaper completely anonymously, provided I can pay in cash. However, if I want to access funds from my bank, my identity must be verified, by my PIN or by the signature on my cheque guarantee card or credit card. The same is true for hosts on a network, businesses must be determine who has access to which applications or data and this can only be done with accurate identity information. Traditionally this is done by building a wall around systems and keeping very tight control on the flow of information into and out of the organisation using a firewall.

Ideally, I'd like to sign on to my computer in the morning and have access to all relevant systems needed to do my job without having to remember another pass phrase. Some people will immediately think of "single sign-on" (SSO) and groan inwardly, in the past, SSO solutions have depended on a single vault of information and all systems relying on this vault for resolving identities. However, such systems tended to lock customers into a limited set of technologies. More modern thinking on digital identity is based on the concept of trust. For instance, I may log on to an airline website and book tickets to New York, The airline may recommend a rental car from a particular company. When I click on the link the rental company will receive an assurance of my identity (an authentication assertion in the lingo) allowing me to use the services of the site without another lengthy sign on. Here, two different companies can set up a trust relationship and their users reap the benefit.

Phillip Windley is well placed to talk authoritatively on these issues, he was the CTO for iMall.com which provides e-commerce services where identity issues are paramount. In 2000 he became the CIO for the state of Utah helping to develop e-government systems. He is currently an associate professor of Brigham Young University. This book is a timely and informative introduction to the burning issue of digital identity. Windley's writing demonstrates his strong grasp of this difficult subject. He introduces each concept and defines it carefully in technical terms. Without this attention to detail, chapters on trust and privacy would be too woolly to be useful. This careful approach means the text doesn't descend into stultifying technical language or impenetrable management-speak. It is clearly laid out and the sections are short and to the point.

This book is not a technical book, it will not tell you how to install and configure an identity management architecture (IMA). In fact, it's almost impossible to find a reference to an existing product. This makes it all the more interesting as the subject is in turns, philosophical (what do we mean by "privacy"?), and pragmatic (how is it implemented?). Each chapter is peppered with relevant examples, many from Windley's personal experience, many from the banking world where issues of digital identity have been paramount for a long time. When discussing related technologies, he is quick to point out interoperability problems between standards and the fact that implementing an IMA is politically challenging to say the least.

I expected the book to be worthy but dry, and this was not helped by the picture on the cover of a couple of girls attending what looks like the most boring fancy dress party ever. I found instead, that the subject was interesting and the text informative. Some of the diagrams did little to enlighten the text and had inadequate explanation. Those looking for technical book with code examples will find little to interest them, instead this is a thorough review at the architectural level of the technology required to implement identity management. Highly recommended.

## **Learning SQL**

**Alan Beaulieu**

**O'Reilly and Associates**

**ISBN 0-596-00727-2**

**306pp.**

**£ 24.95**

**reviewed by Lindsay Marshall**

I have to confess to being sufficiently old that I have never been on a database course. They weren't invented when I was a student, at least not the whole Ted Codd stuff and third normal forms and relational algebras etc. etc. I suspect there are quite a few readers of this journal who also fall into this category. My point is, of course, that I have never officially learned SQL. (And I pronounce it S. Q. L. – probably because I have never been on a database course). Oh, I can rub along in it and mostly get what I want out of MySQL, or at least, get back the stuff I put in. But I can't say that I am entirely conversant with the minutiae of the language, nor indeed some of the big stuff (what is the opposite of minutiae?), particularly since I mostly use systems that don't implement all the language or do it in a non-standard way. What I am getting at, the veritable bottom line, the embarrassing confession, is that I just don't understand joins. I do simple but ones but I really don't get the model at all. So I jumped on this book, turned straightaway to the chapter on joins and read it carefully through. it's got lots of examples and takes it easy, but, no, I still don't understand the more outré joins. My bad I think – I've never been one for joining things, I was never a scout, though my great-uncle George was a joiner.

I didn't let my immunity to joining stop me though and I have to say that I learned loads of really good stuff about SQL that I never knew was there. I'm going through all my database code replacing the clunky things I used to do with shiny new methods. I've learned about unions and case statements. Sub-queries I'm still having trouble with: there is a certain joininess about them that gets in the way, but I am starting to see some light.

So a good book for all you oldies who didn't know about first and second normal form, let alone third. I'm not sure how useful this would be to someone who actually had been taught about SQL and databases though – it might be a good refresher. I did mention above that there were lots of examples, and there are, but I have to confess to finding them a little tedious: all of them are based around a banking example and I just got bored with the whole scenario. But that's my only quibble – I am still dipping in and still learning. Though not about joins.

---

## **XML Pocket Reference**

**Simon St. Laurent and Michael Fitzgerald**

**O'Reilly and Associates**

**ISBN 0-596-10050-7**

**175pp.**

**£ 6.95**

**reviewed by Lindsay Marshall**

See the combined review below.

---

## **XSLT Pocket Reference**

**Evan Lenz**

**O'Reilly and Associates**

**ISBN 0-596-10008-6**

**176pp.**

**£ 6.95**

**reviewed by Lindsay Marshall**

I was going to try to be clever and write this review all formatted up as valid XML, and packed with witty in-jokes. As you will have noticed I decided against it, because, let's face it XML just isn't that funny, and XSLT is not funny at all. Quite the reverse in fact. A very serious matter XSLT and which I have always found to be badly served in the documentation department. Writing XSLT is not the easiest of things I have ever done, indeed it is one of my least favourite activities (though it does rank above writing Java programs), but it is made worse by a shortage of easy to understand examples and explanations. Now, don't get all excited and think that I am about to tell you that all your problems are solved and that O'Reilly have come up with the goods yet again (yawn). This book is only OK. It is much better than anything else I've come across and its cheap, but it is still pretty opaque. All right, it's a pocket reference not a tutorial, it's meant to be opaque, but I would have liked it to be a little more accessible.

I'd still send you out to buy it though if you need to use XSLT and aren't already fluent.

The XML reference is altogether much more friendly - some nice clear examples (though the one I used in a class the other day may not actually have been strictly valid XML....) and a solid information about XML Schema which is useful since we do need to be getting away from DTDs. (And stuff about Relax and Schematron too). You could definitely get going with XML from just this book and then proceed to more official standard material once you had the ideas internalised. Let's face it, basic XML is not hard, though people do seem to be trying to make it so (see XSLT). This is still a good little book, and it goes next to the PHP and SQL references on my desk.

---

## **Classic Shell Scripting**

**Arnold Robbins and Nelson Beebe**

**O'Reilly and Associates**

**ISBN 0-596-00595-4**

**300pp.**

**£ 24.95**

**reviewed by Sam Smith**

Shell scripting is at the heart of UNIX, and almost everything considered UNIXy is built, somewhere, on a shell script somewhere. While, a number of years ago, anyone who used UNIX needed to have some reasonably advanced level of scripting shells, this is becoming less true as user-friendliness gets better.

However, despite all of the pretty interfaces on top of systems, to delve into the core of our UNIX systems. Pipes are the loose joins that hold the many pieces of UNIX together. While they are ubiquitous and flexible when you appreciate their power, getting there takes demonstration.

The book covers all of the default POSIX tools; many of those binaries that live in `/bin` and `/usr/bin`, and why and how they can be joined together. It provides a good grounding in a wide range of tools and notes on why things are the way they are. While it mentions extensions found in popular implementations, while taking care to be clear that they are extensions and,

hence, may not be portable across implementations.

It is not an “Introduction to UNIX” book – it assumes that you know how to log in, what the shell is and some basic knowledge. It takes the user from there onwards, to the point that the appendix discusses how man pages are written, and the basics of troff. While this is not something that you need to know to use UNIX, knowing the basics and the structure makes for a better learning experience on the journey towards expertise.

This book is not targeted at the people who have been using UNIX forever, but at those who have learned what they know, and who want to look at more, with a slightly bigger picture of how it all fits together.

---

## Contributors

**Dave Cross** runs Magnum Solutions Ltd, an Open Source consultancy based in London. He is a well-known member of the Perl community and a frequent speaker at Perl and Open Source conferences. Since 2002, Dave has been the Perl Mongers Group Co-ordinator for the Perl Foundation. He is the author of “Data Munging with Perl” (Manning, 2001) and a co-author of “Perl Template Toolkit” (O’Reilly, 2003). His company and personal web pages are at

<http://mag-sol.com/>

<http://dave.org.uk/>

**Charles Curran** was a UKUUG Council member and director 1998-2004. He worked at University of Oxford 1973-2002, and is now an under-employed consultant pursuing various interests in Europe.

**Leslie Fletcher** works part-time as UKUUG Campaigns Manager, with the brief of improving the visibility and credibility of UKUUG and its mission in key arenas – business, politics, public service, education. His main first-hand involvement with Open Source is as chair of governors at Parris Wood Technology College in South Manchester. He also has some experience in IT management, having been head of the Department of Computer and Mathematical Sciences at Salford University for five years.

**Reyk Floeter** is 25 years old and a co-founder of .vantronix, a company specialising in network consulting and OpenBSD support. He is the chairman of the EICAR Task Force on Wireless LAN Security and works as an OpenBSD hacker on improving the free wireless and wired network support.

**Mark Jones** studied at Manchester University where he obtained a BSc in Mathematics and Physics and a PhD in Nuclear Physics. He spent a year at BNFL developing criticality and shielding models before moving into IT consultancy in 1998. He is currently a Technical Design Authority for LogicaCMG, specialising in JEE and .NET systems.

**Lindsay Marshall** developed the Newcastle Connection distributed UNIX software and created the first Internet cemetery. He is a Senior Lecturer in the School of Computing Science at the University of Newcastle upon Tyne. He also runs the RISKS digest website and the Bifurcated Rivets weblog.

**Greg Matthews** is a Senior Unix and Linux administrator for the Natural Environment Research Council.

**Sam Smith** has been on UKUUG Council for 3 years and is currently the treasurer, with many random interests in addition to OpenBSD and Mac OS X. He’s also active in the UK Online Democracy group mySociety and the Manchester Universities’ Gilbert and Sullivan Society.

**Alain Williams** is a Council member of UKUUG and works as an independent Unix and Linux consultant, running Parliament Hill Computers Ltd.

<http://www.phcomp.co.uk/>

---

## Credits

The article by Dave Cross was first published in Linux Format and is reproduced here by kind permission of the author.

We are grateful to Reyk Floeter for permitting the reprinting of his article which was originally published in the OpenBSD journal

<http://undeadly.org>

---

## Contacts

Ray Miller  
Council Chairman; Events; Newsletter  
Oxford  
Tel: 01865 283284  
[ray.miller@ukuug.org](mailto:ray.miller@ukuug.org)

Mike Banahan  
Ely  
[mike.banahan@ukuug.org](mailto:mike.banahan@ukuug.org)

Sam Smith  
UKUUG Treasurer; Website  
Manchester  
[sam.smith@ukuug.org](mailto:sam.smith@ukuug.org)

Alain Williams  
Watford  
[alain.williams@ukuug.org](mailto:alain.williams@ukuug.org)

Roger Whittaker  
Newsletter  
London  
[roger.whittaker@ukuug.org](mailto:roger.whittaker@ukuug.org)

John M Collins  
Welwyn Garden City  
[john.collins@ukuug.org](mailto:john.collins@ukuug.org)

Dean Wilson  
London  
[dean.wilson@ukuug.org](mailto:dean.wilson@ukuug.org)

Newsletter  
[newsletter@ukuug.org](mailto:newsletter@ukuug.org)

Jane Morrison  
UKUUG Secretariat  
PO Box 37  
Buntingford  
Herts  
SG9 9UQ  
Tel: 01763 273475  
Fax: 01763 273255  
[office@ukuug.org](mailto:office@ukuug.org)