

**Contents**

<b>From the Secretariat</b>	<b>4</b>
<b>Chairman's Report</b>	<b>4</b>
<b>From the Editor</b>	<b>6</b>
<b>John Pinner: 1944-2015</b>	<b>6</b>
<b>OpenTech 2015</b>	<b>7</b>
<b>Dynamic Languages Conference: Encouraging Cross-Language Understanding</b>	<b>8</b>
<b>LDAPCon 2015</b>	<b>9</b>
<b>Advanced Perl Course Review</b>	<b>10</b>
<b>Unconference Review</b>	<b>11</b>
<b>DevOps Spring 2015 Review</b>	<b>12</b>
<b>Low Power ARM Desktop Computing</b>	<b>13</b>
<b>Tablets, not just for consumption?</b>	<b>16</b>
<b>Book review: Data Science at the Command Line</b>	<b>18</b>
<b>Book review: High Performance Python</b>	<b>19</b>
<b>Book review: PostgreSQL: Up and Running, 2nd Edition</b>	<b>20</b>
<b>Book review: Introducing Elixir</b>	<b>21</b>
<b>Book review: Introducing iOS 8</b>	<b>22</b>
<b>Book review: Swift Pocket Reference</b>	<b>22</b>
<b>Book review: iOS 8 Swift Programming Cookbook</b>	<b>22</b>
<b>Book review: Programming iOS 8</b>	<b>23</b>

<b>Book review: Learning PHP, MySQL &amp; JavaScript, 4th Edition</b>	<b>24</b>
<b>Contributors</b>	<b>26</b>
<b>Contacts</b>	<b>28</b>



## From the Secretariat

***Jane Morrison***

Thank you to all the members who have paid their subscription invoice promptly this year. The outstanding subscriptions will be chased at the end of March and any not paid at the end of April will not receive the June Newsletter.

We are pleased to announce that 2nd Quadrant Ltd. have recently taken out Silver Sponsoring membership of FLOSS UK.

If your company is interested in Sponsoring Membership please see all the details and benefits at: <http://www.flossuk.org/Sponsors>

At the time of writing we have everything in place for the forthcoming Spring Workshops and Conference being held in York from 24th - 26th March. The event is being kindly sponsored by Google, ELIGO Recruitment and our Sponsor members SuSE and 2nd Quadrant Ltd. You can read a write-up about the event in this newsletter.

The Unconference this year was held on 7th February at London Hackspace, 447 Hackney Road London E2 9DY. It was well attended and a write-up about the event appears in this Newsletter.

Other forthcoming events include:

- OpenTech: 13th June - London
- Dynamic Languages Conference: 20th June - Manchester
- Unconference: October - London

We also hope to bring you more tutorials later in the year.

And looking ahead - Spring 2016 Tutorials/Workshops and Conference will be held in March - city / venue tbc.

The next Newsletter will be the June issue and the copy date is Friday 15th May.

Any comment about past or future events, or if you have something to say about our User Group or this Newsletter please contact: [newsletter@ukuug.org](mailto:newsletter@ukuug.org)

If you do NOT wish to receive future issues of the Newsletter in hard copy (all issues can be found on our web site in PDF format) please let me know.

On a sad note, John Pinner, ex FLOSS UK Council member, passed away in February. He was such a lovely man and our condolences go out to his family, friends and colleagues. A full obituary appears in this newsletter.

---

## Chairman's Report

***Kimball Johnson***

### Spring Conference

The Spring Conference in York was a great success, I was pleased to see so many people and so many great talks. For those of you that were not able to make it, or who just want to review what was discussed, I am pleased to say that the videos will be made available soon.

---

All the talks for the main conference were recorded by York Student Television, and I am very grateful for the service they have provided. As I write this the YSTV team are in Preston for the National Association of Student Television Awards, and I wish them every success in that. Once they return they will get to the job of editing the videos, and once they have made them available I will upload them to the FLOSS UK YouTube Channel and make an announcement. In addition we hope that they are able to take the videos that were taken at last year's conference and make them available as well.

Thank you once again to everyone that attended the conference and for all the hard work of my fellow council members, and to my wife Stephanie, who once again looked after the reception desk of the conference. I could not do it without you all. Also a very great thanks to the wonderful speakers, especially to John Leach for taking on the Wednesday Key Note slot at such short notice.

All the attendees will have received a short survey, I ask that you please complete this as the feedback is extremely helpful in ensuring future events are as much of a success.

### **Upcoming Events**

We have a number of events coming up in the calendar, firstly on Saturday June 13th at ULU in London OpenTech will return once again. This wonderful one day conference is always full of interesting and exciting talks, and I highly recommend you attend. For more information please see: [www.opentech.org.uk](http://www.opentech.org.uk)

The following week sees the first FLOSS UK Dynamic Languages Conference in Manchester. This is a revival of a conference that was held in 2011 in Edinburgh, where it was put together by the Perl Community, and in particular Mark Keating from Shadowcat Systems as part of the Edinburgh International Technology Festival. FLOSS UK has decided to revive it, helped greatly by the fact Shadowcat Systems are now Corporate Members of FLOSS UK and Mark is on the council. The call for papers is still open, we are looking for talks about any aspect of Dynamic Languages, and look forward to a wide ranging program. For more information please see: [www.dynamiclanguages.co.uk](http://www.dynamiclanguages.co.uk)

Looking forward, in November we are hosting the annual LDAPCon conference, this year it is being held in Edinburgh. More information will follow soon. We are always looking for ideas for tutorials to run for our members, so if you have any topic you wish to hear about, please contact me at [kimball.johnson@flossuk.org](mailto:kimball.johnson@flossuk.org) or Jane on [office@flossuk.org](mailto:office@flossuk.org)

### **Support for local user groups**

The budget is still available for FLOSS UK to assist local user groups by helping them to obtain speakers for their events and assist with travel expenses. In addition we have a small budget for assisting with projects that would benefit the Free Software or Free Hardware communities in some way.

If you have an idea and wish some support, please contact Jane on [office@flossuk.org](mailto:office@flossuk.org) and it will be discussed by Council.

### **Get Involved**

FLOSS UK exists to serve its members and we are always on the lookout for people who are keen to get involved in any capacity, whether that be through volunteering to help with organising events, writing newsletter articles, or entirely new activities which haven't been tried before. If you would like to help out in any capacity please do get in touch via [office@flossuk.org](mailto:office@flossuk.org).

## From the Editor

*Paul Waring*

It is with great sadness that I include an obituary of long-standing member John Pinner, who passed away at the end of February. I have collected contributions from people who knew John during his involvement with FLOSS UK and the Python community which is available on page 6.

On a happier note, we have details of a range of upcoming events as well as reviews from the Unconference and Spring conference. We also have an article from Les Pounder on tablets and Bob Easthope on Low Power ARM Desktop Computing, in addition to the usual collection of book reviews.

We are always looking for newsletter content across the following categories:

- Articles: 500-1000 words on a technical subject. Longer articles can be split over multiple issues if required. Past issues have included topics as diverse as Arduinos, ZFS and software patents.
- Event reports: Both FLOSS UK events and those which are of interest to our membership in general, e.g. Oggcamp.
- Book reviews: 250-500 words on a technical book. We can now provide review copies of O'Reilly books in electronic (PDF, ePub, Mobi) and print formats. Subscribe to the mailing list at <http://lists.ukuug.org/mailman/listinfo/books> to find out which books are available for review.

If you have any ideas for newsletter items, please do get in touch via email: [newsletter@flossuk.org](mailto:newsletter@flossuk.org).

---

## John Pinner: 1944-2015

*Paul Waring*

John Pinner, an active member of FLOSS UK and the wider community, passed away at the end of February after battling cancer for several months.

John started his career as an apprentice for Armstrong Siddeley, a manufacturer of luxury motor cars, aero engines and aircraft. He joined Land-Rover in the 60's, worked on the Range-Rover prototype and became responsible for engine testing. During this time he discovered Python and ultimately set up his own company, Clockwork Systems to develop applications which included engineering control systems, payroll and stores and stock control.

More recently he acquired the Linux Emporium which, as well as goodies like T-shirts and mouse mats, sold boxed software sets and held a repository of most distributions which could be purchased on CD. He pioneered the production of Linux laptops. The typical laptop came from the current Thinkpad range and customers could choose from Debian, SUSE or Redhat. More recently Ubuntu became the standard distribution. The Emporium was well known for its customer support and for making wireless work when it was still something of a black art.

John continued his interest in Armstrong Siddeley, owning, maintaining and exhibiting his own Sapphire. He was a leading light in the Owners Club and at one time or another occupied most positions on the Committee. He edited their monthly magazine for many years. The

---

Club is very unusual in that it managed to acquire all the patents, drawings and specifications and the complete stock of spares at the closure of the company. John held the archive on behalf of the Club and provided and maintained their stores control system and web site.

John served on the UKUUG/FLOSS UK Council for several years and was always willing to get stuck in with helping to organise events, whether that be charring sessions, providing speakers or hosting the conference in 'sunny' Birmingham. He made a special effort to welcome partners of delegates, arranging events around the conference schedule so that they could enjoy a day out and then join us for dinner in the evening.

Despite being a self-confessed 'Unix beardie', John recognised the need to bring new people into the community, and he was particularly keen to broaden UKUUG's membership and conference delegate base. He was a key player in the rebranding of the organisation to FLOSS UK, sourcing a graphic designer for the logo and helping set up a new website.

As well as FLOSS UK, John was known for his active involvement in the Python community, helping to initiate EuroPython and PyCon UK, two popular conferences which continue to sell out every year. He was presented with the Python Software Foundation Community Service Award in 2010 as acknowledgement of his hard work. When the sad news of his passing was announced, tributes flooded in from across the community, showing just how many people he had helped and inspired over the years.

Everyone will have their own special memory of John, whether it be the 'Pinner Reversal' (reversing a string in Python using `string[::-1]`), his alarm clock timer for lightning talks or his tours of Birmingham on the way to conference dinners. Personally, I will never forget the time Council decided to have a meeting in the FLOSS UK office in Buntingford, when John drove the two of us from Birmingham and back, during which time we discussed subjects from HMRC's internal systems (which 'paused' from 1-2pm as the data entry clerk went for lunch!) to classic cars. His presence in the community will be sorely missed, but I am sure that many events will continue to be a success thanks to the groundwork he laid.

*Many thanks are due to everyone who helped provide anecdotes for this piece, especially Quentin Wright who filled in the background for John's early life.*

---

## OpenTech 2015

**When:** Saturday 13th June 2015

**Where:** Student Central (formerly ULU), London

**Cost:** £5 on the door donation

Registration for OpenTech 2015 is now open: <http://www.opentech.org.uk/2015/>

We're starting to get a nice schedule of OpenTech talks. We'll have sessions on:

- Toys
- Transport
- Not your normal databases
- Election - What happened? What's next?
- Data successes (and the other outcome)

- Not the Government Digital Service
- Post-Snowden and beyond
- Selling Ideas
- Blockchain uses and abuses
- Law: making it? breaking it?
- Cities
- New considerations in development (coding)

We look forward to seeing you in June. The first cut of a schedule will be up soon. If you want to get your ticket before the rush, please register at <http://www.opentech.org.uk/2015/> and then tell your friends!

The OpenTech Team (@opentechuk)

---

## **Dynamic Languages Conference: Encouraging Cross-Language Understanding**

The 2015 Dynamic Language Conference is a one day event held on Saturday, June 20th at Studio, Manchester.

Dynamic Languages Conference (DLC) is a cross-language event aimed primarily at, but not restricted to, Open Source Dynamic Languages. The organisers hope to bring together as many languages as possible to an open forum. The aim is to inspire discussion, evaluation and presentations on the manner in which these languages approach and solve tasks at the cutting edge of development technology. As we move towards a world dominated by an increasing number of primarily C-derivative scripting languages, high level languages, we encounter a plethora of approaches to similar issues. The likelihood is that many of these approaches are equally valid and that integration of ideas across the different languages will educate and inspire those practitioners in them.

This conference matches the style of Open Source Developers Conferences (OSDC) but with a specific focus on the dynamic languages.

The Dynamic Language Conference is a joint initiative organised by FLOSS UK and Shadowcat Systems. The hope is to encourage a wide uptake from both business, entrepreneurs, freelancers and interested individuals. We are hoping that people will become advocates for the event and sign up, submit talks and encourage others in their local networks to do the same. FLOSS UK is the United Kingdom's chief advocate of open and free technology.

There is a Call for Papers that has been opened and will run until the 5th June 2015. The organisers will do their best to notify speakers immediately as to who has been chosen but will post the schedule on the 12th June ahead of the event a week later.

The Organisers will be operating a system of announcing some speakers early, this is especially true of well-respected speakers, experienced or talks with a high degree of interest. They are keen for discussion with speakers about their presentations and to any innovative use of the talk time.

If you require more time to organise a schedule we are willing to pre-authorise your talk based on your popularity, standing in the community, knowledge and experience. However,

the organisers would dearly love to see a broad spectrum of talks pitched at different levels of expertise and will accept talks from beginners or new speakers.

Speakers should feel free to contact the event organisers, Mark, Kimball or Jane at the FLOSS UK offices. A list of addresses can be found at <http://www.dynamiclanguages.co.uk>.

---

## **LDAPCon 2015**

The fifth International Conference on LDAP and Directory Services will be held in the UK at the University of Edinburgh School of Informatics Forum.

**Tutorials:** Wednesday 11th November 2015

**Conference:** Thursday 12th & Friday 13th November 2015

### **Call for papers and tutorials**

Topics:

- You are using LDAP in interesting projects?
- You do LDAP client or server development?
- You have used LDAP in a new way?
- You do identity and access management on top of LDAP?

Why not share your ideas and experiences with others?

We are looking for speakers who are willing to talk about any topic related to LDAP and identity management, including:

- LDAP technology implementation (Servers, API, User interfaces etc.)
- LDAP Usage (Schema, Security, Operations, Scaling, big data, etc.)
- LDAP related technologies (PKI, XACML, SAML, etc.)
- LDAP and Beyond (IAM, Identity Federation, Authentication on the web, etc.)
- Best Practices for directory services.

Accepted talks will be grouped into tracks such as a standards/development and deployment/administration.

### **Deadlines and Important Dates**

- Submission Deadline: 28th June 2015
- Author Notification: 10th July 2015
- Final Papers due: 10th October 2015
- Tutorials: 11th November 2015
- Conference: 12th-13th November 2015

### **Talk Submissions**

Main presentations should last about 45 minutes including discussion; we will also provide smaller slots of 15 minutes and 5 minutes for poster presentations or lightning talks. Please tell us which duration you prefer when proposing your talk. The talk must be in English.

---

The one and only way to submit your abstract (approximately 200-800 words, accompanied by your biography of about 100-300 words) is via email to [submissions@lists.ldapcon.org](mailto:submissions@lists.ldapcon.org). Abstracts must reach the Program Committee by 28th June 2015. Early submission is encouraged.

All abstracts will be reviewed by the program committee.

For accepted talks we expect you to submit slides and/or a paper of approximately 2-10 pages (A4 or US Letter format, 25mm borders, preferably L<sup>A</sup>T<sub>E</sub>X source or OpenOffice).

For 5-minute talks, a brief abstract is required. A short paper, slides or a poster should be provided for accepted talks. We will provide display boards for posters throughout the conference.

By submitting a paper you grant the conference organizers the non-exclusive right to publish your paper in the conference proceedings and on the website; you maintain the right to publish it elsewhere at your discretion.

### **Tutorial Submissions**

We are looking for high-quality tutorials on LDAP and related subjects, at any level from introductory to advanced. Tutorial length can range from an hour to a full day. Wireless Internet access will be available if required. The purpose of the tutorials is focussed education, so they should cover established topics and best practice rather than presenting new work.

The Programme Committee has an open mind about the format of the tutorial day, but has a limited number of rooms available. Make your proposal early and we will aim to build an attractive programme for the day.

### **Expenses**

Speakers get free access to the conference, including the social event.

If requested in advance we will provide accommodation for speakers.

Travel expenses might also be covered in special cases. If you need this, please contact us early so we can try to arrange it.

### **Website**

<http://ldapcon.org/2015/>

### **Contacts**

**General enquiries:** [enquiries@lists.ldapcon.org](mailto:enquiries@lists.ldapcon.org)

**Paper/Tutorial submissions:** [submissions@lists.ldapcon.org](mailto:submissions@lists.ldapcon.org)

---

## **Advanced Perl Course Review**

***Quentin Wright***

Yet another successful Advanced Perl Course was held at the Ambassadors Hotel in London last November. The Perl Courses we arrange have endured despite changing trends in web technology and yet again we had a full house of 10 attendees from both academic and commercial environments. The London venue works well and is convenient for the main line railway stations.

---

The course covers Perl 5.1, Advanced Testing, the DBIx::Class and then moves on to using Moose which brings a more consistent approach to object oriented programming in Perl. Additionally time is spent on the Catalyst web framework. To round things off there's a session on benchmarking and profiling. Everyone gets a traditional set of course notes.

Our tutor, Dave Cross, was the founder of London Perl Mongers way back in 1998. Dave is a regular speaker at Perl and Open Source conferences and is often invited to present tutorials alongside the main conference. He is the author of Data Munging with Perl (Manning, 2001) and a co-author of Perl Template Toolkit (O'Reilly, 2003).

If you are interested in participating in any of the Perl courses please email [office@flossuk.org](mailto:office@flossuk.org). As well as the 2-day Advanced course there are 1-day Introductory and Intermediate Courses.

---

## Unconference Review

*Quentin Wright*

### **Report on the FLOSS UK Unconference held on 7th February 2015 at the London Hackspace**

This was our fifth unconference and as something of a change was held at The London Hackspace which is to the East of London near Shoreditch.

A goodly portion of the 40 or so registered attendees were there for a prompt start at around 10am. FLOSS UK Chairman Kimball Johnson launched the event, welcoming everyone and reminding all of the upcoming Spring DevOps Conference.

We then rattled through what is now becoming a familiar format for the beginning of the Unconference. All are encouraged to put up topics that they would like to talk about or alternatively would like to find out about. A brief description is put on a Post-It note and the notes stuck on a whiteboard. After a sentence or two's description of each topic a show of hands determines which are the most popular subjects.

We decided to run with a single stream of topics and to limit presentations and talks to 20 minutes. Each talk was followed by questions and discussion, with a more flexible approach to timing. We did have coffee breaks and a lunch break with a generous allowance for networking.

It's always interesting to discover the topics that are of current interest. Maybe not-surprising security was top of the list and it was clear that many are grappling with the issues in a hands-on way. Tim Fletcher did us proud with a talk on his use of Yubikey, and later on in response to a request gave an impromptu overview of his experience of 'Digital Forensics'. Again on security Hazel Smith described how the Linux capability system can be used to sub-divide root privileges and make remote backups more secure.

A quite different topic of Open Data was comprehensively covered by a presentation from Ben Webb.

The week had also seen the launch of the Raspberry Pi 2 and we had a discussion slot where people talked about their use of the Raspberry Pi and embedded systems. To set the scene we had a presentation from Mark with an overview of his 20 or so years experience.

---

One of the benefits of our using The Hackspace venue was that we could go ‘Walkabout’, and we were given a comprehensive guided tour of the London Hackspace. This included demonstrations of various pieces of equipment and a descent into the bowels of the building to see a robot which they are busy commissioning.

Also the Hackspace’s SpaceShip which was so popular at Brighton was parked at the back of the building. To explain to the uninitiated the SpaceShip is installed in a caravan and is a life-size spaceship simulator. FLOSS UK have some affection for the project as we donated on-board cameras when the project was in its early days. Tom Wyatt gave us a presentation on the ‘Bikeshed’ as he called it. This was followed by an invitation to the more adventurous to take a trip during the afternoon, and a couple of teams re-appeared later on flushed and exultant at their experiences!

In response to a request about open software and music David Geary gave a compelling demonstration of his own use of Superlooper, showing how using a drum machine and synthesiser and support from the audience on the finer points of composition it was possible to quickly create a piece of music. We also were shown a project called gibber, which allows the programmatic use of software to make an audio/visual performance on line.

At the end of the Unconference many adjourned to the Sebright Arms to carry on their discussions.

Thanks to all the speakers and contributors others who might not have been mentioned above and especially to all the participants who made it such a successful event.

Thanks are due to London Hackcentre for the accommodation, endless coffee, their tour and generous hospitality. Hazel Smith who has the distinction of being both a long-standing member of FLOSS UK and London Hackspace acted as a coordinator. Also thanks to my fellow FLOSS UK Council Members: Chairman Kimball Johnson who organised the lunch and Tim Fletcher who in unflappable style was able to talk about a whole variety of topics at the drop of a hat.

---

## **DevOps Spring 2015 Review**

***Paul Waring***

This year’s DevOps Spring conference took place in the city of York, surrounded by historical buildings from castles to merchant halls. As usual, the first day consisted of tutorials, this year covering Large-scale System Design (Google), Practical Digital Forensics (Tim Fletcher) and Zero to Perl (Shadowcat Systems), followed by two days of conference talks. This year there were up to three parallel sessions, so unfortunately I have not been able to do justice to every talk in this review, and it was often difficult to choose which session to attend.

Things kicked off on the Wednesday with a keynote speech on Docker by John Leach, who was kind enough to fill in after a last minute cancellation. Docker is a new way to package up applications and configuration for deployment, and as a ‘hot topic’ there were lots of questions and a lively discussion.

After a coffee break I headed back to hear David Proffitt speak on ‘Enhancing SSH for Security and Utility’. He took us through the wide variety of configuration options available on both the client and server and I picked up several useful tips as a result - in particular the use of the `ssh-copy-id` command for copying a public key to a server with the correct permissions.

David's talk was followed by a non-technical but extremely interesting presentation on 'Job Interview Techniques' by Rick Deller, which contained some useful advice on how to deal with tricky situations in the interview process.

Lunch was followed by Nick Moriarty from the University of York discussing the challenges of upgrading a diverse infrastructure from Puppet 2.7 to 3.x. The final scheduled talk of the day was given by Matt S. Trout on 'Build management with a dash of Prolog'. This was delivered in Matt's customary rapid-slide style, and I could just about keep up with my jotted-down notes. The day was finished off with lightning talks, where speakers have five minutes to talk about anything they want, but the audience counts down when there are ten seconds remaining and no overrunning is permitted! Several speakers managed to cram an enormous amount of information into such a short slot, with some incredibly useful Bash tips and information on using filesystem capabilities for full-system backups.

One of the highlights of the annual Spring conference is the conference dinner, and this year did not disappoint. Held at the historic Merchant Adventurers Hall, the dinner consisted of a buffet with a variety of options, and those of with a substantial appetite were pleased to see that the options were not mutually exclusive! Previous dinners have tended to be formal sit-down three course affairs, but the buffet option seemed to work well as it meant everyone was served quickly, and the usual problem of staff trying to find the right tables for various dietary requirements was avoided.

The second conference day started off with a keynote by Stu Teasdale on how to migrate a legacy application (in this case an energy monitoring system) to Continuous Integration and the Cloud. After a break to take in some additional caffeine, I listened to Mark Cairney talk about Federated Access Management in the UK academic sector and beyond, including some of the difficulties encountered and the way forward for future authentication data exchange.

The penultimate talk was a run down of the extensive list of upcoming features in PostgreSQL 9.5 from Gianni Ciolli, who also won the 'best speaker' award. Last but not least was Peter Tribble discussing his 'Tribblix' distribution based on illumos, a derivative of OpenSolaris. It was interesting to hear about progress made with an alternative kernel, especially as illumos and the BSD family don't get the same coverage as Linux, despite still being used behind the scenes in many organisations.

Of course, it's not just the talks which make a conference - often the most important conversations take place during the coffee breaks, over dinner or at the pub. As well as catching up with conference veterans, I was pleased to see some new faces who I hope to see again at future events.

Overall, this was a great conference and many thanks are due to the people involved in organising it, the speakers who kept us informed and entertained over the three days, and the delegates for a wide range of interesting discussions.

---

## **Low Power ARM Desktop Computing**

***Bob Easthope***

For the past three years I have been experimenting with ARM devices to get full desktop on with everything I need to browse the web, connect to servers at work using VNC and SSH, play games, write documents, receiving email and generally having a great time experimenting with the hardware and software.

Low power has ranged from 500 milliwatts to 7 watts which is virtually nothing in comparison with a desktop that will be consuming 100-150 watts! The amount of power you can go down to depends on requirements, for example a Raspberry Pi will consume 500 milliwatts but some limitations in memory and clock speeds you will get with this set-up, however you may find this is powerful enough as it will handle SSH connections to remote servers fine and the newer version is OK for VNC and VPN type tasks.

There are various boards to choose from with a variety of processors used and clock speeds for any given device. For example the A7 is a popular chip design from ARM that is what the chip used in the Raspberry Pi Version 2 was based on, the older version being based on another Broadcom chip using an ARM 11 design that has much less cache, lower clock speeds and an older instruction set which does place a limit on the versions of Linux available that are compiled to run on this architecture.

### **Development boards and devices I used and some results**

The Raspberry Pi B version 1 was very limited as a full desktop, you can get some browsers running such as Midora, and Libreoffice will run but takes a long time to load. I have used VNC and Rdesktop to get to machines which was OK although redrawing can be an issue.

The Raspberry Pi B version 2 released several months ago is more powerful having a newer processor at a higher clock speed, and 1 gigabyte of RAM instead of 512 Megabytes. Additionally the chip is quad core which gives an improvement when multitasking and running applications that can make use of multiple cores, for example playing video files or running a VPN and VNC session, you can really see an improvement over the original device. Surprisingly the cost is still £35 and the power consumption is still around 500 milliwatts.

The Beaglebone Black is very comparable with the original Raspberry Pi, however it does have one big advantage instead of using a Micro SD card it runs off a 2GB eMMC solid state memory chip. You flash this chip using a Micro SD card and boot time is 30 seconds, however due to the single core A8 processor it can be quite slow loading applications, although faster than the original Pi due to the memory and slightly faster chip. Makes for a nice cheap NFS server with USB 2 hard disk connected for backup machine, but not ideal as a desktop device.

Utilite PC another device that is made by an Israeli company but now sold on Amazon is a really promising device. These come with an A9 based processor the Freescale i.MX6 running at 1 GHz on standard version (dual core) and 1.2 GHz on the professional version (quad core), various configurations of RAM from 1 to 4 gigabytes. I have tested both Pro and Standard and standard having a 32GB SSD makes the whole device far snappier. I often connect to the company network using the Juniper SSL VPN web software, stream internet radio and have VNC and SSH sessions running and everything works just fine, no sudden freeze ups as on the PI devices and enough RAM and CPU cores. You can really see the cores being used well and Ubuntu works really well, you can have 3D graphics support and get very clear output display at 1080p to large screens.

Odroid XU3 available at [www.hardkernel.com](http://www.hardkernel.com) is also worth a look if you enjoy experimenting with kernel modules and recent chips, it uses an octacore chip made by Samsung the Exynos5422, which is used in mobile phones from last year such as the Samsung Galaxy S4. The whole idea is 4 of the cores are low power A7s (like in the PI V2) and 4 cores are higher power A15 cores. If you are using your phone with this chip to make a call you would use the smaller lower power cores and then switch to the higher power A15s if you want to play a game that would need more resources. In testing on Xubuntu I could see the faster cores kick in when starting to watch video clips after browsing the web. A new mode allows all cores

to be used at once if you have many tasks running. Cost is around £100 so you pay for extra performance, but this is more at the cutting edge. There's a very good community forum and many people explain building new kernel modules or you can download pre-built images for Xubuntu and Debian as well as XMBC for connecting to TV.

### **So what is the best device to use?**

What I use for my desktop matches up well with the Utilite Pro very nicely, having a reasonable amount of RAM (2 Gigs) is useful for running Java based applications and Firefox. Boot-up time is very fast and good amount of eMMC storage (32 Gig) so can have full desktop and plenty of room for lots of applications. If like me you like to run desktop across two monitors the Utilite is the best device available as it supports this.

The Odroid XU3 is costlier than the Utilite if you buy the eMMC module, however it's far more experimental and tends to need a lot more work to get drivers working, especially when you want to upgrade Ubuntu! Due to license restrictions some drivers are not in the kernel and you have to compile them, insert into a specific kernel version and run that or wait a while for others to figure out the build problems and release an image that you can put on the device which you take a leap of faith in trusting will work. On the plus side can really learn a lot about the hardware having to piece things together and was great for experimenting and very active community for this device and others that hardkernel have released.

### **Limitations of ARM Linux devices**

There's no official support for Flash, something some websites use. There is an open source flash player, that I've got to work with some flash video content, some sites that need a specific flash version sometimes run but with limited functionality. A lot of Linux software is available for ARM chips but some things still need to be ported, however I have found a lot will compile once you work out the dependent packages required.

Memory is also another consideration: 2GB is low by standards of desktops available today, you can buy development boards with 4GB but cost goes up and 4GB is the absolute limit for 32 bit ARM devices available today for low cost.

### **New devices?**

Towards the end of the year it is likely that we will see 64 bit based ARM development boards becoming available at low cost, some of which probably will be suitable for desktop use. The A53 and A57 chip designs are 64 bit and with bigger caches, able to address more memory and likely more cores per chip, with on-chip gigabit ethernet as well as SATA support!

### **Some conclusions**

In today's world of expensive electricity and cost to the environment, low power ARM computers can really be useful to use at home. The Utilite PC consumes 7 Watts when you are using it to its full potential, just compare this with a desktop which often will be running at 150-200 Watts if you run something without fancy graphics that is! Or if you can drop some horsepower you can go to 500 milliwatts if running a Raspberry PI.

My idea was to have something that I could easily leave on all the time at low cost and also at noise because the power consumption is so low the heat generated is low meaning that a fan isn't required, an idea that worked out well for me and has reduced my power consumption and quietened my office at home significantly.

Linaro, Canonical and Debian project team have done a great job at getting Ubuntu and Debian running on ARM powered devices, much work has been done at both ends of the spectrum, from servers to very small embedded Linux devices. RedHat too have started porting to ARM chips so never a better a better time than now to run the RISC.

---

## Tablets, not just for consumption?

*Les Pounder*

In the open source community there is a strong, and just, focus on the activities of Linux and other free operating systems. Typically I would write about how good "<insert Linux distro here>" would be, but instead I would like to talk about tablet computing. Schools around the world are using iPads as tools to teach Computing. While they are great as consumption devices, they are pretty lousy when used to create anything. When was the last time that you wrote code on a tablet?

In late 2014 I heard about a cheap Windows tablet, called the Linx 7 that came with Windows 8.1, the full desktop not the phone operating system. I was intrigued as to the educational aspects of such a device so I picked one up for £60. While Windows is not an open source platform, the aspect of using a desktop OS on a tablet device piqued my interest, as it enabled schools to try out open source technologies in a safe manner.

The Linx 7 is an Atom powered, quad core CPU with 1GB of DDR3 RAM and comes with 32GB of storage. It can plugin to an external screen using HDMI and has USB OTG (On The Go) to enable peripherals to be attached. For £60 this would be an ideal device for use in many levels of education, but as my focus is on Primary and Secondary level education I wanted to use open source applications on this Windows device as it is notoriously difficult for Schools to adopt Linux and other open source projects into their IT infrastructure due to being unwilling or nervous about how it will interact. In fact only the Raspberry Pi has made such an impact. showing schools how versatile and cost effective open source can be for schools.

I wanted to test a few Windows version of applications that I used with Linux to see if such a cheap device could help schools embrace open source technologies in the same way as the Raspberry Pi.

### **Programming tools**

Children in Primary schools across the world are taking their first steps with Scratch[1] and it is great to see a child building code in such a simple way. I tested Scratch and Scratch 2 with the Linx 7 and it performed really well with both.

Moving to a text based language and Python 3 is well served by IDLE which runs quite well on the low powered hardware. I was able to quickly try out a few demo scripts that I wrote and all worked well, not bad for £60 of hardware.

But Python is not the only method of learning to code, oh no we have they fabulous Sonic Pi[2] which promotes the use of coding as a means to generate music. Using the Ruby language children can quickly pick up programming skills via creating music, in fact the tutorials written for this project mention nothing about Ruby until much later in the learning process. This is a great way to teach "computing by stealth" with children enjoying the process of making something rather than focussing on on coding concepts.

Minecraft has steadily taken over the world and has recently been purchased by Microsoft. Thanks to the Raspberry Pi programming using Minecraft as an outlet has become an interesting and exciting avenue for schools to pursue. On the Raspberry Pi you can easily use Python to alter the Minecraft world and teleport players, create buildings and control external devices with standard coding logic. On the Linx 7 you can surprisingly run Minecraft quite well and replicate the same coding functionality found on the Raspberry Pi version using a custom server.[3]

### **Electronics**

When tinkering with electronics and circuits one of the most labour intensive activities used to be documenting the layout of your project. But in recent years we have seen Fritzing[4] being used more and more in both print and web publications. It produces high quality drawings of projects and contains libraries of common components and libraries provided by online retailers such as Adafruit[5] which enable classes to quickly design projects and generate a shopping list of parts. Fritzing is a large application and runs reasonably well on the Linx 7, granted the more components used, the more memory that is gobbled up by the application.

### **ICT**

Computing is not just about writing a letter or creating a database but at the same time computing is not just about Coding. Computing is an umbrella to so much and that's why the standard ICT skills such as Office skills and e-safety are still relevant to this day. Using LibreOffice[6] on the Linx 7 I was able to write this very document and keep my accounts in order.

### **Conclusion**

So I learnt that the applications that I know and love on Linux work remarkably well on even the lowest specification of hardware. I also learnt that this small device could make a big difference to schools as a low cost portable solution that offers so much more than an iPad. It might not have the same appeal as a Raspberry Pi, but this tablet really does provide a decent user experience for schools wishing to dip their toe in to open source applications.

Now all we need to do is hack the device to run Linux and we have the perfect portable Linux device.

[1] <http://scratch.mit.edu/>

[2] <http://sonic-pi.net/>

[3] <http://www.stuffaboutcode.com/2014/10/minecraft-raspberryjuice-and-canarymod.html>

[4] <http://fritzing.org/home/>

[5] <http://www.adafruit.com/>

[6] <http://www.libreoffice.org/>

---

## **Data Science at the Command Line**

**Jeroen Janssens**

**O'Reilly Media**

**ISBN: 978-1-4919-4785-2**

**212pp.**

**£ 22.52**

**Published: October 2014**

**reviewed by Rob Charlton**

Data science is a hot topic by any measure. A popular online shop lists 146 books on Data Science published in 2014 which is twice the number in 2013 and five times the number from 2012. The author of Data Science at the Command Line is a Senior Data Scientist at YPlan in New York and clearly knows his subject well.

This book is very much a hands on guide. It is packed with practical examples, all of which can be easily worked through while reading the book. It is aimed primarily at Data Science practitioners, or at least those with some appreciation of the subject, who are looking to augment their tools and workflow using the command line. These individuals would normally reach for a programming language like R or Python for their work, but Jeroen Janssens sets out to demonstrate that the old but powerful Unix tools philosophy can offer a lot to this new subject. Further, it is likely that command line experts with no Data Science experience will still find many useful tips and new ways of using the commands they are already familiar with.

The book starts with a good introduction to the command line and popular command line tools. Data Science is presented as a 5 step process: Obtain, Scrub, Explore, Model and iNterpret (referred to by the mnemonic OSEMN which doesn't exactly trip off the tongue). A useful package which is presented in the book is the Data Science Toolbox. This is a virtual machine image which contains all the tools and utilities mentioned in the book, as well as all of the test data sets and examples. It can be run locally using vagrant, or online using Amazon Web Services. It is a really handy addition, otherwise a lot of time working with the book would be wasted finding, building, installing and configuring the many tools used.

Most of the remaining chapters correspond to steps in the OSEMN process. 'Obtaining Data' focusses on getting data onto your system from web pages, web APIs, databases and spreadsheets. 'Scrubbing Data' gives many worked examples on slicing, filtering and processing data, especially CSV data, to get it into the format you need. In fact this is the biggest chapter of the book. It introduces a lot of useful tools, many of which were new to me and some of which were written by the author himself. 'Exploring Data' shows several methods to visualise your data, both in text as well as graphs. This chapter introduces Rio, which is a tool provided by the author to allow R scripts to play nicely with command line workflows. Finally 'Modelling Data' shows how to use Tapkee, Weka, Scikit-Learn and BigML all from the command line to process the data to make predictions or find patterns. Both exploring and modelling assume some knowledge of Data Science concepts.

There are three chapters dedicated to ideas to help get the best from the command line. There is a chapter on turning BASH one liners into shell scripts, for those unfamiliar with how to do this. The workflow tool Drake is introduced, which offers some of the same benefits of Gnu Make to data processing pipelines. Finally Gnu Parallel is introduced in order to show how large workflows can be parallelized on a single machine or a bank of servers.

One topic which is not covered, which you might expect in a book of this nature was Hadoop

or any of the popular 'Big Data' tools and technologies. I think it would be safe to say that this book was aimed more at working on smaller data sets, not multi-terabyte files.

I can thoroughly recommend this book. It is well written, contains many practical examples, and I have learned about many new command line utilities which I can apply in my daily work.

---

## **High Performance Python**

**Micha Gorelick, Ian Ozsvald**

**O'Reilly Media**

**ISBN: 978-1-4493-6159-4**

**370pp.**

**£ 16.65**

**Published: September 2014**

**reviewed by Rob Charlton**

High Performance Python is a comprehensive manual for anyone interested in improving the performance of their Python code. It also shines a spotlight into some Python internals that any Pythonista should be interested in exploring. The book is primarily aimed at users with CPU bound problems and mostly those in the scientific computing arena. Don't be put off by this though, Micha and Ian do a great job of covering I/O bound, memory bound and other optimisation problems. You will find the book's advice useful for just about any Python application. You will need solid Python skills to be able to make use of the knowledge presented though - this isn't a book for Python beginners. There are a lot of examples, with assertions all backed up by measurements and charts so you can be confident of the results. Finally the text is full of battle-hardened advice drawn from the experience of the authors and other industry experts.

The book starts with an introduction to computer architecture from the perspective of where bottlenecks can occur. It then looks at the yawning gap of performance from C code to any given Python script. By the final chapter you should have all the tools you need to be able to get your Python code a lot closer to the C code's performance though there are a number of tradeoffs that you will learn about along the way.

The second chapter looks at profiling. It provides a deep dive into a lot of different profiling techniques and these are returned to and augmented in later chapters. This is absolutely essential advice and something I will apply in my daily work from now on. The next three chapters look at some Python primitive data types and patterns (lists, tuples, sets, dicts, iterators and generators) giving lots of useful insight into how they are implemented, how they compare and how they perform in terms of RAM and CPU usage. This whole section would be useful to any Python developer.

Moving more into the territory of numerical computation, Chapter 6 looks at Vector computation. It introduces the perf command line tool, then takes an example to show how successive optimisations using numpy can make a dramatic difference - a 40x speed improvement in this case. It also gives a great illustration of the performance improvement lifecycle.

The next chapter covers compiling Python code to C. This is a very active area of work in the Python community, and a range of different projects (Cython, shed skin, numba and pypy)

are looked at in detail and compared. The end of the chapter covers the various techniques of calling out to C (and Fortran!) from Python.

The next 3 chapters focus on doing more work in parallel. Chapter 8 looks at speeding up I/O bound problems using `gevent` and `tornado` to do event based programming. Chapter 9, one of the largest, looks at the multiprocessing module to show how you can break free of the Global Interpreter Lock (GIL) and parallelize CPU bound problems. Chapter 10 looks at clustering, and a range of different tools and techniques to partition jobs over multiple machines and communicate with them.

Chapter 11 starts to look at RAM performance and shows how to measure RAM use and pick data structures to use less memory. It also gives worked examples of using probabilistic and approximate data structures such as `hyperloglog++` and bloom filters.

The final chapter is called 'Lessons from the field'. It is a really worthwhile read all of its own. It provides 6 separate essays from start-ups and other Python users in cutting edge fields, each discussing how they use Python, what their technology stack is, how they have optimised performance and what special advice they would give.

Overall I would say this book would be a very worthwhile purchase. It has a lot to offer to a wide range of readers, provides a lot of detail and examples, and is clear and easy to understand. Performance tuning covers a very wide subject area but this book manages to cover a very impressive range. I only have one minor criticism which is that some of the figures are not reproduced well in greyscale, as they have details that were clearly meant to be in colour.

---

## **PostgreSQL: Up and Running, 2nd Edition**

**Regina O. Obe, Leo S. Hsu**

**O'Reilly Media**

**ISBN: 978-1-4493-7319-1**

**234pp.**

**£ 16.58**

**Published: December 2014**

**reviewed by Mike Fowler**

PostgreSQL: Up and Running has been updated for a second edition covering some of the new features introduced in PostgreSQL 9.3 and 9.4. The authors are well known and active in the PostgreSQL community having written *PostGIS In Action* and the *PostgreSQL OnLine Journal*. Aimed at new and intermediate PostgreSQL users but there are tips for experts too, I was pleased to learn a few tricks myself. The book assumes a certain level of knowledge from it's reader, namely that you can install software and understand basic SQL. This has allowed them to skip a lot of tedious detail that often feels patronising to more experienced users. Examples are plentiful and well explained often being numerically annotated making the commentary even easier to follow.

A wide range of topics are covered and of course given the size of the book it is not exhaustive and it thankfully doesn't try to be. Early chapters focus on core SQL concepts and highlights where PostgreSQL does things differently. Data types are given a decent chapter with particular focus on the JSON data type and it's supporting functions. Later chapters look

at PostgreSQL specific features including an extensive look at EXPLAIN and a brief coverage of replication. Personally I would've preferred a greater focus on replication however the material is detailed enough to get a working cluster up and running.

One drawback to the physical version of the book is that it is riddled with links. Fortunately the links are usually where you can go for greater detail on the topic in question but as a curious reader you would start to tire of keying these in. It's picky, but you can't help but feel this is a digital book at heart.

PostgreSQL: Up and Running will serve as a good PostgreSQL primer and can often be used in place of the often more complex manual. I'd recommend this book to any new user of PostgreSQL and to seasoned PostgreSQL users who aren't up to date with the latest features.

---

## **Introducing Elixir**

**Simon St. Laurent, J. David Eisenberg**

**O'Reilly Media**

**ISBN: 978-1-4493-6999-6**

**210pp.**

**£ 14.18**

**Published: September 2014**

**reviewed by Charlie Harvey**

Elixir (<http://elixir-lang.org/>) is a dynamic, functional language that runs on the Erlang virtual machine. The language goals are productivity, extensibility and compatibility. It has a nice syntax that looks a lot like Ruby, though it feels quite like Scheme in some respects - for example in its extensive meta-programming support which is provided by hygienic, scoped macros. Elixir provides excellent support for concurrency, scalability and resilience, all of which it inherits from Erlang.

Introducing Elixir is a short book. O'Reilly seem to be producing more short books than they used to back in the days when their output was mostly big, manual-like tomes. Obviously this means that the book functions as a brief introduction to the language and not a comprehensive reference.

The intended audience for Introducing Elixir isn't completely clear to me. The introduction identifies it as "people who've been programming in other languages but want to look around" and not for people "in a hurry to get things done". Which is fine, but the pace in the first chapters seems rather too slow for experienced programmers even given that caveat. It is, nevertheless, solidly written and covers a fair amount of material. I imagine that folks with little exposure to functional programming will appreciate the coverage of higher order functions, pattern matching and recursion though it would have been good to see more in-depth treatment of some of the basic functional tools like folds.

The book's examples are all variants of solutions for a single problem - calculating fall velocities on various planets. I found this approach to be reasonable enough, though a little repetitive. Other readers may find it bores or irritates. Some more guidance on the idiomatic ways to solve the problem in Elixir might have been a good addition to the book.

There are no exercises included in the book - I find these a much better way to learn a language than reading (or even typing) examples, but there is a set of exercises designed to go along

with the book and available for free from <http://chimera.labs.oreilly.com/books/1234000001642>. I haven't reviewed the material, but I rather like the approach of pulling out arguably non-essential content to the web.

Overall, I felt that this was a solid, though short, language introduction to an exciting and very interesting new language. If the repetitive examples and the fairly slow pace are not to your taste, the resources section of the book mentions Dave Thomas's Programming Elixir as a possible alternative introduction. That book may appeal more to those who are interested in something faster paced with more emphasis on the unique aspects of Elixir.

---

## **Introducing iOS 8**

**Steve Derico**

**O'Reilly Media**

**ISBN: 978-1-4919-0861-7**

**334pp.**

**£ 22.52**

**Published: December 2014**

**reviewed by Lindsay Marshall**

See combined review below.

---

## **Swift Pocket Reference**

**Anthony Gray**

**O'Reilly Media**

**ISBN: 978-1-4919-1542-4**

**186pp.**

**£ 9.99**

**Published: December 2014**

**reviewed by Lindsay Marshall**

See combined review below.

---

## **iOS 8 Swift Programming Cookbook**

**Vandad Nahavandipoor**

**O'Reilly Media**

**ISBN: 978-1-4919-0869-3**

**902pp.**

**£ 31.99**

**Published: November 2014**

**reviewed by Lindsay Marshall**

See combined review below.

---

## **Programming iOS 8**

**Matt Neuburg**

**O'Reilly Media**

**ISBN: 978-1-4919-0873-0**

**1018pp.**

**£ 29.20**

**Published: December 2014**

**reviewed by Lindsay Marshall**

A quartet of books and over 2000 pages. All with the usual high standard that we expect from O'Reilly books: good content, good production etc.

The Swift Pocket Reference is like the other pocket references: an excellent summary of the language, small and convenient. In some ways it doesn't seem vastly different from the free iBook that you can get from Apple, just shorter, and of course printed. This will be useful for someone who knows Swift well, but the language has some hard and probably unfamiliar bits in it too and getting to grips with those is going to take a lot more than this.

Introducing iOS 8 has the subtitle "Swift programming from idea to App Store" and it seems to deliver on this. Having been looking at quite a few books like this recently I can definitely say that this is one of the best in this category, though I do worry about its claim to be useful even if you have absolutely no programming experience: there is some tricky stuff in here. Though not as much as in the next two.

Programming iOS 8 is like the grown up version of the previous book, hugely detailed with lots and lots of useful examples. It covers parts of the system that most programmers will never use. It's comprehensive and is nearly 1000 pages long. You could do someone a serious damage with this book, it's heavy. There do seem to be omissions though, presumably because this is most likely a revamp of the equivalent iOS 7 book. Not much on extensions, nothing about adding keyboards (that I could find - it's a big book!!), no mention of the Watch, HealthKit, HomeKit, CarPlay etc.

This Swift programming cookbook is exactly what it says: A collection of recipes that solve problems that app developers may run into, all newly translated into Swift. Another massive book, full of useful information. And here's why this all so brief, even for me, but...

One book says 'covers Xcode 6', another says 'covers Swift, iOS 8 .1, Xcode 6.1'. The day these four books arrived, Apple released the betas of Xcode 6.3 and iOS 8.3. They also released Swift version 1.2 which has significant language changes. So not only will they be missing information but some of the examples may have code that is now syntactically wrong, that uses deleted/renamed methods or that changed semantics. These are good books but about fast moving targets and they are not cheap. I can recommend them wholeheartedly but I think you would be mad if you bought them.

---

## **Learning PHP, MySQL & JavaScript, 4th Edition**

**Robin Nixon**

**O'Reilly Media**

**ISBN: 978-1-4919-1866-1**

**812pp.**

**£ 22.78**

**Published: December 2014**

**reviewed by James Roberts**

Any book covering PHP, MySQL, JavaScript, and jQuery with CSS and HTML despite having nearly 750 pages of content, can certainly only devote a limited amount of attention and detail to any of these topics; my usual MySQL reference alone has over 1200 pages.

I had not read previous editions of this book, but comparing contents with a PDF version of the previous edition it appears that a major change in this edition is the introduction of an extensive treatment of jQuery, with a whole new chapter dedicated to it, and the dropping of the PHP mysql access extension for mysqli; there are various other changes, but much of the rest of the content would appear to be similar.

So: who is the book targeted at, and how good a job does it do in introducing so many topics?

After reading the book I am still not clear on the target audience. The author states that the book is 'for people who want to learn how to create effective and dynamic websites', and goes on to list web-masters and graphic designers, school and graduate students, and individuals: the only requirement is some knowledge of static HTML-with-CSS-based websites. On page 28 we see that the 'book is aimed mostly at PC and Mac users', and so it spends some time in guiding the reader through installing WAMP or XAMPP. Of course all of the tutorials and exercises will work, perhaps with a little prior setup, on at least Linux and the BSDs; while LAMP installation for Linux is mentioned it is not covered. I was disappointed but not surprised to see that in 2015 we are still given plain FTP as a file transfer method.

Continuing, we are guided to select a programmer's text editor - Editra, which is freely available across all the targeted platforms, is suggested. The Windows-oriented background reveals itself with the recommendation of the phpDesigner IDE - which will not run natively under Linux/BSD (though it seems to work fine in Wine at least on Kubuntu), and which requires a license purchase to use beyond evaluation, though some alternative IDEs are indicated.

In Chapter 3 we have finished the preliminaries and launch into PHP. This introduction assumes very little, and on page 40 we are introduced to variables, with the nice image of a matchbox which can contain a value. This also illustrates my difficulty in conceiving the exact target audience ; this sort of introduction is for the complete beginner, either an adult with zero programming experience or, say, a 11-14 year old school student in an introductory course (but would they know what a matchbox is?). Similarly on pages 46, 47, 77 et seq. there are very low-level introductions useful to the complete beginner - but would this beginner have written a static website? I will come back to this question at the end of the review.

We move through the material at a fair old rate and in a logical order. I was interested in how the vexed issue of PHP security could be handled in such a compact introduction, and by page 61 there are clear warnings about the use of superglobal variables - this only covers a paragraph and I would worry about someone putting together a PHP site live on the web after

reading just this book without having it audited by someone who knew PHP security well, as I myself do not.

The introduction to PHP continues through chapters 4,5,6 and 7, with some more warnings about input validation on page 158, but this is only introduced after an insecure approach has been presented.

Chapter 8 introduces MySQL. We start out on the command line, and necessarily the MySQL introduction is a little less gentle with the complete beginner; by this stage such a beginner who had worked through the previous chapters probably could handle this material. The examples I tried also worked under MariaDB.

Very rapidly we arrive at Chapter 9, 'Mastering MySQL', where database design and data normalisation are briefly introduced before the reader is led into transactions and the different storage engines available in MySQL. There is no space in the book for the sort of exposition that would enable the reader fully to understand the reasons behind the selection of InnoDB, and that is only to be expected in this concentrated format. Basic MySQL data backup and recovery are introduced.

Chapter 10 puts together the material learned so far, with MySQL being accessed via PHP. Page 255 adds more security warnings, this time about user input validation of data collected via PHP, with 6 pages covering this in reasonable detail.

By Chapter 11 we move to JavaScript and the pace is more swift for the absolute beginner, as relevant programming topics were introduced in the PHP chapters and the examples of course are related to web pages and HTML. Quite soon browser 'gotchas' are described, and ways of getting debugging output while developing JS scripts are explored and evaluated.

In Chapters 12 to 15 JavaScript is further explored, and we then arrive at Chapter 16 which covers using JavaScript for input validation and error handling together with PHP, this chapter apparently expanded over the treatment in earlier editions.

Chapter 17 introduces Ajax and gives clear use cases. Chapter 18 introduces CSS, while Chapter 19 is entitled 'Advanced CSS' and recommends CSS3 over JavaScript where possible. Where this is not possible, Chapter 20 offers a similar Advanced CSS access from JavaScript.

Chapter 21, new to this edition, explains the basics of jQuery after explaining why this tool-kit was chosen.

Chapter 22 brings us to HTML5, with HTML5 Canvas being treated in Chapter 23. Chapter 24 covers HTML5 Audio and Video. This is quite an important topic with the increasing deprecation of Flash, but still mentions Flash fall-backs for IE5/6, presumably required if you are developing for the government or councils (sigh). Chapter 25 covers other HTML 5 features including a rapid pass over geolocation, GPS, Local Storage and Offline web applications.

Finally we arrive at Chapter 26 where all the threads are pulled together in the form of the development of a very simple social networking site that incorporates aspects of everything covered so far in the book.

So, back to my original two questions; who do I see as the target audience? And how well does the book achieve its aims?

The book covers a great deal of material in a very compact form. It cannot also go into depth, but enough detail is given to explain the exact reasons for the existence of each of the components described and demonstrated. This book would work well without further

support as a rapid introduction to dynamic web programming for someone with prior notions of another programming language and some HTML/CSS experience.

With support in a live or virtual classroom, it would work well for complete beginners at adult, intermediate, or junior levels; I believe quite a few questions would come up for that beginner that are not answered in the book.

This is not a flaw as such in the book; no book is perfect or complete, but this one does a good job at what it sets out to do, which is to walk a naïve user through to a point where they will have a good notion of how modern feature-rich websites are put together, and they will have built a basic example themselves. With some more study and mentoring readers could eventually readily produce such sites.

As a final very general comment, dynamic web page design now requires learning at least a server-side and a client-side language; experience and expertise in tool-kits and frameworks to simplify application of these tools; and further expertise in a database and its language, together with a presentation formatting language and a text description language, just to provide what could be provided by a single language tool in a desktop application: is this really the right way to go about it?

---

## Contributors

**Rob Charlton** is Chief Software Engineer for Vertu Corporation Ltd where he runs the Services team. In his day job he does Python, Ansible, AWS, Splunk and dabbles in data science as well as managing the team. Outside of the day job he dabbles with Raspberry Pi and Arduino. Over the past 20 years he has started startups, built products, and written code for servers, PCs, phones and embedded devices.

**Bob Easthope** is based in Cambridge and works for ARM as a developer in the internal IT group, developing and building new applications and services using Linux and Open Source tools and applications.

**Mike Fowler** is a Senior Software Engineer working in the Midlands. He is a firm open-source advocate and has helped companies move from Microsoft based solutions to complete Linux/PostgreSQL stacks. He is a major contributor to YAWL, an open source BPM/Workflow system. When not melled with a keyboard, he spends his time playing as many different stringed instruments as he can.

**Charlie Harvey** is IT Director of New Internationalist workers co-op, who produce an award-winning magazine, books about global justice and run ethical online shops. He has been working with free software since the late 90s. He enjoys cider and blogs at <http://charlieharvey.org.uk>.

**Kimball Johnson** has been programming since a very early age, starting with BBC Micros, then MS DOS and Windows Systems, however was enlightened with a copy of Debian GNU/Linux Woody at university. From this developed an affinity to Systems Administration, but like all good admins he is lazy and so tries to automate as much of his job as possible, and now follows the way of the Infrastructure as Code.

**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 Newcastle University. He also runs the RISKS digest website and the Bifurcated Rivets weblog.

**Jane Morrison** is Company Secretary and Administrator for FLOSS UK, and manages the FLOSS UK office at the Manor House in Buntingford. She has been involved with FLOSS UK administration since 1987. In addition to FLOSS UK, Jane is Company Secretary for a trade association (Fibreoptic Industry Association) that she also runs from the Manor House office.

**Les Pounder** works closely with North Western Linux and Free Software groups to promote the use of Open Source software as opposed to proprietary software. He is also the organiser of UCubed, a free Linux and open source event in Manchester, and an organiser of Barcamp Blackpool and Blackpool Geekup, and has been head of crew at Oggcamp. He writes for Linux Format magazine and contributes to Linux podcasts including Fullcircle, Ubuntu UK Podcast and Linux Outlaws.

**James Roberts** was one of the last people in the UK to learn to program on 80-column cards at Control Data, and ran into Linux SLS in 1992. He runs a small company/consultancy that sneaks Linux and \*BSD into SME Windows shops wherever it fits. He has inter alia taught micro-computing and MIDI and audio recording to young adults.

**Paul Waring** currently works in teaching and software support at a large UK university. Outside of work he can usually be found filing documentation bugs against various open source and free software projects. He also edits the FLOSS UK newsletter.

**Quentin Wright** is a Director of Sunfield Technology working with Linux and Ruby and system integration projects. In his spare time in between playing with reluctant motor vehicles he is pre-occupied with Web and Javascript programming.

---

## Contacts

Kimball Johnson  
Chairman  
Preston

Gavin Atkinson  
Council member  
York

Tim Fletcher  
Council member  
Manchester

Ian Norton  
Council member  
Morecambe

Stephen Quinney  
Council member  
Edinburgh

Quentin Wright  
Treasurer  
Warwick

Paul Waring  
Newsletter Editor  
Manchester

Alain Williams  
System Administrator  
Watford

Holger Kraus  
Website  
Leicester

Sam Smith  
Events and Website  
Cambridge

Jane Morrison  
FLOSS UK Secretariat  
The Manor House  
Buntingford  
Herts  
SG9 9AB  
Tel: 01763 273475  
Fax: 01763 273255  
**office@flossuk.org**