



Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

ISSUE #95 - March 2015



BQ AQUARIS E4.5 THE UBUNTU PHONE HAS ARRIVED

Full Circle Magazine is neither affiliated with, nor endorsed by, Canonical Ltd.



HowTo

| | | |
|---|---|------|
| 1 | 6 | 1/10 |
| 2 | 7 | 4/15 |

Program in Python p.15




LibreOffice p.19



Use LaTeX p.22



Blender p.XX



Inkscape p.25



Graphics



Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY



Linux News p.04

```
#An alias to make the ls
command more detailed
alias ls = "ls -la --
color=always --classify"
```

Command & Conquer p.13




Arduino p.28




Linux Labs p.32



LinuxLabs p.XX



Review p.37



My Story p.47



Blender p.XX




Letters p.XX



Tuxidermy p.55



Q&A p.56



Security - SSH p.58



Ubuntu Games p.62



Ubuntu Games p.64



The articles contained in this magazine are released under the Creative Commons Attribution-Share Alike 3.0 Unported license. This means you can adapt, copy, distribute and transmit the articles but only under the following conditions: you must attribute the work to the original author in some way (at least a name, email or URL) and to this magazine by name ('Full Circle Magazine') and the URL www.fullcirclemagazine.org (but not attribute the article(s) in any way that suggests that they endorse you or your use of the work). If you alter, transform, or build upon this work, you must distribute the resulting work under the same, similar or a compatible license.

Full Circle magazine is entirely independent of Canonical, the sponsor of the Ubuntu projects, and the views and opinions in the magazine should in no way be assumed to have Canonical endorsement.



WELCOME TO ANOTHER ISSUE OF FULL CIRCLE.

He's back! Fear not programmers. Our Python series (and Greg) has returned. As ever the LibreOffice series continues, and to complete the trinity we have a piece on LaTeX. For those of you who like the idea of Dropbox, but don't trust it (for whatever reason) we have an article on Syncthing. It looks quite like BitTorrent Sync which I wrote about a while back.

Last month the big news was, of course, the Ubuntu Phone. Unfortunately, those of you who purchased one of said phones had to wait almost six weeks for it. While I applaud Canonical and BQ for bringing an Ubuntu phone to the market having a flash sale but only saying at the end of the sale that the buyer has to wait almost a month for the unit is not the done thing. If the stock wasn't there then BQ should say so up front before the buyer can even hit a 'buy' button. I hope that by the time you read this you'll have your phone. If you do have it, send me an email telling me what you think about it. I'm keen to hear what other folks have to say about it. If you want my thoughts on it, head over and read my review of it.

Kevin O'Brien returns this month to give us an intro, and the basics, on SSH. While it's not something I've used, I know many of you are fans of SSH. Speaking of something I've not used, Alan Ward has sent in a really interesting piece on Ubuntu (and Linux in general) being used in a professional setting for the 3D viewing of molecules.

And there's more! We have a competition this month with five copies of Able2Extract Professional up for grabs. I've written up a review of it this month and it's got some pretty impressive OCR capabilities.


All the best, and keep in touch!

Ronnie

ronnie@fullcirclemagazine.org



This magazine was created using :



Full Circle Podcast

Released monthly, each episode covers all the latest Ubuntu news, opinions, reviews, interviews and listener feedback. The Side-Pod is a new addition, it's an extra (irregular) short-form podcast which is intended to be a branch of the main podcast. It's somewhere to put all the general technology and non-Ubuntu stuff that doesn't fit in the main podcast.

Hosts:

- Les Pounder
- Tony Hughes
- Jon Chamberlain
- Oliver Clark



<http://fullcirclemagazine.org>



Download



CANONICAL STARTS TO BUILD IoT ECOSYSTEM

Canonical, the company behind the Ubuntu distribution of Linux, is pulling together an ecosystem of Internet of Things partners that promises to make it easier to both build and manage IoT applications.

Maarten Ectors, vice president of connected devices for Canonical, says the challenge with IoT these days is that rather than being an actual platform for developing applications, the IoT as developers see it as more akin to the Internet of isolated things.

To address that issue, Canonical has created Ubuntu Core, an implementation of Linux for embedded systems that makes it possible not only to run multiple IoT applications on the same device, but also to continuously update them.

As part of that effort, both Amazon Web Services and

Microsoft have agreed to publish their cloud APIs on Ubuntu Core.

Ectors says Canonical is trying to extend the presence that Ubuntu has into the cloud out to IoT devices. About 70% of the servers running inside the AWS cloud, for example, are running Ubuntu Linux, says Ectors. In addition, 62% of the implementations of the open source OpenStack cloud management framework are running on Ubuntu Linux, he adds.

Source: <http://www.programmableweb.com/news/canonical-starts-to-build-iot-ecosystem/2015/02/27>

Submitted by: Arnfried Walbrecht

OPEN SOURCE: MORE EYES, FEWER VULNERABILITIES, GREATER SECURITY

Future posts will dive deep into open source and its relationship to autonomous

devices, but first, let's take a few paragraphs to level-set why open source might be an ideal option. First, full disclosure: I'm an advocate of open source software, so I've seen proof that a community of shared ideas and projects that can be modified, improved, and distributed freely can be a better way to develop technology. Being able to see the code, learn from it, ask questions, and offer improvements is the open source way.

While it might seem counterintuitive, open does not mean less secure. In fact, the opposite is often true. Because the development process is collaborative, bugs, flaws, and vulnerabilities can be found sooner, and more often, and fixed more quickly. By granting access to the code, more people can work to solve issues. It's been said about open source that "given enough eyeballs, all bugs are shallow." More eyes and greater transparency can lead to fewer vulnerabilities and greater security.

Source: <http://mil-embedded.com/guest-blogs/open-source-more-eyes-fewer-vulnerabilities-greater-security/>

Submitted by: Arnfried Walbrecht

RASPBERRY PI ENTHUSIASTS CAN NOW SHARE THEIR LITTLE WORKHORSE WITH WINDOWS, MAC AND LINUX USERS

RealVNC, the developer of VNC remote access and control software, is allowing Raspberry Pi enthusiasts to share their PC with Windows, Mac and Linux users.

RealVNC released its flagship VNC product for the Raspberry to coincide with the device's birthday at the end of last month, enabling users to connect to their Pi from any Windows, Mac or Linux computer. Once VNC has been downloaded to the Pi, users can apply for a free license or, for an enhanced set of features, a



personal or enterprise license.

Raspberry Pi users will also be able to connect to their Pi from an iPad, iPhone, Android or Chrome device by downloading RealVNC's free VNC Viewer app, which is available in the Apple App Store, Google Play or the Chrome web store.

VNC can be used with Raspberry Pi in two ways. If the Raspberry Pi is connected to a monitor or TV and is running a graphical desktop, the connected user can see exactly what a user sitting in front of the Pi would see – to support generic remote access and also good for sharing a screen with a friend, for instance. If the Pi is not running a graphical desktop, then graphical remote access can still be obtained by creating a virtual desktop instead - "ideal for users who have built their Pi into a robot", said RealVNC.

Source: <http://www.techworld.com/news/operating-systems/raspberry-pi-enthusiasts-can-now-share-with-windows-mac-linux-users-3600302/>

Submitted by: Arnfried Walbrecht

COMPANIES REALLY WANT LINUX-SAVVY EMPLOYEES AND THEY WANT THEM NOW

According to the Linux Foundation and tech job company Dice, in the 2015 Linux Jobs Report, "Nearly all hiring managers are looking to recruit Linux professionals." While programmers and Linux system administrators are in high demand, your chances of landing a great job are greater if you have cloud, security, and/or software defined networking (SDN) skills.

In particular, "42 percent of hiring managers say experience with or knowledge of OpenStack and CloudStack are having a big impact on their Linux hiring decisions" while "49 percent of Linux professionals believe open cloud will be the biggest growth area for Linux in 2015."

Container technologies, such as Docker, get a lot of ink but aren't as important for finding a job. Only five percent of hiring managers say that container knowledge has made a difference in their hiring decisions. 19 percent of Linux professionals see containers as the

top growth area.

Hiring managers are looking for developers and network administrators who know their way around SDN. Specifically, "19 percent are looking for Linux talent with SDN skills."

Source: <http://www.zdnet.com/article/companies-really-want-linux-savvy-employees-and-they-want-them-now/>

Submitted by: Arnfried Walbrecht

MEERKAT LINUX MINI PC WITH INTEL PROCESSOR UNVEILED BY SYSTEM76

A new Linux Mini PC called the Meerkat has been unveiled this week by PC builder System76, which has been equipped with an Intel Broadwell processor and will be available to purchase later this month around March 19th 2015.

The Meerkat will be launched offering different processor options in the form of either an Intel Core i3-5010U supported by Intel HD 550 graphics or a Core i5-

5250U processor supported by Intel HD 6000 graphics depending on your requirements and budget.

Other features of the new Linux mini PC that is supplied running the Ubuntu Linux operating system and measures just 4 x 4 x 2 inches, include the ability to install up to 16 GB of RAM. Together with up to 2TB of internal storage using the supplied M.2 for solid state disk (SSD) storage, or the 2.5 inch drive bay for the inclusion of a more traditional hard drive or other SSDs.

Connections on the small desktop mini PC include 4 x USB ports, Gigabit Ethernet, HDMI and DisplayPort allowing users to connect two monitors to the system if required.

Unfortunately, pricing for the new Linux mini PC system has not been revealed as yet, although the system is expected to sell from around \$500 and when it launches later this month.

Source: <http://www.geeky-gadgets.com/meerkat-linux-mini-pc-with-intel-broadwell-processor-unveiled-by-system76-05-03-2015/>

Submitted by: Arnfried Walbrecht

VMWARE SUED FOR FAILURE TO COMPLY WITH LINUX LICENSE

In 2007, top Linux contributor Christoph Hellwig accused VMware of using Linux as the basis for the VMware ESX bare-metal hypervisor, an essential part of VMware's cloud offerings.

Years went by and the Software Freedom Conservancy, a non-profit organization that promotes open-source software, claims to have negotiated with VMware for the company to release ESX's code, and its successor ESXi. That way, argued the Software Freedom Conservancy, these programs would legally comply with Linux's Gnu General Public License version 2 (GPLv2). VMware refused in 2014.

Now, Hellweg and the Software Freedom Conservancy are suing VMware in the district court of Hamburg in Hamburg, Germany.

The group explains that they see this as a "regretful but

necessary next step in both Hellwig and Conservancy's ongoing effort to convince VMware to comply properly with the terms of the GPLv2, the license of Linux and many other Open Source and Free Software included in VMware's ESXi products."

What's surprising about VMware's stubbornness is that there's never been much question that VMware had used Linux in ESX and ESXi. As Hellwig wrote in 2007, "VMware uses a badly hacked 2.4 kernel with a big binary blob hooked into it, giving a derived work of the Linux kernel that's not legally redistributable."

Source: <http://www.zdnet.com/article/vmware-sued-for-failure-to-comply-with-linuxs-license/>

Submitted by: Arnfried Walbrecht

GRAB YOUR PITCHFORKS: UBUNTU TO SWITCH TO SYSTEMD

The Ubuntu Project is set to move forward with a plan to make a controversial system

management tool a key part of Ubuntu Linux.

On Monday, March 9, the Ubuntu maintainers will reconfigure code base for the forthcoming version of the OS so that it uses the much-debated systemd suite of tools to handle core initialization tasks and manage system daemons.

That means that when Ubuntu 15.04 ships (presumably in April), all new Ubuntu installs will be running systemd by default.

It's a move that's sure to annoy some. When the Debian Project announced that it was switching to systemd last year, it sparked an angry protest by old-school admins who wanted Debian to stick with the familiar Unix System V-style init software to handle its startup and management chores.

That didn't stop Debian from following through on its plan, though, and the systemd naysayers were forced to take their metaphorical ball and go home – to a new, systemd-free fork of Debian, known as Devuan.

Ubuntu already doesn't use old-

style init; it uses an alternative called upstart. But that won't prevent a few grumbles from those who'd rather just leave things as they are – to say nothing of those who just plain don't want systemd.

Source: http://www.theregister.co.uk/2015/03/07/ubuntu_to_switch_to_systemd/

Submitted by: Arnfried Walbrecht

GNOME 2 IS BACK: UBUNTU MATE IS NOW AN OFFICIAL FLAVOR

Ubuntu MATE is now an official flavor of Ubuntu.

Yes, that means Ubuntu is giving a stamp of endorsement to GNOME 2 once again. You don't need to switch to Linux Mint – just install the Ubuntu MATE disc and get a desktop that works like it did before Ubuntu's Unity and the GNOME Shell came along.

If you're using Linux, there's a good chance you've heard of the MATE desktop.

It's a continuation of the old GNOME 2 desktop code. The developers are continuing to improve how it works with new technologies without dropping everything and starting from scratch. This makes it more controversial than projects like Linux Mint's Cinnamon desktop, which takes modern code and try to make it behave more like a traditional desktop, instead of bringing the old GNOME 2 code forward. GNOME 3 now also offers a Classic Mode to appeal to users who want a more traditional desktop experience.

But never mind the competitors, and never mind which project will have the easiest time coping with new technology in the future. The MATE desktop works well today. If you don't want to play with new stuff, if you're comfortable with GNOME 2, or if modern desktops just seem too unwieldy, it's a great option. This is a core strength of Linux: When users don't like desktop changes, they can bring the old desktop forward. Windows users are stuck with however Microsoft wants the desktop to work this year.

Source:

<http://www.pcworld.com/article/2893647/gnome-2-is-back-ubuntu-mate-is-now-an-official-flavor.html>

Submitted by: Arnfried Walbrecht

LINUX ADOPTS CONFLICT RESOLUTION CODE

If you can't take the heat, get out of the kitchen" could be the unofficial motto of the Linux kernel community. Over the years, there has been one conflict after another in the heart of the the Linux development community, the Linux Kernel Mailing List (LKML). Now, in order to make the LKML more peaceful, the group has adopted a Code of Conduct.

There's no question that Linux is the most successful operating system and open-system project. But it's also true that if you watch it closely, you'll see a lot of conflict within the community. In particular, Linus Torvalds, Linux's founder, doesn't suffer fools gladly, and he's never afraid to let other developers know when he thinks they're wrong.

There's nothing new about this, but recently, Lennart Poettering, a Red Hat engineer and one of the creators of the controversial systemd system and service replacement for Unix and Linux's sysvinit daemon, called out Torvalds publicly. Poettering accused Torvalds of encouraging hate speech and attacks on him. While Poettering's reputation suffered more from this episode than anyone else's, it sparked a discussion on how better to address conflicts within the Linux kernel development community.

Source:

<http://www.zdnet.com/article/linux-adopts-conflict-resolution-code/>

Submitted by: Arnfried Walbrecht

OPEN COMPUTE PROJECT (OCP) FORMALLY ACCEPTS OPEN NETWORK LINUX (ONL)

Big Switch Networks, the company bringing hyperscale networking to data centers worldwide, today announced that the Open Compute Project (OCP)

has formally accepted Big Switch's contribution of Open Network Linux (ONL) as its reference Network Operating System (NOS). ONL is a Linux-based open source network operating system to bare-metal and branded white-box ("brite box") switches.

In addition to supporting commercial products, ONL was initially created as the base hardware-testing reference platform for DIY hobbyists and researchers. ONL currently supports 12 different open switch hardware platforms and basic L3 routing, and also has an OpenFlow agent based on OF-DPA and the Indigo project in progress.

Big Switch Networks' contribution -- Open Network Linux (ONL) -- is the Linux distribution for bare-metal switches that runs underneath Big Switch's commercial Switch Light OS™. ONL's goal is to give people deploying OCP switches a simplified experience with a standard Linux distribution that comes prepackaged with all of the relevant drivers, loaders, and platform-independent benefits.

Source:

<http://www.marketwired.com/press-release/open-compute-project-ocp-formally-accepts-open-network-linux-onl-1999108.htm>

Submitted by: Arnfried Walbrecht

VMWARE BRINGS HORIZON 6 TO LINUX

Desktop virtualization has always been a necessity for companies and organizations who need a mobile, connected workforce. And it is true, simultaneously, that many organizations rely on the flexibility and security that Linux workstations offer over commercial solutions such as Microsoft Windows. It is, therefore, a good news for companies, that VMWare, the world-famous virtualization solution, is bringing its latest offering, Horizon 6, to Linux.

According to a recent official announcement, VMWare will soon be porting Horizon 6 (released around a year ago) to official Linux repositories. Although advanced users choose SSH (Secure Shell) for

remotely working on their machines over networks, newcomers to Linux, especially in the corporate sector, will find this move incredibly effective. This is partly due to the ease of use that VMWare offers, in the form of GUI solutions over conventional, CLI-based approaches in Linux, which often baffles new users.

Already, VMWare has released a test version of its virtualization suite, that is available to install and test on Redhat and Ubuntu based servers. Among several features, the suite offers virtual GPU support for NVidia cards, to enable smooth 3D graphics over the network.

Source:

<http://www.pcworld.com/article/2896232/vmware-expands-desktop-virtualization-to-linux.html>

Submitted by: Anirban Chatterjee

MAJOR MILESTONE IN LINUX GAMING: STEAM ON LINUX CROSSES '1000 GAMES' MARK

Gaming on Linux has often been a bit of an oxymoron, until even a few years back. Most users, who chose Linux over Windows or Mac as their primary workstations, relied on secondary OSes or platforms for satisfying their gaming needs. Of course, Linux has prestigious and open-source titles, including Warzone2100 and Battle for Wesnoth; but PCs dominated the market when it came to games.

Thankfully, now those days are history. Steam for Linux has recently crosses the '1,000 games' mark, offering more than 2,000 items in total (including demos, videos, etc). Although the number is still less than that offered to Windows and Mac users (4,814 and 1,613 respectively), this can be easily seen, and regarded, as a major milestone in the history of Linux.

Quality differences between the same games, when run on Direct3D and OpenGL in Windows and Linux OSes respectively, were quite notable at the beginning. But slowly, that is starting to minimize, as more developers devote time to make games specifically targeting the Linux users, rather than simply

porting existing titles to the platform. Things are going to improve, even more, once the developers turn to Steam's next-generation Vulkan API.

Source:

<http://www.extremetech.com/gaming/201055-potent-penguinistas-steam-for-linux-crosses-1000-game-threshold>

Submitted by: Anirban Chatterjee

IS THE LINUX FOUNDATION TRYING TO GAG LINUS TORVALDS?

Linus Torvalds has taken a lot of criticism for his...er...blunt responses to Linux developers over the years. But now the Linux Foundation has set up a "code of conflict" that might change the way Linus interacts with developers.

Julie Bort reports for Business Insider: *On Monday, the Linux Foundation kinda sorta slapped him on the wrist when they issued a new "Code of Conflict" policy that declared "personal insults or abuse are not welcome."*

It says that if "anyone feels personally abused, threatened, or otherwise uncomfortable" while working on Linux, they should report the situation to the Technical Advisory Board who will step in and mediate.

Torvalds was not the one to write this policy. His right-hand man, Greg Kroah-Hartman, wrote it and cutely submitted it as a "patch" to the Linux system. That meant that the ultimate keeper of Linux, Torvalds, had to see the "patch" and approve it, which he did, adding the public comment, "Let's see how this works."

Source:

<http://www.itworld.com/article/2895396/is-the-linux-foundation-trying-to-gag-linus-torvalds.html>

Submitted by: Arnfried Walbrecht

OPEN SOURCE ON THE AGENDA AS LMAX BECOMES MEMBER OF LINUX FOUNDATION

LMAX Exchange, the world leading FCA regulated MTF for

global FX trading and the UK's fastest growing technology firm, today announced it has become a member of the Linux Foundation, the nonprofit organisation dedicated to accelerating the growth of Linux and collaborative development.

LMAX Exchange will be a part of the Linux Foundation's new Core Infrastructure Initiative, a multi-million dollar project bringing together more than 20 major international technology companies, including Microsoft, Intel, Google and Cisco, to identify and fund critical open-source infrastructure projects in need of investment.

Dr. Andrew Phillips, Director of Technical Operations at LMAX Exchange commented, "At LMAX Exchange we are firm believers in the power of open source technology to upend markets and disrupt traditionally opaque and complex systems, especially in the financial services. Linux is a speed and performance workhorse that is pervasive in financial services today, and as Linux Foundation members, we're excited to collaborate on this transformative technology and help spread the

benefits of open source technology."

"We are proud to welcome LMAX Exchange as the newest member of the Linux Foundation" said Amanda McPherson, Chief Marketing Officer at The Linux Foundation. "LMAX Exchange is leading the transformation in financial trading with its unique approach to exchange style trading for FX and its innovative use of open code. Their work is sure to be of great benefit for the future of Linux."

Source:

<http://leaprate.com/2015/03/open-source-on-the-agenda-as-lmax-becomes-member-of-linux-foundation/>

Submitted by: Arnfried Walbrecht

AUTOMOTIVE BUS OPEN SOURCED WITH LINUX-BASED DESIGN

A German university is open sourcing a secure, two-tier Automotive Service Bus for car computers, available on a control unit running Linux on a

PandaBoard.

Technische Universität München (TUM) has open-sourced an automotive computer bus design developed as part of its "Visio.M" (Visionary Mobility) electric car project, according a Mar. 10 press release by TUM. Next week at the CeBIT show in Hanover, Germany, TUM will demonstrate the carbon fiber Visio.M prototype, which was backed by the German government with 7.1 million Euros, as well as the car's newly open "Automotive Service Bus."

The system is controlled by a cross-platform central control unit built by IAV. A separate, web-enabled control unit responsible for driver and Internet communications communicates wirelessly with a touchscreen, which in the case of the Visio.M is an Apple iPad. According to TUM's announcement, Visio.M's OSGi hardware platform is based on a hardware design that runs Linux on an open-spec PandaBoard, which in turn is equipped with a Texas Instruments 1GHz, dual-core, Cortex-A9 OMAP4430 system-on-chip. However, TUM's announcement does not specify

which block in the architecture diagram below contains the PandaBoard.

Source:
<http://linuxgizmos.com/automotive-bus-open-sourced-with-linux-based-design/>

Submitted by: Arnfried Walbrecht

BLUEBERRY, LINUX MINT'S ELEGANT SOLUTION TO MANAGING BLUETOOTH DEVICES

In what should be termed a good news for heavy bluetooth users and enthusiasts alike, Linux Mint has announced the upcoming release of Blueberry, an elegant, minty front-end to gnome-bluetooth. Aimed to make the task of managing multiple bluetooth devices simpler and hassle-free, BlueBerry is known to provide a unified destination for interacting with wireless mice, keyboards, speakers and the type that rely on bluetooth technology.

Interestingly, although the software has been designed in a

Mint-specific way, it will be independent of any Mint-exclusive dependencies. So, it will run flawlessly on other distros, even those that run different desktop environments, such as Unity, GNOME 3 and Xfce. Since BlueBerry is a GUI front-end, it is being designed to automatically recognize the underlying desktop environment, and switch to the necessary back-end tools without needing the user to specify them.

Blueberry is expected to ship with LMDE's (Linux Mint Debian Edition) second release.

Source:
<http://www.omgubuntu.co.uk/2015/03/linux-mint-bluetooth-set-up-tool-blueberry>

Submitted by: Anirban Chatterjee

VMWARE WANTS AMICABLE END TO 'MERITLESS' LINUX-LIFTING LAWSUIT

VMware thinks it will be possible to find an amicable resolution to the lawsuit alleging it has pinched parts of the Linux kernel.

The lawsuit was brought two weeks ago by kernel developer Christoph Hellwig, who set the ball rolling in his native Germany. Hellwig's complaint alleges VMware has combined code issued under GPLv2 with its own code into products "without providing nor offering complete, corresponding source code for that combined work under terms of the GPLv2."

VMware's now responded to that allegation, saying "We believe the lawsuit is without merit, and we are disappointed that the Software Freedom Conservancy (SFC) and plaintiff have resorted to litigation given the considerable efforts we have made to understand and address their concerns."

Virtzilla's post on the matter goes on to say that "VMware has worked in earnest with the SFC to understand and address their concerns. We did so out of respect to the free and open source software community and we are optimistic that this can be resolved amicably."

Source:
http://www.theregister.co.uk/2015/03/16/vmware_wants_amicable_end_to_meritless_linuxlifting_lawsuit/

Submitted by: Arnfried Walbrecht

GNOME SHELL 3.15.92 IMPROVES GNOME'S CLASSIC THEME

GNOME Shell 3.15.92 was released today (17 March 2015) and among other minor enhancements it improves the GNOME classic theme.

Florian Müllner announced the GNOME Shell 3.15.92 release today and among its enhancements are an improved classic theme, fixed ordering of calendar events, pointer barriers for the legacy tray, menu fixes, GDM changes, and various other items. The classic theme changes come down to items that looked "bad" or improper up to now.

Source:
http://www.phoronix.com/scan.php?page=news_item&px=GNOME-Shell-3.15.92

Submitted by: Michael Larabel

GTK+ WAYLAND NOW SUPPORTS HiDPI CURSORS

Matthias Clasen has released the latest GTK+ 3.15 development release that's near final and about ready to be named GTK+ 3.16.

Yesterday's GTK+ 3.15.12 release brings improvements to the GtkCellRendererPixbuf, GtkEntry and GtkMenuItem alterations, and the GTK+ Wayland back-end now has support for HiDPI cursors. There's also various bug-fixes in GTK+ 3.15.12 like crash fixes and Pixman warnings. The official changes can be seen outlined via this Git commit.

Source:
http://www.phoronix.com/scan.php?page=news_item&px=GTK-3.15.12-Released

Submitted by: Michael Larabel

STATE OF VOIP IN LINUX

Like most people, I find myself using the same VoIP options everyone else is using. Thankfully, these days there are far more options available than what we might think.

One of the popular VoIP applications in Linux is Skype which coming from any other platforms, Linux VoIP clients often find themselves being compared to Skype. Foss advocates are usually quick to point out the flaws in trusting Skype with your voice calls, yet the fact is that this is what most people use. There are more than one alternative applications for VoIP communications in Linux.

Today, I'll look at these options and also explore up-and-coming alternatives as well.

Source:
<http://www.datamation.com/open-source/state-of-voip-in-linux.html>

Submitted by: Matt Hartley

SAVE AND RECOVER DATA FROM CRASHED DISKS WITH THE DDRESCUE COMMAND

Horrible event that really want to avoided is data loss because of broken harddisks. But, you still can do something with your harddisks if that event occurs. By utilizing ddrescue, a good tools for save your data, you still can get back your data.

GNU ddrescue is a program that copies data from one file or block device (hard disk, cd/dvd-rom, etc) to another, it is a tool to help you to save data from crashed partition i.e. it is a data recovery tool. It tries to read and if it fails it will go on with the next sectors, where tools like dd will fail. If the copying process is interrupted by the user it is possible to continue at any position later. It can copy backwards.

Source:
<http://www.cyberciti.biz/tips/how-do-i-save-recover-data-from-crashed-disks-with-dd-and-ddrescue-command.html>

Submitted by: NixCraft

MAKULINUX CINNAMON 8.1

MakuluLinux Cinnamon is a freely distributed, easy-to-use, easy-to-install, portable and open source desktop-oriented operating system based on the award-winning Debian GNU/Linux distribution and built around the beautiful, lightweight and modern Cinnamon desktop environment.

Its claimed as a very first x64 Edition for Makulu Linux family. This release is special for so many reasons, It is sure to mark a major milestone, not just for Makulu, but considering what is inside, the whole of the linux world.

Source:
<http://linux.softpedia.com/get/Linux-Distributions/MakuluLinux-Cinnamon-103650.shtml>

Submitted by: Marius Nestor

OPEN SOURCE VS. PROPRIETARY FIRMS ON THE IOT BATTLEGROUND

The Internet of Things is driving an abundant amount of investment to the middleware tier. This involves activity among larger companies and numerous start-ups in developing platforms. The good news is that many of these platforms are using some common standards, noted Ian Skerrett, vice president of marketing and ecosystem at the Eclipse Foundation.

Technology wars are predictable. Every new wave of gadgetry brings a fight over who will be the next king of the software hill. The next big battle is brewing over control of the Internet of Things marketplace.

The IoT is quietly gaining momentum as companies develop software to connect all sorts of consumer products to the Internet. Consumers see only convenience and extensions to their always-on mobile devices. Product makers see a pathway to streaming data that can be monetized from buyers' connections.

Could history be ripe for repeating itself as open source begins to take on the current, yet unsustainable, walled-garden core of the IoT? Based on the victories in some early skirmishes, innovations developed by open source start-ups may be the David in the here-again fight against proprietary Goliaths.

Source:

<http://www.linuxinsider.com/story/81741.html>

Submitted by: Jack M. Germain

RICHARD STALLMAN'S GNU MANIFESTO IS NOW 30 YEARS OLD

Richard Matthew Stallman is a stubborn man. And it is this trait of his above all else that has spawned the wonderful world of free and open source software, a world that was barely hinted at when he wrote the first document about his intentions.

In true leftist style, Stallman called it the GNU Manifesto. It was published in March 1985 in Dr

Dobb's Journal of Software Tools, a venerable technology publication that shut shop in December last year after 38 years of publishing.

But Stallman's manifesto remains. It encapsulated a desire to create a free operating system, for use by all, one that the users could control. This desire grew out of the fact that UNIX source code was not being released after A. T. & T. was broken up and the anti-trust decree under which it was operating became void. Stallman was not inspired by the path proprietary software development was taking.

He was working at the Artificial Intelligence Laboratory at MIT when he wrote the manifesto. "GNU, which stands for Gnu's Not Unix, is the name for the complete Unix-compatible software system which I am writing so that I can give it away free to everyone who can use it. Several other volunteers are helping me. Contributions of time, money, programs and equipment are greatly needed," was how he began.

Source:

<http://www.itwire.com/opinion-and-analysis/open-sauce/67342-richard-stallmans-gnu-manifesto->

[is-now-30-years-old](#)

Submitted by: Arnfried Walbrecht



In staying with the theme of the last few articles of C&C, I will be introducing two new JavaScript frameworks for web development. AngularJS, which is all client-side, and Meteor - which can be both client-side and server-side. After all this focus on content management systems, I felt it was necessary to mention a few smaller tools which can help add features to static web pages, without having to use a CMS, or creating massive amounts of custom code.

WHAT IS ANGULARJS?

AngularJS is a framework designed to make the creation of client-side web applications easier. It integrates into HTML by adding to its syntax. It allows things such as replacing content without reloading a page, simplifying testing, and offering optimizations to how JavaScript is loaded.

WHAT IS THE DIFFERENCE BETWEEN ANGULARJS AND SOMETHING LIKE NODEJS?

NodeJS is a framework that acts both as a server, as well as a web application in general. This means many things are possible in both systems, but NodeJS won't integrate into a normal static HTML page. As AngularJS is included into a web page in the normal fashion (with a `<script>` tag), it can be added into any web project.

WHERE CAN I START LEARNING ANGULARJS?

AngularJS is included in Code School's list of courses. For anyone who doesn't know Code School - it's a website that offers videos and exercises for learning programming languages. As everything is done in the browser in real-time (including the programming exercises), you're able to easily check your work, or re-watch sections of the videos to better understand concepts. The AngularJS course is free, and an excellent introduction. For those of you who learn best by doing, you can follow along with the

videos in a text editor, and then simply answer the exercises once you reach that stage.

WHY MENTION METEOR THEN?

Meteor is an open source system for creating web and apps in pure JavaScript, and offers libraries to make designing the application faster and more efficient. The main difference is that Meteor also offers a NodeJS-based server for running your code, which is integrated with other applications to make deploying a production system easier, among other things.

WHY CHOOSE ONE OR THE OTHER?

It depends on the scope and aim of the project. If you have a static HTML page, and want to simply add some modern features to it (no refresh, or making certain content dynamic), then AngularJS is best. This is because it can be

added into the existing HTML page, and is relatively lightweight.

If you're writing an entire application from scratch, and need systems in place on both the server and the client side, you'll be better suited to Meteor. The same is true if you're deploying a web application on various devices (phones, tablets, desktops), and it will be more than just an HTML page with some dynamic content.

Lastly, it boils down to personal choice. If both AngularJS and Meteor can do the things you need, pick the one that you like the look of best. If the extra features of Meteor appeal to you, use it.

WHAT ABOUT A CMS?

As AngularJS is not a server-based framework, however, there are some CMS that integrate it into their systems. As AngularJS integrates easily into existing webpages, you could technically use it with any CMS you may already know. In the case of Meteor, there are some CMS that

I've seen: OrionJS, Meteor-Admin, or Azimuth. I have not used any of them, and they seem to vary in size and complexity. Depending on the size of the project, and on what features of a CMS you may need, simply creating a custom system in Meteor may be easier and faster.

WILL YOU WRITE A TUTORIAL ON EITHER ANGULARJS OR METEOR IN A C&C ARTICLE?

As AngularJS is featured in Code School, I feel an article focused on this would be a duplication of effort. However, if there are specific questions, I will be happy to answer them. Also, if there is enough interest in an article on Meteor, I will happily write that too.

Hopefully this has enthused some readers to give AngularJS and Meteor a shot in some of their existing (and/or future) projects. If you run into any specific issues, or are interested in a proper tutorial on Meteor, feel free to send me an email at the following address. Also, if anyone has any questions, suggestions, or requests for C&C articles in general, feel free to

email me at lswest34+fc@gmail.com.

FURTHER READING

<https://angularjs.org/> - Official homepage of AngularJS

<http://campus.codeschool.com/courses/shaping-up-with-angular-js/> - Code school free AngularJS course

<https://www.meteor.com/> - Official homepage of Meteor



Lucas has learned all he knows from repeatedly breaking his system, then having no other option but to discover how to fix it. You can email Lucas at: lswest34@gmail.com.



EXTRA! EXTRA! READ ALL ABOUT IT!

Our glorious news reporters are now posting regular news updates to the main Full Circle site.

Click the NEWS link, in the site menu at the top of the page, and you'll see the news headlines.

Alternatively, look on the right side of any page on the site, and you'll see the five latest news posts.

Feel free to discuss the news items. It's maybe something that can spill back from the site into the magazine. **Enjoy!**



The Ubuntu Podcast covers all the latest news and issues facing Ubuntu Linux users and Free Software fans in general. The show appeals to the newest user and the oldest coder. Our discussions cover the development of Ubuntu but aren't overly technical. We are lucky enough to have some great guests on the show, telling us first hand about the latest exciting developments they are working on, in a way that we can all understand! We also talk about the Ubuntu community and what it gets up to.

The show is presented by members of the UK's Ubuntu Linux community. Because it is covered by the Ubuntu Code of Conduct it is suitable for all.

The show is broadcast live every fortnight on a Tuesday evening (British time) and is available for download the following day.

podcast.ubuntu-uk.org



First, let me thank all the readers who sent me emails of hope and wishes for a quick recovery. They were very kind and helpful. I also want to thank Ronnie, our wonderful editor, for his support and patience during that painful period. I still have issues with sitting for long periods of time, so this is being done over the course of a number of days, so I hope the continuity that I'm trying for works. Now on with "the show"...

Not too long ago, I was walking to the time clock and the General Manager of my "day job" called me into his office. Hoping it was just a "how's it going" talk, I went in and sat down. He then started the meeting with "I'm having a problem with my spreadsheet program, and was hoping you could help me".

As my vision darkened and the three-note ominous orchestral string hits "Da Da DAAAAAAAAA" that we all know from the horror flicks of the 70's and 80's rang through my mind, rather than

running screaming from the room, I innocently asked what was wrong. He responded that there was something wrong with one of the macros and "the thing just quits in the middle of the calculations". As I whipped out my white cowboy hat, I said in my best hero voice "Don't worry citizen. We'll have you up and running in no time." Within a short while, I discovered the reason the spreadsheet was unceremoniously crashing was that one cell in one of 35 workbooks was getting a divide by zero error due to an expected value not being entered in another cell in yet another one of the 35 workbooks. Let me make this perfectly clear, it was not my boss's fault. All he had asked for was a simple way to get the higher-up values from the data. (The previous two sentences have absolutely nothing to do with the fact that my boss may read this article! Or maybe it does.)

As I walked back to my work area, brushing the spurious bits of computer code from my white hat, I realized that this would be an excellent teaching moment. So,

here we are. But first, let's revert back to 1979 when Apple introduced Visicalc. That was the first "Free Form Calculation type system" to really make a hit in the marketplace. While there were many bugs in the software, the world loved the idea and clones (bugs and all) began to pop up on other computer systems, like the Commodore Pet and other Apple competitors (including Microsoft in 1981 with a program called Multiplan). Finally, in 1983, a company called Lotus Development Corp. introduced Lotus 1-2-3. While very close to Visicalc in many aspects, including the menu structure, it was written completely in x86 assembly language, which made it very fast, and many of the bugs of Visicalc were fixed. Lotus 1-2-3 was so popular that it became a common benchmark to test a machine for "PC Compatibility".

The advent of the Free Form Calculation systems, allowed the "normal" person to deal with numbers in a way that previously was in the realm of the

programmer. Almost anyone could, in a few hours or so, make sense of numbers, create charts and graphs, and share that information with coworkers. Shortly after that, the ability to automate some portions of the spreadsheet through Macros and Basic-like embedded languages gave these non-programmer users even more power over their destiny. They could get the answers themselves, and pretty charts and graphs as well, without having to wait in the queue for I.T. assistance. However, as we all learned from Peter Parker's uncle Ben...

WITH GREAT POWER, COMES GREAT RESPONSIBILITY.

Soon the spreadsheet was taken into areas that were better suited for databases than spreadsheets. We now had workbooks upon workbooks that relied on other workbooks, and if one little number along the way didn't happen to get updated... well, we had the old "house of cards" effect.

While I don't think that every spreadsheet is evil, there are some (read this to say 'many') that should have been converted to databases many years ago. They just became too large and unwieldy for their own good. If someone had just sat down with the programmers and said, "Please help", the world would be a kinder, gentler place.

Now as I step down from my soapbox, we come to the real reason for this month's article. Every good Python programmer should have a way to deal with spreadsheets in their arsenal of tools. You never know when you will be called upon to pull data from a spreadsheet and manipulate it. While there are many ways to get data from spreadsheets like using CSV files, which has its own drawbacks, sometimes you need to read and write directly from and to a 'live' spreadsheet. After looking around, I settled on a very nice library to access my boss's problematical spreadsheet.

We will be adding the library called XLRD, which one might imagine stands for eXcel ReaD.

This library allows us to easily read data from Excel files (.xls, .xlsx and .xlsm) from versions 2.0 onward.

Let's create an excel spreadsheet that we can use to examine the functionality of XLRD. Either open excel, or openoffice or libreoffice calc. In the first column (A), enter the numbers 1 to 5 going down. In the next column (B), enter 6 to 10. It should look something like this:

| | A | B |
|---|---|----|
| 1 | 1 | 6 |
| 2 | 2 | 7 |
| 3 | 3 | 8 |
| 4 | 4 | 9 |
| 5 | 5 | 10 |

Now save the spreadsheet as "example1.xls" in the folder you will use to save the test code. This way, we won't have to worry about paths.

Now download and install XLRD (<https://pypi.python.org/pypi/xlrd>). We can use it like is shown below.

Save the file as example1.py in the same folder as the spreadsheet. Since the code is so short, we will simply discuss it here. Of course, the first line imports the library. Then we create a function called OpenFile and pass the name (and path if needed) of the spreadsheet to the function.

Now we call the open_workbook method and get back a 'book' object. Then we use the nworksheets attribute to return the number of ACTIVE workbooks. We can also get the name of the workbooks. In this case, they are the default. We use the

sheet_by_index method to get Sheet1 into the first_sheet object. Now we can start getting data. We get the information from the cell at position (1,1) which translates to cell position B2 (it's Zero based, so cell A1 would be (0,0)). We print the data from there, both what the cell contains and the value, so we could use it in a calculation if we wish.

That was really easy, wasn't it? Now, let's do something a bit more useful. Enter the code shown on the next page (top right) and save it as 'example2.py'. This example will print out the contents of the workbook.

Since we already used the first four lines of code in the first

```
import xlrd
def OpenFile(path):
    # Open and read excel file
    book = xlrd.open_workbook(path)
    # Get number of active workbooks
    print "Number of workbooks: ",book.nworksheets
    # Get the names of those workbooks
    print "Workbook names: ",book.sheet_names()
    first_sheet = book.sheet_by_index(0)
    cell = first_sheet.cell(1,1)
    print "Cell at 1,1: ",cell
    print "Cell Value at 1,1: ",cell.value

if __name__ == "__main__":
    path = "example1.xls"
    OpenFile(path)
```

example, we'll skip them. By using the 'sheet.nrows' and 'sheet.ncols' attributes, we get the number of rows and columns. This can be helpful, not only so we know what we are dealing with; we can write "generic" routines that use those values in our calculations as you will see. In fact, we use 'rows' in a for loop to obtain each row's information.

Notice the line that has 'first_sheet.row_slice'. This gets a block of cells of a given row. The syntax is as follows:

```
X = first_sheet.row_slice(RowInQuestion, Start_Column, End_Column)
```

So we have used the number of rows and the number of columns in calculations. The output from our program should look something like this...

```
There are 5 rows in this workbook.
There are 2 cols in this workbook.
[number:1.0, number:6.0]
[number:2.0, number:7.0]
[number:3.0, number:8.0]
[number:4.0, number:9.0]
[number:5.0, number:10.0]
Press any key to continue . . .
```

```
import xlrd
def OpenFile(path):
    book = xlrd.open_workbook(path)
    first_sheet = book.sheet_by_index(0)
    # Get the number of rows in this workbook
    rows = first_sheet.nrows
    # get the number of columns in this workbook
    cols = first_sheet.ncols
    print "There are %d rows in this workbook." % rows
    print "There are %d cols in this workbook." % cols
    for r in range(0,rows):
        cells = first_sheet.row_slice(rowx=r,start_colx=0,end_colx=cols)
        print cells

if __name__ == "__main__":
    path = "example1.xls"
    OpenFile(path)
```

We'll do one more example before we end this month's article. Go to the spreadsheet and in column C put some dates. Here's what my spreadsheet looks like now:

| | | |
|---|----|-----------|
| 1 | 6 | 1/10/2014 |
| 2 | 7 | 4/15/2015 |
| 3 | 8 | 6/24/1986 |
| 4 | 9 | 9/30/1963 |
| 5 | 10 | 3/3/2000 |

You can use any dates you like. Now let's re-run our example2.py program. Here is the output from mine.

```
There are 5 rows in this workbook.
There are 3 cols in this workbook.
[number:1.0, number:6.0, xldate:41649.0]
[number:2.0, number:7.0, xldate:42109.0]
[number:3.0, number:8.0, xldate:31587.0]
[number:4.0, number:9.0, xldate:23284.0]
[number:5.0, number:10.0,
```

```
xldate:36588.0]
Press any key to continue ...
```

Well, that's not what we expected. It seems that excel holds dates as a value that is simply formatted for whatever we ask it to. This might be helpful for sorting and calculations, but, for showing the actual data, this won't do. Luckily, the writers of the library already thought of this. Delete the line that says "print cells" and replace it with the code shown below.

```
for c in cells:
    if c.ctype == xlrd.XL_CELL_DATE:
        date_value = xlrd.xldate_as_tuple(c.value,book.datemode)
        dt = str(date_value[1]) + "/" + str(date_value[2]) + "/" + str(date_value[0])
        print dt
    else:
        print c.value
```

Here, we go through each cell in the cells list and check the type of the cell to see if it is considered a XL_CELL_DATE. If it is, then we convert it to a tuple. It is stored as YYYY,MM,DD. We simply pretty it up to print it as MM/DD/YYYY. Here is the output of our new program...

There are 5 rows in this workbook.
There are 3 cols in this workbook.

```
1.0
6.0
1/10/2014
2.0
7.0
4/15/2015
3.0
8.0
6/24/1986
4.0
9.0
9/30/1963
5.0
10.0
3/3/2000
Press any key to continue ...
```

Just for your information, there is a library from the same wonderful people called XLWT, which allows you to write to excel files. There is a wonderful tutorial and documentation on these two libraries at <http://www.python-excel.org/>.

The source code for example3.py is on pastebin at <http://pastebin.com/bWz7beBw>.

Hopefully, I'll see you next month.



Greg Walters is owner of RainyDay Solutions, LLC, a consulting company in Aurora, Colorado, and has been programming since 1972. He enjoys cooking, hiking, music, and spending time with his family. His website is www.thedesignatedgeek.net.

PYTHON SPECIAL EDITIONS:



<http://fullcirclemagazine.org/issue-py01/>



<http://fullcirclemagazine.org/issue-py02/>



<http://fullcirclemagazine.org/python-special-edition-issue-three/>



<http://fullcirclemagazine.org/python-special-edition-volume-four/>



<http://fullcirclemagazine.org/python-special-edition-volume-five/>



<http://fullcirclemagazine.org/python-special-edition-volume-six/>



HOW-TO

Written by Elmer Perry

LibreOffice Pt.47 - Linking To The Web

Spreadsheets are a great place to collect data. The web is a great resource for data and much of that data is dynamic. You may even want to collect this data into a spreadsheet and make the spreadsheet mirror the dynamic nature of the web data. Calc allows us to link to external data sources, including web pages, and it will surprise you how easy it is.

Many people trade stocks. Trading requires diligently tracking the stocks. In the past, I have shown this done manually, but today, I am going to show you a way to use a Calc spreadsheet to track the top 100 stocks as compiled by the web site [barchart.com](http://www.barchart.com).

Like so many things, there is more than one way to do this. I'm going to show the method I feel is the easiest for discovering and getting exactly the information

you want. You will import the web page using the Web Page Query filter. Once imported, you can examine the elements of the page to find the data you want to extract. You will then create your sheet of 100 stocks by pulling from your imported source file.

SETTING UP THE SOURCE SHEET

When you import the source using the Web Page Query filter, you are not actually working with a local copy, but a link to the page on the Internet. You will not save this import (but you could). Instead, you are using it as a reference to the actual page on the [barchart.com](http://www.barchart.com) web site.

From the menus, File > Open. In the Open dialog, you will find the file type filter drop-down list at the bottom of the dialog. All the


different file types recognized by LibreOffice are listed here. Scroll through and find the filter named "Web Page Query (Calc)". This will help you create your link to the web page. In the "File name" text box, enter (or copy and paste) the following URL:

<http://www.barchart.com/stocks/signals/top100>


Click the Open button. Calc will take a few seconds then come up with the Import Options dialog. The Automatic option will import the web page "as is." Custom allows you to select another language for importing the page. I recommend using Automatic as it has the greatest chance for proper interpretation. The "Detect special numbers (such as dates)" checkbox

does what you think it does. When checked, it will detect dates, times, etc, and format them appropriately in the sheet. We do not have special numbers we are concerned about, so we don't need it. Click OK to accept the options. Depending on your Internet connection, computer speed and memory, Calc will take a few seconds to import the web page. Notice that it imports the entire web page minus the images.

NOTE: This is a link to the actual web page. Nothing has been saved on your local computer at this point.



THE NAVIGATOR



We will use the Navigator to examine the imported page and link the data to our new document. The Navigator shows you the elements and structure of a document and changes depending on the type of document you have open. Not only is it good for our purposes here, but for navigating and manipulating a document, especially a large document. Currently, there are four ways to open the Navigator:

- Press F5 on your keyboard
- View > Navigator
- Navigator tab on the sidebar
- Navigator button on the standard toolbar




EXAMINE THE SOURCE

When a web page is imported, several range-names are created, as well as importing named tables on the web page. The prefix HTML_ is added to any names imported from the page. The filter creates two special ranges, HTML_all and HTML_tables. HTML_all allows you to select the entire document. HTML_tables allows you to select all the tables. The problem you run up against is the creator of the web page probably wasn't thinking about you importing his page into Calc. Imported table names were for the creator's reference not yours.

If you open the Navigator using one of the methods listed above, in our source document, you will discover a list of names under the "Range-names" item in the Navigator. Double-click on a range-name to highlight it in the document. As you go through the list, you soon discover that the names HTML_4 and HTML_dt1 both contain the table with our stock data. You will use one of these to create your stocks document.

IMPORT TABLE TO NEW SHEET

Now that you know what range-name to import for the stocks data, it's time to create the sheet you want to save and keep. You will import from the source document into a new document. Remember that your source document is just a link to the actual web page, nothing has been saved to your local computer yet.



From the menus, File > New > Spreadsheet. Use any of the four methods discussed above to open the Navigator. From the documents list at the bottom of the Navigator window/panel, select the source document, top100. Click on the drag-mode icon in the Navigator toolbar and change the setting to "Insert as link." Expand the entries under the "Range names." Select either HTML_4 or HTML_dt1, drag it to cell A1 in the new document, and release. After a few seconds, depending on your internet and computer speeds, the stocks data will appear in your new sheet. Save the new document. You can now close the source document. There is no need to save it. Your new document is actually linked to the page on the web site and not the source document.

You can close your new document, and when you open it, you are prompted about whether to update the links in the file. If you answer Yes, Calc will retrieve a fresh version of the page and update the data in your sheet. You can set the file to update periodically when it is opened. In the menus, Edit > Links brings up the Edit Links dialog. With the link selected, click on the Modify button to bring up the External Data dialog. Check the "Update every..." check box, and fill in the number of seconds between every update. For example, every five minutes is 300 seconds, ten minutes is 600 seconds, and an hour is 3,600 seconds. Click OK to save your changes and Close to close the Edit Link dialog.

Importing dynamic data from a web page in Calc is pretty simple. Use the Web Page Query filter to link the page to a sheet. With the Navigator, you can examine the page to determine which range-name contains the data you need. Once you know the range you need to use, you use the Navigator to drag the range-name into a new sheet and save the new document. The new document will prompt you to update every time you open

the document or periodically, should you assign an update period to it.



Elmer Perry's history of working, and programming, computers involves an Apple][E, adding some Amiga, a generous helping of DOS and Windows, a dash of Unix, and blend well with Linux and Ubuntu. He blogs at <http://eeperry.wordpress.com>



LaTeX (pronounced laytech) is an enhancement of the TeX document creation software. TeX is the creation of Donald Knuth who is a computer programmer. TeX was released in 1978 and later TeX was improved upon by Leslie Lamport in 1984. That is why it is called LaTeX. LaTeX has been around for a long time, and I call LaTeX the Great Grandma of desktop publishing. LaTeX is more structured than modern desktop publishing applications. There are many editor applications for creating LaTeX documents, and it can be easy to call LaTeX a word processing program. Nothing is further from the truth – LaTeX is a language and it is a language that has to be learned in order to be successful at creating documents. LaTeX has a steep learning curve, and some people claim that it is so much faster and easier to use a word processor. Except that: *"Many people discover LaTeX after years of struggling with word processors and desktop publishing systems, and are amazed to find that TeX has been around for nearly 25 [+] years and they hadn't heard of it. "It's not a*

conspiracy, just a well-kept secret known only to a few million people", as one anonymous user put it." – Peter Flynn

People who have had enough of the problems encountered with word processors, especially Microsoft Word, have found LaTeX to be much easier to use, and a more productive, stable method. At the same time that I am writing this, two researchers from Germany published a paper comparing Latex to Microsoft Word. Their findings favoured Microsoft Word for efficiency, but not only that: the two authors of the paper suggested that researchers at universities be prevented from using Latex. Many scholars weighed in on this strange conclusion and found the study to be flawed, if not manipulated, to favour Word.

As you can imagine, there was a firestorm over this one; I thought that this was one of the best comments: *I've shown people LaTeX documents of mine, the content of which they are completely ignorant,*

and they always remark how pretty it looks. I've never gotten a compliment like that about a Word document! – David Kotschessa, University of South Florida, Tampa.

See this site if you want to read more about this debate: <http://lemire.me/blog/archives/2015/01/14/knauff-and-nejasmic-recommend-banning-latex/#comments> (the site has a link to the original article).

Forty years ago when I was at university, I would have thought I had died and gone to heaven if I had LaTeX to write my papers with. Over the years, I have taken some stabs at LaTeX but it was after I completed a beamer slide project using LaTeX on my Raspberry Pi that I realised that LaTeX has a lot of power with very little overhead. Now I use LaTeX as often as I can, almost abandoning LibreOffice entirely. I have to admit that there are some things that are very easy to do in LibreOffice compared to LaTeX, but the end product from LaTeX is so much better.

The Latex community is a good example of what the free software movement is all about. It has, in a way, made this article difficult to write as there is so much good help out there on the Internet.

WHY BOTHER TO LEARN LATEX?

- There is no Meta-data in a Latex file. In some work environments, that can be critical.
- I have seen many testimonials in favour of LaTeX by students who, after years of using MS Word, switched to LaTeX to write their thesis at university.
- Forget Powerpoint; the Beamer class of document in LaTeX creates excellent slides. Save as a PDF document and present it anywhere. I have seen enough Powerpoint presentations get messed up by the presenter not really knowing how to use the application.
- Latex can be run on very limited hardware. You can run Latex on an old XP laptop with 500 megs of RAM or on a Raspberry Pi

computer.

- It is a new skill, probably much more valuable to a young person's life than top scores playing computer games. Give a young person a "LaTeX challenge" to improve their knowledge and get them off of the joystick for a change.
- It is so cool that we can create all kinds of fancy stuff with an ASCII file.

GETTING STARTED

Anyone who has done coding in HTML will be able to pick up on LaTeX much easier than someone who has not, or has created HTML only in a WYSIWYG editor.

The best method to install would be to go to the Ubuntu Software Centre and install TeX Live. If you want to be on the cutting edge of Tex, you could install TeX Live from the Tex Users Group. Either way will give you the files required. The instructions are here:

<https://help.ubuntu.com/community/LaTeX>.

SELECTING AN EDITOR

It is necessary to select a LaTeX text editor for creating LaTeX documents. There are lots to choose from within Synaptic or the Ubuntu Software Centre and recommending one is the stuff that flame wars can be made of.

I prefer any editor that has pull-down menus listed in a logical way – the formatting code required to create documents. Latex is a huge environment with lots of code out there. As we are learning, it is helpful to be able to see the code we need. It is also helpful to have a built-in viewer to display our document as it will be printed.

There are three LaTeX editors that I have liked: LaTeXila, TexMaker, and the text editor Geany with LaTeX extensions. LaTeXila would be a good choice to get your feet wet as it looks like many text editors that you have used before. TexMaker looks confusing and complicated, but after you get adjusted to the wide selection of buttons and menus, it is an easy program to use as well. Geany worked well on my Raspberry Pi. You may have to

install "latexmk" to get LaTeXzila to run.

After you have downloaded the software, I would recommend watching this video: 'Learn Latex in Five minutes'

<https://www.youtube.com/watch?v=Y-kXtWdjt看> – now I have let the cat out of the bag! You can now see that there are lots of Latex HowTo videos on YouTube.

Before you write your first LaTeX document, one initial bit of advice: always create a folder for the document that you are working on and keep all of the files related to the document in it. This will serve you well when you try more complex documents. Go get your feet wet, and we will come back to this next issue.

At least now you are in on one of the world's well kept secrets.



John Eddie Kerr is a Law Librarian at a county law library in Guelph, Ontario, Canada. Ubuntu powers his desktop at work and at home. He is a member of the Kitchener-Waterloo Linux Users Group and the WFTL-LUG.

Get *unlimited* access to a cutting-edge technology and business library with **Apress Access!**

For **\$199**

YOU GET:

- Unlimited access to Apress titles for a full year
- Instant access to each new Apress publication
- Compatibility with any device—desktop, laptop, or mobile
- Use of our new exclusive-to-Apress reader with unparalleled search functions
- Option to download any eBook for just \$4.99 for a limited time



www.apress.com |  @apress

Want more info? Check out www.apress.com/subscription





HOW-TO


Written by Mark Crutch

Inkscape - Part 35

Cast your mind back to part 30 of this series where I introduced the notion of an “unset” fill, which allows each clone to have its own color that is independent of the parent object. It's a handy trick for creating a collection of similar-but-not-identical objects, such as a crowd of characters with differently colored hair or clothes. You can use this same mechanism with the Tiled Clones dialog to produce arrays of clones whose colors differ from the parent object either in subtle drifts of shade and hue, or in big, bold steps.


As usual, we'll start by drawing a simple parent shape – our familiar round-cornered square. But, rather than fill it with color, we'll unset it either by using the “?” button in the Fill tab of the Fill & Stroke dialog, or by right-clicking on the color swatch in the bottom-left of Inkscape's status bar and selecting “Unset fill”. We'll also use the Reset button at the bottom of the Tiled Clones dialog to get back to a sensible set of defaults, regardless of your

experimentations as a result of the previous two articles. With all the preparations in place, let's take a look at the “Colour” tab (in my British English installation) of the Tiled Clones dialog.




The general layout of this should be familiar by now, but the details differ a little when compared to the tabs we've looked at previously. The first change is the addition of the “Initial colour” field at the top. Clicking the swatch there opens a color picker from which to choose the initial

color that your clones will take. It's “initial” because the rest of the fields can subsequently change the color quite drastically. With everything else in this tab set to zero, clicking the Create button will produce an array of clones, all taking on that initial color. The visual effect will be no different to cloning a solid-colored parent object, so, in this case, we end up with an array of red squares. I've moved the parent out from under the first clone a little, so you can see that its own color remains unset.



The remaining fields in the tab allow us to change H (Hue), S (Saturation) and L (Lightness) for each row and column, with the usual options for randomise and

alternate. If you're not very familiar with the HSL color model, it's perhaps best explained by looking at the “Wheel” tab on any of Inkscape's color pickers. Yes, there's also a dedicated HSL tab, but although I find it to be the more useful for day-to-day use, the wheel view is a better explanatory tool.



Hue, the first of our three values, represents a position on the outside circle. You might expect a value from 0° to 360° – or

the equivalent in radians if you're more mathematically inclined – but that would be too sensible. Instead the range of values available varies in different parts of the Inkscape interface. Within the HSL tab, for example, the numbers run from 0 to 255. Within the Tiled Clones dialog, however, they run from 0% to 100%. In either case, 0 represents pure red with increasing numbers progressing anti-clockwise through yellow, green, blue and purple before the end of the scale brings you back around to red.


Having picked a base Hue, the triangle in the center is used to select a combination of the Saturation and Lightness values. With the hue set at 0 (pure red), the triangle is oriented as shown in the screenshot. Now imagine a pair of axes, one running from the pure colored corner of the triangle to its opposite edge (a horizontal line in this case) and another running along this edge between the two other corners (a vertical line). Saturation is the position along the first line, and defines the amount of the pure color that's present in the final swatch – how “washed out” it appears. Lightness is the position on the second line,

representing how dark or light the color is. When Saturation is zero there is none of the pure color present, so the result is a shade of gray that can run from pure black (when Lightness is zero) to pure white (when Lightness is at its maximum). The ranges for Saturation and Lightness also run from 0 to 255 on the HSL tab, or 0% to 100% in the Tiled Clones dialog.

The important thing to realize is that the Hue can wrap round – a value of 50% gives you exactly the same pure cyan as 150% or 250%. Saturation and Lightness don't wrap: values above 100% won't suddenly wrap round to lower values, but neither will they result in extra saturation or extra lightness. Values less than 0% behave similarly.

With all that in mind, let's put a value of 25% in the Per Column “H” field. We've got four columns, so the colors will be picked from our color wheel at positions of 0, 25%, 50% and 75%, working anti-clockwise from your selected Initial Color – pure red in this example. It should be easy to see that this gives us red, green, cyan and purple for the columns of our

clones.




Can you work out what will happen if we change the number of columns to 8? Remember that the Hue value can wrap round. How about if we use a value of 33.3%, 50% or something else entirely?

Now try putting a value of -50% into the Per Row “S” field. With each row you'll get less and less of the pure color included. Given that our starting color is already pure red at 100% saturation, this gives us values for our three rows of 100%, 50% and 0%, resulting in rows that are pure colors, washed out colors, and completely gray. Given that the Saturation value doesn't wrap, can you guess what the result would be for more than three rows? Also try picking an initial color with low saturation and then putting a positive value in the

field instead.


Finally, let's reset our initial color to pure red, and play with the Per Row “L” field. You might expect that putting -50% in here would have a similar effect to the Saturation, giving values of 100%, 50% and 0% for rows that are bright, dark, then black. Instead you get this:



The issue is that the Lightness scale runs from 0% (black) to 100% (white) – pure red, of course, has neither too much white nor too much black, so its value is actually 50%. Thinking of Lightness as running along a vertical line in the earlier color wheel image, it's easy to see that the red corner of the triangle lies 50% of the way up. Checking the HSL tab will also show that your pure red color has a Lightness of 128 (out of 255). Now you should be able to see that a

HOWTO - INKSCAPE

value of -50% in the field leads to rows of 50%, 0%, 0% (Lightness doesn't wrap either). -25% will give us the expected outcome.




Try creating a larger array of clones with small values in the fields to get gently sweeping changes in color or tone. Or use bigger values – especially in the “H” field – to get bold differences between the clones. Finally, try drawing a simple leaf with veins but an unset fill. Group the parts, then use the tiled clones dialog to create an array of them. With a little use of the Random fields in each of the tabs we've discussed so far – plus some negative offsets in the Shift tab to pull everything together a little – you can quickly and easily create an autumnal forest floor background.

You may recall that it's possible to unset the stroke of a parent



object as well as its fill. This also works with the Tiled Clones dialog, but, as there is only one Colour tab, there's no way to use different generated colors for the fill and stroke: you can unset the fill color, the stroke color, or both, but anything that's unset will be given the same generated color. There's also no way to set any of the other stroke parameters through this dialog – although you can manually set them on each clone afterwards. This means that the Tiled Clones dialog isn't a great help if you want to create hundreds of clones whose stroke width or line style

varies. In this final example, I've cloned a yellow rounded square with an unset stroke, but the



different widths and dash styles of the strokes all had to be manually set afterwards via the Fill and Stroke dialog.

I had promised to cover the Trace tab in this instalment, but the Color tab ended up being a more nuanced topic than I had previously expected, so the Trace tab has been postponed until next time.



Mark uses Inkscape to create three webcomics, 'The Greys', 'Monsters, Inked' and 'Elvie', which can all be found at <http://www.peppertop.com/>




HOW-TO

Written by Ronnie Tucker

Several months ago, I bought a couple of cheap screens for use with my Arduino. I got them from Banggood for a couple of pounds/dollars each. One was a Nokia 5110 style screen, the other was a 1.8" Color TFT screen which does 160 x 128 pixels:

<http://www.banggood.com/1-8-Inch-Serial-SPI-TFT-LCD-Display-Module-With-Power-IC-SD-Socket-p-909802.html>



It's the color screen I'd like to discuss this month.

Coming from cheapo sellers such as Banggood (or if you use some eBay/Amazon sellers), most stuff has scant information. Having

scoured around Banggood comments, and on Google, I eventually managed to find the right library, adjustments, and code, to get my screen to work.

First, the libraries. You can grab them (and some documentation) from:

http://devacron.com/QDtech_TFT180A_S6D02A1%20LCD%20Module.zip

To install the libraries: go to the Arduino programming interface; in the menu, go to Sketch > Import Library > Add Library; and point it to the ZIP file that you downloaded. If you go to File > Examples > TFT > Arduino, you'll see some example code.


Some screens that I've bought in the past (such as the Nokia 5110) come with the header pins (the pins you stick into the breadboard) as a separate piece that you need to solder on. This isn't much of a problem, but this color screen came with the header pins attached. So, plug the screen into your breadboard, wire up your

Arduino 5V and GND pins to the relevant strips on your breadboard, and let's get the screen on.

First, the important part that's not mentioned in most documentation: you must use 1K resistors on all the data lines. In other words, all but the 5V, BL, and GND. You also need to connect the BL pin to the 5V line.

Now, with it all wired up, we need some code. This code I'm going to show you is really for next month (where I'll add an ultrasonic sensor), but it is still a good foundation of how to get the screen to show you something.

Code is at: <http://pastebin.com/UYQe58xB>



A couple of quick pointers and notes on the code:

- You can ignore the <NewPing.h> – that's for next month. Also ignore the #define lines and the NewPing line below them.
- The commented lines (beginning with //) are the pinouts for the screen.
- The extern lines are for the font(s) used on the screen. SmallFont() is about the best, but you can chop and change between it and, say, BigFont() in the code.

The setup() is exclusively for the screen. All commands with myGLCD are for the screen. Let's look at a couple of them:

We begin by initialising the screen:

```
myGLCD.InitLCD();
```

Telling it we want to use the small font:

```
myGLCD.setFont(SmallFont);
```

Next, we clear the screen:

```
myGLCD.clrScr();
```

Pick a color to use in the next command:

```
myGLCD.setColor(255, 255, 255);
```

A very important thing to note here is with the color values. Normally you would use RGB, but this screen uses BGR. Using 255,0,0 won't get you red, it'll get you blue.

Then we print 'Distance in cm:' one pixel down and in the center of the screen.

```
myGLCD.print("Distance in cm:", CENTER, 1);
```

The rest of the code is mainly for the sonar ping and displaying the distance, but we'll get to that next month.

For now, feel free to read through the PDF's in the ZIP you downloaded. It has some interesting commands in there that will let you draw shapes and even set the screen display to landscape/portrait mode.



Ronnie is the founder and (still!) editor of Full Circle. He's a part-time arts and crafts sort of guy, and now an Arduino tinkerer.



GUIDELINES

The single rule for an article is that **it must somehow be linked to Ubuntu or one of the many derivatives of Ubuntu (Kubuntu, Xubuntu, Lubuntu, etc).**

RULES

• There is no word limit for articles, but be advised that long articles may be split across several issues.

• For advice, please refer to the **Official Full Circle Style Guide:** <http://url.fullcirclemagazine.org/75d471>

• Write your article in whichever software you choose, I would recommend LibreOffice, but most importantly - **PLEASE SPELL AND GRAMMAR CHECK IT!**

• In your article, please indicate where you would like a particular image to be placed by indicating the image name in a new paragraph or by embedding the image in the ODT (Open Office)

document.

• Images should be JPG, no wider than 800 pixels, and use low compression.

• Do not use tables or any type of **bold** or *italic* formatting.

If you are writing a review, please follow these guidelines :

When you are ready to submit your article please email it to: articles@fullcirclemagazine.org

TRANSLATIONS

If you would like to translate Full Circle into your native language please send an email to ronnie@fullcirclemagazine.org and we will either put you in touch with an existing team, or give you access to the raw text to translate from. With a completed PDF, you will be able to upload your file to the main Full Circle site.

REVIEWS

GAMES/APPLICATIONS

When reviewing games/applications please state clearly:

- title of the game
- who makes the game
- is it free, or a paid download?
- where to get it from (give download/homepage URL)
- is it Linux native, or did you use Wine?
- your marks out of five
- a summary with positive and negative points

HARDWARE

When reviewing hardware please state clearly:

- make and model of the hardware
- what category would you put this hardware into?
- any glitches that you may have had while using the hardware?
- easy to get the hardware working in Linux?
- did you have to use Windows drivers?
- marks out of five
- a summary with positive and negative points

You don't need to be an expert to write an article - write about the games, applications and hardware that you use every day.





- ➔ Access all your data in one de-duplicated location
- ➔ Configurable multi-platform synchronization
- ➔ Preserve all historical versions & deleted files
- ➔ Share folders instantly in web ShareRooms w / RSS
- ➔ Retrieve files from any internet-connected device
- ➔ Comprehensive 'zero-knowledge' data encryption
- ➔ 2 GBs Free / \$10 per 100 GBs / Unlimited devices

<https://spideroak.com>

Online
BACKUP

Secure
SYNC

Easy
SHARING

Whether you need to access a document you have stored on a remote server, synchronize data between a Mac, Windows or Linux device, share important business documents with your clients, or just rest easy knowing all of your data is safely, securely, and automatically backed up - SpiderOak's free online backup, online sync and online sharing solution can handle all your needs!

SpiderOak offers a different approach to online backup by combining a suite of services into one consolidated tool - free online backup, synchronization, sharing, remote access, and storage. This difference is further measured in our zero-knowledge privacy policy - the first one ever employed in this setting. Our flexible design allows you to handle data from any operating system (Mac, Windows and Linux) or location (external drives, network volumes, USB keys, etc...) using just one centralized account.

Download mobile clients
for **iOS & Android**

JOIN SPIDEROAK NOW
Get 2 Free GBs

Get 25% off any SpiderOak package
with the code: **FullcirclemagFans**



Nowadays it is the time of the Cloud. What is the Cloud? It is not pretty clear. By the way, for home users, the most common meaning of the Cloud is file synchronization between devices, like Dropbox, Google Drive, and so on

If you don't trust these solutions, e.g. you have to store sensitive data, or if you need to sync files only inside your network (you don't need these files to be accessible from the Internet), and also while achieving better network performance, or you want as much storage space as possible without spending a cent, then you have to install some software on your own, on your own home machines.

An interesting and pretty new software to perform file synchronization is Syncthing (<http://syncthing.net/>). It is fully Open Source. It is pretty simple to install and easy to use. It is a very promising piece of software, useful to share and replicate your files between your devices using a



Syncthing

P2P architecture, and since it is multi platform, you can synchronize your files among Windows, Linux, OS X, Android, Raspberry-Pi, Solaris and openBSD.

It doesn't require any dependency, so you don't have to install and configure Apache, PHP and MySQL like you have to do with Owncloud or Seafile. On the other hand, Syncthing doesn't provide a plethora of features like those programs: it simply synchronizes files between your machines.

HOW TO INSTALL SYNCTHING ON UBUNTU

Installing and configuring Syncthing on Ubuntu is pretty straightforward. Go to

<https://github.com/syncthing/syncthing/releases> and download the release suitable for your architecture (32 or 64 bits).

```
cd /var/tmp/

wget
https://github.com/syncthing/syncthing/releases/download/v0.10.21/syncthing-linux-amd64-v0.10.21.tar.gz

tar xzvf syncthing-linux-amd64-v0.10.21.tar.gz

cd syncthing-linux-amd64-v0.10.21/
```

In this directory, you can find some documentation (FAQ and Getting started PDF documents), and some sample startup scripts.

There is not a default place where to put the executable file. But in order to perform automatic

updates, the directory (and the executable itself) should be writable by the user running the service. Yes, the service is not something system wide. The Syncthing process is related to a single user. So, potentially, if your PC is shared between multiple users, every user should run a separate instance (or copy) of the service. But, here, let's suppose that you are the only user of your PC.

```
sudo mkdir /opt/syncthing
sudo chown youruser /opt/syncthing

cp syncthing /opt/syncthing
```

Once that is done, you can delete the TAR file and the directory in /var/tmp.

Now, in order to start the Syncthing process at every reboot, you need to configure some Linux initialization script: systemd, runit, rc.local. You can find various examples on the Syncthing forum (<https://discourse.syncthing.net/>).

Here below I will show you how



to configure the default init daemon on Ubuntu 14.04.

Let's create a new configuration file

```
sudo vi /etc/init/syncthing.conf
```

containing these lines shown right.

Please use your username, your home directory, and the path where you have placed the syncthing executable.

Then start the service:

```
sudo initctl start syncthing
```


To configure it, Syncthing offers a clear and clever web GUI. You can eventually find the configuration files in the following directory: `/home/youruser/.config/syncthing/`

So open a web browser and connect to <http://127.0.0.1:8080>

Delete the default shared folder. Then add a new one: give it a Folder ID (it should be the same on every node), and select a path; if it doesn't exist it will be automatically created.

Here you can enable file versioning, and you can decide if the folder will be shared as read-only to the other nodes. Now click Save, then click on restart.

At this point you should change some settings. Click on the gear wheel in the upper right corner, then select "Settings".



```
description "Syncthing service for user youruser"

start on (local-fileSYSTEMS and net-device-up IFACE!=lo)
stop on runlevel [!2345]

env STNORESTART=yes
env HOME=/home/youruser
setuid "youruser"
setgid "youruser"



exec /opt/syncthing/syncthing

respawn
```

Here you can set a different "Device Name," and you can configure the web GUI to listen on the Ethernet interface (i.e. use the 0.0.0.0 address) rather than on localhost only, and on which port (default is 8080). If you enable such options, even if you are on a private network, it is a good thing to set up authentication (User and

Password) and HTTPS. Then you can define some bandwidth limits related to this device (Upload and Download rate limit).

Since I've not tested it yet, and it is out of the scope of this article, remember to disable Global Discovery.



Global Discovery should be useful in order to access your private cloud from the Internet, and in order to share the files with your friends around the world, but, as stated already, this functionality will not be covered here.

When requested, click again on the restart button.

Now it is time to install Syncthing on another node. Follow the same steps also on the other PC in your network (if they run Ubuntu Linux), otherwise download and install the Windows package, or the Mac one. Delete the default shared folder and stop here.

OK. Now, click on "Show ID".

Here you will see an ID and a QR-Code. You can scan the QR-Code with your smartphone in order to configure the Android app, very handy! But for now, take note (copy it in the clipboard) of the very long ID string.

Go back to the first configured node, click on "Add Device" and write the ID of the second node in the Device ID text box. Then check the box of the folder you want to share with this device.

On the other node web GUI, you should see a request that the first node wants to connect and that it wants to share a folder. Click on "Add" and then specify the folder path.

You are done!

Put a file on the shared folder on one of the two nodes and magically it will be replicated on the other one.

Following the above instructions, at this point you can add as many nodes as you want. And you can share a folder between some nodes, but not with others: for instance, if it is not necessary to have a folder shared with a specific node, you can prevent useless bandwidth wasting.

WHAT'S NEXT


Well. We have looked at the

hard way; now we can look to the unofficial Syncthing PPA repository:
<https://launchpad.net/~ytwld/+archive/ubuntu/syncthing>

And you can try the GTK GUI that you can find here <https://github.com/syncthing/syncthing-gtk> (also in this case there is a PPA repository).

This is very interesting: it provides configuration capabilities, like a first run wizard, and more.

It will integrate with the desktop notification area, and with Nautilus as well. In addition, it watches the filesystem in order to perform instant synchronization after changes on the synchronized



directory.

MULTI PLATFORM

As said, Syncthing is multi platform: Windows, Mac, Solaris (and derivatives, like SmartOS), openBSD. There is also a version for Linux on ARM platforms, so you can install it on devices like the Raspberry Pi. And on the Google Play (or on the FOSS F-Droid market) you can find an app for Android. There is not an iOS app right now.

CONCLUSION

As you have seen, Syncthing is very simple to install and configure. With this software you can only share and replicate files among your devices, a job that it performs pretty well.

On the other hand, other than the configuration web page, there is not a web interface that allows you to access and manage files via the web: so you can't access or upload your files from a web browser, like with Dropbox, Google Drive or Seafile and Owncloud.

And you don't have other

facilities, like the ones you can find on products like Owncloud: shared calendar, address book, task scheduler, and so on.

But if all that you need is to share your documents and pictures between all your devices (and eventually share files with your friend without relying on a third party infrastructure), Syncthing is what you need.



Alessio is an unpretentious sysadmin at FTGM, a specialist cardiopulmonary health foundation in Italy. Linux and FOSS are not his hobby: they are a job. Sometimes he blogs at <http://blogoless.blogspot.it>

[PACKT] CLASSICS

DID YOU MISS ANY?

New to Packt? Check out some of our all time classics to build your essential learning library and make sure you start 2015 ahead of the curve.

<https://www.packtpub.com/books/packt-classics>



REVIEW

Written by Ronnie Tucker

BQ Aquaris E4.5 : Ubuntu Edition


Well, I threw out a rope asking for questions, and you guys threw me back a noose. So many questions! I'll try to answer as many as I can throughout the review, and I'm hoping I can get an interview with someone from Canonical to answer the more technical questions.

I'm purposely going to skim through the obligatory first switch on and basic usage as these things are mentioned in almost every review. I want to get to the real meat of things which is what you folks have asked.

TURNED ON

The first thing to do with the phone is to insert a SIM card. This is the first unusual thing about the Ubuntu Phone; it can take two SIM cards. Also, since many of you asked, yes, it is unlocked. When you first switch on the phone, don't panic about the long boot time, it's much faster thereafter. You are

asked to sign in using Ubuntu One. That confused me. I know Ubuntu One only as the cloud hosting that was closed down a while back. Anyway, I signed up for an account and was signed in to my phone. You can also add accounts from Evernote, SoundCloud, Flickr, Facebook, Vimeo, and Google.




asked to sign in using Ubuntu One. That confused me. I know Ubuntu One only as the cloud hosting that was closed down a while back. Anyway, I signed up for an account and was signed in to my phone. You can also add accounts from Evernote, SoundCloud, Flickr, Facebook, Twitter, Fitbit, Instagram, Vimeo, and Google.

SWIPING

Since using Ubuntu on a phone is different from Android and iOS, you're greeted with a tutorial that tells you all about the swiping. Ubuntu Phones use swiping from all sides of the screen and can have either a short swipe, or a long swipe. For example: a swipe for a short distance from the left to the right will bring in the side menu, but swiping longer will wipe away the current screen and return you to the Today scope. A short swipe from the left will switch to the previous app, a longer swipe will show a carousel of previously run apps to switch to.

SCOPES?

Speaking of scopes. Scopes are effectively screens that you can swipe between. There's no desktop on Ubuntu Phones so




you're greeted with the Today scope which shows you basic info such as weather, missed calls, news, etc. It's known as an aggregated scope as it pulls its info from elsewhere. Swipe right to left and you'll see the next scope, and so on. Swiping up from the bottom of the

REVIEW - BQ AQUARIS E4.5


screen will get you a list of possible scopes which you can add by pressing the star icon, or hold down on a scope in the list to rearrange your scopes.

Available scopes range from Amazon to eBay, BBC News to Flickr, Grooveshark, Soundcloud, Vimeo, YouTube, and everything in between. Just remember though: most of these scopes are really just




links to their mobile website in a pretty webapp.

THE APPS




The phone comes with apps for phone calls, SMS, contacts, camera, gallery, playing media, reminders, tasks, but also webapps for Amazon, eBay, Facebook and Twitter. It even comes with Cut The Rope, the

first well known title to get an Ubuntu edition.




For messaging, a lot of you asked about WhatsApp. While there's no version, as of writing; the phone does have Telegram which lets you message in a similar vein to WhatsApp, but it does require that your friends also have Telegram installed.

Regarding contacts. While I did add my Google accounts, it didn't sync my contacts across. It did say something about syncing, but my contacts were (and still are) empty. You do need to enable access to G+, Gmail and Contacts in the account settings. Even with them on, I got no contacts.




What about you, travellers... does it have maps to stop you

from getting lost? Yes it does. It uses HERE Maps by default. Of course, you can install a Google Maps webapp if you like. HERE (from Nokia) has similar functionality to Google Maps - with satellite imagery and navigation. You can have map/satellite views - with public transport and live traffic layers if required. Navigation can be via car, walking or public transport.



ADDING APPS

In the Apps scope, at the bottom, there's a nice big red button to take you to the Ubuntu Store. It's from here that you install new apps. Initially you see an 'app of the week'; then scrolling down you see top apps, game of the week, etc. At the top is a drop-down menu with categories, or you can press the magnifying glass icon to search for apps.



Having found your app (reviews and ratings are shown on the app page), you simply click the 'Install' button and wait for your app to install. Easy as that.

Many of you asked about Google apps. Well, the good news is that there are apps to get you Gmail, Maps, Drive, and the like. One particularly handy app is simply called Gmail (the app is by Canonical), and it's a webapp that grants you access to Gmail, G+, YouTube, et al, and lets you sign in with multiple accounts and so on. Very handy. There's even an app for Google Drive. You can't edit documents with it, but you can view them.

Speaking of cloud storage. Dropbox? Yes, there's an app for that too. Several of you asked about OneDrive. Yes, there's one for that too, and, while on the subject of Microsoft, there's also a webapp for Outlook. As the creator JoshStrobl says: any missing functionality is the fault of Microsoft. These webapps are, again, mobile web pages. What you get here is what you

get in a mobile web browser.



Several of you crazy people asked about the availability of a terminal within the Ubuntu Phone. I'm happy to report that there is a terminal app that can be installed. Now, I'm no expert (not by any means) at the terminal, so I tried only some basic listing and directory commands, but they all worked as they would in Ubuntu (desktop). I typed apt-get and it

gave me the help for it. Same with ssh. So it would seem that it's a fully functioning terminal.

I've yet to find a document viewer app that will display .DOC or .ODT files, but there are several text and PDF viewers that work just fine and at least one app that lets you create markdown documents in ODT/PDF.

If you're curious about what's in the Ubuntu Store, there's an unofficial site here: <https://appstore.bhdouglass.com/apps> which shows the apps available.

SETTINGS

SYSTEM SETTINGS > NETWORK

This lets you turn on/off WiFi and, of course, connect. Mobile lets you turn on/off mobile data and roaming. It's here that you see the first mention of 2G/3G. That's right. No 4G. But that doesn't worry me as I have no 4G reception where I live. Bluetooth. This is where it's a

bit odd. It seems that while the Bluetooth is, technically, functioning, it won't accept incoming files/connections. I got it to recognise my Nexus 5, but it wouldn't accept anything I sent over. Looking at the specs, it says (and I quote): "software not currently integrated". I can only hope they integrate it soon. I hardly use Bluetooth (if ever) but it seems a lot of you like it.

SYSTEM SETTINGS > PERSONAL

In here, you can change the default background and set your sounds for ringtone and messages. There are pages of languages to choose from. I'm pretty sure you're covered here. Accounts is where you add your Facebook/Google/UbuntuOne, and notifications is where you can turn on/off popups from Telegram, Gmail, Twitter, etc.

SYSTEM SETTINGS > SYSTEM

Pressing the Battery icon gives you a nice graph of your

battery usage while Brightness lets you alter just that. Phone lets you turn on/off button sounds, while time & date let you alter that too. Security & Privacy is certainly something a lot of you were concerned about and in here is where you can set a lock type (code or passphrase) and time until lock. Set a SIM pin, show/hide stats on the welcome screen, and where to search (phone and/or internet). Location Access shows you which apps would like that property and whether it is on/off for that app. Diagnostics lets you send (or not) crash data to Canonical. Lastly, Updates lets you check for software updates to install.

The only two remaining options are 'About this phone' (which gives serial, IMEI, storage, software and update info) and 'Reset phone'.

HARDWARE

Quite a few people are doubtful of the hardware. While the phone is mid-to-low range with only 1GB of RAM, it can still handle Ubuntu. Swiping is

smooth and while there is sometimes a swirling circle when something is loading, it's a second at very most. Hardly the end of the world!

Battery life is as good as any other phone. I spent several hours tinkering with it, installing and uninstalling apps, keeping the screen on while typing this review and, after three hours, it's barely dropped by 20%.

Some folks were asking about storage. While it does only come with 8GB, Ubuntu uses 2.5GB and, after installing quite a few apps, I still have 4GB free on the device. The phone does accept microSD cards so storage shouldn't be a problem.

Don't expect DSLR photographs with the camera, but it's certainly more than capable of taking a good photo. Indoors it does look grainy, but outdoors it looks OK. The back camera is 8MP with the front one being 5MP, and video being in full HD quality. To get your media to/from the phone, you can use either something like a Google Drive, or DropBox app, to send



Updates are done, as stated previously, using the System Settings. No sooner did I have mine out of the box and booting than it was telling me there was an Ubuntu update and several app updates, so here's hoping the updates continue often.

So far, I'm really impressed with the Ubuntu Phone. It's quick, smooth, cheap, uses Linux, and has good app support - even this early. I like it. What annoys me from other reviews and naysayers is that

it to the cloud for later retrieval. Or, if you're old school and want to use a cable, then it showed as a portable media player in my Kubuntu 14.10. I dragged photos (taken with my Nexus 5) to it and they showed up just fine in the Gallery app.

Some people are also confusing the Kickstarter Ubuntu Edge phone with this BQ one. This one certainly, not sure about future ones, does not boot Ubuntu to a desktop PC/monitor from a dock.




REVIEW - BQ AQUARIS E4.5

those people seem to forget that this is a phone that's coming in at under €200 (€169.90 as I write). It's great value for money.

AVAILABILITY

This is the tricky bit. Since demand is unknown, the phone is available only through so called 'flash sales' on the BQ.com website. You need to keep an eye on various Ubuntu and BQ sites and social media pages to check availability. Sorry non-Euro's, the sales are in the EU only.



MY QUICK DEMONSTRATION

<https://www.youtube.com/watch?v=tRI7G1rx90g>

USER MANUAL:

<https://static-bqreaders.s3.amazonaws.com/file/Ubuntu-Aquaris E4 5/Manual Aquaris E4.5 ubuntu EN.pdf>

UBUNTU ON MY DEVICE?

Several folks asked about the possibility of installing Ubuntu Phone on their existing device. While I know very little on the subject, it seems it can be done. A compatibility list is being kept at:

<https://wiki.ubuntu.com/Touch/Devices>

SPECS:

Screen

Dimension: 4.5"

Technology: IPS multi-touch screen, 5 capacitive points. Protective Dragontrail display

Layer Resolution: qHD 540 x 960 - 240 ppi (HDPI)

Aspect ratio: 16:9

Dimensions and weight

Dimensions: 137 x 67 x 9 mm

Weight: 123 g

Processor

CPU: Quad Core Cortex A7 up to 1.3 GHz MediaTek

GPU: Mali 400 up to 500 MHz

Memory

Internal memory: 8 GB

Ram: 1 GB

Battery

LiPo 2150 mAh

Connections

dual micro-SIM

micro-USB OTG slot, Bluetooth® 4.0 hardware compatibility (software not currently integrated).

3.5 mm TRRS headphone jack (CTIA)

MicroSD slot, up to 32 GB

Connectivity

Wi-Fi 802.11 b/g/n

Bluetooth® 4.0, Bluetooth® 4.0 hardware compatibility (software not currently integrated).

2G GSM (850/900/1800/1900)

3G HSPA+ (900/2100)

GPS and A-GPS

Interface

Operating System: Ubuntu

Languages: Spanish, English, French, Portuguese, German, Italian, and many others

Camera

Rear: 8 Mpx (Dual-flash and autofocus)

Video resolution: Full HD (1080p)

Frontal: 5 Mpx

Sensors

Brightness sensor, proximity sensor, accelerometer, eCompass, gyroscope

Other system functions

LED notification, microphone, noise canceller



REVIEW

Written by Ronnie Tucker

Able2Extract Professional 9

While it is possible to extract text from a PDF using a selection with copy/paste, it doesn't always work as planned. Also, you can lose formatting. Able2Extract Professional 9 can do all of that and more. Built into the Pro version is a rather impressive OCR feature which can extract text from images.

INSTALLING

Installing Able2Extract is easy enough. You download the


Ubuntu/Debian .DEB file, double-click it and let it install. If you have a key to unlock it then you can enter that after the install.

USAGE

On first use you're taken step-by-step on how to open a file and convert it to text. In short, you're working across the menu from left to right.

In steps:

- Open a file (PDF or text)



COMMAND & CONQUER

Written by Lucas Westermann

The last few issues of Command & Conquer have focused on programming aspects. However, I realized that I haven't talked about development environments. For a long while, I generally used a terminal and Vim, but for a few months now, I've been using Atom. In this article, I'll focus on what Atom is, how you can install it, and what the benefits are to using it.

WHAT IS ATOM?

Atom is a text editor created by GitHub. It's hackable (meaning you can configure it however you like), and is based on web technologies. This means you can adjust the look and feel of Atom by editing CSS files and adding features with HTML and JavaScript. If anyone has used Brackets or Adobe Edge CC, Atom should be familiar to you.

INSTALL ATOM?

Homepage: <https://atom.io/>

rpm package). It can also be built from source, but it's easiest to just download the .deb file from the homepage, and install it in Ubuntu. Unfortunately, as it's not in a PPA, it will not update automatically. You will need to re-run these steps to update it, or use an unofficial PPA such as <https://launchpad.net/~webupd8team/+archive/ubuntu/atom>.

WHY SHOULD I TRY IT?

Atom offers a large collection of plugins - ranging from themes,

to syntax highlighting, to plugins that will compile and execute code directly from Atom. Due to its hackable nature, you can install exactly what you want, and configure it to run however suits you best.

The features integrated into its core (a file tree, tabbing, file management directly from within Atom, etc) are features almost every heavy-duty IDE has. However, not every text editor offers these same features out of the box while remaining relatively

lightweight. Furthermore, support for things such as Emmet (a plugin for generating HTML using css-like selectors) can simplify your workflow.

- A small list of plugins I use:
- Stylus - syntax highlighting and snippets for Stylus files
 - web-browser - a browser that opens and runs directly in atom
 - atom-terminal - opens a terminal in the current directory
 - build - build your current project from within Atom
 - script - runs your code in Atom
 - color - css color viewer
 - color-picker - allows you to select a color from a pallet.
 - emmet
 - project-manager - allows you to save open folders/paths into a project for easy access later.

Of course, there are many, many more to choose from. Depending on what languages you program in, or your personal workflow, you may find packages I have never heard of.



The last few issues of Command & Conquer have focused on programming aspects. However, I realized that I haven't talked about development environments. For a long while, I generally used a terminal and Vim, but for a few months now, I've been using Atom. In this article, I'll focus on what Atom is, how you can install it, and what the benefits are to using it.

WHAT IS ATOM?

Atom is a text editor created by GitHub. It's hackable (meaning you can configure it however you like), and is based on web technologies. This means you can adjust the look and feel of Atom by editing CSS files and adding features with HTML and JavaScript. If anyone has used Brackets or Adobe Edge CC, Atom should be familiar to you.

INSTALL ATOM?

Homepage: <https://atom.io/>


rpm package). It can also be built from source, but it's easiest to just download the .deb file from the homepage, and install it in Ubuntu. Unfortunately, as it's not in a PPA, it will not update automatically. You will need to re-run these steps to update it, or use an unofficial PPA such as <https://launchpad.net/~webupd8team/+archive/ubuntu/atom>.

WHY SHOULD I TRY IT?

Atom offers a large collection of plugins - ranging from themes,

to syntax highlighting, to plugins that will compile and execute code directly from Atom. Due to its hackable nature, you can install exactly what you want, and configure it to run however suits you best.

The features integrated into its core (a file tree, tabbing, file management directly from within Atom, etc) are features almost every heavy-duty IDE has. However, not every text editor offers these same features out of the box while remaining relatively



REVIEW - ABLE2EXTRACT PROFESSIONAL 9

- Select an area (all or an area)
- Select an output format (HTML, image, and LibreOffice Calc and Writer are supported)
- Save.

Using the OCR took a little while to figure out, but you just convert the image to a PDF, or print the image to a PDF.

PDF TO TEXT

Upon opening FCM#94 (previous page, top right image) I skipped to page 13 and selected the first three columns of text. This also made it select the image, so I went with that and clicked the OpenOffice (surely it should say

LibreOffice?) button. From the popup I clicked the 'Convert' button below Writer to get an ODT file.

The ODT file was saved then automatically opened in LibreOffice Write.

While the output (previous page, bottom right image) isn't identical to the PDF it has kept the header, and text colours, which is nice. Even the dotted vertical lines were kept. The 'drop cap on' did knock it out of whack for those two lines, but the output as a whole is still very usable.

One thing I did notice is that even with small PDF files, like FCM

(10MB) it takes a few seconds to skip through the PDF.

Anyway, getting text from a PDF isn't that impressive. Time to give the OCR a run for its money.

IMAGE TO TEXT

Seeing that it could do Calc, I decided to get a bit cheeky and convert a table from an image to Calc format.

Would it be able to read the text from the image, make it editable and keep it within a

decent table format?

The answer is a resounding yes! While some text is a bit off, it has to be said that the original was a PDF printed, scanned, and turned into a PDF again, so the quality was a bit ropey.

It would certainly be easy to convert that Calc output into a table that would resemble the original.

What about an image of text to editable text?

Yep!

AERODROME CHART - ICAO

| RUNWAY/TAXIWAY/APRON PHYSICAL CHARACTERISTICS | | | |
|---|-----------------|------------------|-----------|
| APRON / RWY / TWY | SURFACE | BEARING STRENGTH | ELEVATION |
| RWY 05/23 | Grooved Asphalt | 65/R/B/W/T | - |
| Main Apron | Concrete | - | 22ft amsl |
| Main Taxiway | Asphalt | - | - |

GUND (Geoid Undulation) =
The height of the Geoid (MSL) above the Reference Ellipsoid (WGS 84) at the stated position.

BEARINGS ARE MAGNETIC
ELEVATIONS AND HEIGHTS ARE IN FEET

| | |
|--------------------------|-------|
| ELEVATIONS IN FEET AMSL | 136 |
| HEIGHTS IN FEET ABOVE AD | (110) |

| A | B | C | D |
|----|--|-----------------|------------------|
| 1 | | | |
| 2 | | | |
| 3 | AERODROME | | |
| 4 | CHART - ICAO | | |
| 5 | RUNWAY/TAXIWAY/APRON PHYSICAL CHARACTERISTICS | | |
| 6 | | | |
| 7 | APRON/RWY/TNY | SURFACE | BEARING STRENGTH |
| 8 | HWY 05/23 | Grooved Asphalt | ELEVATION |
| 9 | Mon Apron | Concrete | 2211 amsl |
| 10 | Main Taxiway | Asphalt | |
| 11 | | | |
| 12 | GUNS (Geoid Undulation) = | | |
| 13 | The height of the Geoid (MSL) above the | | |
| 14 | Reference Ellipsoid (WGS 84) at the stated position. | | |
| 15 | | | |
| 16 | BEARINGS ARE MAGNETIC | | |
| 17 | | | |
| 18 | ELEVATIONS AND HEIGHTS ARE IN FEET | | |
| 19 | | | |
| 20 | ELEVATIONS IN FEET AMSL | 136 | |
| 21 | HEIGHTS IN FEET ABOVE AD | -110 | |
| 22 | | | |

I like how it converts it to editable text, does a good job of it, and even keeps headers in bold. It's not just a dump of plain text. It really does try to copy the format of the original.

CONCLUSION

Of course, it's not infallible. Give it a coloured background with white text and I'm pretty sure it'll fail, but so will the vast majority of OCR applications. I was particularly impressed with how few errors there were in a good quality image to editable text.

If you have high quality images that you need converted back to text, then this application is definitely one to consider, and

kudos to Investintech for making a Linux version of their app available.

Linux System Requirements

OS: Linux Fedora 20 or newer, Ubuntu 13.10 or newer, 32-bit edition

RAM: 512+ MB of free memory available for the software

Hard Drive Space: 250 MB of disk space for the program components

Monitor: 1366 (Width) x 768 (Height) screen resolution

Download trial from:
http://www.investintech.com/productdownloads/able2extract_pro.htm

COMPETITION:

To win one of five life-time keys to Able2Extract Professional 9 all you have to do is answer the following question:

What does OCR stand for?

Email your answer to:
misc@fullcirclemagazine.org


Deadline for entries is Sunday 19th April. Five winners will be drawn at random.

In the top section, we have the pilot briefing and procedure notes
The first box tells us a number of things, such as the localizer frequency, final approach course, runway landing lengths, touchdown zone elevations, and the airport elevation.

The middle box will tell us the special notes, such as if the ASOS/AWOS is inop, special alternate and takeoff minimums/procedures, along with the approach lightning system and a text description of the missed approach procedure.

The bottom box tells us all of the frequencies that we will need to know, such as weather, approach, and CTAF. The white L with a black circle around it next to the 122.7 (Unicom) frequency tells us the lightning is pilot controlled with that frequency.

Plain View



Able2Extract PDF Converter 9

All-in-one PDF solution

- ✓ Convert PDFs to Microsoft Word, Excel, PowerPoint, Text, Images, OpenOffice and more with precision.
- ✓ The most accurate converter for PDF tables to spreadsheet format.
- ✓ Generate industry standard PDFs with powerful PDF creations options.
- ✓ Protect and Secure your PDFs.
- ✓ Resize, rescale, delete and move pages inside PDF.

Works with:



Ubuntu



Fedora



@able2extract



www.investintech.com

 **INVESTINTECH.COM**
PDF SOLUTIONS



MY STORY

Written by Alan Ward

Ubuntu has become well-known as a distribution designed for normal users, with an emphasis on ease of use. However, its open-source nature also makes this kind of work environment especially useful for the scientist.

Investigators form a rather specific category of computer users. They tend to have quite precise needs, that overlap those of "normal" users only up to a certain extent. For example, a statistician may in some cases have a use for the very same spreadsheet as a manager, though put to other tasks. However, at some point, the statistician will require a more powerful number-crunching environment, such as R (also available in the Ubuntu repositories).

However, the number of potential users for specific, scientific, programs is naturally much smaller than that for more ordinary tasks. Many of those developing programs for scientific purposes are actually scientists themselves, as the specialization

of modern science makes a background in each specific field at the very least a certain advantage. Having an open-source operating system makes building programs easier for these people who may not be computer engineers. At the same time, having a software management tool such as the apt and repository system at your disposal makes distributing your program much easier than without them. All of this has contributed to making a wide range of scientific applications available both for Ubuntu and for the distribution it is based on, Debian.


To illustrate this topic, in this piece I would like to show you some of the options to display chemical molecules in 3D on your computer, with an emphasis on organic chemistry. Applications include not only teaching chemistry in itself, but also learning more about biology and, to some extent, genetics. For example, we could view a 3D model of hemagglutinin (PDB code 1RUZ) that viruses such as the infamous Influenza A virus use to

bind to the host's cells - the "H1" part of e.g. H1N1 representing the specific type of hemagglutinin contained by that virus.

GETTING MOLECULES

Several file formats are currently in use, but perhaps the most extended are the MDL Molfile format (extension: .mol) and the Protein Data Bank format

(extension: .pdb). Most molecule viewers are capable of handling both, or even of converting a molecule between formats. It is worth noting that both formats are originally text-based files with a well-documented structure, which illustrates how open data formats are helpful to share data in the scientific world. Compressed versions may also be found, generally using standard gzip compression.



An example of the molecule of glycerol (glycerin) in the Molfile format is shown below.

NOTE: the distances between atoms are completely off, this is just an example).

There are several good sources for molecule files on the Internet. One of the better known is the Research Collaboratory for Structural Bioinformatics (RCSB) Protein Data Bank (PDB), at <http://www.rcsb.org/pdb/home/home.do>. This has a comprehensive collection of molecules contributed by many teams from all over the world. Of special interest for the layperson such as myself is their PDB-101 primer http://www.rcsb.org/pdb/101/structural_view_of_biology.do with a structured presentation by topics.

RCSB PDB Deposit Search Visualize Analyze Download Learn More MyPDB Login

Summary 3D View Sequence Annotations Seq Similarity 3D Similarity Literature Bot & Chem Methods Links

T STATE HUMAN HEMOGLOBIN [ALPHA V96W], ALPHA AQUOMET, BETA DEOXY 1VWT

DOI: 10.2210/pdb/1vwt/pdb

Primary Citation

Novel water-mediated hydrogen bonds as the structural basis for the low oxygen affinity of the blood substitute candidate rHb(alpha 96Val→Trp).

Puius, Y.A.P., Zou, M.P., Ho, N.T.P., Ho, C.P., Almo, S.C.P.

Journal: (1998) Biochemistry 37: 9258-9265

PubMed: 9849306

DOI: 10.1021/bi9727287

Search Related Articles in PubMed

PubMed Abstract:

One of the most promising approaches for the development of a synthetic blood substitute has been the engineering of novel mutants of human hemoglobin (Hb) A which maintain cooperativity, but possess lowered oxygen affinity. We describe here two crystal structures...

[Read More & Search PubMed Abstracts]

Biological Assembly

3D View More Images...

Symmetry: C2 view

The “Molecule of the Month” section contains a large assortment of articles on specific molecules that can certainly fill us in on the salient points of how the biology works.

By searching for different keywords, I was able to find a

specific molecule of interest: hemoglobin (PSB code 1VWT) from human red blood cells. Each molecule is described, the team that announced it is given, as is the citation to the scientific publication it initially appeared in. A download link is also provided (to the right of the PDB code in

large letters), by which we can download the corresponding file in the PDB format.

VIEWING MOLECULES

There are quite a few programs available in the Ubuntu repositories to view the file we have just downloaded. One of the oldest and best known is Rasmol, that now sports a GTK interface.

The window itself is very simple: all options are accessible through the menu bar at the top. The user can rotate the structure with the mouse within the main window itself, so spatial relationships that cannot be shown on a printed page become much more clear.


When we load up a file we see it by default in stick-form representation, where bonds between atoms are represented by short sticks, color-coded by atom type (white for carbon, red for oxygen, yellow for iron, etc.) Hydrogen atoms are usually not shown directly, though this option can be set if desired. This is the molecule of hemoglobin, with its four main structures (alpha and

Molecule exported from chemtool

```

6 5 0 0 0 0 0 0 0 0999 V2000
    2.7500      2.7763      0.0000 C    0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
    5.1316      4.2632      0.0000 C    0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
    7.0658      2.8684      0.0000 C    0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
    1.1184      4.1974      0.0000 C    0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
    5.1316      5.5132      0.0000 C    0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
    9.1053      4.2895      0.0000 O    0 0 0 0 1 0 0 0 0 0 0 0 0 0 0
  1  2  1  0  0  0  0
  2  3  1  0  0  0  0
  1  4  1  0  0  0  0
  5  2  1  0  0  0  0
  3  6  1  0  0  0  0
M  END
    
```



File View Display Colours Options Settings Help



beta units) surrounding a central space.

Other viewing options allow us to show atoms represented as filled spheres (Display > Ball and Stick, or Display > Spacefill), which can be useful for the smaller molecules or to see the complete volume a molecule occupies. However, for large molecules with several hundreds or thousands of carbon atoms, it may be more clear if we hide individual atoms and


bonds, and instead move to a view based on strands (Display > Strands) or the cartoon view (Display > Cartoon). In this screenshot, the strands view was colored by functional units (Colours > Chain) so we can distinguish the alpha and beta chains by color. We can also activate stereoscopic vision (Options > Stereo) to see a separate view for each eye if so desired.



MY STORY

JMol is a more recent offering. Written in Java, it is available for different platforms such as Windows and OS-X as well as GNU/Linux, and should be readily portable to others. It has similar options to Rasmol, though the interface is different. Some tools are available to slightly edit the molecule (add or delete atoms) and to connect to other programs. However, some of these are unfortunately no longer easily

available on Ubuntu, such as the Povray raytracing environment. The default representation in JMol is sufficiently clear for easy viewing of biological models, and can be rotated using the mouse as before. This is JMol's view of the hemoglobin model from the PDB file. Two complexes that imprison iron (Fe) atoms (in yellow) are quite visible in the lower part of the foreground:



The newer PyMOL Molecular Graphics System is one of the more recent applications available. Written in the very same Python modern interpreted language that has often been seen in the pages of Full Circle, its presentation revolves around not one but two windows. One of these is used as a combination of log viewer and general input dialog, while the other holds the molecule view proper and its associated options.


PyMOL has the richest collection of options of all the applications presented here - although the interface is perhaps not very intuitive. Just to set you on course, the "A" button is to add elements to the molecule, the "S"

button serves to set (activate) view options, and the "H" button to unset (hide) features. The "C" button changes between colorizing schemes.

There are also more options to see molecules in several types of stereo, and some options to build videos of the molecule, that I have not played with too much.

ROLL YOUR OWN

Playing around with existing module files is interesting not only in themselves, but also as a way to appreciate the actual amount of useful information (accent placed on "useful") found on the Internet. However, at some point we may






MY STORY

wish to start drawing up our own molecules.

A simple place to start is the molecule of propane-1,2,3-triol - perhaps better known as glycerol or glycerin. It can be found not only in soaps, but also in foodstuffs and even in electronic cigarettes. Basically, we have a chain of three carbon atoms (the propane skeleton), with a hydroxyl (-OH) group hanging off each carbon.

There are in fact several applications in the Ubuntu repositories that draw planar representations of organic chemical molecules, and more can



be found in various places on the Web. One of the easiest to use is Chemtool.

Drawing tools in the upper toolbar allow us to place various chemical bonds at specific angles to each other, to form the carbon skeleton for our molecule. When done, a text tool can be used to add the functional groups at various places. Double and triple bonds are naturally also available. Once drawn, elements can be moved, erased, flipped horizontally and vertically, etc.

The finished molecule can be exported in various flat graphical formats such as PNG, but also in Molfile format. This can then be read in by PyMol or any of the other viewers. In PyMol, the missing hydrogen atoms can easily be added to our geometry.

However, we can see that something weird has happened to the center carbon atom: the supplementary hydrogen needed to complete its bonds has somehow grown out at a strange angle. This is not what we expected, and may be attributed to the fact that Chemtool is basically a molecule sketching

application for 2D. Its output is fine for publishing on paper, but lacks the information about depth needed to draw realistic 3D models of molecules.

This is where another program, Avogadro, comes in handy. This is more of a molecule builder than just a sketching tool. In much the same way as Chemtool, Avogadro comes with an interface that allows us to build up the carbon skeleton of our molecule, and then add oxygen atoms where needed to form the additional functional groups. In this case, however, additional hydrogen atoms are adjusted dynamically during construction, helping us see exactly what we are constructing.

Once the molecule has been built, we can choose Extensions > Optimize geometry and the program will calculate the most realistic positions for the atoms (using "realistic" in the sense of positions with the least potential energy). This can then be exported as a flat image file, or as a file in PDB format that can then be opened with PyMol.

The end result is much more satisfactory, as we can observe

that each of the three carbon atoms has a tetrahedral structure, not planar. This can be even better appreciated as we rotate the molecule with the mouse.

The software reviewed in this article is from the apt packages named: *rasmol*, *jmol*, *pymol*, *chemtool* and *avogadro*. They can easily be found and added to an existing *buntu installation using the usual tools, such as apt-get, synaptic, Ubuntu Software Center, etc.



Alan teaches computer science at Escola Andorrana de Batxillerat (high-school). He has given GNU/Linux courses at the University of Andorra and taught GNU/Linux systems administration at the Open University of Catalunya (UOC).



MY OPINION

Written by Kavinda Bandara

I have been using linux for the last 7 years, the first five as a dual-boot with Windows, and the last two solely dependent on Ubuntu. In the early days, I tried Redhat, Fedora and with the live CDs, I have now completely shifted to Ubuntu.

Being a doctor by profession, this was an odd combination (as many others found it nauseating to take the bitter pill and stay on the learning curve), but I always managed to find an alternative to the Windows software others

used. The simplicity, stability, and good reliable updates, and the trustworthy community, always were the blessings for me to stay on the track.

As my previous laptop died after 6 years of service (HP 9000 series), I bought a new one last year. This is a Dell 3521, coming with Intel core i3, 500GB HDD, 4GB RAM, hybrid graphics with AMD RADEON 7670M and Intel HD Graphics 4000 and Dell wireless 1704 (manufactured by Broadcom). This came preinstalled with Ubuntu



12.04 LTS so I kept it that way.

From the beginning, the wireless didn't seem to function well. But I usually used a 3G dongle to connect to internet. So the wireless didn't give me a problem in the early days. Once I bought my new Smartphone, I couldn't connect it to the laptop via Bluetooth or wireless. Then I started to troubleshoot.

After doing a thorough literature search in the community support, I started with reinstalling network manager, tried with different managers and then handled the drivers. I tried with reinstalling the existing ones, and then tried different suggestions given by the community.

On one fine day, when I was trying to fix it with drivers, the wireless suddenly disappeared from the list. (I am so sorry I cannot post the link to the thread here, as I lost it with system reinstallation). I tried with "lspci" several times, and with reinstallation of default factory drivers. And then I tried

with a live boot, but the result was the same. I decided to hand over the job to the shop where I bought the laptop, as the laptop was still in the warranty period.

What they did was installation of a trial version of Windows 7 and installation of drivers! And the wireless is working!! And I could do all the stuff I wanted.

I felt bad about the situation. Why that happened and why we couldn't solve the situation by ourselves is the main question I have so far. Why did we have a long term support system if we couldn't tackle problems for the next four years? And it would have been easy if we categorized and ranked the community threads.

I'm still using the trial version, and waiting impatiently to install 14.04. I wish everything would go fine and I could get rid of Windows. It's always a bad dream to go back to Windows.





LETTERS

If you would like to submit a letter for publication, compliment or complaint, please email it to: letters@fullcirclemagazine.org. PLEASE NOTE: some letters may be edited for space.

Join us on:



goo.gl/FRTMI



[facebook.com/fullcircle
magazine](https://facebook.com/fullcirclemagazine)



twitter.com/#!/fullcirclemag



[linkedin.com/company/full-
circle-magazine](https://linkedin.com/company/full-circle-magazine)



[ubuntuforums.org/forum
display.php?f=270](https://ubuntuforums.org/forumdisplay.php?f=270)

FULL CIRCLE NEEDS YOU!



Without reader input **Full Circle** would be an empty PDF file (which I don't think many people would find particularly interesting). We are always looking for articles, reviews, anything! Even small things like letters and desktop screens help fill the magazine.

See the article **Writing for Full Circle** in this issue to read our basic guidelines.

Have a look at the last page of any issue to get the details of where to send your contributions.

Tuxidermy

OK, JONES! YOU'RE HERE BECAUSE PEOPLE SAY YOU'VE LOST YOUR MIND! YOU STARTED RANTING ABOUT NONSENSE THINGS, LIKE A NEW WORLD ORDER, FREEDOM AND SOME KIND OF MAGICAL SIGN WHO'D LEAD PEOPLE TO A BRIGHT NEW FUTURE! EXPLAIN YOURSELF! NOW!



WELL, SIR. IT ALL BEGUN WITH A STRANGE BOOK I'VE FOUND BY CHANCE IN A LIBRARY.

IT SEEMED LIKE A NORMAL BOOK, BUT ONCE I'VE OPENED IT, I COULDN'T STOP READING! IT SAVED MY LIFE AND I HAD TO TELL THE WHOLE WORLD ABOUT IT! EVERYONE MUST KNOW ABOUT THE BOOK, SIR!



WHAT BOOK WAS IT?

I CAN'T TELL YOU, SIR. YOU'LL HAVE TO READ IT YOURSELF BUT, ONCE YOU SEE THE SIGN, IT'LL CHANGE YOUR LIFE TOO!

THE SIGN? THE YELLOW SIGN, YOU MEAN? POOR BOY, THEN ALL IS LOST FOR YOU! IT DRIVES PEOPLE INSANE!



YELLOW SIGN? YELLOW SIGN? NO, MAN! YOU GOTTA STOP DIGGING INTO THESE OLD HORROR NOVELS!



THIS SIGN ACTUALLY MADE ME FEEL VERY SANE!





Q&A

Compiled by Gord Campbell

If you have a Linux question, email it to: misc@fullcirclemagazine.org, and Gord will answer them in a future issue. Please include as much information as you can about your query.

Q What is the easiest way to get rid of old kernels?

A (Thanks to *Bashing-om* and *morgaes* in the Ubuntu Forums) Run these commands:

```
sudo apt-get update
sudo apt-get dist-upgrade
```

Reboot.

One more command:

```
sudo apt-get autoremove
```

It's a new use for "autoremove."

Q I went and bought this WNDA3100v2 N600 Wi-Fi adapter; now I need help setting it up.

A If you can follow these instructions, you should be OK:
<https://help.ubuntu.com/community/WifiDocs/Driver/Ndiswrapper>

Generally, it's a good idea to Google "[device name] linux"

before buying hardware. Many USB Wi-Fi adapters work as soon as you plug them in and reboot, but printers, webcams, and Wi-Fi adapters, have varying levels of support.

Q I use VLC media player to watch TED Talk videos, but I would like to download them instead.

A (Thanks to *mc4man* in the Ubuntu Forums) The TED site supports downloading, but it's not obvious. When you select a video in your browser, there is a "download" button. That starts the video playing, but you can pause it, then right-click in the video and select "save video as".

Q I added a ppa and installed simplescreenrecorder. How can I completely remove it?

A (Thanks to *deadflowr* in the Ubuntu Forums) Simplescreenrecorder installs two

packages, so use these commands:

```
sudo apt-get purge
simplescreenrecorder*

sudo add-apt-repository --
remove ppa:maarten-
baert/simplescreenrecorder

sudo apt-get update
```

Q How do you delete Windows software installed in 14.04 using the program's native installer (not through WINE)?

A The uninstaller is also a Windows program, so you may be able to uninstall using an uninstall program UNDER WINE.

Many applications do not include an uninstall program, so all you can do is delete the files.

The general approach is this: in your file manager, set it to display "hidden" files and folders. In your home folder you will find a folder called .wine, and under it is drive_c, which corresponds to a C: drive in Windows. It contains Program Files, among other things.

My own approach is to leave things alone until I do a re-install of a new version, such as Ubuntu 16.04 LTS. Deleting a single movie saves more space than dozens of programs.

Q Today, Xubuntu 14.04 software update installed the same kernel 3.13.0.46 for the third time. Why?

A It was actually 3.13.0-46.75, and there was a regression requiring the updates.

Q Is there any way to synchronize Evolution Mail and my Hotmail account?

A (Thanks to *sandyd* in the Ubuntu Forums) See this web page for IMAP and SMTP settings that can be added to evolution: <http://windows.microsoft.com/en-us/windows/outlook/send-receive-from-app#>

Q&A

Q My system sometimes locks up. I would like to know the exact time my system locks.

A Run conky, and include the time in the conky display.

Q In a text file, once the line containing "Heading" is found, how can I display the next 20 lines?

A (Thanks to **Lars Noodén** in the Ubuntu Forums) Grep can do this:

```
grep -A 20 Heading  
/path/to/somewhere/somefile.txt
```

TOP QUESTIONS AT ASKUBUNTU

* What is the relation between ALSA and PulseAudio sound architecture?
<http://goo.gl/y365oR>

* How can I get the monitor resolution using the command line?
<http://goo.gl/k6Ji5B>

* 10 million files in one folder
<http://goo.gl/1kfp1A>

* What should I do straight after installing Ubuntu? [on hold]
<http://goo.gl/kFb7WK>

* What happens after 14.10 expires? Can I upgrade to 15.10 or do I have to reinstall?
<http://goo.gl/rSvWpg>

* How do I check whether a module is installed or not in Python?
<http://goo.gl/JOdUyf>

* Is it safer to install two OSes on different hard drives or just different partitions? [on hold]
<http://goo.gl/u3xFJj>

* Why doesn't Ubuntu remove old kernels automatically?
<http://goo.gl/SpLcPU>

* Grep searching two words in a line
<http://goo.gl/ZF9UaT>



Gord had a long career in the computer industry, then retired for several years. More recently, he somehow found himself "The IT Guy" at a 15-person accounting firm in downtown Toronto.



One of the interesting things about the Internet is that it was not designed for security, and kind of happened almost accidentally. The early researchers, people like JCR Licklider and Vint Cerf, were mostly interested in facilitating communication between university researchers, and assumed that anyone they were in contact with was another "good guy". It was only when the Internet started to take off in the 1980s and 1990s that anyone began to pay attention to this stuff. Technologies like Javascript and SSL were introduced by Netscape, for instance, though they have since been adopted by the whole Internet.

INTRO

In the early days, if you wanted to log on to a remote computer to do some work, you would probably use a program called Telnet to do this. Telnet would let you log in to a remote Unix machine if you had an account on it, and, once logged in, you could do anything your

account was authorized to do. And if you had root access, it could be just about anything. I recall my first introduction to this in the 1990s when I was managing the website of my university. I was given a shell account on our Red Hat server that also housed the web site, and sternly warned not to do anything to screw it up. I would Telnet into the server from my desktop to do things like chmod the CGI scripts and similar things we needed to do with websites in those days. This worked reasonably well, but the problem with Telnet is that it was never designed to create a secure connection, and since this server had our Website, it was perforce exposed to the whole Internet. If you combine an unsecure connection with a man-in-the-middle attack, or with vulnerabilities in the Telnet program that allowed an attacker to get in, and even escalate their privileges, you can see why this became an issue.

In 1995 there was a password-sniffing attack on the network of the University of Helsinki in

Finland, and this led a researcher there, Tatu Ylönen, to create the first SSH implementation. SSH is an acronym for Secure Shell, and expresses the idea that you can securely log in and get a shell on a remote server. This was initially released as free software, but in later versions he took it proprietary. But the developers at OpenBSD decided that a free software implementation was needed, and they created OpenSSH, which is the basis for most implementations today. And while SSH was initially intended for Unix-like environments (which came to include Linux), the OpenBSD developers created a portability branch to make it available to any OS. So, if you are running a Windows environment, you can use SSH just as easily. If you want a little more detail on this you can consult Wikipedia.

DESIGN CONSIDERATIONS

SSH was designed to do several things:

- Create secure and private communication between two

different machines. This means that the connection must use encryption.

- Establish the integrity of the communications to ensure that messages have not been altered en route. Again, encryption makes this possible.
- Authenticate both parties to the conversation to prove their identities. Again, encryption makes this possible.

As regards authentication, SSH can be done using passwords, but that is a weaker form of authentication. If you are serious about security, you should authenticate using a key. Our old friends the public/private key pair come in at this point, and, as you might imagine, the outlines of how this is done are pretty similar to what we saw in both E-mail and in Certificates. This is not perfect, of course, but it reduces the chances of an attack very significantly. Note that even if you are using password authentication, there is encryption and key-exchanging going on. That is necessary to accomplish the first two of the



above goals.

ENCRYPTION AND TUNNELING

The basic idea in SSH, just as with Virtual Private Networks, is to use encryption to create secure communications between two different systems. It has become customary to refer this type of connection as a “tunnel”. This is a metaphor, and like all metaphors it both illuminates some things and obscures others. The idea of a tunnel does help to get across the security of the connection, in that if done properly the outside world cannot see what is going on. The encryption does work if you do it right, and as Bruce Schneier famously said in the wake of the Snowden revelations, you can trust the mathematics. However, the metaphor somehow implies to people that this traffic is flowing in some other place than the rest of the Internet, and that is simply untrue. All “tunnel” traffic flows through the same routers and switches as all other Internet traffic, and is made up of the same kinds of packets. If you are on a network where someone is using SSH, you can “see” the packets

using Wireshark or other similar software. But you would not be able to see anything in the payload of each packet other than random noise because of the encryption.

It is important to understand the mechanisms and how they work if you are going to be secure, since otherwise you might make a mistake and open yourself to an attack.

SSH USES

Although SSH originally was developed to simply provide a secure shell session on a remote server, it has been extended in a number of interesting ways which we will look at in upcoming tutorials. For example, SSH can be used to:

- Create tunnels
- Forward TCP ports
- Create X11 connections
- Securely transfer files (SFTP)
- Securely copy files (SCP)
- Securely mount a remote file system (SSHFS)

WHERE DO YOU GET SSH?

SSH uses the client-server

model. Generally, you are coming from a desktop and want to connect to a remote server. If the server is Unix or Linux, it should have SSH installed and properly configured if the sysadmins are any good. On Windows servers, it may need to be installed, but this is not difficult. For any Windows sysadmins out there, here is an article that explains the installation on a Windows 2008 server.

As for desktop clients, again all Unix-like systems have the SSH client installed by default. This includes Unix, Linux, MacOS, all flavors of BSD. For Windows users, I recommend installing PuTTY. This is free software, is distributed under the MIT license, and complies with the Debian Free Software Guidelines. There are several PuTTY applications, since they use different applets for each of the capabilities, so one is for Secure Shell, another for SFTP, and so on.

There is a useful manual for OpenSSH and for the various applications that make up all of OpenSSH which can be found at <http://www.openssh.com/manual.html>. From here you can see that

the applications which make up all of OpenSSH include:

- ssh – The basic rlogin/rsh-like client program.
- ssh_config – The client configuration file.
- sshd – The daemon that permits you to login.
- sshd_config – The daemon configuration file.
- ssh-agent – An authentication agent that can store private keys.
- ssh-add – Tool which adds keys to the above agent.
- sftp – FTP-like program that works over SSH1 and SSH2 protocol.
- scp – File copy program that acts like rcp.
- ssh-keygen – Key generation tool.
- sftp-server – SFTP server subsystem (started automatically by sshd).
- ssh-keyscan – Utility for gathering public host keys from a number of hosts.
- ssh-keysign – Helper program for hostbased authentication.

BASICS

So, as we saw in the last tutorial, SSH uses the Client-Server model. Now, technically a server is

just the machine you are connecting to, and there is no reason in principle that it could not be another desktop, a laptop, or even a telephone if it has the appropriate software. So, the model really reduces to you as the client, and the other machine as the server. As with all Internet connections, there are standards and protocols involved. The original Telnet communicated over TCP through port 23. Because SSH was conceived as a replacement, it used the same TCP protocols, and was assigned the adjacent port number of 22. This is the standard, but it is not set in stone. Indeed, one of the ways you can improve security is to use a non-standard port. This requires that the server is configured for a different port. Servers monitor ports using daemons, so the server administrator would need to configure the daemon to listen to the alternate port, such as 16180, for SSH traffic, and then communicate this to the potential clients. This is a good thing to know if you are using SSH to log in to a remote server that you administer (for example, you may have a server co-located in a data center, or a Virtual Private Server that you control and administer). If

a vulnerability in the SSH protocol were to be discovered, you can bet that the bad guys would immediately start to hammer on port 22 of every IP address on the Internet looking to take advantage, but if your server does this on a non-standard port it gives you that much more protection. Still, if you are connecting to a server that you do not control, you probably will connect to port 22, and chances are your client software is already configured to do that by default.

How It Works

To begin with, all SSH connections are initiated by the client. You, as the client, are going to the server and asking "Please, sir, may I have a shell connection?" And you will generally do this on port 22. The server will have a daemon listening to that port, and it should respond to your request. If you have the same account name on both your client and the server, you can just log in to the server. If they are different you should put in your account name. These examples assume you are using a terminal.

Example one:
`ssh 192.168.1.24`

Example two:
`ssh myserver.host.com`

Either of these would work if your account name is the same as on your local client. If they are different, you could speed things along by adding the account name:

Example three:
`ssh phred@myserver.host.com`

The server should then send back to your client the protocol version, which should be SSH 2.0. If your client also supports SSH 2.0, the connection proceeds. Otherwise, the connection is dropped. All modern clients support SSH2.0, however. If you really want to dive into the details, you can start with RFC 4253, written by the previously mentioned Tatu Ylönen. If you want to see everything going on, use the `-v` switch with the `ssh` command to turn on verbose mode. This will show you all of the communication going on between your client and the server.

Then the Binary Packet Protocol kicks in. This specifies each of the

fields in the packet sent using SSH. If you need the full details on this, consult RFC 4253, but this is probably something you can skip if you are not going to be writing your own client.

Next is the Server's turn to identify itself, which it does by transmitting its public key. If this is your first time attempting to log in to this server, you will get something like this"

The authenticity of host 'myserver.host.com' can't be established.

RSA key fingerprint is d8:09:f4:42:....

Are you sure you want to continue connecting (yes/no)?

Since this is the first time you have tried to connect, you don't know for sure if this really is the server you want. This is where a man-in-the-middle attack could take place. If you were sitting in a coffee shop using the free wifi, for instance, someone could be at the next table intercepting the traffic and could respond with their own public key, thus hijacking your session. This should only happen the first time you connect with this

particular machine, because the public key should be stored for future reference in what is called the “known_hosts” file. On a Linux machine, this is usually found at `~/.ssh/known_hosts`. On Windows 7 generally you can find it at `%USERPROFILE%\ssh` or `%USERPROFILE%\ssh`. But if you get a new laptop you have to go through this initial connection again with each of the sites you connect to.

This is definitely a weakness, so how can the server admins prevent this? Your login attempt got the fingerprint returned to you, so that is the key to this, if you will pardon the pun. You probably don't want to post it publicly on an unsecured website, for instance, since the bad guys might figure out a way to spoof it. And email carries the same risk. In a corporate environment, you might post it on an encrypted website behind the firewall, and require the employee to use their credentials to access it.

The next step depends on whether you are using SSH v.1 or SSH v.2. Since the vulnerability in SSH v.1 was found long ago in Internet time, you probably should

question what is happening if you see it used today. SSH v.2 was adopted in 2006, which makes it comparatively old and stable. As Michael W. Lucas put it in his book *SSH Mastery*: “SSH-1 permits man-in-the-middle attacks and session hijacking, as discussed in Chapter 1. If someone insists on using SSH-1, practice saying “I told you so.” His book is excellent, and the e-book is only \$10, and worth every cent. In particular, if you need to set up an SSH server, which I won't really cover much of, this book is mandatory reading. Among the changes introduced by SSH v.2 are:

- improved encryption standards, including 3DES and AES
- Public key certification for clients (we will discuss later on)
- The use of sound cryptographic Message Authentication Code (MAC) algorithms for integrity checking
- SSH v.1 was monolithic, meaning that all of the protocols needed for encryption, authentications, etc. were part of a single large protocol built into SSH v.1 in SSH v.2 each protocol is split out into its own and defined in a separate RFC such as
 - Transport Layer Protocol
 - Connection Protocol
 - Authentication Protocol

I don't want to take the time to dig into how SSH v.1 works, and there are a few differences, so I will focus on only SSH v.2 here.

Once you have accepted the Public Key of the server, it is up to you as the client to respond. You do this by first generating a symmetric key (called the session key) which will be used to encrypt all traffic. Remember from our previous tutorials that Asymmetric key pairs carry a very large computational overhead, so they are generally used only to set up a connection and exchange the symmetric key. The client creates this key, then using Diffie-Hellman-SHA1 key exchange sends it back to the server.

NOTE: There is a provision in the protocol for a Certificate Authority to attest to the validity of the public keys used. This would help immensely in preventing the man-in-the-middle attack you are vulnerable to at your first connection since the CA would give you confidence in the validity of the server's public key. But not all servers use this at present, in part because Certificates are expensive.

We are not done yet. You should only have access to the server in accord with the rights you were given when your account was created. This means it is time to authenticate. And that is the topic for our next tutorial.



Kevin is a geek, Linux enthusiast, a certified Project Manager by day, and in general a tech lover. His blog is at: <http://www.zwilnik.com>



Amnnesia: The Dark Descent is still considered to be one of the best horror survival games ever released. A third-party total conversion modification for Amnesia has recently been released and it rivals Amnesia's spine chilling scare score. Penumbra: Necrologue was originally released for Microsoft Windows on Halloween night, October 31, 2014. Then, on February 6, 2015, Penumbra: Necrologue was released for Linux. Penumbra: Necrologue is available as a free mod, but only if you already have Amnesia: The Dark Descent installed in your computer. You can obtain Penumbra: Necrologue direct from the <http://penumbranecrologue.ru> website or from Steam depending on your preference. Penumbra: Necrologue is a third-party modification to Amnesia: The Dark Descent from a dedicated team of Russian fans who call themselves CounterCurrent Games. A third-party total conversion modification is a mod to a game which completely changes the game while at the same time still using

its original game engine. This is exactly what Necrologue is. However, although Necrologue is based on Amnesia: The Dark Descent, there are also elements borrowed from the Penumbra video game series. In fact, it can be argued, and I actually believe, that Penumbra: Necrologue is the fourth installment of the Penumbra game series.

At its core, Penumbra: Necrologue is a puzzle, first-person, survival, horror game. It picks up where Penumbra: Requiem left off. You play the role

of Philip as you awake in a sort of underground lab called The Shelter. The only thing you need to know is that your goal is to survive. You can basically interact with almost everything in your environment. Anyone who's played any of the games from the Amnesia series (The Dark Descent or Machine of Pigs), or the Penumbra series, will feel right at home playing Necrologue. When you first awake, you have no memory of anything nor do you know how you got there or even where you are – you don't know if you're alone or accompanied. In

the first set of rooms when the game begins, there isn't much for you to do other than learn how to navigate around, how to examine objects (by the way, you can examine pretty much anything), how to open doors, etc. It's especially important that you learn how to save your game progress and how to use the journal as well as your inventory.

I've been playing the game with mouse/keyboard setup though you can also use a gamepad controller. I would have liked playing Necrologue with my controller but unfortunately there was a minor bug affecting my controller. Everything works fine except that, for some reason, I can't move sideways (left/right) with the gamepad so instead I've chosen to play with my mouse/keyboard. I use the standard WASD to move and the mouse to look around with a few other keys for other actions, such as picking up objects, rotating them, turning the flashlight on/off, etc.

As you're first setting up the



game, one of the first screens informs you that this game is better enjoyed while using headphones (instead of speakers). I completely agree. I also recommend turning the lights off and isolating yourself from others while you play. By using headphones, and playing in the dark, you are cranking up the game's fear factor beyond 10. The eerie soundtrack that ominously plays while you aimlessly wander the dark halls of this decrepit facility keeps your body tense and alert as if you were really physically walking in The Shelter and your very life were at stake. At the end of the dark hall, as it intersects another hallway, you notice someone's flashlight approaching from the right which makes the hair on the back of your neck stand up from the chills. Is it friend or foe? You've spotted them but have they spotted you? Do you keep walking in that direction or open the door to your left and hope they haven't noticed you? These are the kinds of decisions you must make while playing the game.

Now might be a good time for me to say that you have zero weapons in the game. This is a

survival game to the core, all you can do is run, hide, or run then hide. Often you must solve puzzles in order to get to a new section or to get out of your current predicament. If you're looking for a good scare, if you're a fan of survival games, or maybe if you're looking for something a little different, then you should try Penumbra: Necrologue.

Of course, you've got to keep in mind that it is absolutely necessary that you already own Amnesia: The Dark Descent and that it be installed on the same machine in which you're planning on playing Penumbra: Necrologue. If by chance you're interested in playing Necrologue but you've never played Amnesia: The Dark Descent, then you're definitely missing out on perhaps the scariest horror



survival game ever made and without a doubt you should buy Amnesia and play it either before or after playing Penumbra: Necrologue. You can get Amnesia: The Dark Descent from Steam, The Humble Bundle, or Desura, for \$19.99 and, as I've stated earlier, you can then add Penumbra: Necrologue for free.

CONCLUSION

I've enjoyed playing Penumbra: Necrologue very much. For being a fan-made modification, it has exceeded my expectations. The gameplay is nothing new, especially for having already played games from the Amnesia and Penumbra games series. I didn't expect it to be drastically different, after all it is a mod. Think of it as a continuation of Penumbra

and Amnesia combined into one. The map is completely new, the characters are new, and the monsters are new. The frightful mood maintained throughout is really this game's strong point. However, the minor glitches I encountered, specifically with the gamepad controller lack of left-right movement, prevent me from giving it a perfect score.



MINIMUM SYSTEM REQUIREMENTS:

OS: Major Linux Distribution from 2010, Debian should work but is not supported.
Processor: High-range Intel Core i3 / AMD A6 CPU or equivalent.
Memory: 2 GB RAM
Graphics: Mid-range NVIDIA GeForce 200 / AMD Radeon HD 5000. Integrated Intel HD Graphics should work but is not supported; problems are generally solved with a driver update.
Hard Drive: 2 GB available space



I'd been searching high and low for a good Linux native golf game when suddenly Perfect Golf was released on Steam Early Access.

EARLY ACCESS

As the title suggests, this is an early access game. In other words: beta. So, when you're about to click the BUY button, just remember that the game is still in the works and will have some features missing, broken, or both. Obviously, you also need Steam installed. This review is for v1.2.1.1 from the end of January 2015.

PLAYER CREATION

Before stepping onto the

course, you need to create a new player. You choose a name, shirt colour and some default clubs to take with you. There's no tweaking of the player gender, clothing, or features. It's 'one character fits all' at the moment.

THE LOBBY

Before going anywhere, you're in the lobby area of Perfect Golf. Here you can join the chat with others, create an online game (yes, you can play your friends, or random people, online), create a local game (yes, you can play with offline, real world, friends too), or just practice.

I haven't tried the online mode as yet, so I can discuss only the local play in this review.



PRACTICE

Before hitting the fairways, I'd recommend going for a few practice swings to hone your game.

There are two main modes to choose from in how you swing. One is the good old three-click method where you click to begin swinging, click at the top of your swing, then click at the end of the swing. With this method, your second click determines the power of the shot and the last click determines the trueness, or not, of your shot. This is the method I

prefer. The other is using the mouse to pull back, then push forward to swing. This is a far more precise method as it uses the actual angle and speed of your movements to determine the shot.

Ball dynamics is really what makes this game shine. It's pretty unforgiving for you. Once a shot



UBUNTU GAMES

has been taken in practice, you will see some pretty detailed results on how you did.

It's this attention to physics and detail that makes Perfect Golf highly regarded.

SETTINGS

Before starting your first game I'd recommend visiting the settings to edit, if nothing else, the weather.

There's not a lot to tweak, initially, in graphics as I find it best to get a game going, and then tweak the graphics to get you as

many frames per second (FPS) – as a lagging game will not help your swing any.

The weather settings is where Perfect Golf shines. You can change the day length, and time, as required. Dynamic sky looks great but, if you're on a slow machine, you might need to disable this. RealTime Weather is great and you can set it to your locale to get your actual current weather into game. With me being in Glasgow, that usually means rain and high winds. Not always a great thing!

LOCAL PLAY




Finally. A game of golf!

There are (as of writing) five courses to choose from. Most are easy to medium rated, but, with specific/local weather (more on this later), you could make the courses harder.

First, you choose your course, how many players, holes, and various other rules, and, finally, go.

The screen layout is pretty familiar to anyone who's ever played a golf game. Top left is the course title, hole, and par. Below that is the player name and hole



info. Top right are six icons you use to tune your shot (we'll come to them in a moment). Bottom right is the swing meter. Below the swing meter is a yellow button. When this is on, you're taking a shot. Click it off and you're taking a practice swing. Right click in the swing meter and you get a choice of club. Right click a second time and you can adjust where you'll hit the ball for back spin, etc.



The icons are (top to bottom) settings, map, fly by, editor, ground grid, flag on/off, and show pin. I mostly use the map to select where I'd like the ball to land, and to see distances, and the ground grid when putting.

Initially, the camera is pretty static, but you can change that in settings to have it follow the ball,

or switch to the ball as it lands. Panoramic camera zooms out to let you see a bit more of your surroundings rather than zoomed in on the player.

You're on the green. Your putter is out. I normally switch the ground view on to see the lay of the land. Clicking in the grid will show you a pin which will display the distance to the pin and whether it is above/below the ball.

The one thing here to get used to is that, when putting, the small flag icon on the swing meter is where the flag is from you – ignoring the distance. It's up to you to give it more/less power depending on whether the hole is above/below the ball.

CONCLUSION

Even in its current early access state, this is the golf game to beat. I like their way of thinking in that they won't hold your hand and give you a ton of help or be unrealistic. Games in the past let you adjust the ball in the air and told you where the 100% power level was, and where to click for a 50-yard hit. Not here. The swing meter is blank

for percentage, and the realism is there in that, if you are just a touch off perfect, then the ball will swing to the left/right.

My only criticism is that when you are a way off from the hole, you are working in yards, which is fine, but the closer you get it changes to feet. This sometimes catches you off guard when you look and see 50 and try to go half-way on a 100 yard club. The ball takes off and you realise, too late, that it was actually 50 feet – which is about 17 yards. Say bye-bye ball and hello trees. For reference: there are three feet in a yard.



MINIMUM SPECS:

Processor: core 2 duo E8400 @ 3.00 GHz.
Memory: 4 GB RAM
Graphics: video card with 512MB video RAM
Hard Drive: 1 GB available space

RECOMMENDED:

Processor: Quad Core
Memory: 8 GB RAM
Graphics: GeForce GTX 770
Network: Broadband Internet connection
Hard Drive: 1 GB available space

Steam page:

<http://store.steampowered.com/app/288140/>





HOW TO CONTRIBUTE

FULL CIRCLE NEEDS YOU!

A magazine isn't a magazine without articles and Full Circle is no exception. We need your opinions, desktops, stories, how-to's, reviews, and anything else you want to tell your fellow *buntu users. Send your articles to: articles@fullcirclemagazine.org

We are always looking for new articles to include in Full Circle. For help and advice please see the **Official Full Circle Style Guide**: <http://url.fullcirclemagazine.org/75d471>

Send your **comments** or Linux experiences to: letters@fullcirclemagazine.org
Hardware/software **reviews** should be sent to: reviews@fullcirclemagazine.org
Questions for Q&A should go to: questions@fullcirclemagazine.org
Desktop screens should be emailed to: misc@fullcirclemagazine.org
... or you can visit our **forum** via: fullcirclemagazine.org



FCM#89

Deadline:
Sunday 05th Apr. 2015.
Release:
Friday 24th Apr. 2015.



Full Circle Team



Editor - Ronnie Tucker
ronnie@fullcirclemagazine.org
Webmaster - Rob Kerfia
admin@fullcirclemagazine.org
Podcast - Les Pounder & Co.
podcast@fullcirclemagazine.org

Editing & Proofreading

Mike Kennedy, Gord Campbell, Robert Orsino, Josh Hertel, Bert Jerred, Jim Dyer and Emily Gonyer

Our thanks go to Canonical, the many translation teams around the world and **Thorsten Wilms** for the FCM logo.

Getting Full Circle Magazine:



EPUB Format - Recent editions of Full Circle have a link to the epub file on the downloads page. If you have any problems with the epub file, you can drop an email to: mobile@fullcirclemagazine.org



Issuu - You can read Full Circle online via Issuu: <http://issuu.com/fullcirclemagazine>. Please share and rate FCM as it helps to spread the word about FCM and Ubuntu Linux.



Google Play - You can now read Full Circle on Google Play/Books. Either search for 'full circle magazine' or click this link: <https://play.google.com/store/books/author?id=Ronnie+Tucker>