Full Circle

Phablet

Password

THE POCKET SERVER
FROM AN OLD NEXUS 7 TABLET
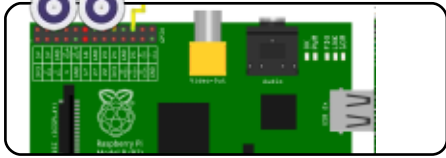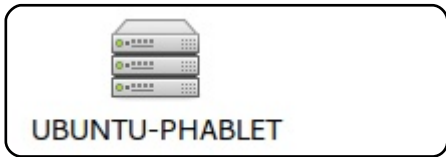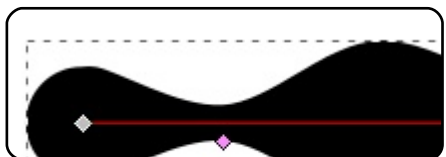
Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

EDITORIAL

WELCOME TO ANOTHER ISSUE OF FULL CIRCLE.

It's a full house this month. We have Python, LibreOffice, part three of the VAX-VMS tutorial, Inkscape, Arduino, more 3D printer building, and (as the cover shows) an extra HowTo article showing how you can use an old (well, 2013) Nexus 7 tablet, with Ubuntu Touch, as a pocket server. And, of course, Gord answering your questions.

Our good friend Art returns with a review of Deepin OS. A distro I certainly hadn't heard of before. And speaking of distros, Steve is looking at Apricity OS. Meanwhile, I'm reviewing Able2Extract 10 which has some nice additions since I last reviewed it almost a year ago, and in games Oscar is reviewing Saints Row IV and I'm reviewing a new Early Access game called The Kindred. Think Minecraft meets The Sims.

It's also that time again where I have to beg for you, the readers, to send in some articles. I'm running low again. Your article can be about anything as long as it has something to do with Ubuntu/Linux in some way. It can be a review (hardware/software), a helpful HowTo, an opinionated article where you want to get something off your chest (but no Microsoft bashing), anything. Remember: no articles means no magazine. ANYONE can write for FCM. Not just the good folks you see here each month. Email your articles to me, Ronnie, at: ronnie@fullcirclemagazine.org. Don't forget to include screenshots/photos in your document.

If you want to be kept abreast of the latest Ubuntu/Linux news then be sure to add the Full Circle Weekly News RSS feed (shown right) to your favourite audio downloading software. At the moment we're receiving around 15,000 downloads each week! I'm also hoping to get the show on Stitcher radio for your convenience.

All the best, and keep in touch!
Ronnie
ronnie@fullcirclemagazine.org

This magazine was created using :

SCRIBUS

gimp

LibreOffice The Document Foundation

creative commons

Find Full Circle on:

goo.gl/FRTMl

facebook.com/fullcircle magazine

twitter.com/#!/fullcirclemag

http://issuu.com/fullcircle magazine

http://www.magzter.com/publishers/Full-Circle

Weekly News:

http://fullcirclemagazine.org/feed/podcast

http://www.stitcher.com/s?fid=85347&refid=stpr

http://tunein.com/radio/Full-Circle-Weekly-News-p855064/

THE FULL CIRCLE WEEKLY NEWS

A short podcast (<10min) with just the news. No chit-chat. No time wasting. Just the latest FOSS/Linux/ Ubuntu news.

RSS: http://fullcirclemagazine. org/feed/podcast

AUDIO MP3

LUBUNTU-BASED LXLE 14.04.4 "POSH PARADIGM" LINUX OS COMING SOON WITH NEW UI

Dubbed Posh Paradigm, LXLE 14.04.4 promises to be a great release that brings all the cutting-edge GNU/Linux technologies, along with a new user interface to its dedicated users.

Best of all, it would appear that the major changes prepared by the LXLE team for the next major version, LXLE 16.04, will, in fact, be implemented in the soon-to-be-released LXLE 14.04.4 version.

"*So instead of back tracking, the distribution tippy-toed to the edge with the latest LXDE core to bring the features necessary to implement the new LXLE user interface the 'Posh Paradigm,*" developers have explained.

And this move is one of the best the LXLE team could have done in the entire history of their GNU/Linux operating system, as users will now have all the features that were planned for the next version.

Source: http://news.softpedia.com/news/lubuntu-based-lxle-14-04-4-posh-paradigm-linux-os-coming-soon-with-new-ui-50109

ANYONE CAN NOW PORT UBUNTU LINUX FOR RASPBERRY PI 3 WITH UBUNTU PI FLAVOUR MAKER

Martin Wimpress, the leader of the Ubuntu MATE Linux operating system, had the great pleasure of informing Softpedia, just a few minutes ago, that he updated the Ubuntu Pi Flavour Maker tool introduced by us in December 2015 to support porting of Ubuntu MATE, Xubuntu, Lubuntu, and Ubuntu Server OSes for the Raspberry Pi 3 Model B.

It appears that Martin Wimpress had early access to the Raspberry Pi 3 board, thanks to the awesome people at the Raspberry Pi Foundation, so he had three days to update his Ubuntu MATE 15.10 (Wily Werewolf) operating system for Raspberry Pi 3, as well as to the powerful Ubuntu Pi Flavour Maker tool.

"*The Raspberry Pi 3 is available. The Raspberry Pi Foundation sent me a unit on Friday night and I've been working with them over the weekend to make sure Ubuntu MATE, and other flavours, have day-1 support,*" said Martin Wimpress.

Source: http://news.softpedia.com/news/anyone-can-now-port-ubuntu-linux-for-raspberry-pi-3-with-ubuntu-pi-flavour-maker-501140.shtml

AVNET SIGNS EUROPEAN LINUX SOFTWARE DEFINED NETWORKING DEAL

Avnet has announced a pan-EMEA distribution agreement with Cumulus Networks to offer "the industry's first true full-featured Linux OS for networks" to

customers and partners.

The agreement will drive innovation in open data centre networking, making software-defined networking (SDN) a reality, said Avnet.

Avnet will offer Cumulus Networks' data centre network operating system software, Cumulus Linux, which can be integrated with hardware from multiple leading switch vendors. "This flexible, robust approach to data centre networking significantly reduces costs and complexity, which is paramount in the data centre environment," said Avnet.

Open networking solutions based on Cumulus Linux are designed to allow MSPs and organisations of all sizes to leverage efficient data centre technology, providing customers with a choice of hardware and software elements to fit their individual needs.

Source:
http://www.channelbiz.co.uk/2016/03/03/avnet-signs-european-linux-software-defined-networking-deal/

CONVERGED PC AND SMARTPHONE IS THE FUTURE, SAYS CANONICAL'S MARK SHUTTLEWORTH

Shuttleworth was at Mobile World Congress (MWC) in Barcelona last week, where Ubuntu exhibited to show off its phones, tablets, and IoT (Internet of Things) initiatives.

Among the displays was BQ's Ubuntu tablet, which has both a tablet mode and a windowed mode that can be enabled when attached to an external display, keyboard and mouse.

"We're showcasing Ubuntu as a converged OS," said mobile product manager Richard Collins. "It's one codebase that can go onto any device that has some kind of need for a display, one SDK that developers can use. We'll do the magic to make sure that apps scale for the display. All of that has been engineered into our first tablet product. If you connect a mouse and keyboard, it is running a full Ubuntu PC."

Ubuntu smartphones were also

on display, destined for the Chinese market, from brands well-known in China, such as Meizu. Collins said that Chinese vendors wanted to build their own ecosystem, and that Ubuntu gave them freedom to do that.

Source:
http://www.theregister.co.uk/2016/03/03/converged_pc_and_smartphone_future_mark_shuttleworth/

PERL VULNERABILITIES CLOSED IN ALL SUPPORTED UBUNTU OSES

Perl doesn't get all that many fixes for vulnerabilities, but it happens from time to time. The Ubuntu maintainers have been quick to update the package, and it is now ready in the repos. Fixing the vulnerabilities and updating the system is a really good idea.

"Markus Vervier discovered that Perl incorrectly handled nesting in the ata::Dumper module. An attacker could use this issue to cause Perl to consume memory and crash, resulting in a denial of service," reads the security notice.

This is just one of the

vulnerabilities found and fixed. For a more detailed description of the problems, you can see Canonical's security notification. Users have been advised to upgrade their systems as soon as possible. If you haven't upgraded in a while, you might pull more than just the Perl patch.

The flaws can be fixed if you upgrade your system to the latest libpng12-0 package specific to each distribution. To apply the patch, users will have to run the Update Manager application.

Source:
http://linux.softpedia.com/blog/perl-vulnerabilities-closed-in-all-supported-ubuntu-oses-501302.shtml

LINUX DESKTOP GETS MIPS PROCESSOR

Baikal-T1 has two MIPS P5600 cores clocked at at least 1GHz.

"P5600 currently holds the highest CoreMark/MHz score among 32-bit CPU IP processors, making it one of the most powerful processors when it comes to performance per MHz and per watt," claimed

Imagination Technologies, owner of MIPS intellectual property.

The PC is called Tavolga Terminal TB-T22BT, runs the Debian 8 operating system, and can be configured as a workstation or a thin-client terminal.

Free software including LibreOffice - with similar function to Microsoft's Office - and the Firefox browser are available for Debian 8.

Source:
http://www.electronicsweekly.com/news/products/microprocessors/linux-desktop-gets-mips-processor-2016-0

PINGUY OS 14.04.4 IS NOW BASED ON LINUX KERNEL 3.13.0-79, DROPS CLOSED REPOS

Pinguy OS is a Linux distribution based on Ubuntu, and the developers are using only the LTS version of the Canonical-built operating system. Since Canonical released their latest point release, Ubuntu 14.04.4, two weeks ago, it's now time to upgrade Pinguy OS as well.

It's interesting to see just how fast the Pinguy OS developers have moved after the launch of Ubuntu 14.04.4, especially since they have made some changes of their own. Canonical was a little bit late, and the debut of their point release was delayed, but that shouldn't be a problem.

Pinguy OS is built by the same people who are also making Pinguy Builder, an amazing tool that can be used to build your own distro from the version that you have installed. It's not all that well known, and the devs have been considering dropping it.

According to the changelog, HandBrake has been updated to version 0.10 to take advantage of the H.265 codec, the Linux Kernel has been updated to version 3.13.0-79, and the closed repos have been removed from the source list.

Source:
http://news.softpedia.com/news/pinguy-os-14-04-4-is-now-based-on-linux-kernel-3-13-0-79-drops-closed-repos-501345.shtml

IF YOU LIKE FEDORA, YOU'LL LOVE KORORA

The Korora Linux desktop is a solid computing platform that's loaded with options to suit a wide range of user needs.

Korora 23 final was released last month. It's a Fedora-based distribution featuring many user-friendly enhancements as well as a choice of five desktop environments. Fedora is the community version of Red Hat Linux Enterprise.

Unless you are looking for bleeding-edge desktop environments, Korora gives you all the top contenders. You can run Cinnamon, GNOME, KDE, MATE and Xfce.

Korora's use-it-out-of-the-box philosophy is one of the reasons the distro keeps getting better. If you want a better, more user-friendly Linux distro that reaches beyond Fedora's enterprise appeal, you can't go wrong with any of Korora's five desktop versions. It leaves little for users to desire and makes choosing another distro unnecessary to get your preferred interface.

Source:
http://www.linuxinsider.com/story/83193.html

DENVER POST PITCHES CANONICAL'S UBUNTU LINUX OPEN SOURCE OS

Ubuntu and GNU/Linux got some good press recently from the Denver Post, which pitched the open source operating system as a healthy alternative to Windows and OS X on desktop computers.

The article, written by Tamara Chuang, focused on PCs from System 76, one of a handful of companies that sells laptops and desktops with Ubuntu pre-installed. It described the Ubuntu experience in exceedingly positive terms.

"*With Ubuntu's user-friendly graphics, it looks a lot like Windows or Mac OS," Chuang wrote. "Software programs like imitation Microsoft Office-like tools are free downloads. And you're not forced to upgrade, though Ubuntu does push out security updates when needed.*"

The article also said Ubuntu "*can be a better alternative to Windows computers.*"

Source:
http://thevarguy.com/open-source-application-software-companies/denver-post-pitches-canonicals-ubuntu-linux-open-source-o

SOLUS: A WELCOME GROUND-UP BREAK FROM THE LINUX HERD

The Linux world is full of spin-offs, clones and branded distros. The most famous lineage is Mint, based on Ubuntu, which is in turn based on Debian. What's less common is entirely new distros, starting from zero and building their own stack.

That's exactly what Solus has done. That's not to say Solus doesn't, like any software, stand on the shoulders of giants, it's just that Solus stands on a few less shoulders than others.

The Solus project recently hit the 1.0 milestone and then quickly followed that up with a 1.1 release. If you're looking for something new, including a desktop that's entirely its own, Solus is well worth a look.

Solus is designed to be user friendly, but minimalist. You won't find tons of graphical bling or any experimental "HUD" or "Shell" interfaces. Instead, Solus offers a clean, well-designed desktop experience that sports a flat, modern look and a healthy dose of its own take on the desktop experience.

Source:
http://www.theregister.co.uk/2016/03/08/solus_linux_review/

CANONICAL WANTS YOUR FEEDBACK ON SNAPPY AND SNAPCRAFT IN UBUNTU 16.04 LTS

Canonical is preparing to unleash the major Ubuntu 16.04 LTS (Xenial Xerus) operating system this spring, on April 21, 2016. Snappy/Snapcraft are an important part of the upcoming long-term support (LTS) release, so they would like to get the pulse of the community.

"Snappy and Snapcraft are shaping up very nicely for their 16.04 release and quite a few of you have been working with them already. While the team is still busy getting everything ready, we would like to get your feedback and hear what your experience was like and how you feel about 16.04."

Therefore, if you have five minutes to spare on answering a few questions about your experience with the latest Snappy and Snapcraft technologies in the Ubuntu 16.04 LTS (Xenial Xerus) operating system, go ahead and take a look at Canonical's new survey and help the devs figure out what else they need to improve for the final release.

Source:
http://linux.softpedia.com/blog/canonical-wants-your-feedback-on-snappy-and-snapcraft-in-ubuntu-16-04-lts-501467.shtml

FSF DEMANDS THE W3C REJECT ENCRYPTED MEDIA EXTENSIONS PROPOSAL

For years, our Defective by Design campaign and the anti-DRM movement have been fighting media and proprietary software companies who want to weave Digital Restrictions Management into the HTML standard that undergirds the Web. Winning this is a top priority for us -- the DRM proposal, known as EME (Encrypted Media Extensions), would make it cheaper and more politically acceptable to impose restrictions on Web users, opening the floodgates to a new wave of DRM throughout the Internet.

The battle is coming to a head as EME approaches a final vote by the Web's standardization organization, the World Wide Web Consortium (W3C). We need to make our voices heard now -- the W3C is convening March 20-22 and is scheduled to discuss the proposal.

Source:
https://my.fsf.org/civicrm/profile/create?gid=183&reset=1

CANONICAL IMPRESSES WITH UBUNTU CORE AND LOW-POWER MOBILE DEVICES

We saw three surprising things at Canonical's booth at Mobile World Congress 2016, and each has to do with conversion. By "conversion," I mean a mobile device equipped with Ubuntu Core – a shared code-base that enables apps to run on mobile and desktop, whether designed for touch, or mouse and keyboard input – that can connect to an external monitor and see the apps optimized for the larger display.

If that sounds a lot like Microsoft's Universal Apps and Continuum, that is because it's the same concept. But as Microsoft is struggling to finish and polish its software, Canonical went ahead and did it better.

What Canonical is doing here is all the more impressive because of the issues Microsoft appears to be having in implementing its own convergence experience for mobile devices using Windows 10 and Continuum. Windows 10 for phones is still not finished, and as Collins noted, "The mobile version [of Windows 10] has the ability to display certain apps to a desktop, but still remains a mobile OS at its core."

Canonical itself still has some work to do. More apps need to get the convergence treatment, and we need to see more available devices, Further, getting video-out capabilities on these higher-end Ubuntu phones is crucial--there's no two ways about that--and adding dual-display support for both phones and tablets will take the experience to a more compelling level.

Even so, Canonical has proven that its work with Ubuntu Core and the "convergence" paradigm is impressive and promising.

Source:
http://www.tomshardware.com/news/canonical-ubuntu-core-convergence-mwc,31359.html

ONEPLUS LETS UBUNTU TAKE A STROLL ON THE

ONEPLUS ONE

OnePlus seems to have achieved something rather cool on the software front. They've partnered with Canonical, the people who make the popular Linux desktop distribution Ubuntu, in order to bring the mobile version to the original OnePlus One. Neat! OnePlus made the announcement on Google Plus.

If you'll recall, there was a big buzz around a mobile version of Ubuntu about two and a half years ago, made more notable by the quite wise decision to build user-flashable files for popular Android phones in the same manner as Android ROMs. (Ubuntu is also available on a few commercial phones, most notably the upcoming Meizu Pro 5.) Now there's a public build available for the OnePlus One for those who want to try out the Ubuntu take on mobile computing.

Source:
http://www.androidpolice.com/2016/03/09/oneplus-plays-nice-with-canonical-and-lets-ubuntu-take-a-stroll-on-the-oneplus-one/

CANONICAL RECOMMENDS OPEN-SOURCE AMDGPU AND RADEON DRIVERS FOR UBUNTU 16.04 LTS

Ubuntu 16.04 LTS will not support the AMD Catalyst proprietary graphics driver for AMD GPUs (Graphics Processing Units), which is known to the community as the fglrx driver. Instead, Canonical recommends users using the open-source Radeon and AMDGPU alternatives.

According to the Ubuntu developers, who go to all the trouble of backporting the latest kernel code from the upstream Linux 4.5 kernel branch for the open-source AMDGPU and xf86-video-ati AMD Radeon graphics drivers into Ubuntu 16.04 LTS' Linux kernel packages, AMD has put a lot of effort into the video drivers.

Source:
http://news.softpedia.com/news/canonical-recommends-open-source-amdgpu-and-radeon-drivers-for-ubuntu-16-04-lts-501556.shtml

MARCHER TROJAN USES NEW TACTIC TO INFECT ANDROID USERS

A fraudulent Adobe Flash installer package is a pathway to infection and potential financial losses by way of the Marcher Trojan.

Security firm Zscaler is warning about a new variant of the Android Marcher Trojan that is using Adobe Flash and adult content sites as a way to trick users into becoming infected and giving up financial information.

There are a lot of different vulnerabilities in Adobe Flash—in fact, Adobe just updated for 23 new vulnerabilities this week—but the new Android Marcher Trojan isn't using an authentic version of Flash or exploiting vulnerabilities that Adobe has already patched. Rather, the Android Marcher Trojan uses a fake version of an Adobe Flash Player installer to infect users.

"*The majority of the Marcher Trojan downloads that we are blocking in the cloud are from porn sites,*" Deepen Desai, head of security research at Zscaler, told eWEEK. This appears to be a popular social engineering tactic where the user is prompted to install the Flash Player update to view the porn video and the attack cycle can start with an email or SMS."

Source:
http://www.eweek.com/security/marcher-trojan-uses-new-tactic-to-infect-android-users.htm

FIGHTING CYBER ATTACKS WITH ARTIFICIAL INTELLIGENCE

The next frontier of anti-virus software is leveraging artificial intelligence (AI) to not only predict what threats are out there, but to also actively fight back before they strike.

This is according to American-based Cylance's chief marketing officer, Greg Fitzgerald, speaking at the NetEvents Press and Analyst Summit in Rome, Italy.

The company says it is "*revolutionising cyber security through the use of AI and machine learning to proactively prevent advanced persistent threats and malware*".

Cylance announced it is expanding into the Europe, the Middle East and Africa (EMEA) with the establishment of a London-based team led by Evan Davidson, former enterprise sales director at FireEye. It also established a channel partnership with CoreSec Systems, which supplies cyber security and networking solutions in Sweden and Denmark.

Source:
http://www.itweb.co.za/index.php?option=com_content&view=article&id=150902

HP LINUX IMAGING AND PRINTING 3.16.3 ARRIVES WITH SUPPORT FOR UBUNTU 16.04 LTS

The development team behind the HPLIP (short for HP Linux Imaging and Printing) project, an open-source initiative to bring the latest HP printer drivers to GNU/Linux distributions, has released the third maintenance build in the 3.16 series.

HP Linux Imaging and Printing 3.16.3 has just been announced today, March 18, 2016, and it includes initial support for Canonical's upcoming Ubuntu 16.04 LTS (Xenial Xerus) operating system, thus providing drivers for the latest HP printers and scanners to Ubuntu Linux users installing the new OS.

Another interesting change in HPLIP 3.16.3 is the fact that the RPM packaging for the Red Hat Enterprise Linux 5.x series of operating systems has been discontinued, but recent releases of RHEL are still supported. Detailed installation instructions are provided by the project's maintainers on the official website.

Source:
http://news.softpedia.com/news/hp-linux-imaging-and-printing-3-16-3-arrives-with-support-for-ubuntu-16-04-lts-501900.shtml

THE LINUX FOUNDATION AND EDX OFFER FREE CLOUD INFRASTRUCTURE MOOC

Understanding cloud technologies tops the list of most important skills for any developer, sysadmin or emerging DevOps professional. If you don't believe me, just try to hire a cloud expert. Good luck!

This course won't turn you into an expert, but it will at least get you on your way. It provides a primer on cloud computing and DevOps software. It will cover next-generation cloud technologies like Docker, CoreOS, Kubernetes, and OpenStack. The course will also provide an overview of software-defined storage and networking solutions and a review of DevOps and continuous integration best practices.

The class is not exclusively focused on open-source cloud software. Mark Hinkle, The Linux Foundation's VP of marketing and a cloud expert in his own right, told me at the Great Wide Open conference, that it covers the fundamentals of cloud computing irregardless of platform.

Source:
http://www.zdnet.com/article/the-linux-foundation-and-edx-offer-free-open-source-cloud-infrastructure-course/

AMD RELEASES RADEON GPU-PRO BETA DRIVER FOR UBUNTU LINUX WITH VULKAN SUPPORT

The new driver has been named AMD Radeon GPU-PRO by AMD, and it is currently targeted at AMD R9 series of GPUs (Graphics Processing Units), such as AMD Radeon R9 380X, AMD Radeon R9 380, AMD Radeon R9 285, AMD Radeon R9 M395X, AMD Radeon R9 Fury X, AMD Radeon R9 Fury, and AMD Radeon R9 Nano.

Being the first public Beta release and all that, the AMD Radeon GPU-PRO driver software has been made available only for the Ubuntu Linux platform. The package contains an installation script and Debian packages (.deb) for the Ubuntu 14.04.4 LTS (Trusty Tahr) operating system.

"This driver is intended as beta level support for use solely with Vulkan applications and as such some driver functionality has been removed. This is including, but not limited to, support for other Graphics APIs, Radeon Settings and other Radeon Software driver features," says AMD in the release notes.

The driver currently supports the Vulkan 1.0, VDPAU, OpenGL 4.5, GLX 1.4, and OpenCL 1.2 APIs (Application Programming Interfaces), offers basic display and power management features, supports both the ADF (Atomic Display Framework) and KMS (Kernel Mode Setting) technologies, and includes a GPL-compliant kernel module.

Source:
http://news.softpedia.com/news/amd-releases-radeon-gpu-pro-beta-driver-for-ubuntu-linux-with-vulkan-support-501921.shtml

UBUNTU BEING PULLED INTO "GAME CHANGING" AREAS – CANONICAL CEO

Mark Shuttleworth, founder of Canonical, said developers are now targeting the company's Ubuntu platform for "game changing" areas such as Network Function Virtualisation and Internet of Things.

Last month, the company significantly boosted its convergence strategy, unveiling the first Ubuntu-powered tablet, from European vendor BQ, following earlier launches in the smartphone space.

Speaking to Mobile World Live, Shuttleworth said it was a "nice step" to include a tablet as part of its convergence story, as he opened up on how the demand for Ubuntu was beginning to grow.

Source:
http://www.mobileworldlive.com/featured-content/home-banner/ubuntu-is-beginning-to-branch-out-canonicals-shuttleworth/

EDUBUNTU LINUX WILL SKIP THE UBUNTU 16.04 LTS RELEASE

Edubuntu is a version of Ubuntu Linux designed for schools, students, and folks generally interested in education. Formerly known as Ubuntu Education Edition, the operating system is based on Ubuntu, but includes a suite of apps aimed at teachers and students.

The first version of the operating system was released in 2005, and the last major release came in 2014, when the developers decided to only offer new versions alongside Ubuntu's LTS (Long Term Support) releases every two years instead of the more frequent releases which come out every six months.

Now it looks like the Edubuntu team has decided to slow the pace even further: Ubuntu 16.04 is coming in April, but Edubuntu will stay on version 14.04 indefinitely.

The developers plan to continue supporting Edubuntu 14.04 through April, 2019, which is when the "long term support"

promise would normally end anyway.

Source:
http://liliputing.com/2016/03/edubuntu-linux-will-skip-ubuntu-16-04-lts-release.html

MONEY TALKS: LINUX FOUNDATION POURS FUNDS INTO R LANGUAGE

The Linux Foundation-backed R Consortium, created to support the math-and-statistics-centric R language, will soon put its money where its mouth is.

The consortium is prepping several new projects designed to advance the language, its implementations, and the culture of development around it. In addition to defining standards and practices for R, this includes funding R-related projects the consortium believes will be a boon to the community.

Seven projects got the thumbs-up from the consortium to receive a total of $200,000 in grant funding. Most prominent among them is a project to develop a unified framework for distributed

computing in R -- a common method to run R applications across compute clusters.

Source:
http://www.infoworld.com/article/3047181/application-development/money-talks-linux-foundation-pours-funds-into-r-language.html

After having set up my intel NUC (see last month's article for more information), I started using NGINX and Apache to serve my in-progress web projects. However, setting up various virtual hosts, managing it, and a complicated series of dependencies, had me abandon that plan relatively quickly. Instead, I switched to Vagrant for a few projects. Unfortunately, while setting up a Vagrant system is pretty quick, it's also heavier than it needs to be resource-wise. Instead, I looked into Docker. This month, I'd like to cover what Docker is, and how I've set it up to manage my various web projects.

WHAT IS DOCKER?

Docker is a way of creating virtualized containers for running software and services. The main difference between Docker and Vagrant is that Docker shares the base of the virtualization with all containers. So while Vagrant creates a standalone full-fledged VM, Docker instead creates a VM that uses a shared kernel between all VMs, and is based on LXC.

While it's difficult to explain the difference between Vagrant and Docker in layman's terms (as much of the differences are in the details), there are some simple things to note. For example, Docker is generally faster to start up, and less resource intensive (as it shares some of the host's kernel).

It does have to run on a Linux machine, but there are tools offered by Docker to easily set it up on Windows or Mac (using a base virtual machine to supply the linux kernel).

SETUP

Docker containers can be assigned to various virtual networks, can expose ports, can run off specific images, and can share folders/files between the host and the guest. My current system is as follows:
• Custom network (called "webnet")
• One Docker container running HAProxy, and exposing port 80. Also includes a static IP on webnet.
• Multiple docker containers running the nickistre/ubuntu-lamp image, and sharing a host folder with the guest at /var/www/html (default Apache folder). Each has a static IP on the webnet network, but doesn't expose any ports (as communication is funneled through HAProxy).

I have set up dnsmasq with a catch-all DNS entry for all subdomains of home.lan. These get forwarded to the local machine at port 80 (which is, in turn, passed to the docker HAProxy machine). HAProxy is then used to check which subdomain it is, and (if configured) points it to the correct docker container's static IP.

To illustrate (pelican is a static site generator): pelican.home.lan points to 192.168.1.16 (local machine) via dnsmasq. Once the request arrives, HAProxy checks the configuration file, and then passes the request through to 172.18.0.4:80 (the docker instance running my pelican site). If HAProxy doesn't have an entry for the domain, it just ends at 172.18.0.2 (the HAProxy machine).

WHAT FIRST?

First, you'll need to create the network you want to use.

```
docker network create --
subnet=172.18.0.0/16 webnet
```

This creates a new network called webnet, with a possible IP range of 172.18.0.0 to 172.168.255.254. I won't go into the specifics of the format used here. If you want to use a different IP range, just adjust the 172.18 part - it's unlikely you'll ever need more than 65534 possible addresses for Docker containers. I chose 172.18 because the default docker IP range (which is dynamically allocated in the default network) is 172.17. That way, I should be able to always tell that the IP corresponds to docker. Do not choose the same IP range as your actual local network (typically 192.168).

The reason why we need the custom network, is simply because the default networks from docker don't allow assigning static IPs.

How do I go about creating my machines?

The standard docker command will pretty much always be the same.

```
docker run –d –v {SHARE} ––
net webnet ––ip 172.18.0.X ––
name {NAME} {IMAGE}
```

What it does:
• 'docker run' fires up a container.
• -d detaches the created container (otherwise all output from the container is automatically printed to the terminal, and closing the terminal will close the docker instance).
• -v {SHARE} - specifies the shared folder, in the format /local/path:/remote/path. For example: -v /home/lswest/web/pelican:/var/www/html
• --net webnet - configures which network the container should use.
• --ip 172.18.0.X - this is the static IP I'm assigning. I like to keep them in a basic order, in order to make adding HAProxy entries easier. If you chose a different IP range in the step above, adjust accordingly.
• --name {name} - this is the name the docker instance will be known as. For example --name pelican. This can be used in the docker start/stop/restart/rm commands, and appears in the listing of docker ps.
• {IMAGE} - this is the image to use for the base of the container. I like the nickistre/ubuntu-lamp image, which contains Ubuntu 14.04 and LAMP. There are other images (such as ArchLinux), but since my Digital Ocean servers typically run on Ubuntu, I stuck as close as possible to the real-world environment. The HAProxy image I use is HAProxy:1.5 (official HAProxy image, version 1.5).

In the case of the HAProxy image (which should be created first), the command will look like this:

```
docker run –d –v ~/docker-
config/haproxy/haproxy:/usr/l
ocal/etc/haproxy/haproxy.cfg:
ro ––net webnet ––ip
172.18.0.2 –p 80:80 ––name
proxy haproxy:1.5
```

Main differences:
• -p 80:80 - exposes the guest port 80 to the host port 80 (so visiting 172.18.0.2 in a web browser should spit out the typical 503 error from HAProxy).
• -v - in the command above, I just like the actual HAProxy.cfg file into the location for the config for HAProxy. NOTE: editing this file with some text editors (such as vim) will result in HAProxy not accepting the changes. This is due to the inode changing. To fix, just restart the machine with 'docker restart proxy'. Nano appears to avoid this problem.

Autostarting

If you want the docker containers to always run (after crashing, reboots, or restarts), you can add --restart=always to the run command. This must be done when creating the container - so if you've already created one, you'll need to stop it, remove it (rm), and then recreate it.

Does this work for only port 80?

The settings for the HAProxy container should never need to change (during docker run).

However, if you want to point a domain to, for example, a NodeJS application running on 8000, just adjust the IP in the HAProxy configuration (see below). For example, 172.18.0.6:8000.

HAProxy Config

I've saved an example of my configuration here: http://pastebin.com/1M5DMkF3

Basically, adding new containers is as simple as copying the acl line (#33), the use_backend line (#36), and the backend block (#39-43). Be sure to keep increasing the index numbers (host_test1 would become host_test2, etc.). And give each backend a unique name (and adjust the IP!).

How do I make a wildcard dnsmasq entry?

You can add a single line to your dnsmasq.conf file (typically found at /etc/dnsmasq.conf). It looks like this:

```
address=/home.lan/192.168.1.1
6
```

Replace the IP with your own, and the home.lan portion with the domain you'd like (minus any subdomains - so instead of www.google.com, it would be google.com). Note, also, that using a real domain here isn't a good idea, as any requests will be directed to your local machine. So if using google, you'll never reach the Google homepage again. Hence why I use home.lan.

I hope this article was helpful to anyone wondering about Docker, who has been looking for a better way to host local projects while in development. If you have any questions, or want to share a docker trick of your own, feel free to email me at lswest34+fcm@gmail.com.

FURTHER READING

https://en.wikipedia.org/wiki/LXC
LXC

https://www.quora.com/What-is-the-difference-between-Docker-and-Vagrant-When-should-you-use-each-one
Docker VS Vagrant

https://www.docker.com/enterprise
Docker

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 reporter (Arnfried) is 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!**

Welcome back to the crazy world of Python Programming in the real world. Before we get started, I need to make a confession. Last time I goofed. The images in part 63 are wrong. The LEDs are backwards from what they should be. Brian Kelly noted this and was brave enough to point out the old man's errors. Thank you Brian. If you follow the text, you should be good to go.

Secondly, I have to apologize for not making it last month (FCM#106). I'm having more medical issues that are keeping me from sitting for too long. Hopefully this will be taken care of soon.

Enough of that. Now for this month's offering.

THE MYSTERY LED

In the last two articles, we learned how to turn on and off LEDs programmatically. That was simple enough. This is digital output as opposed to analog output. The RPi, unlike the Arduino, cannot do analog I/O. So we are limited to turning a GPIO pin (and in this case, a LED) either on or off. This time we will be using that knowledge to do something pretty interesting.

So get your Pi and your breadboard and we'll start working.

THE WIRING

You will need a Raspberry Pi, a breadboard, two LEDs - one Red and one White, two 220 Ohm resistors and 3 jumper wires.

I've used the original Pi for this wiring image example. If you have a Pi B+ or 2B (or even the brand new 3), the pins at this point are exactly the same.

Just to avoid confusion (on my side), the Cathodes (Negative side) of the LEDS are connected to the resistors going to ground, and the Anodes (Positive side) are connected through the jumper wires to the Pi pins. The positive
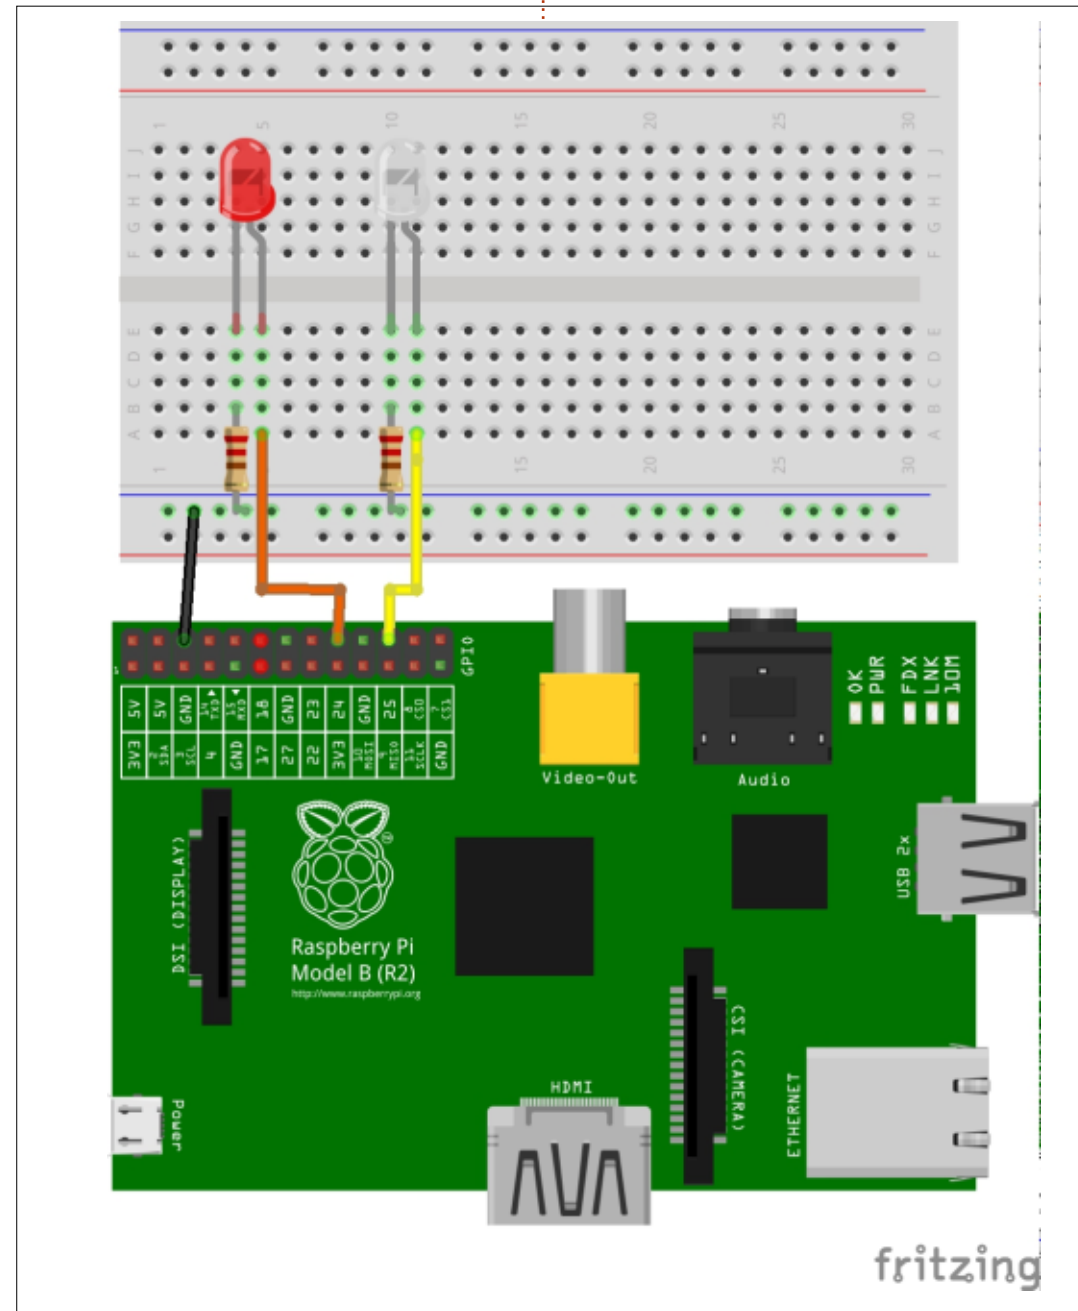
side of the LED is usually marked by the longer lead and the negative side is the one that has the flat spot on the base of the LED.

THE CODE

I won't explain the code just yet. Just put into the editor as it is. We will discuss it in a bit.

Once you have the code entered correctly, then run it and see what happens.

THE REVEAL

If you have been paying attention over all these years, you probably have figured out what the code is doing. If you can't figure it out, don't feel bad. We'll jump into the explanation.

Instead of the LEDs being either on or off, they pulse on and off. Since I said earlier, we can only send out (or read) a On/Off voltage (or 1/0, or High/Low), so how can this be?

We are using a trick called PWM or Pulse Width Modulation. We are still living with the rules, but we

```python
import RPi.GPIO as GPIO
from time import sleep
GPIO.setmode(GPIO.BCM)
GPIO.setup(25,GPIO.OUT)
GPIO.setup(24,GPIO.OUT)
white = GPIO.PWM(25,100)
red = GPIO.PWM(24,100)
white.start(0) # start white led on 0 percent duty cycle (off)
red.start(100) # red fully on (100%)
pause_time = 0.05
print("Program Starting...Press CTRL+C to exit")
try:
    while True:
        for i in range(0,101): #101 because it stops when it finishes 100
            white.ChangeDutyCycle(i)
            red.ChangeDutyCycle(100-i)
            sleep(pause_time)
        for i in range(100,-1,-1):
            white.ChangeDutyCycle(i)
            red.ChangeDutyCycle(100-i)
            sleep(pause_time)
except KeyboardInterrupt:
        white.stop()
        red.stop()
        GPIO.cleanup()
```

are bending them to our benefit. The pictures below, taken from my oscilloscope connected to the project, should help explain a bit clearer. We will be concerned with only one LED at this point.

If we send out a Low to the GPIO pin to the LED it's zero volts. The LED is getting nothing on the Anode, so it is off. In the last two articles, when we turned the LED on by sending the Anode of the LED a High So we have in the first instance a zero, and in the second

a 1. Just like we have assumed... either Off or On.

This time we vary the amount of time that the GPIO signal is high and low. If we do it slowly, the LED would simply flash on and off in response to the voltage. In the case of this version, we are switching it on and off very quickly and at the same time, changing the amount of time it is on compared to off, which is called the duty cycle.

You can see that the signal is on for about 80% of the time and off for about 20%, which would be a 80% duty cycle. By doing this quickly, the LED reacts by dimming
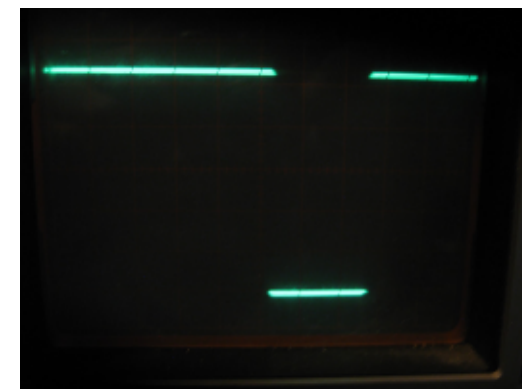
a bit from the 100% on all the time. As the program does its loop, it changes the duty cycle and makes the high longer or shorter depending on what part of the loop it is.
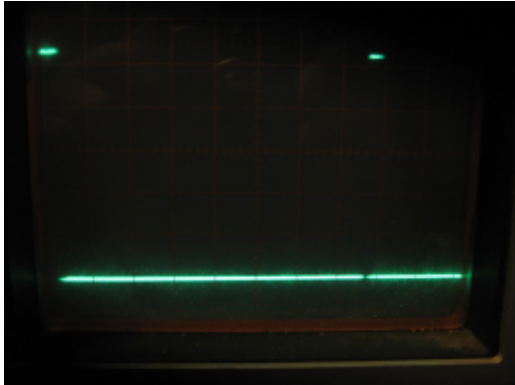
In the picture above, we have a duty cycle of about 5%. In this case the LED is turned on for such a short time, that it is extremely dim and for all intents and purposes it is off.

Now, let's start taking apart the code.

```
import RPi.GPIO as GPIO
from time import sleep
```

As always, we start with our imports. We import the GPIO library, and this time, we import the sleep function from the time library. You will understand the reason for that shortly.

```
try:
    while True:
        for i in range(0,101): # 101 because it stops when it finishes 100
            white.ChangeDutyCycle(i)
            red.ChangeDutyCycle(100-i)
            sleep(pause_time)
        for i in range(100,-1,-1):
            white.ChangeDutyCycle(i)
            red.ChangeDutyCycle(100-i)
            sleep(pause_time)
except KeyboardInterrupt:
        white.stop()
        red.stop()
        GPIO.cleanup()
```

```
GPIO.setmode(GPIO.BCM)
GPIO.setup(25,GPIO.OUT)
GPIO.setup(24,GPIO.OUT)
white = GPIO.PWM(25,100)
red = GPIO.PWM(24,100)
```

In these five lines, we set the GPIO mode to BCM, and set the GPIO pins 24 (physical pin 9) and 25 (physical pin 11) to be output pins. We have done this before. Now we set the values for the PWM to 100% duty cycle.

```
white.start(0) # start white
led on 0 percent duty cycle
(off)
```

```
red.start(100) # red fully on
(100%)
```

We next turn the Red LED on (100%) and the white LED to 0 volts.

```
pause_time = 0.05
```

```
print("Program
Starting...Press CTRL+C to
exit")
```

We set the pause_time variable to 0.05 seconds. This makes it fast enough to (hopefully) not allow for a flicker.

In the next block of code, we do our loops. The first loop is to make the white LED get "brighter" and the red LED to get "dimmer". The second is to reverse the process. Just using the first loop as an example, we use a FOR LOOP to set the value of i and then we set the duty cycle for the white LED to i and that of the red LED to 100-i.

Notice that we have wrapped this with a TRY...EXCEPT set. This allows us to continue to run until the user enters CTRL+C. When that happens, we fall out of the TRY side so we can do our clean up code.

So now you know that we can bend the rules to our use.

Next time, we will start to examine a different GPIO library. Until then, have fun.

Alan holds a PhD in Information and the Knowledge Society. He teaches computer science at Escola Andorrana de Batxillerat (high-school). He has previously given GNU/Linux courses at the University of Andorra and taught GNU/Linux syadmin at the OU of Catalunya.

When you're writing in LibreOffice, do you ever repeat yourself? I said, "Do you ever repeat yourself?" Oh, sorry, I'm repeating myself. In Writer, especially in business writing, there is text you repeat often: greetings and closings in letters, address headers, enclosure statements, and disclaimers, are all examples of repeated text. At other times, you're just trying to design a layout for a newsletter or template, and you need some text to fill in the space so you will know what it looks like. This is where Auto Text is useful. This handy little utility makes it easy for you to repeat yourself without having to type it over and over again.

You can open the Auto Text dialog with Edit > AutoText.

PREDEFINED AUTO TEXT ENTRIES

Auto Text comes with many predefined entries. There are four default categories, but you can add more. The three that already have entries are Business Cards, Work (3 ½ x 2); Only for Templates; and Standard. You will find a few gems in each of the categories.

The business cards have several different predefined card styles. The different styles have placeholders for you to fill in with your information. Using fields, like placeholders, makes these very handy to use. The topic of placeholders is beyond this article, but you will learn about them soon. There are three elegant styles, and two modern styles.

The templates category contains mostly layouts for different headers. The category contains both simple and complex header styles. There are also two for meeting minutes. These can give you a head start on setting a template for a newsletter or meeting notes.

The Standard is a group of general inserts into a document. You'll find default greetings and closings, business document markup (like Attention, Enclosure, CERTIFIED MAIL), and a formula-numbering entry. Included in this category are several text placeholders. These are good when you're trying to design a newsletter or document layout. There is the aptly named Dummy Text, which is a 290 word paragraph. There is also the traditional Lorem Ipsum (a paragraph of "Latin" text). Finally, there are some starter texts for different situations, including job applications and debt-collection reminders.

INSERTING AUTO TEXT

When it comes to inserting Auto Text, there are two methods: from the Auto Text dialog and using the shortcut. The shortcut method is the quickest by far, but let's take a look at both.

When you have the Auto Text dialog open, you can select any of the entries and see the text in the preview window. Once you find the text you want to insert, just click the Insert button to insert the text into your document at the current
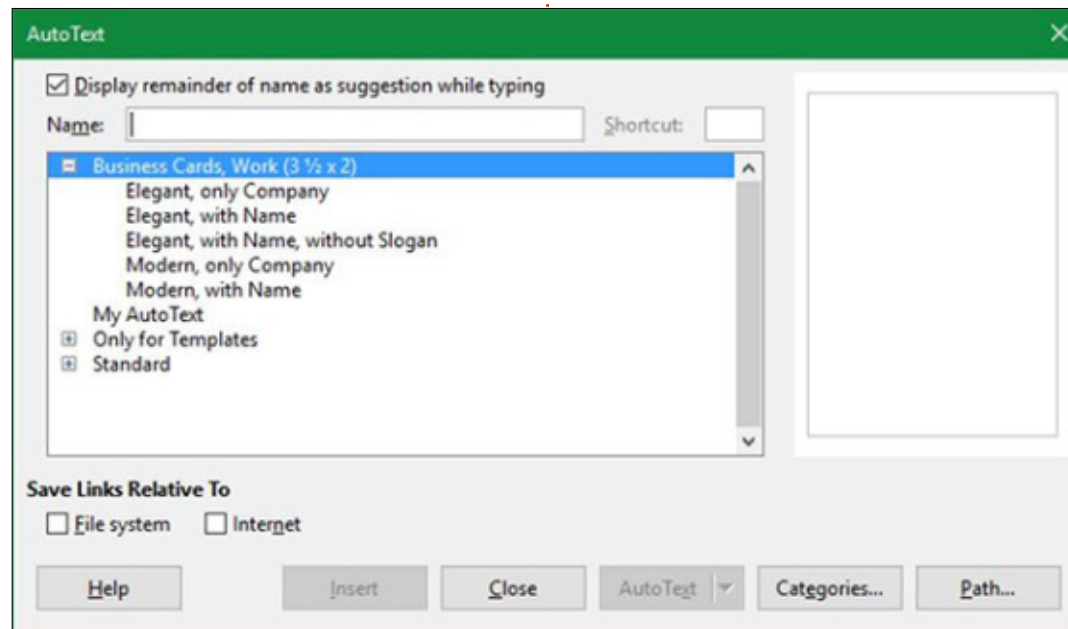
cursor location.

But who wants to open a dialog, find, and select a text entry? You will notice that each Auto Text entry has a name and a shortcut. You can quickly add the text by typing the shortcut for the entry and pressing the F3 key on the keyboard. For example, the shortcut for Lorem Ipsum is LOREM. If you type in LOREM and press the F3 key, Writer will insert the Lorem Ipsum text into your document. Note that while the shortcuts show in all capitals, Auto Text will recognize the shortcut even when in lower or mixed case.

CREATING AUTO TEXT ENTRIES

You can create your own Auto Text entries, so you can repeat yourself using your own words. You probably don't want to get too

carried away here, but if you have something you type almost daily, you should consider creating an Auto Text.

A good example is my journal entries, which I make every day. At the beginning of the day, I start the entry with a date and time stamp to identify when the entry was created. If I make another entry on the same day, I just throw in a time stamp. This involves going through the Insert menu to select the fields, but if I used Auto Text, I could enter a shortcut, press F3, and start typing my entry, which is what I like to do. Let's create the journal entry and time-stamp Auto Text entries, and then we will modify the journal entry to show how you can change an existing entry.

In a new or existing document, Insert > Fields > Date and press enter. Insert > Fields > Time and

press Enter. Highlight the date and time entries you just created in the document. Edit > Auto Text. Select the category My Auto Text. This is the only default category you can add entries to without adding your own categories. In the Name text box, type in Journal Entry. The shortcut should automatically set itself as JE. By default the shortcut is created by the first letter of each word in the name, but you can edit the shortcut should you want to change it from the default. Click the Auto Text button and select New from the drop-down menu. This will create an entry in the My Auto Text category named Journal Entry. Click close to exit the Auto Text dialog.

Now, test your new Auto Text

entry. Type JE on a new line and press F3. You should get your new journal entry header. Because we used fields, the date and time will reflect the date and time of the moment you inserted the entry.

To create the time-stamp, we just need the time, Insert > Fields > Time and press enter. Highlight the time entry you just created. Edit > Auto Text. Make sure the My Auto Text category is selected, and enter Time Stamp in the Name text box. The shortcut will set itself to TS. Click the Auto Text button and select New from the drop down menu. The Time Stamp entry is created in the My Auto Text category.

Test the Time Stamp entry in
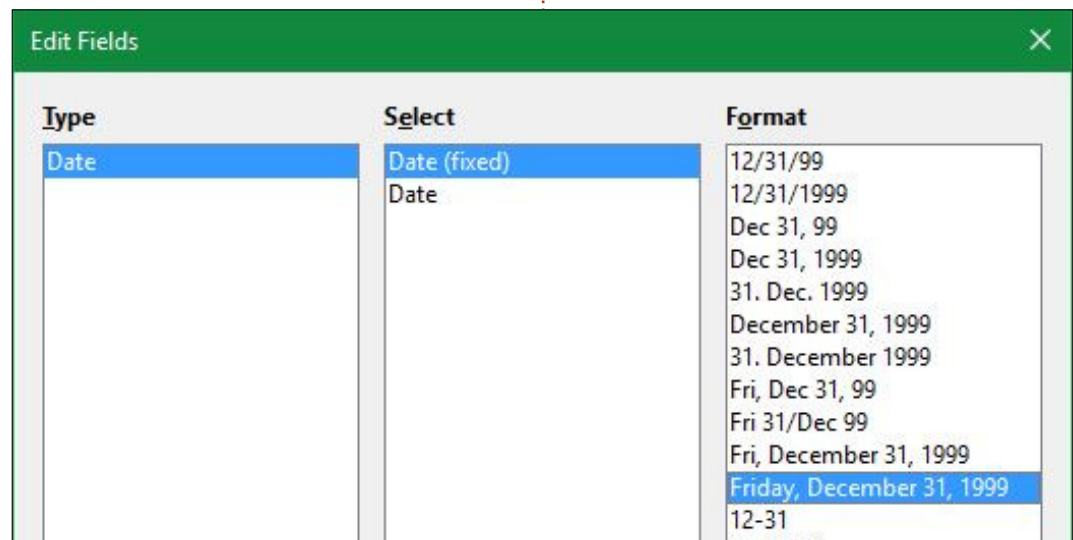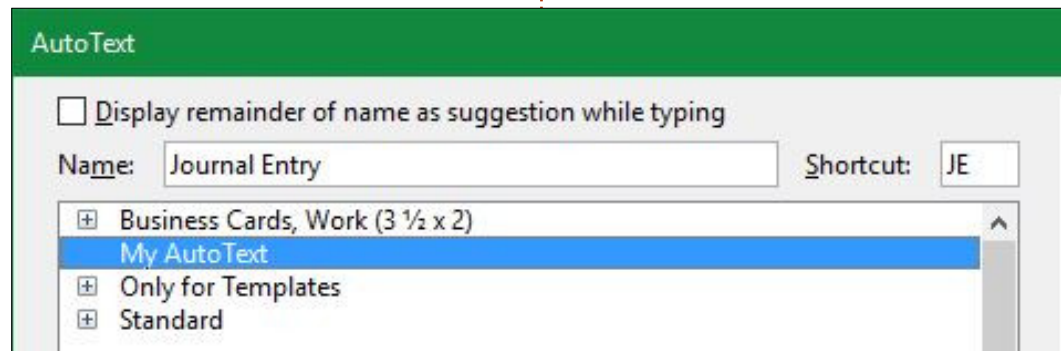
the same manner as for the Journal Entry.

About a week after creating my Auto Text entries, I was reviewing some past entries and thought, "It would be nice if the date showed the day of the week as well." I decided I needed to change the Journal Entry Auto Text to include the day of the week. Let's see how I made the change.

In a document, create a new Journal Entry (JE and press F3). Double-click on the date field to bring up the Edit Fields dialog. In the Format column, select the format Friday, December 31, 1999, which gives you the full date with the day of the week. Click OK. Highlight the changed date and

time. Edit > Auto Text. Under the My Auto Text category, select Journal Entry. Clicking on the Auto Text button, select Change from the drop down menu. Click on Close and test your modified Journal Entry header (JE and press F3).

I know it may not seem like much, but for me, it makes my journal entries quicker, so I can get straight to writing. I'm sure you can think of similar snippets to make your own work go quicker.

PATHS, CATEGORIES, AND LINKS

There are some buttons and check-boxes in the Auto Text

dialog I haven't mentioned yet. They are the Path and Categories buttons, the two "Save Links Relative To" check-boxes, and the "Display remainder of name as suggestion while typing" check-box. Let's take a quick look at these.

If you click on the Path button, it brings up a dialog with a list of paths where LibreOffice stores the Auto Text files. You should see two paths by default. One is a system folder, where the defaults that you can't change are stored. The other is to a user path, where your personal data is stored. This is the path you want to back up in order to preserve the Auto Text entries you have created. You can Add a

new path, like a path to a server on your network, by clicking on the Add button and browsing to the location. The default two are usually enough for most people unless you are in an enterprise or company setting where company Auto Text entries are stored on a server. When you are finished making changes, click on OK.

Clicking on the Categories button brings up a list of the categories and their paths. You can't delete any of the default categories except for My Auto Text. The others are fixed. But you can create new categories. To create a new category, enter the name in the Category text box and select a non-system path from the
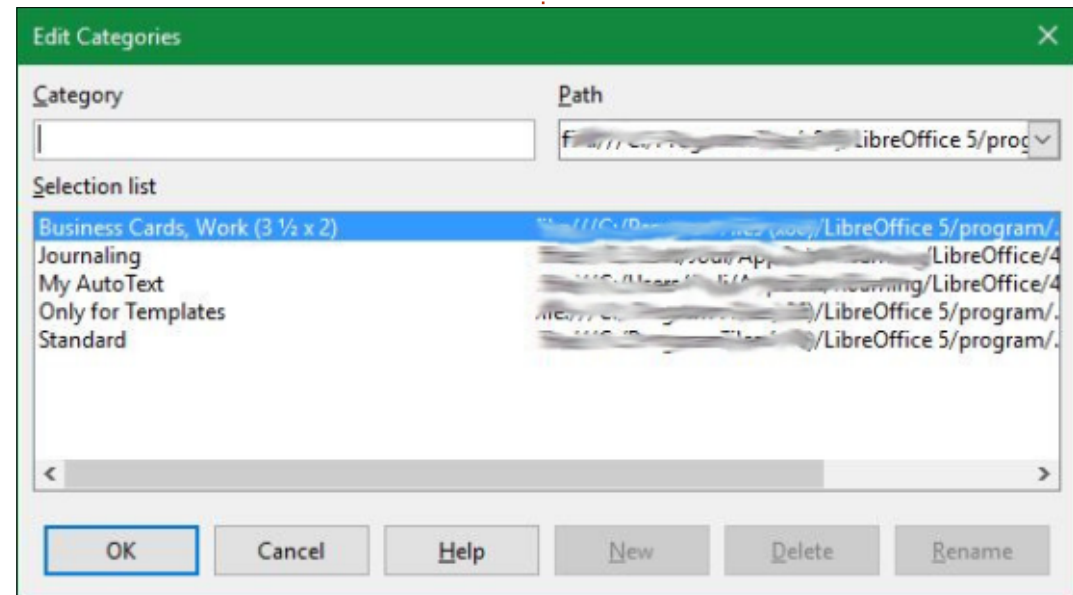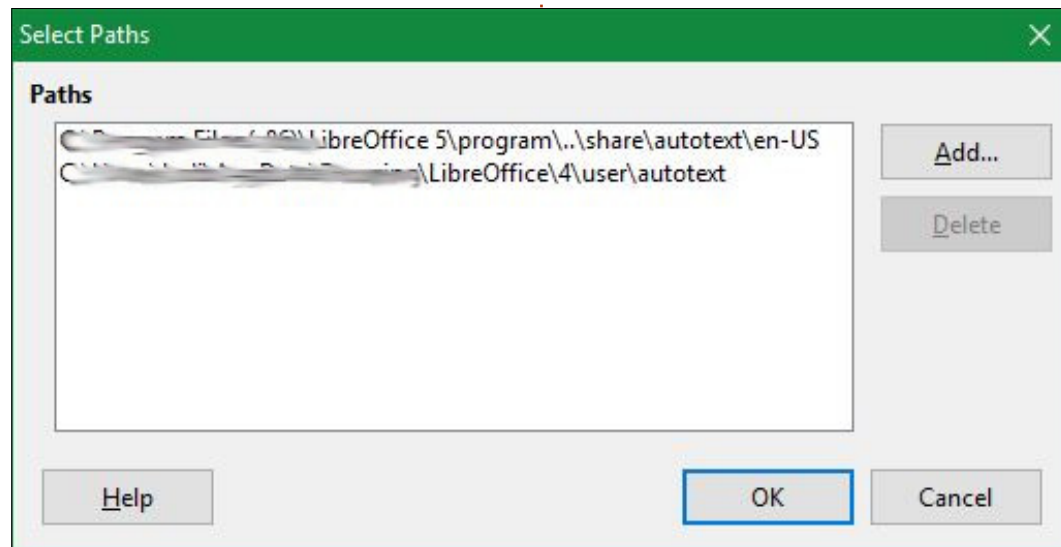
Path drop-down list. Click the New button to add it to the list of categories. I created a new category named Journaling in this manner. The Rename button allows me to edit the name of any of the categories. Just select the category from the list and type the new name in the Category text box. Click the Rename button to change the name. Click the OK button when finished making changes to the categories. You can move entries from one category to another by clicking on the entry and dragging it into the other category. I moved my two journal entries into the new Journaling category.

The "Save Links Relative To" check-boxes control whether the paths to the Auto Text files are relative or absolute. One check-box is for the local file system and the other is for the Internet. For most people, leaving these unchecked is the way to go. In some bigger deployments, the IT staff might decide to check these. The reasons for checking these are beyond the scope of this article, but I wanted to make you aware of them.

The last check-box is the

"Display remainder of names as suggestion while typing" at the top of the dialog. This check-box works in conjunction with the Auto Complete settings. If checked, and the Auto Text shortcut has more that three letters, the name of the Auto Text entry will display in the same manner as other Auto Complete entries. You can use the Auto Complete accept key to insert the Auto Text into your document. A good example is Lorem Ipsum Auto Text. The shortcut is LOREM. With this, and Auto Complete turned on, when you type the R, it shows Lorem Ipsum as an Auto Complete. Pressing the accept key (Enter in my case) inserts the Lorem Ipsum text into your document.

Auto Text is a feature that allows you to repeat information without having to enter it each time. Besides the number of entries already created, You can add new ones of your own. Those entries can contain fields as well as other text formatted the way you want. You can create your own categories for storing and organizing your Auto Text entries.

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

In the early days of computers, a company called Digital Equipment Corporation (DEC) created its 32-bit VAX computer using openVMS as its operating system. Because a VAX/VMS computer is so reliable, there are today - after more than 25 years - still a large number of them in use. But in the end, even these reliable computers will have to be replaced. As described in part 1, you could migrate from VAX/VMS to Linux because the way Linux works is largely compatible with VAX/VMS. If you use Pascal as your programming language, you will find that Lazarus/Free Pascal is a good replacement. But there are technical functions used in VMS with no apparent replacement in Linux. In this article, I will describe how I replace logicals.

FILE-SYSTEM PHILOSOPHY?

To understand logicals, I have to explain the philosophy behind the file-system of VMS. In Linux there is one "device" ('/') and all physical devices are mounted to become a virtual folder (the mounting points). VMS uses a totally opposite approach: Devices are separate items, and you are allowed to create a new (virtual or pseudo) device pointing to a folder.

As an example, this is how you would address a file in VMS:

```
DEVICE:[DIRECTORY.SUBDIR1.SUB
DIR2]FILENAME.EXT;123
```

The same file in Linux would look like:

```
/mountingpoint/directory/subd
ir1/subdir2/filename.ext
```

Note all the uppercase characters in VMS, and the number at the end (the file version, described in part 4).

WHAT ARE LOGICALS?

A logical is a name/value pair maintained by the system, consisting of strings of up to 255 characters. As with eventflags, the reason why they are important is they are used consistently throughout the entire system in VMS. Whenever you access a device, the system will first look if the device name you specified is a logical. If so, the device name is replaced with the value of the logical (in VMS we say: the logical is translated), and that process is repeated up to 10 times (to prevent loops).

This mechanism is very versatile. You can influence what device is accessed by changing the corresponding logical (give it another value), and the value can contain more information than just a device: it can also contain the path to a folder (directory), or other information like TRUE, FALSE, OPEN, CLOSED, etc, etc.

MORE THAN ONE TRANSLATION.

The value is not limited to one string. You are allowed to specify more than one string, separated by a comma to get a search list. If the system does not find the file in the folder resulting from the first "translation", it will continue the search by following the second "translation", and the third, and so on. Though very useful, be warned: this can have a nasty side-effect. If you access a file through a search list, and you don't have permission to access the file (due to file protection, root is the owner, etc), the system will continue the search, but does not find it in the other folders. Then it reports to you the FINAL error of "file not found". But you can see the file and you will think: is that computer crazy or am I? Luckily, neither; the fault is in the message. It should be "access denied" (VMS: "file protection violation"), and I thought VMS was perfect!

In part 1, I mentioned that I have had bad experiences with frequent, but random, errors on Charon-VAX. These errors involved the translation of a logical which contained a search list. One in every hundred or so translations failed. It was impossible to find what was wrong, because the next time the same logical was translated, there was no error. It

🏠 contents ^

was because of the randomness of the error that we decided not to use Charon-VAX on our production computers, only the development computers were emulated.

NOT AS SIMPLE AS IT SOUNDS...

The implementation of logicals in VMS is very complex: a logical is part of a table and there are many different tables. Which tables and the order in which the system looks for the logical is specified in yet another logical. I will not try to explain all of this – it would be outside the scope of this article. Let me only mention that the most used tables are process, job (session), group (GID) and system tables, and they are searched in that order. This makes it possible to create a pointer to a folder for your group in the group table, and temporarily create the same pointer to a different folder for test purposes in your job or process table for your own use - without disturbing the other users or services, as would be the case when using hard links.

EXAMPLES OF USE

To clarify this complex instrument I will show some examples of how logicals are or could be used:

• In Linux, the system is located on '/'; in VMS the system is on disk "SYS$SYSDEVICE".

• In VMS, the logical "TT" points to the terminal you work on. This could be a physical terminal (TTAxxx), a network terminal using the LAT protocol (LTAxxx), a telnet terminal (TNAxxx), an Xterm terminal (FTAxxx), or a remote terminal using the DecNet protocol (RTAxxx). Your program does not need to know which device started the program; it just reads from and writes to "TT:" (note the colon, this signals that TT is a device, not a file). Do not confuse this with stdin and stdout. These are file handles and they open the file "TT:".

• Suppose you made a process that reads files with file-type XML from 3 folders (placed there by 3 other process), and put the result in another. You could define 2 group logicals pointing to these folders:

```
define/group INPUT_FOLDER
SYS$DISK3:[QUEUES.FROM_PROCES
S_1],SYS$DISK3:[QUEUES.FROM_P
ROCESS_2],SYS$DISK3:[QUEUES.F
ROM_PROCESS_3]

define/group OUTPUT_FOLDER
SYS$DISK4:[PROCESSED_DATA]
```

Then your process would look for input using the file specification INPUT_FOLDER:*.XML, and write the result to OUTPUT_FOLDER:RESULT_'DATE'_'TIME'.DAT

Now you made a change to your process and want to test it. Reading files from the input folder(s) would interfere with your system, so you define 2 new logicals at job level:

```
define/job INPUT_FOLDER
MYDISK0:[TEST_QUEUE]

define/job OUTPUT_FOLDER
MYDISK0:[RESULTS]
```

Now you can run the new process and test it by copying files to MYDISK0:[TEST_QUEUE] and examine the resulting files in MYDISK0:[RESULTS].

• Suppose you measure the weight of your product with an old scale, and then you purchase a new scale. In the first few weeks, you want to be able to switch between the scales. You could have 2 different executables, and stop the process that reads the scale and restart it using the other executable when you want to switch scale, or you could put the following code in your program:

```
new_scale := $TRNLNM
("NEW_SCALE"); (TRNLNM =
translate logical name)

if new_scale = 'ACTIVE' then
weight := read_new_scale()
else weight :=
read_old_scale();
```

You could use the following terminal command to use the new scale:

```
define/system NEW_SCALE
ACTIVE
```

and, to use the old scale:

```
define/system NEW_SCALE
INACTIVE (or 'FALSE' or
'USE_OLD_SCALES' or anything
else)
```

To use version control, I create a new folder tree for every version and use a logical to specify which version is active. As an example: My current version is 7.2, so I

create the folder APP$DISK:[V7-02] and the sub-folders DATA and EXEC: APP$DISK:[V7-02.EXEC] and APP$DISK:[V7-02.DATA]

The logical KW21 used as pointer to the current version is created with:

```
define/group KW21
APP$DISK:[V7-
02.]/translation=concealed
(note the dot at the end!)
```

The switch "translation=concealed" makes the logical KW21 a pseudo device, so the sub-folders can be addressed with KW21:[EXEC] and KW21:[DATA].

To use version 7.3, I would create APP$DISK:[V7-03.EXEC] and APP$DISK:[V7-03.DATA] and define KW21 as APP$DISK:[V7-03.] – without changing anything in my applications.

The only drawback of this method is that you have to use KW21:[000000] if you want to address the folder where the logical KW21 is pointing to.

There is a lot more to tell about logicals. If you want to read more about logicals and have some time on your hands, try google with "VMS logicals".

HOW TO HANDLE LOGICALS?

I tried replacing the logicals using hard links, but that works only for logicals with only 1 translation. And - as with eventflags - you can ask the system what the value of a logical is – a feature that cannot be implemented with hard links. This way, you could also define a kind of "system parameters" to be used by running processes. This makes it possible to implement software changes that can be switched on and off by changing the value of a logical instead of stopping a process and restarting it using a different executable. Or changing to or from daylight saving time without setting the clock. The list of other possible uses in VMS is long.

For the search list, you can use environment variables (PATH), but these are valid within only one process and cannot be changed by another process. Also not an option...

MY SOLUTION

To use logicals on Linux, I created a few procedures in my library to replace the VMS functions to define, delete ("deassign") and translate logicals. In VMS, the logicals are part of the kernel and located in the kernel's "paged memory". This part of the kernel's memory can grow in size and can be temporarily swapped out (VMS: paged out). As my solution is created in user-space, without the kernel controlling which process has access to which table, I decided to take a different approach: I placed the tables in separate data files.

Upon program start, the shared tables (system, group and job) are opened (or created if necessary), and mapped into shared memory, so they can be accessed directly. This happens only as soon as any of the logical-library routines is called for the first time. Any changes due to (re)defining or removing of logicals are written back to the corresponding file and are immediately visible by other processes as it's mapped into shared memory. The process local table (LNM_PROCESS) is created and also mapped into memory, but not shared. And here lies a little problem:

LNM_PROCESS is not a table. It's a logical pointing to another table containing the process ID (pid) in its name to separate the many process tables from each other. The behavior of Linux concerning process ID's is different from that of VMS (I will go into more detail on this in part 5 about DCL). DCL (Digital Command Language) is a real shell: it wraps around a process, and every time an executable ("image" in VMS) is started, it runs within this shell, within the same memory, with the same logicals and symbols (environment variables) and with the same process ID. Bash is called a shell, but it is NOT! It is a CLI (Command Line Interpreter). Every time an executable is started in Linux, it's started in a new subprocess. Besides the greater overhead of the creation of a new process, it also gets a new process ID. And therein is the problem: A different process ID means a different process table. In VMS, you can define a process logical and then start an executable that would use this logical. In Linux, the defined process logical will be gone on exit of the program defining the logical.

To circumvent this problem I decided to use the process ID of the parent instead. This will be the same every time you start a new executable from the same terminal. Because this causes a different problem (having the same process table) when starting several background ("detached" in VMS) programs from a single terminal or program, I created the program "RunDetached" to do this for you. This again causes all programs to use process ID 1 from the init program, but that is easily detected. In that case its own process ID minus 1 is used.

As mentioned before, logicals are used consistently throughout the entire system in VMS. This means that whenever you want to access a device, VMS will first try to translate the name of the device. To get the same behavior, I created a library function named OPEN to substitute the file access function of VAX-pascal, but this works only for TEXT handles. If another type of file handle is used, than it will be used to determine the file attributes (like record size), and I cannot mimic that. In this case the best solution is to manually replace the OPEN statement by a sequence of:

• Call to "translate name" to get the VMS behavior;
• Call to ASSIGN to open the file, while specifying the correct file handle.

I am hoping that I can incorporate this into my conversion software later on.

The above mentioned functions ("define", etc.) can also be accomplished from the terminal, so I created the three corresponding programs too. They are – of course – available as open source just like RunDetached.

Next month: In the next article I will go more in-depth about other functions such as mailboxes (IPC), what you MUST know about the difference in the way VAX-pascal and Free Pascal handle "PACKED ARRAY OF CHAR" (strings), and how to deal with file version numbers.

After keeping VAX/VMS systems running for 30 years, **Theo** needed a new challenge and started from scratch with Linux to re-do every mistake he made all over again. You can email Theo at: info@theovanoosten.nl

THE OFFICIAL FULL CIRCLE APP FOR UBUNTU TOUCH

Brian Douglass has created a fantastic app for Ubuntu Touch devices that will allow you to view current issues, and back issues, and to download and view them on your Ubuntu Touch phone/tablet.

INSTALL

Either search for 'full circle' in the Ubuntu Touch store and click install, or view the URL below on your device and click install to be taken to the store page.

https://uappexplorer.com/app/fullcircle.bhdouglass
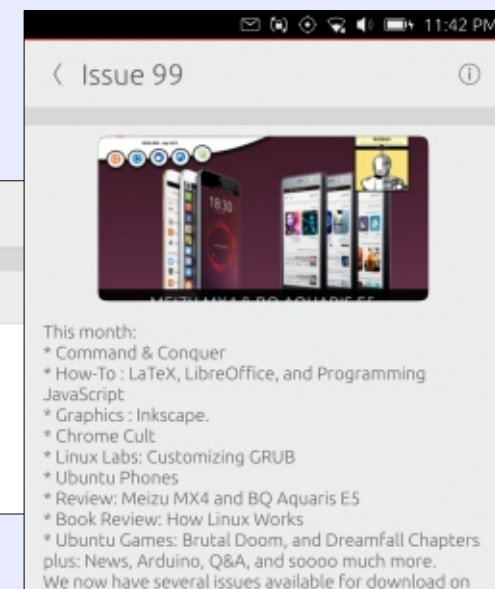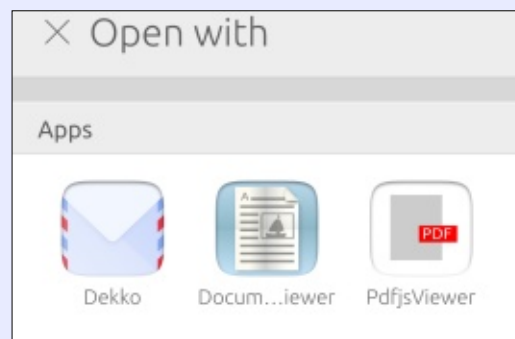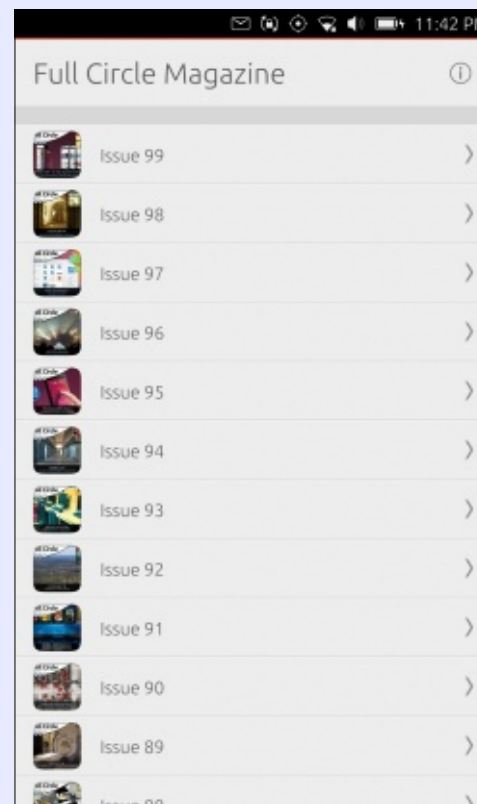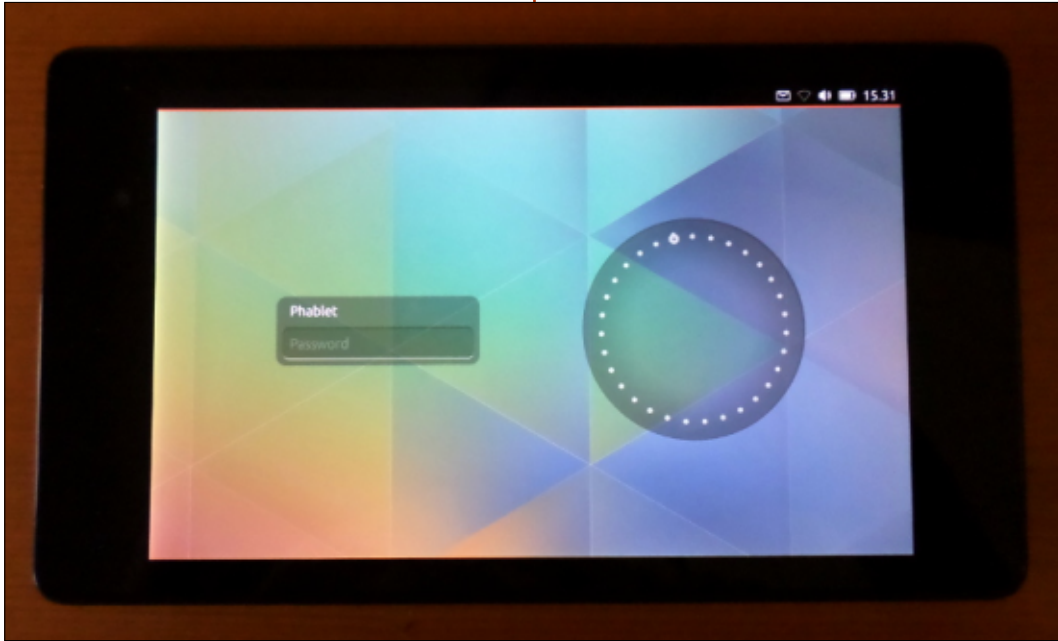
I would like to start off this article by presenting my new server to the reader.

Yes, this is it (shown above). A 2013 version (actually 2014) Nexus 7 tablet, with the Ubuntu Touch operating system. Naturally, it is no news today that mobile phone and tablet hardware is up to some creative use (light use, actually) as a server - the Raspberry Pi took care of that. A point could even be made that ARM processors are ideal for servers with light or sporadic use, since the very same power management characteristics that are so useful making the most of your mobile phone's battery can actually be put to use, putting such a server to sleep when not processing requests, thus saving on electrical bills.

What was lacking was in fact the operating system. Neither iOS nor Android can be seen as server-grade operating systems, they are just not built that way. Their paradigm is rather that of a single-user system with a graphical desktop, which is fine since most devices are used precisely for that purpose. But what about Ubuntu Touch? Since the tablet version of Ubuntu is a direct derivative of the desktop offering, there are many characteristics of the system that can be used to set up a pocket server. We will consider some of them in the following lines.

WE HAVE ROOT

To operate a server, having root access to the underlying system is a must. Software needs to be installed and configured, and indeed some services need root access just to start up - mainly those using privileged TCP/IP ports in the 1-1023 range (think Web servers).

We all know about the hoops iOS and Android make users run through just to gain root access. Special programs need to be installed, that basically use much the same tools a hacker would need to escalate privileges and become administrator. This is obviously a bit problematic, on two counts. In the first place, the very existence of rooting applications means that there are well-known defects in system security. So, what is to avoid other applications or malware using the very same defects for nefarious purposes? On the other hand, the user installing such a kit should always ask him- or her-self whether the person making such an app available could eventually be tempted to include a backdoor, making the entire device controllable from elsewhere without the legitimate proprietor's knowing anything about it. Somebody who has the knowledge to root a device's operating system will probably also have the information to set up such a trojan.

In Ubuntu Touch, on the other hand, we can just fire up the terminal - and yes, there is a terminal available as standard. Much as you would on any Ubuntu computer, the default user (who goes by the login "phablet") just needs to issue a sudo bash, and there one is with root privileges. The password is the same

password or PIN code used to set up the device.

This can be seen both as a good, and a bad feature. On one hand, there is no futzing about with software of dubious origin. On the other, any software that is well-enough thought out could eventually persuade the user to enter his password, thus gaining root access. If Ubuntu Touch were to gain a considerable market share, then such phishing attempts would unfortunately become more a probability than a mere hypothesis.

THE PACKAGE MANAGER

Since we have root, it should be a relatively simple affair to install our favorite software packages from the repositories, and off we go. Unfortunately, things are not quite that simple.

In the first place, Ubuntu Touch

has gone the Snappy way. This is a new way of distributing the actual files that a software package contains into our computer's filesystem. In the regular versions of the *buntu distributions, the apt package system is used. In each package file (actually a .DEB extension file), many individual files are contained. On installation, each file will be written to the appropriate directory on our computer. So configuration files go into /etc, binary (program) files into /usr/bin, libraries into /usr/lib, and program data into /var. Snappy packages work in a different way:

"Each snappy package is installed into its own directory. snappy packages will never overwrite files that belong to different packages or older versions of the same package. A normal snappy package can read only its own space and write to a special writable area. This is enforced via the apparmor profile for ubuntu-core apps."

Source: Ubuntu Snappy Filesystem Layout Guide https://developer.ubuntu.com/en/snappy/guides/filesystem-layout/.

Instead of using the Snappy default location for applications

(which is /apps), Ubuntu Touch uses directory /custom/click. For example, the Terminal application that comes with Ubuntu Touch can be located in three versions, all in subdirectories of /custom/click/com.ubuntu.terminal:

```
# find / -name terminal
```

```
/custom/click/com.ubuntu.terminal/0.7.121/lib/arm-linux-gnueabihf/bin/terminal
```

```
/custom/click/com.ubuntu.terminal/0.7.121/lib/i386-linux-gnu/bin/terminal
```

```
/custom/click/com.ubuntu.terminal/0.7.121/lib/x86_64-linux-gnu/bin/terminal
```

Ubuntu Touch also has a different software manager. Both graphical (the "Ubuntu Store" app) and console ("pkcon") versions are available. Unfortunately, neither of them has complete access to the vast range of software available inside the apt repositories. Even worse, some packages would seem to be available using pkcon, but cannot actually be installed in this way.

So, what can we use to install our favorite server software? The answer is naturally the very same apt-get, aptitude, etc, commands

we are used to on Ubuntu Server. Now, before going any further, let us stress that this is NOT something Canonical approves of, which is quite understandable in a way since there is no guarantee that the Snappy/Click packages and apt packages will play well together going forward. So please proceed with caution - and be prepared to reinstall the system from scratch if everything should go belly-up. (I do not think anything very bad will happen, but it COULD, so...)

Before proceeding, we should note that the root filesystem is mounted read-only by default on Ubuntu Touch. So the first thing we will need to do is make it remount read-write. To do so, start on the device by going to Settings > About this device > Developer Mode, and turn the Developer Mode on. Now, with the USB cable, connect the device to the computer used to install Ubuntu Touch - or any computer with the phablet-tools package installed. As root, issue the following command on the computer:

```
# phablet-config writable-image
```

You should eventually see the device reboot, now with the filesystem in read-write.

Inside the Terminal app on the device itself, or through a terminal from the computer (try command "adb shell" on the computer with the device connected), we can now issue a series of commands:

```
phablet@ubuntu-phablet:~$
sudo bash

[sudo] password for phablet:

root@ubuntu-phablet:~# apt-
get update
```

and we should see the tablet making its connection to the Ubuntu repositories in the usual manner.

MAKING USE OF AVAILABLE TECHNIQUES

Once the apt system is up and running, we can start installing the software to turn our tablet into a server. Just to make things clear: we will be installing software that is meant to work in the background, with, at most, error messages on the console or in log files. There will be no graphical programs on this one, basically because most graphical software for Ubuntu is still compiled for the X server. Ubuntu Touch runs Mir, which is not compatible.

Perhaps a first step could be to install an SSH server, to enable us to SSH in from another computer. Actually, this is not necessary since the openssh-server package is already installed in Touch. However, for some reason it is not started automatically on boot. A quick fix is to edit the /etc/rc.local file and insert the appropriate command at the end of this file. It should now end like this:

```
service ssh start

exit 0
```

Meanwhile, the SSH service can be started manually at any time using the service command:

```
service ssh restart
```

When SSH-ing in from another computer, please remember the default user is "phablet", so - assuming 192.168.0.117 is the IP address of the tablet, try:

```
ssh phablet@192.168.0.117
```

Root entry through SSH and password-less access can be configured in the usual way.

Since SSH has support for file copying, the scp and rsync commands will work to transfer files to and from the tablet. SFTP will also work, enabling most desktop managers to mount the device's filesystem over the network.

A second service that may be of use is a web-server. Apache is a likely candidate:

```
# apt-get install apache2
```

Once working, we could investigate options such as using Apache for webDAV. This would mean that once up, any other devices (or computers) on the same network could access files on the device, and if using webDAV-enabled software (such as Cadaver) could upload files to the device. Some calendar systems like to use webDAV to synchronize items.

Windows file-sharing is easily enabled. Just install Samba, and the configuration file /etc/samba/smb.conf. Do not forget to add a samba password to user phablet! So:
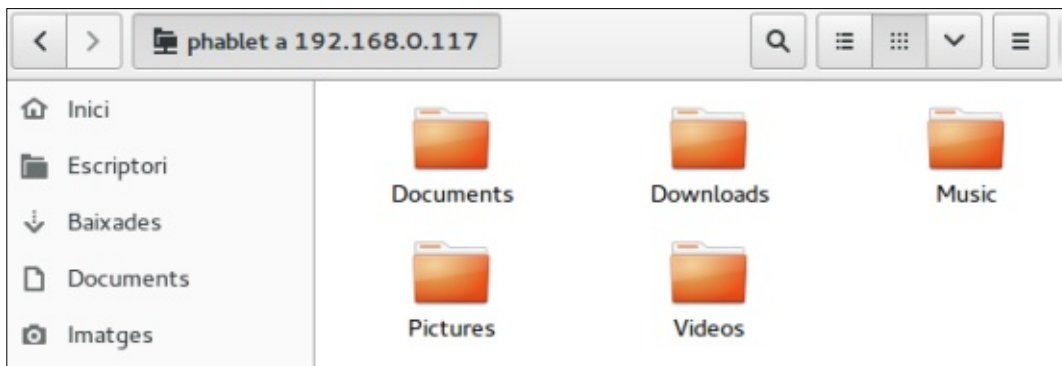
```
# apt-get install samba

# vi /etc/samba/smb.conf

# smbpasswd -a phablet

# service smbd restart
```

From another computer, we can now navigate through the network and log into our device. Try connecting to address "smb://phablet@192.168.0.117"

HOWTO - POCKET SERVER

(substituting your device's IP address). It is always nice to see a mere mobile device recognized as a full-blown server.

For extra Geek points, install a git repository and use it to work collaboratively on a project with other people. Doing a commit to your phone is probably not within the bound of normal users' experience.

Finally, an easy - and elegant - way of sharing the music and/or videos from your device to the local network is using a Universal Plug-'n-Play software tool to export your media library. Unfortunately, a complete Home Theatre such as Kodi/XBMC cannot be installed, since it is much too reliant on having a graphical display available. Even standalone mode does not seem to work on Touch. However, smaller

UPnP/DLNA tools such as miniupnp do work well. Installation is simply:

```
# apt-get install sqlite
minidlna
```

Then, edit files /etc/default/minidlna and /etc/minidlna.conf with appropriate configuration stanzas, such as:

```
media_dir=A,/home/phablet/Music
media_dir=V,/home/phablet/Videos
```

and

```
network_interface=wlan0
```

Reboot the server:

```
# service minidlna restart
```

If there are any problems, you may find some indications on what is happening in the log files:

```
# tail /var/log/minidlna.log
```

From any other computer or tablet on the same network, the tablet's contents should now be available. For example, on VLC:

SOME FINAL WORDS

Just to conclude, it may be well to consider some security aspects. Configuring a server securely is supposed to be a complex endeavour - and it actually is. The techniques shown above are in essence opening up doors to the world, such that those outside could conceivable use to get in. If you store sensitive information on the device, this could eventually be compromised.

So it is perhaps best to consider using a mix of security techniques such as strong passwords and encrypted protocols (HTTPs) where available. It is also good practice not to leave less secure services such as Samba (Windows file sharing) and UPnP open on networks you do not fully control. If you do wish to use them, then perhaps it would be best to make sure these servers are not broadcasting on a public network.

Some possibilities are to tie them to a particular IP address - one that the device uses on your home network, but not on others -, or simply to have these services off by default and turn them on only when required.

With this in mind, there is no lack of interesting projects that can be investigated with an Ubuntu Touch device. Basically, if a Raspberry Pi can handle it, chances are the 'phone in your pocket can also do so. In any case, it is nice to know that such possibilities are now available to make your device a little more than just a window for browsing the Internet.

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

Able2Extract 10

Create, Convert and Edit PDF

- ✓ Convert PDFs to Microsoft Word, Excel, PowerPoint, CSV, AutoCad, Text, Images, OpenOffice etc.

- ✓ Convert any file format to Excel.

- ✓ Edit PDF text right on the spot.

- ✓ Reassemble, merge and split PDFs.

- ✓ Protect and secure your PDFs.

Works with:

Ubuntu   Fedora

Able2Extract 10
Create, Convert and Edit PDF

Convert PDF to Excel, Word, PowerPoint and more...

Compatible with:

INVESTINTECH.COM
PDF SOLUTIONS

INVESTINTECH.COM
PDF SOLUTIONS

This month, we'll conclude our tour of LPEs by looking at the two new additions to Inkscape 0.91: Power Stroke and Clone Original Path. One limitation artists often find with Inkscape is its inability to produce variable thickness strokes. There are ways to fake it, which pretty much all rely on the "stroke" actually being a filled path in its own right. That inevitably leads to the follow-up problem of how to fill a shape drawn using such fake strokes. These two LPEs are Inkscape's answer to those problems.

POWER STROKE

This effect lets you vary the width of your path by adjusting "stroke knots" along its length. For cases where you might otherwise use PAP to provide some variability to the stroke width, Power Stroke will often achieve a similar effect but with more flexibility. Let's start with the simple example of applying this LPE to a straight line:

The top line, in red, represents the original skeleton path. On

applying the LPE the first result you'll see is that the skeleton is replaced by a closed path that encloses the original shape. The closed path takes on the attributes of the original skeleton path, so, in this case, it has a red stroke and a transparent fill, resulting in the second shape in the image. Usually you'll want your Power Stroke to be filled with no outline, so in the third image I've changed the style to a black fill with no stroke. You would be forgiven for thinking that we've just gone through a convoluted way to produce a slightly thicker black line with rounded end caps, but the fourth image shows the real secret to Power Stroke: this is the same as

the third object, but with the Node tool selected (F2), revealing not only the normal start and end nodes, but three additional pink or purple colored handles (the "stroke knots") sitting on the periphery of the shape.

The effect automatically adds these nodes at the start and end of the path, and somewhere towards the middle. Using the Node tool they can be dragged perpendicularly to the skeleton path to set the stroke width at that point, but can also be dragged along the path to change the location at which the thickness changes. Taking the previous example; dragging the handles around a little lets us easily produce this result:

As you can see, the thickness of the line is set by each of the nodes, with sections in-between ramping linearly from one node to the next. Looking at the LPE's UI you'll

notice a pop-up menu for the Interpolator Type. This is what is producing the linear change; pick another value to alter the way in which the width of the path is modified from one node to the next.
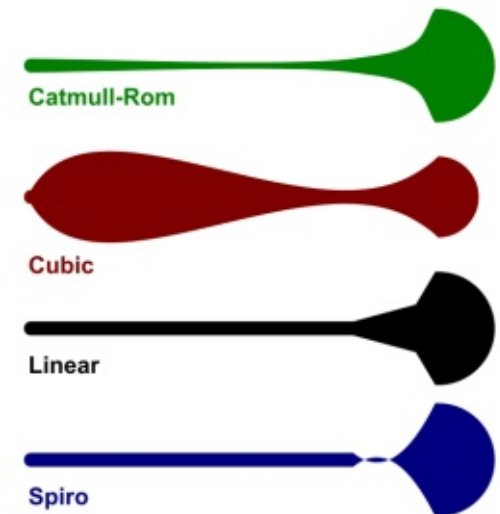
You can't fail to have noticed the rounded ends to the line. They weren't present on the original skeleton path, so where did they come from? A quick scan of the effect UI will show that several of the controls from the Stroke Style tab of the Fill and Stroke dialog are replicated as part of the LPE. From there, you can set start and end

caps (the source of the rounded ends in this case), as well as the join type and miter limit for paths with angles in them. These all operate in a similar way to the equivalent controls in the Style and Stroke dialog, except that the LPE offers more options.
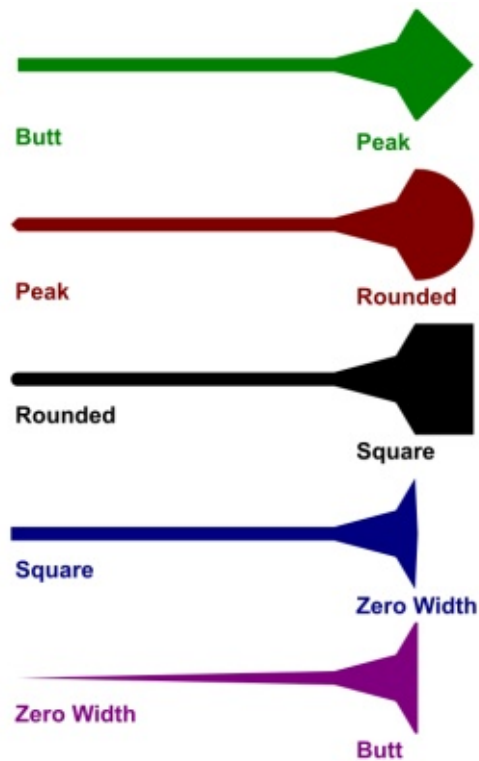
With three handles for manipulating the Power Stroke, you can already produce some useful results, but the real power comes when you add even more handles. Unfortunately, this is done via a rather clunky approach that can sometimes be a bit

unstable, particularly on more complex paths. The first step is to select your path and switch to the Node tool, so that the handles are visible. Next you have to hold the Control-key whilst clicking on a handle. This will create a second handle, directly on top of the first one, which you can then drag to its new location. If you do experience problems, such as the handle becoming detached from the path and having no effect, undo your changes, and then try again, duplicating a different handle instead. Our simple line with three handles quickly turns into something more bumpy once a few more are added:

To delete a handle you have to click on it whilst holding the Control and Alt keys. There's no way to select multiple handles in order to move or delete several at a time. You've probably noticed that you can drag handles past one another with ease – that's thanks to the Sort Points checkbox in the effect's UI. Uncheck that for a different behaviour in which the shape is drawn from handle to

handle based on their original ordering rather than their final position along the line. It's useful for some effects, but generally it's better to leave the box checked.

With all that background out of the way, let's take a look at this LPE when used on a more complex path. It's time to return to my efforts to manually trace "Frankie" (see parts 16-21 for my previous attempts at this):
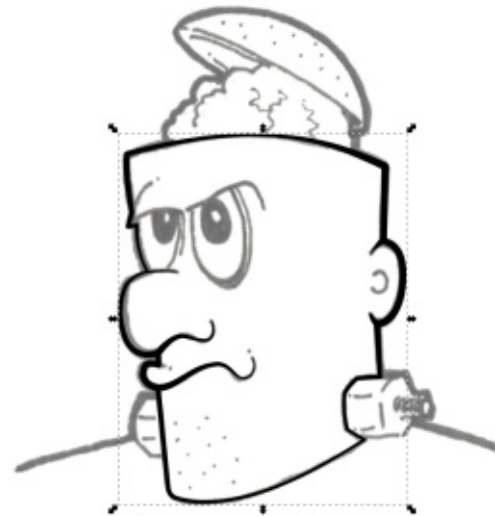
The parts of the path that stretch inside the outline – the mouth, nose and ears – were particularly fiddly to get right. In practice it's probably faster to just use a simple path for the outline, then draw those parts separately. But you can clearly see how the

LPE allows the path to thicken and thin to give more of a dynamic feel to the character than a simple fixed width line could achieve.

Clone Original Path

Our final LPE follows on directly from those efforts to manually trace Frankie. One problem with the Power Stroke or PAP effects is that they draw strokes as filled paths, so setting a fill color on them actually changes the color of the "stroke", not the area inside it. To clarify, if I were to select the Power Stroke path in the Frankie example and set its fill color to red, the result would just be a red Power Stroke, not a black Power Stroke with a red fill inside the face area.

Previously, I've shown you how to work around this limitation using the Bucket Fill tool or by manually drawing a second path that you can fill and then send below the outline. Neither of these solutions is ideal, and both can take a lot of manual tweaking to get right. Wouldn't it be easier if you could just fill the original skeleton path with a different color, without affecting the fill

that's used for the Power Stroke? Effectively that's what the Clone Original Path effect lets you do.

There are a few ways to use this effect, but we'll start with the long-winded method, to give you a better understanding of what's happening. First you'll need a sacrificial skeleton. Don't worry, we're not heading into Voodoo territory, but rather you'll need a skeleton path that will completely disappear once you use the LPE. Its only purpose is to serve as an object to apply the effect to, so a simple straight line will suffice. Select the path and add the Clone Original Path LPE to it, then gasp in amazement as… nothing happens. There are a few more steps to go through before the effect has any visible impact.

Having applied the effect, you next need to select your Power Stroked path, then copy it to the clipboard. Re-select the sacrificial skeleton and use the first button in the effect's UI to paste the path from the clipboard. Your skeleton will disappear, and it will seem that the Power Stroked path has been selected instead. Appearances can be deceptive, however – use the arrow keys to move the selected

item and you'll realise that you've actually got a clone of the skeleton path used in the Power Stroke. Fill it with a color and send it back in the stack and you've achieved in seconds what would have taken several minutes to do manually.
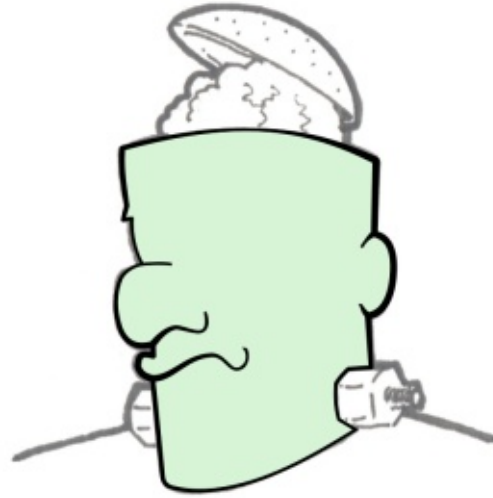
You can fill your clone with a gradient or pattern, if you prefer, or change the stroke style. In principle you can also apply other LPEs – although they don't always chain as well as you might hope. And because it's a clone, you can change the shape of the original skeleton path (the one used for the Power Stroke), and your filled version will automatically update to match it – comic creators rejoice! No more does every tweak to an outline have to entail a

corresponding manual update to the fill shape!

With the cloned path selected, the second button in the UI will select the original – though the classic Shift-D shortcut or Edit > Clone > Select Original menu entry both also work. Cloning a path in this way isn't restricted to the Power Stroke LPE, so if you need a copy of the skeleton that you used with any other effect, just follow the same steps.

You can streamline the creation process a little by copying the Power Stroke path to the clipboard before you create your sacrificial skeleton rather than afterwards. But the Inkscape developers are nicer than that, and have added an option that will automatically create an infinitesimally small sacrificial path that has just a single node, add the Clone Original Path effect to it, and connect it to your original path, all from a single menu entry. Just select the Power Stroked path, then use Edit > Clone > Clone Original Path (LPE), and then set the fill and stroke you want for your clone. The only way it could be any faster is if there was a keyboard shortcut for the menu option.

But wait! Inkscape 0.91 does include a keyboard shortcut editor. Click on the Edit > Preferences menu item, then in the dialog, drill down to Interface > Keyboard Shortcuts. Expand the Edit section of the pane on the right and you should find Clone Original Path (LPE) in the list. Click in the Shortcut column for that entry, then press the new keyboard shortcut you wish to use (I went for CTRL-ALT-SHIFT-D to keep it in line with the other cloning shortcuts).

There's one final trick up the sleeve of this LPE. Back in part 30, I introduced the use of "unset" fills and strokes to allow different clones to have different styles and colors. This effect offers another way to achieve a similar result. You can either select an original object that you wish to clone, and then use the Edit > Clone > Clone Original Path (LPE) menu option, or you can select a clone that you've already created and just click the "+" button in the Live Path Effects dialog. Now you can change the color and style of your clone with impunity, safe in the knowledge that changes to the original shape will still be

reflected. For obvious reasons this works only where the parent object is a path or can be trivially converted into one – so it does work with stars, spirals and even text objects, but doesn't work with groups or bitmap images.

It's worth noting that this approach does result in a second copy of the path data being stored in the clone's XML (see part 43 for details), unlike a normal clone which just hold a reference to the original. This means that not only is the file size a little larger, but any edits made to the original outside of Inkscape won't be reflected in the LPE clone, whereas they would with a "real" SVG clone. One advantage of this approach, however, is that you don't have to unset the fill and stroke on the original, so you won't be left with a black shape that you might have to hide under another object or by placing it off the side of the page.

Conclusion

The Power Stroke and Clone Original Path LPEs are worthy additions to Inkscape 0.91. For a comic artist they could be reason enough to upgrade from an older version. Using little more than

these two effects produced the best manual trace of Frankie so far:
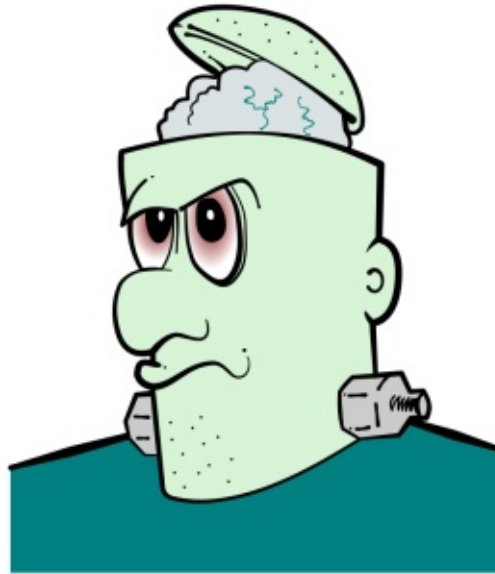
We've now reached the end of our tour of LPEs. There are more being added to the development builds all the time, so do check out the dialog with each future version of Inkscape. They represent perhaps the most common way in which the developers have broken through the limitations of the SVG format in order to add functionality that far exceeds what any normal SVG editor could offer. It's true that the UI for some is confusing, they're sometimes a little unstable, and that they often don't chain as well as they should, but it's well worth spending some time to play around with them as

they can open a path to drawing possibilities that would be impossible or, at least, impractical to produce any other way.

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

IncrediBuild
BEYOND ACCELERATION

ACCELERATE LINUX AND ANDROID DEVELOPMENT

❝

Being able to directly visually audit the build process to look for bottlenecks whilst reducing execution time is wonderful.

Richard Trotter
Geoteric

HAVE YOU EVER THOUGHT WHAT WOULD HAPPEN IF YOU MAKE ALL YOUR DEVELOPMENT PROCESSES RUN 10 OR 20 TIMES FASTER?

Slow builds, long running tests and scripts, compute intensive development processes **delay continuous delivery**, leading to **longer release cycles**, missed deadlines, broken builds, overworked develops, and insufficiently tested software.

INCREDIBUILD ACCELERATES BUILDS, COMPILATIONS, TESTING, AND ANY OTHER DEVELOPMENT PROCESS
WE SPEED UP YOUR DEVELOPMENT LIFECYCLE

Once thought a reality of every development process, **make slow builds a thing of the past.**

Increase your development productivity, accelerate your build lifecycle, and enable truly Agile development.

Realize the premise of faster Continuous Delivery and get your Continuous Integration to perform.

WARNING: This month we're heading into the big league with controlling household voltage devices. It should go without saying, but I'll say it anyway, you must make sure your 110/240V devices are unplugged before chopping cables, and double-check your wiring before you do put that plug back into the wall outlet.

In previous articles, I used an LCD screen which, while it worked, would need about a dozen wires and a potentiometer to control the screen brightness (see photo, top board). This time, I've managed to get some much simpler LCD screens that require only four wires (same photo, bottom board). The wires are VCC, GND, SDA and SCL. These should be marked on your Arduino, but on UNO boards, I believe it's A4 and A5. On my MEGA, it's 20 and 21 and they're marked as such. But you need to put both those lines through a 4.7K resistor from 5V.

These newer LCD screens are I2C which means they have a little controller board on the back. It also means they require a newer LCD library (https://bitbucket.org/fmalpartida/new-liquidcrystal) in the code.

Before adding LCD code, we need to do a scan of the LCD screen to get its I2C address. Different models have different addresses. So, grab the code from: http://arduino.cc/playground/Main/I2cScanner and run it to see your LCD I2C address in the serial monitor. Mine is 0x3F.

LCD CODE

I need to include that new library:
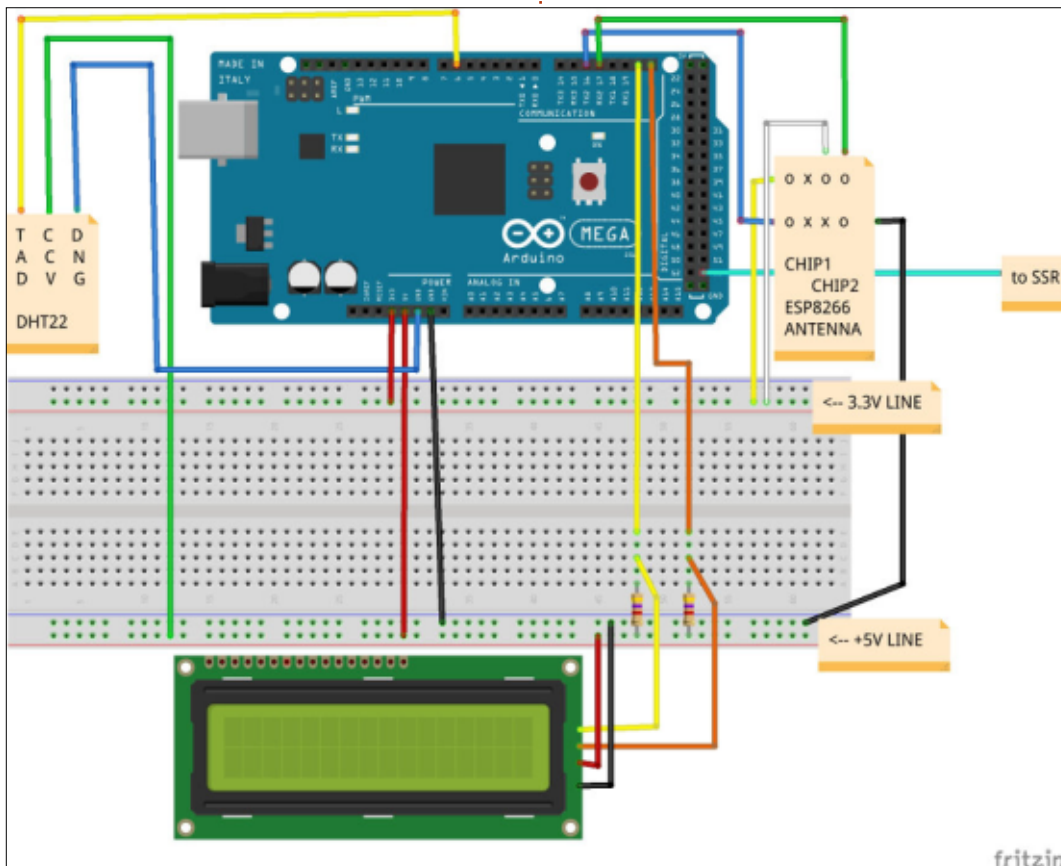

```
#include
<LiquidCrystal_I2C.h>
```

Define various pins (I2C stuff that you shouldn't need to touch), but insert the correct I2C address:

```
#define I2C_ADDR    0x3F
// you shouldn't need to
touch the pin numbers below

#define BACKLIGHT_PIN      3
#define En_pin   2
#define Rw_pin   1
#define Rs_pin   0
#define D4_pin   4
#define D5_pin   5
#define D6_pin   6
#define D7_pin   7

LiquidCrystal_I2C
lcd(I2C_ADDR,En_pin,Rw_pin,Rs
_pin,D4_pin,D5_pin,D6_pin,D7_
pin);
```

In my setup, I initialise the screen as 16X2, turn on the backlight, print a message, and use a delay to allow it to be read.

```
lcd.begin (16,2);

lcd.setBacklightPin(BACKLIGHT
_PIN,POSITIVE);

lcd.setBacklight(HIGH);

lcd.home();
```

```
lcd.print("Screen OK");

delay(2000);
```

Most of those commands are pretty self explanatory.

I now have my temperature, humidity, and some of my serial monitor outputs going to the LCD screen using the lcd.print() command.

```
lcd.setCursor(0,2);

lcd.print("Checking... ");
```

lcd.setCursor() is for placing the text at a particular position on the screen.

NOTE: if you don't see anything on your screen on first run, try turning the potentiometer on the back of the LCD screen to adjust the brightness of the screen.
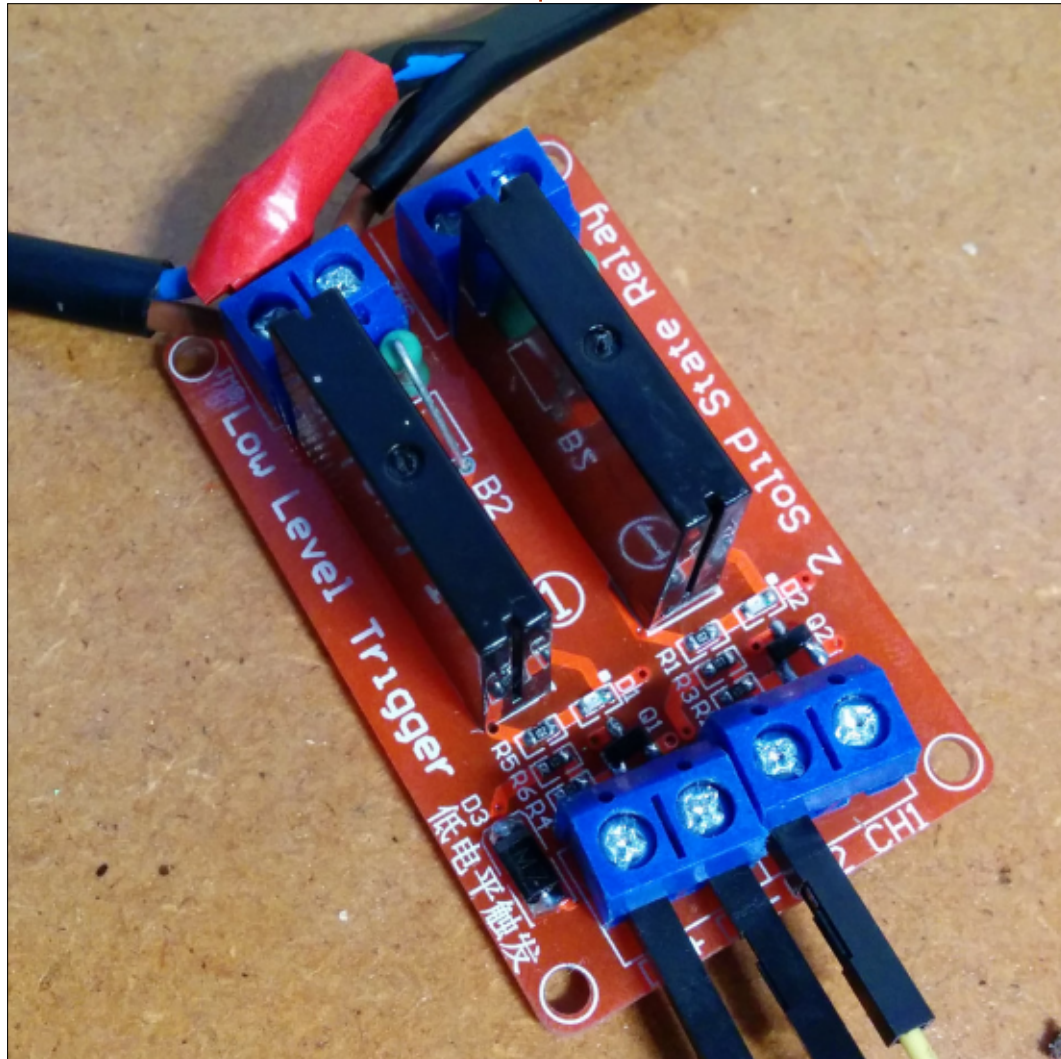
SOLID STATE RELAY

First, you can use a relay switch with this circuit, but you'll get a physical clicking sound as the relay switches back and forth. No big deal, but a solid state relay is digital and without sound. Be absolutely sure that the relay switch (solid state or not) is rated

to do 110/240V.

Before I chop into the cable of my beloved heat mat, I want to make sure this circuit will work. So, instead, I'm going to chop into an unused desk lamp.

With it unplugged, I chopped

through the cable and paired off the two inner cables (live and neutral). I then twisted the two live cables together and left the two neutral cables detached. I also put some electrical tape over that twist to cover it up and keep things somewhat safe.

The solid state relay (SSR for short) has two DC inputs (in the photo, at the bottom left, is the red +5v and blue ground). It also has (depending on your SSR) one or more channels. My SSR has two channels. This is where your Arduino (yellow wire in the photo) will tell the SSR to go HIGH or LOW. The final two inputs are for the cable you wish to open/close. In this case the live (top wires in the photo) from my lamp. Now, it didn't help that my inputs are labelled wrongly. The one marked 'Ch1' is actually for SSR channel two. So if you try this and it doesn't work, try switching your wire from Ch1 to Ch2. You should see a light come on on the SSR to show which channel is HIGH/LOW.

With those things in place it's time to write some code.

To test the SSR I'll flash the lamp on/off. The first new lines of code are:

```
#define ssr1 53

int state = LOW;

unsigned long previousMillis=
0;

const long interval= 1000;
```

This defines pin 53 on the Arduino as the control pin for the SSR. The state is to keep track of whether the lamp is currently on/off. The previousMillis and interval are for the flash. I'm going to try and use millis rather than delays as millis is more preferable to the program halting delay command.

I begin with the usual pinMode and then set the SSR to the initial state of LOW. In other words, off. This SSR is 'low level trigger' which means LOW is on, and HIGH is off.

```
pinMode(ssr1, OUTPUT);

digitalWrite(ssr1, state);
```

In the setup I create:

```
unsigned long currentMillis =
millis();
```

as this will be used in the if/then to switch on/off the lamp.

The only other additional code of note is the switching (top right). CurrentMillis and previousMillis will keep track of how many milliseconds the loop is running

```
if (currentMillis - previousMillis >= interval) {
    previousMillis = currentMillis;
    // flash light
    if (state == LOW) {
      state=HIGH;    // want to turn the SSR off
    } else {
      state=LOW;     // want to turn the SSR on
    }
    digitalWrite(ssr1, state);  // turn the SSR on/off
}
```

for. When it reaches the interval (set up at the start) then the loop ends. Note that this means the code never grinds to a halt like it would with a delay. I can still put stuff to do inside that HIGH/LOW if statement.

The inner if/else is just checking to execute:
• if it's currently off, let's put it on,
• if it's currently on, let's put it off.

And, finally, do the on/off.

With the interval set at 1000 this means the lamp will turn on/off every second.

The only downside of removing the delays is that I'm now hammering ThingSpeak every minute or less, but I'll fix that later. It's working. That's the main thing!

The code for this is a gist on my Github at:
https://gist.github.com/ronnietucker/7fc62df161107116cf93. The code referred to above is marked
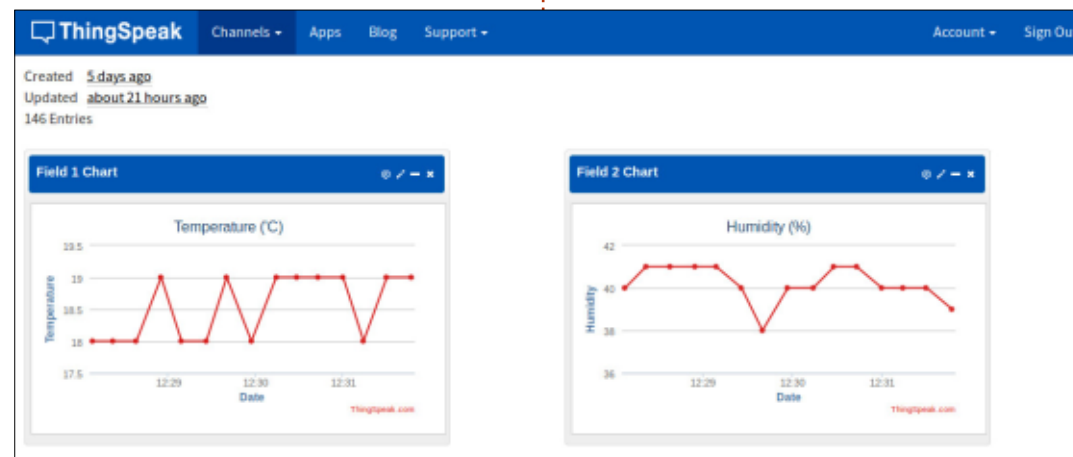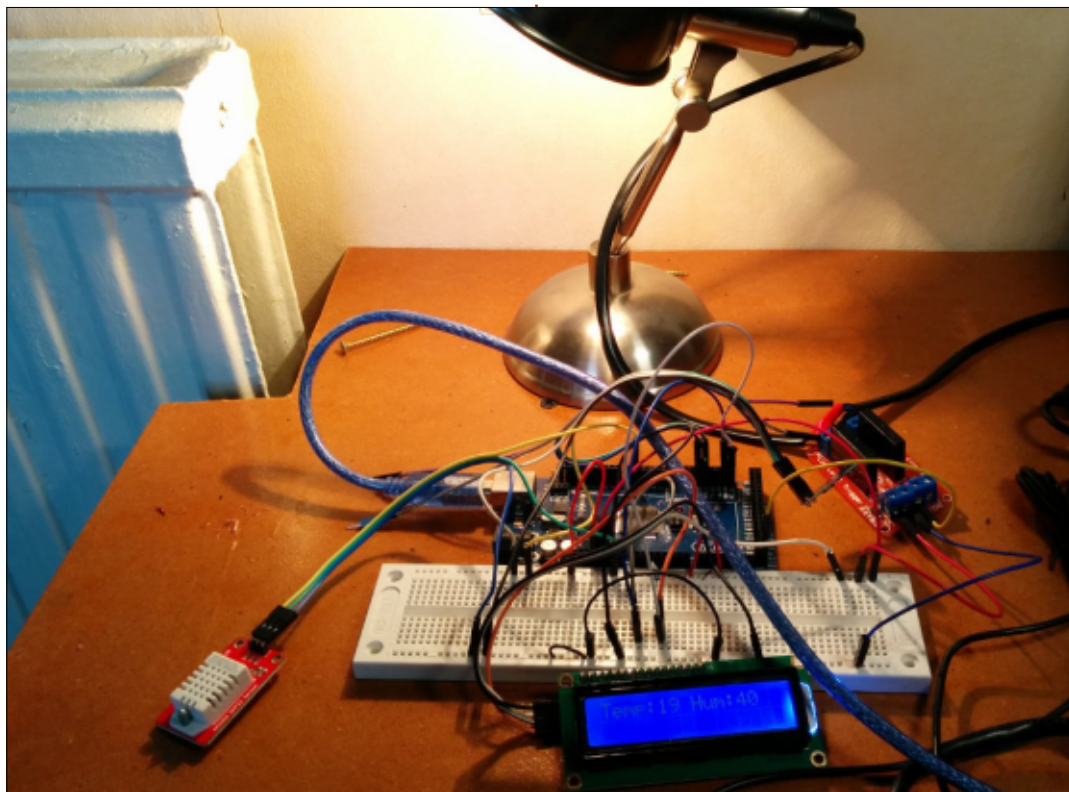
as revision number 4.

To check for a specific temperature, it's as simple as putting in two if statements.

```
if (DHT.temperature > 24 ) {
state=LOW;
}

if (DHT.temperature < 24) {
state=HIGH;
}
```

If the temperature goes above 24, the lamp comes on. If the temperature falls below 24, the lamp goes off.

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.

- -

Mike K. - qualified electrical engineer (true!); and who has an uncle who was electrocuted notes:

Ronnie says: listen to this man. He knows FAR more than me!

Timer-overflow: In the real world, you'd might need to find out the maximum value of that "currentMillis" variable, because, sooner-or-later, it'll reset back to 0. It's defined here as an "Unsigned Long", but I don't know how many bytes is assigned (2/4/8/?), and I don't know if each unit in that var is an actual millisec, or a multiple of one (eg, each unit might be a microsec?).

You may need to allow for that event in your code, otherwise, your code will stop working when it occurs. For example. if a "long" is 4 bytes, and if the time unit is actual millisecs, then this code will fail after 49 days operation (at most!).

Temperature triggers: In the real world, your temp sensor might be "noisy", and, in your case, when the temp is "around 24", the reading might be hovering around 23/24/25, but changing a lot. If that happens, the code might get into a frenzy - switching the device on/off rapidly. To avoid that, one would use two thresholds just outside this frenzy-range, maybe 22-degrees, and 26-degrees. Then, the general approach would be:
• In the initialisation code,set the Device to OFF, and set the "Device_State" to OFF (for simplicity).
• If the new temperature is at/below the lower value, and the Device_State is currently OFF, then set it ON, and set Device_State to ON.
• If the new temperature is at/above the upper value, and the Device_State is currently ON, then set it OFF, and set Device_State to OFF.
• If we can read the current state of the device, then the above "Device_State" variable is not needed. And, reverse the above references to ON/OFF, if appropriate.

Wiring: When cutting the mains cable, strip back the outer insulation very slowly and carefully, and do not damage or cut into the wires which are not used by the relay circuit. If there is an earth cable, and if there is an earthing connector in the SSR circuit, then, do not cut the earth cable, but strip about 1-2 cm of it, fold it, and insert it into the earth connector.

For safety, the SSR circuit must be enclosed in a suitable protective box. If the relays get hot when being used, then ensure there is adequate ventilation in the box, and ensure the box is not placed on soft surfaces such as carpets, duvets, etc - which would inhibit ventilation.

A user might install two relays - one on the positive and one on the negative lines. If a single relay is used - as in this project - it must be connected into the LIVE wire, so that, when the device (lamp, etc) is "off", no mains voltage is reaching the device.

Rating of Relays: Ensure that the relay can easily handle the mains voltage (110V AC or 240V AC), and the maximum current that will be used by the devices. For example, a light bulb might use 1.0 Amp current, whereas an electric room heater might use 20+ Amps.

Overall, we should expect that Arduino fans might have kits in their bedrooms, and might decide to control the mains heating, lighting, ventilation, etc, with the Arduino. In this case, they are messing with lethal devices - which could easily lead to major fires, or injury, or death.

Chromixium OS is no longer available. Per the developing team, Google contacted them and requested a name change. Chromixium was very similar to Google trademarks. It is now called Cub Linux, but it should still carry on as usual.

Now I am an avid fan of Arch-based Linux distros. Therefore, I am greatly enthused by the development of Apricity OS. In fact, in the dictionary, apricity is defined as the warmth of the sun. The development team is in Chicago, IL. This distro has a dedicated blog and forum. The forum is active with quick answers to problems. There will be an online store for Apricity dedicated goods.

Apricity is based on the rolling release of Arch Linux, therefore the OS has the most up-to-date packages and security updates.

The main goal of Apricity is: "*Welcome to Apricity OS, a modern, intuitive operating system for the cloud generation of computing. We believe that an operating system should be easy to use, easily accessible, and beautifully designed.*

We have also incorporated software in order to minimize security risk as much as possible."

I downloaded the iso directly from www.apricityos.com. The iso can be obtained via torrent too. They do not support 32-bit machines. The bootable USB was easy to create and there were no issues in the install. The installer is Calamares. It is a refined product that is on the same level as the Ubuntu Installers.

I have the Gnome desktop installed. Their Gnome desktop is streamlined and responsive. They recently released the Cinnamon desktop. The developers recommend Cinnamon for older hardware or for a more traditional desktop environment. The default backgrounds are wonderful.

The OS uses the Pamac and AUR for software management and updates. It comes installed with great default programs such as PlayonLinux, GIMP, Libreoffice, Caffein, and ICE. Caffeine prevents the computer from going to sleep. ICE is a great option from Peppermint Linux.

CHROME CULT
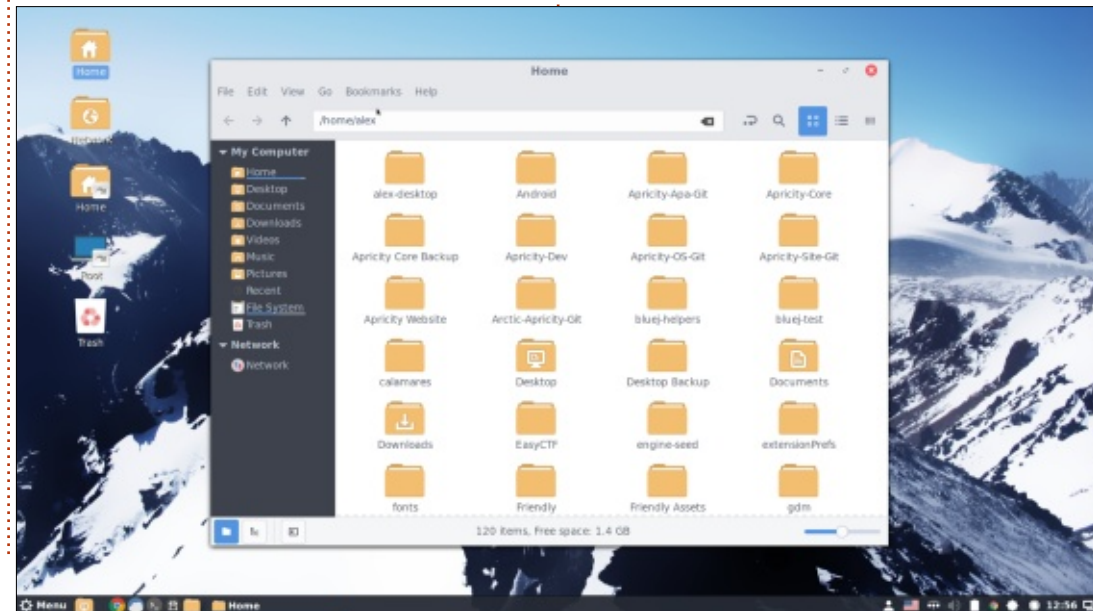
The default browser is Chrome. The Chrome browser has the following extensions installed by default: Ghostery, Great Suspender, and Pushbullet. These extensions have different purposes. Ghostery prevents online tracking. The Great Suspender streamlines the browser efficiency by killing the background processes on inactive Chrome window tabs. Per the Apricity developers Pushbullet, is a tool that lets you send links and small files between your devices, and lets you receive mobile notifications, texts, and phone calls across all your devices.
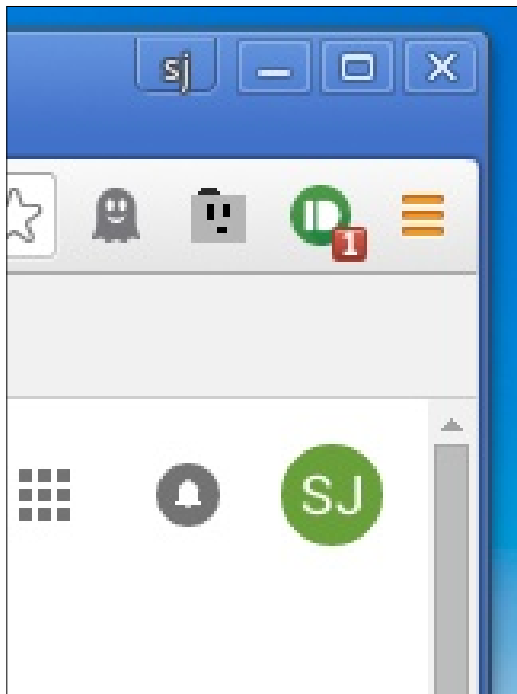
Other features to Apricity are Syncthing and Sbackup. Per the website:
"*Syncthing allows you to share large amounts of data across multiple platforms allowing you to be as connected as possible. Syncthing can transfer large files even faster than cloud-based alternatives. Furthermore, your files and personal information are stored only on local devices, not in the cloud, and are encrypted along every step in the file transfer process.*"

SBackup, is one of the easiest ways to backup your computer. While remaining simple and easy to use, it has many notable features which make it stand apart from the competition. SBackup is able to create both compressed and uncompressed backups, able to split uncompressed backups into multiple parts, supports scheduled backups and manual backups, can back up files to both local and remote destinations, and has a variety of advanced options such as logging, and email notification.

I have not yet used Syncthing and Sbackup on my computer. So I cannot comment on these features. I have used the ICE application to access Full Circle and my Google Drive. I have added these to my Gnome Dock for quick ease. You can see these favorited icons next to Firefox.

I am impressed with Apricity. My only complaint was that a previous update broke Pamac in early February. This was an upstream problem from the developer. Apricity cited that a fix will come soon from the upstream developer. This fix took over 2 weeks to be installed. Antergos and Manjaro fixed the broken Pamac update within 24 hrs. by rolling back to the previous version.

In my humble opinion, the Apricity developers should be proactive in fixing vital programs such as Pamac. Granted there was a bandaid solution offered in the bug forums. This fix was posted only 4 days after the Pamac breakage. However these issues should be addressed quickly if Apricity wants to be a known and respected cloud distro. Despite this criticism, I still like Apricity and the direction it is going.

Next month, I will offer Crouton Ubuntu and the Gallium OS as possible cloud distros directly on the Chromebook hardware.

SJ Webb is a Linux Hobbyist and Research Coordinator. He enjoys fishing, hot rodding, and spending time with his kids and wife. He thanks Mike Ferarri for his mentorship.

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

H ere's where we start to sweat and get shaky hands. Everything previously was done without power. Now we need to apply power to the printer.

On first power-up, the fan above the effector must start up. If it does not then you need to immediately power off. The nozzle should also remain cold. If it heats up, turn off. If everything is good so far then it's time to connect.

Well, almost.

First, you need to put the contents of the zip file (from RepRapPro) onto the (supplied) SD card.

Now, my Fisher is connecting directly to my laptop, not to a router as most people would. This meant some extra networking steps for me.

I had to go into my ethernet connections and delete the existing network connection that was there. My laptop uses wifi for internet. Next, I created a new

ethernet connection, clicked the IPv4 tab and for IP I entered 192.168.2.10 (since my network is normally on 192.168.1.1) and for the subnet I entered 255.255.255.0. Gateway was left blank.

This now puts my laptop's ethernet port on a different network from everything else.

On the SD card, I open the /sys/config.g file and set the IP to 192.168.2.14 and save it.

Now I pop the SD card into the printer, power up the printer, and hook up the ethernet cable from the printer to my laptop.

In a terminal I enter:

```
ping 192.168.2.10
```

Getting a reply to that means the Duet board is up and running. Doing:

```
ping 192.168.2.14
```

and getting a reply means that the printer server is up and running.

Putting 192.168.2.14 into a browser address bar gets me the printer's server page (below). I'm in!

TESTING, TESTING...

Clicking the link on the left that says 'G-Code Console' will display a text entry box that you can think of as being like a terminal to issue commands. Entering:
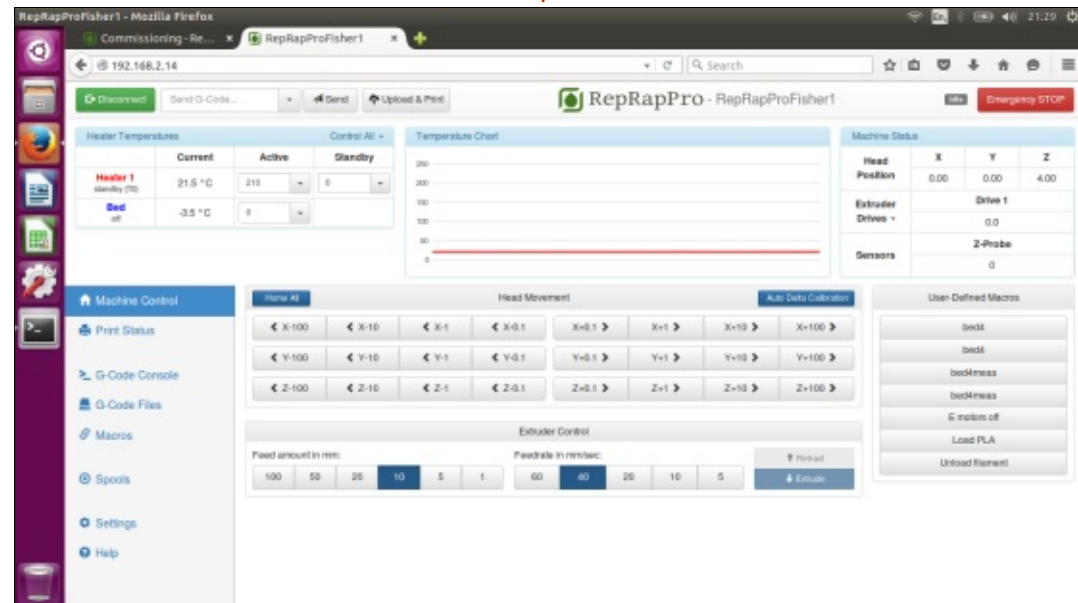
```
G1 S2 X10 F500
```

Will move the X carriage up by

10mm.

After several other commands to check each axis, it's time to go... home.

Going back to the previous page, we click the 'Home' button and all three carriages should move up to touch the microswitches, move down and up a bit, and then return to leave the nozzle a few millimeters above the bed.

Now the scary bit.

Hot End

It's time to switch on the heater and see if we can extrude melted filament. This can be done using the 'Heater' link at the top left. When it reaches the temperature on the dropdown menu (190, maybe 200+), we can try using the manual winder on the extruder to get some melted plastic.

In my case it worked!

I did have to tinker with the screw to give the extruder more grip.

Printing

To actually print something, you load a .g file from the SD card. To do this, click the 'G-code File' link on the left of the screen, and this lists all .g files that are on the card. Click the .g file and you'll be asked if you really want to print this. Click OK and wait.

The RepRapPro zip file (from earlier) usually has a little MakerBot file on there to get you started.

Conclusion

From several weeks of printing random things for practice, I have noticed a glitch in my setup. Ideally when you do HOME and ALIGN (in the main tab), the nozzle should be less than a millimeter away from the bed. Mine is about 10mm away. I end up having to manually move the nozzle down. Not a big deal, and it's something I'll look into later. So make sure your nozzle is right near the bed before printing.

The printer is quick at whooshing back and forth to print plastic, but expect some prints (even small ones) to take quite a while. The little figure of the Kerbal Space Program character is only 12cm high and took just over one hour to print. Almost half of that time was me holding the base in place as I could see the nozzle was going to end up knocking it over.

Most will print without being held, but there'd be nothing worse than seeing a one hour print ruined by it tipping over. By rights, it shouldn't happen, but we know that it inevitably will.

I have to say. This was money well spent. The instructions are

LINUX LAB
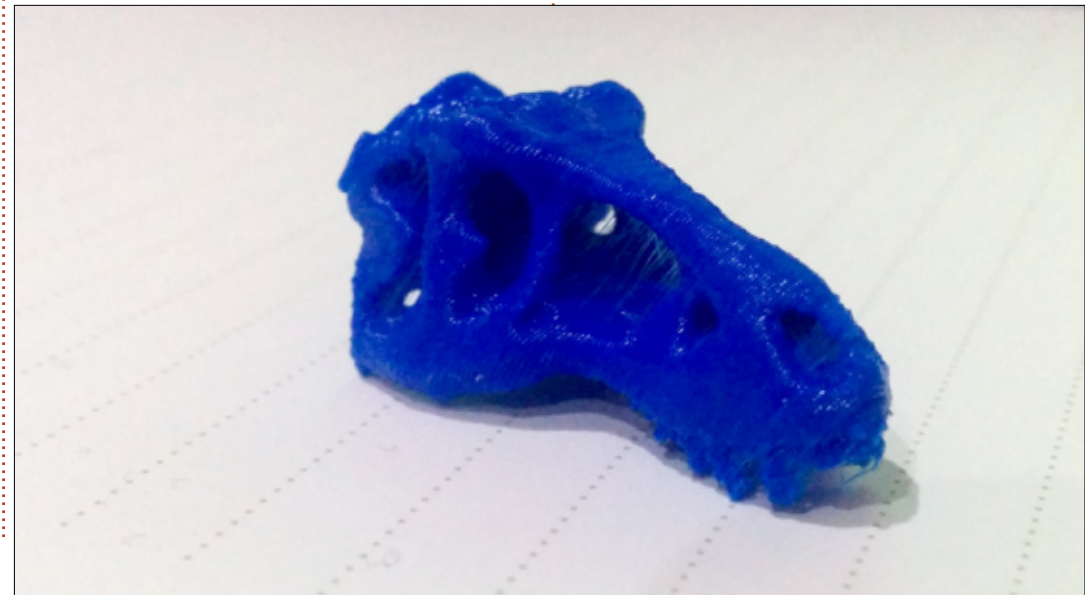
excellent – with plenty of photos for every step of the procedure. Honestly, if you mess this up, you've no one to blame but yourself for not studying the photos, part lists, etc.

There are a ton of good things out there to print, but you need to convert them from STL format to the RepRapPro G format. How's that done? I'm glad you asked.

CONVERTING

Most models out there are in STL format which is no good for the RepRapPro Fisher. Thankfully we can use Slic3r (http://slic3r.org) to convert to the native G format that we need. Click DOWNLOAD

and then, on the download page, click Linux. Follow the instructions there and you'll have Slic3r almost ready to go.

Next, following the instructions at: https://reprappro.com/documentation/commissioning-introduction/printing-duet/#Slic3r_profiles will mean that Slic3r will use the Fisher by default with all the settings done for us.

Now, grab a model from Thingiverse.com. That's my favourite place for models. Save the file somewhere you can remember.

Load up Slic3r and at the top

left, click 'Add…' and point it to the STL file

Slic3r will show you a 3D model of what we're going to print. As long as the Fisher settings have been imported, we don't need to touch the tabs at the top. See how it says 'Fisher' beside 'Printer' at the top right? That's good.

Normally the model is shown in yellow. Click it, it'll turn green, and display the info shown at the bottom right of the screen. Size is in millimeters, so this Lego block will be 31.8mm by 15.8mm by 11.4mm high. That's about all we need to know. If I wanted this as a jumbo Lego, then I could click 'Scale' at the top and increase it to
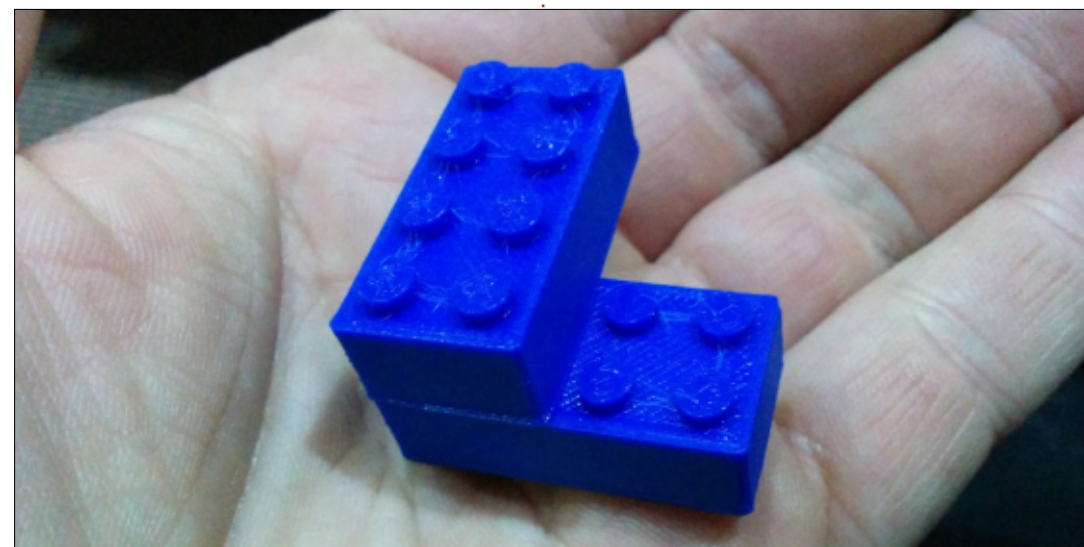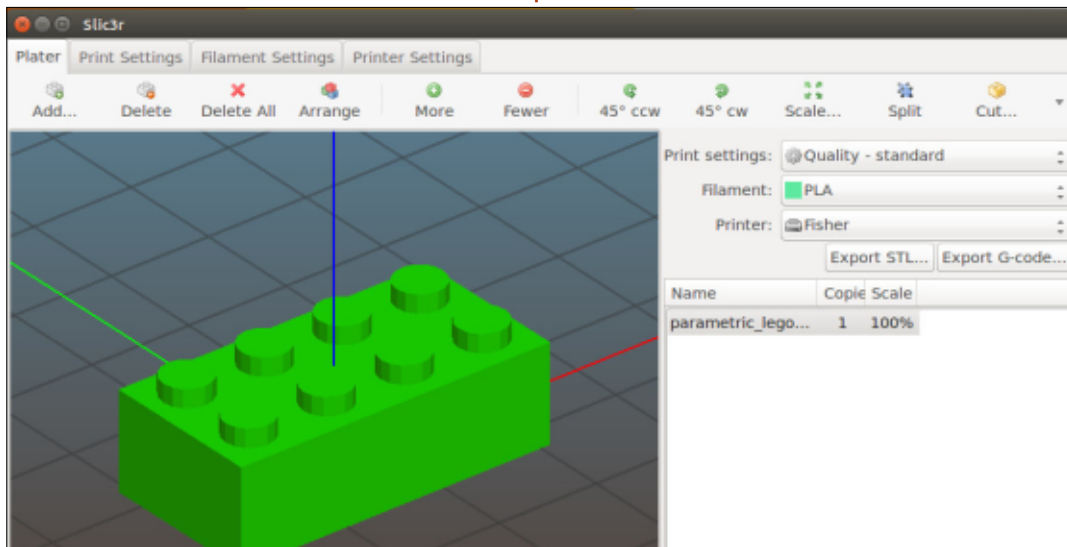
200%, or shrink it to 50% to make a mini-Lego.

The last thing to do is click the 'Export G-code' button and save the .g file somewhere on the SD card.

Pop the SD card into the printer, power it up, network into it, and get printing!

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.
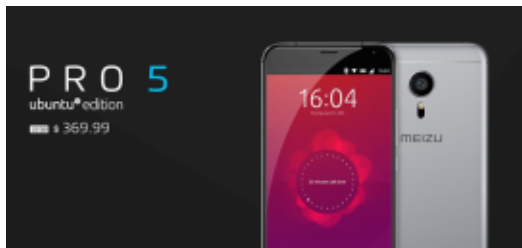
Meizu PRO 5 Ubuntu Edition available for pre-order now!

The most powerful Ubuntu phone was just announced last week, Meizu PRO 5 Ubuntu edition is now available to pre-order from Meizu retailing at USD$369.99 - http://www.meizu.com/en/products/pro5ubuntu/summary.html

The Meizu PRO 5 Ubuntu Edition was showcased for the very first time at Mobile World Congress 2016 in Barcelona. The device being demonstrated at MWC was running a demo version of Ubuntu touch, which will be updated with improved features through the next OTA update.

When the device goes on sale, it will be exclusively available to buy through Meizu and Canonical's joint partner JD.com's global e-commerce platform en.jd.com. JD.com is China's largest online direct e-retailer, offering a seamless online shopping experience for customers around the world. Meizu PRO 5 Ubuntu Edition will be shipped worldwide on its global site.

The Meizu PRO 5 Ubuntu Edition is the fifth Ubuntu device and the second joint launch between Canonical and Meizu in the last twelve months. As VP at Meizu, Li Nan, puts it "We have some interesting plans and so we're looking forward to further deepening our partnership with Canonical in the future." We look forward to an exciting future together.

Community collaborates on new convergent Ubuntu phones

We're celebrating the arrival of two new Ubuntu phone community ports! If you're the lucky owner of a Sony Xperia Z1, you will soon be able to download an image to turn it into an Ubuntu phone and PC! Alternatively, if you're the owner of a OnePlus One, you will be able to flash and use Ubuntu on your phone.

We owe a big thanks to the Ubuntu community for porting Ubuntu to these devices, and especially Marius Gripsgard for his work on the OnePlus One! More widely, this is a meeting of the minds with phone vendors who have been pushing open source in their devices and their developer community for a while. The Sony open source efforts in particular have been essential to this work. Similarly, OnePlus' commitment to open source has helped tremendously towards this port becoming available.

But with any porting work, we know this is just a beginning. We would like to help make sure that these ports are finalised, remain up to date, and bugs are addressed. For this reason, we will be collaborating with OnePlus and participating and contributing to Sony Open Devices community to make sure that the open source community continuously gets all the information, code and tools they need to keep this exciting project going!

Get ready to flash!

Here's a link to the ports https://developer.ubuntu.com/en/start/ubuntu-for-devices/devices

I previously reviewed **Able2Extract** way back in FCM#95. Since then several new things have been added, so it's time to give Able2Extract another whirl, this time version 10, and try out the new features.

CONVERSION FROM IMAGE

The main thing that Able2Extract (A2E) does is convert image text to either a document, a spreadsheet, or a presentation. It does this by saving to either Microsoft Office, or LibreOffice formats. It says OpenOffice in the toolbar – but the formats are the same for LO and OoO.

The first thing I tried was to take a PDF that was comprised of images and test the OCR functionality within A2E.

While the conversion wasn't perfect, I purposely chose this text because it was flawed. It was highly compressed, and not your usual serif style font. I wanted to see how A2E would cope with it. And it coped extremely well, considering...

It even managed to recognise, and keep, the italic words, and where words were hyphenated at the end of a line.

TEXT FROM A PDF

In most cases you can copy/paste from a PDF. But with some PDF's (unfortunately, FCM being one of them), you get less than perfect results.

For this test, I converted a page from FCM to a PNG image. With the PDF opened I did a copy/paste of the first column of text.

Next, I opened the PNG page image into A2E and selected only the first column to convert it to LibreOffice Write (ie: text). In the image, you can see, on the left, the

BITMAP IMAGE OF PDF - ABLE2EXTRACT OCR - COPY/PASTE FROM PDF
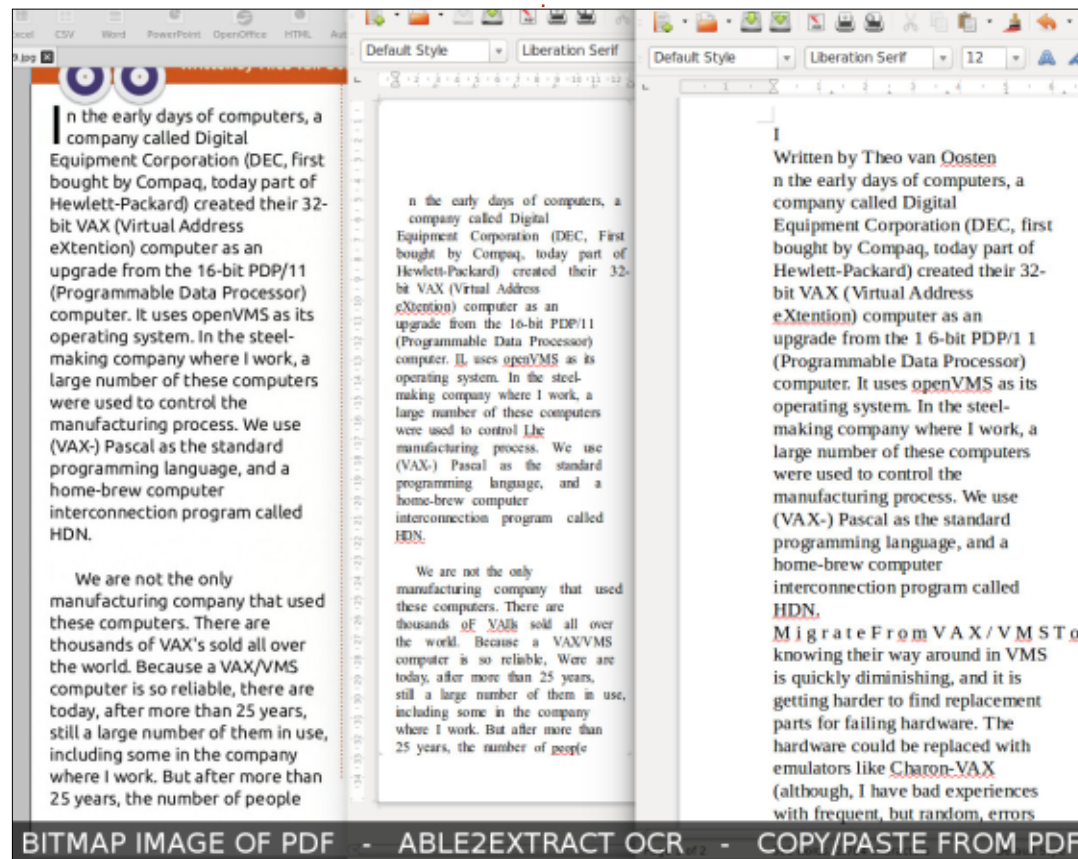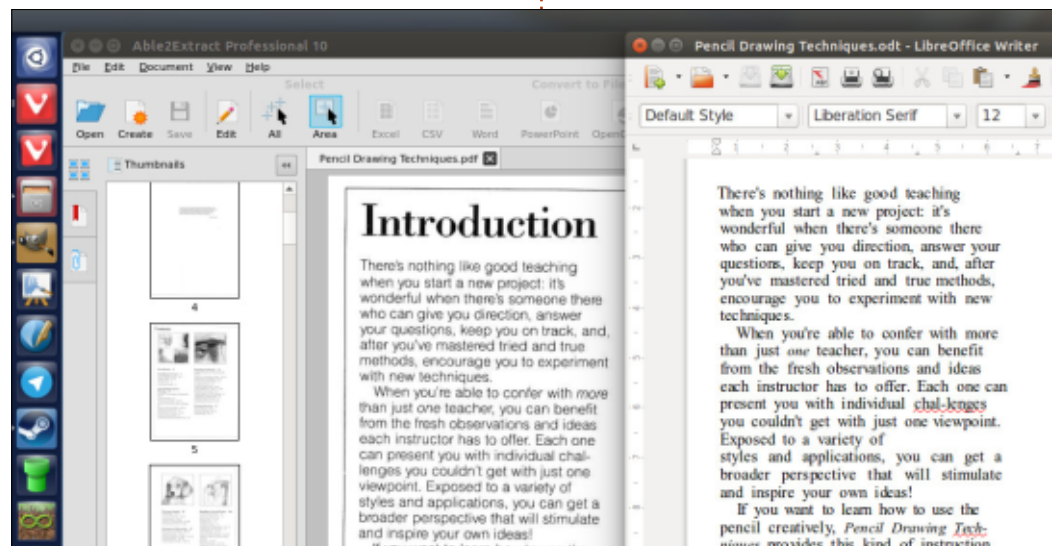
PNG image of the column. In the middle you see the A2E conversion. And, on the right, you see the copy/paste from the PDF to LibreOffice.

The clear winner is the A2E conversion. Again, not perfect, but certainly more usable than the copy/paste.

I should also mention that the OCR conversion from image to text is extremely quick with very little waiting time.

EDITING PDF FILES

A new feature of A2E is being able to edit PDF files. Either by adding/deleting pages, rearranging pages, but also by live editing the text.

Here, I took a PDF that was just text, loaded it into A2E and tried editing the text. Worked like a charm!

In the example image shown below, I have the original PDF loaded in A2E (top) and the edited PDF (bottom). As you can see the edited version has kept the same font, and the edit has nicely blended in the changes.

Clicking the Edit button at the top of the window brings this new feature into play. It's within the edit section; you can add/remove/move pages within a PDF.

OTHER FEATURES

I won't bore you with the details of the other features again. If you'd like to read about them then feel free to check FCM#95 for a full run-down of the other features.

But, in short, A2E can do conversion to text by either a full page, full document, or a selected area. Like I mentioned earlier, it can convert to MS Office, LibreOffice, but it can also convert to HTML, AutoCAD, and various image file formats.

CONS

The only one minor downside I have to report is that a couple of the keyboard shortcuts didn't work. Especially CTRL++ (to Zoom In). CTRL-- does work fine, however.

Live editing of PDFs within A2E is a bit slow, but it does work very well. It's not 100% though, as it couldn't edit text within any of the FCM PDF files. I'm not sure if that's a flaw in A2E or in the PDFs exported by Scribus.

CONCLUSION

In short: if you have images/PDFs that need to be converted, then A2E is definitely something I'd recommend. Certainly quicker and probably more reliable than manually typing

them in – if you require text. But, only if the images you have are high quality scans. Low quality scans will give errors in the resulting text, but that's not really a flaw within A2E. OCR software has always required high resolution scans.

ABLE2EXTRACT:
http://www.investintech.com/prod_featuresList_a2epro.htm
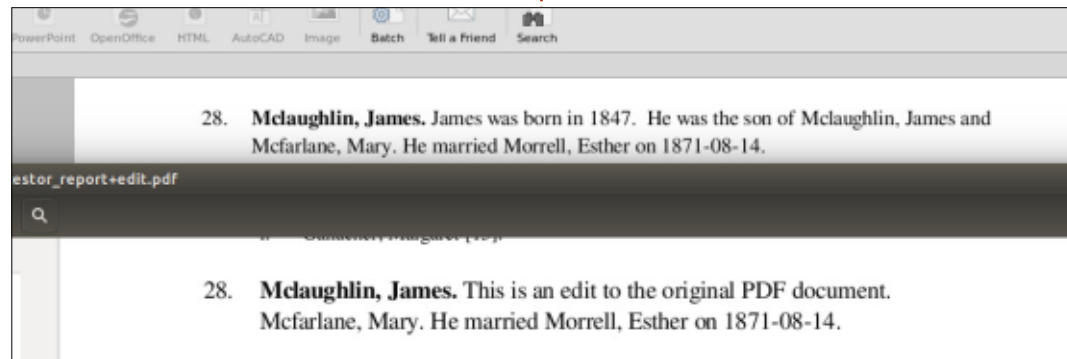
Ronnie is the founder and (still!) editor of Full Circle. He's a part-time arts and crafts sort of guy, and now an Arduino tinkerer.

Ever hear of 'deepin OS'?

Join the crowd. In the space of roughly 12 years, it has been known as Hiwix, Hiweed (I can hear the comments already), Linux Deepin, and possibly a couple more names I missed. Give it time and they'll probably rename it again (the DVD I used for this review was titled Linux Deepin but once it loaded and updated, it was 'deepin OS').

Oh, did I fail to mention the dizzying array of desktop environments? Take your choice of Gnome, IceWM, Xfce, LXDE, Unity, and who knows what else in between (the current one is deepin DE which appears to be a modified LXDE variant but your guess is as good as mine).

In fact, when I started writing this review in 2015, it was still Ubuntu-based, and then suddenly switched to Debian with the release of version 15 in late 2015 (the current 15.1 was released in January, 2016).

Wanna lay odds it changes later this year?

And then there's the website, https://www.deepin.org/?language=en, in which it appears that English is a second language.

Well, guess what -- it is... deepin OS is from China.

You have your choice of 32 or 64-bit variants, with either one running just a tick shy of 2GB, which places this is the middle of the pack size-wise. For version 15.1, you get the newest deepin DE (Desktop Environment if you haven't caught the inference yet), plus an obviously Ubuntu inspired but Debian-based chassis.

Take my advice and pass on the website's download offerings. Slow is the rule there and you'd find it quicker to hop a flight to China for the DVD. Just visit sourceforge.net or do like I did and bum a DVD.

Now, here is where a tick shows up early on. Care to test-drive it first just like Ubuntu? Well, call somebody who cares – because no test-driving is allowed. You either install it or you don't.

Hmmm, guess that's what virtual machine setups are for!

As you might expect, setup is something of a snap, although it's not Ubiquity-based like Ubuntu. If you forget to use a virtual machine, or partition your hard drive ahead of time, you'll have to sacrifice you whole HDD.

Otherwise, it's the usual 20 to 30 minute affair with another hour for the obligatory 400MB or so of updates. Take my advice and use a flash or USB drive, and that installation time will be cut in half.

No hardware issues of note, which is odd given I'm pushing 10 year old hardware with just 3GB RAM. Invariably something goes wrong with wi-fi, video or audio, but not this time. All of it worked just fine without having to add

drivers.

Cold boots on my 10 year old Dell averaging roughly 20 seconds; not bad given the outdated specs, and I will admit I love the opening music that sounds like a Oriental beat-box mix.

The desktop design is pleasant enough; a Docky-like bar at the bottom and the usual offerings elsewhere, but there is a twist. Right-click the dock and users have a choice of two docks similar to OS X (Efficient and Fashion), along with Classic (stretches to fill the bottom screen). Any choice is cosmetic and does nothing to hinder hardware usage.

Although it may look like Gnome, the desktop really isn't. It's a morph that deepin created, and there are differences. Much like previous Apple variations, clicking on the left side rocket ship brings up the program menu, while clicking in the system information (battery, time, etc) panel brings up the Control Panel.

And each of these functions is unique.

For example, the rocket ship

brings up a menu of program icons that looks Unity-like but with a unique design feature -- right-clicking allows for adding the icon to a desktop or panel.

Big whoop, you say? I would agree except for the final function -- uninstall.

Yep, you uninstall by right-clicking. No going through

Synaptic or the Software Center. You just right-click, hit uninstall, and watch the program and icon disappear.

Neat, eh?

Built in programs range from the useful to the downright bizarre. For example:

Forget LibreOffice and

substitute WPS Office. Not a bad substitution, but odd. I know of no other OS that currently pegs this program for its office functions. Of note, the one included with even the newest deepin version is outdated and you'll be forced to update quickly.

Looking for Firefox? Nope, you get Chrome (not Chromium) complete with codecs installed. Adobe Flash is already there, too. Yeah! I can watch videos without getting the black screen of death.

Instead of the usual multimedia offerings, you get deepin Movie, deepin Music, deepin Terminal, deepin Software Center, deepin Media Player, and probably a couple others I missed. What's the main difference between these and the usual Ubuntu or Debian choices? Mainly the name from what I could see.

Like unusual, exotic games? Then install deepin Game (yes, they forgot the "s"). Here you'll find an odd mix of familiar games with Chinese to English subtitles. Unfortunately, this doesn't always come out as intended, especially the running game (seemingly based on our single player run-and-

miss-obstacles fare) called "Crazy Flasher". I do believe that's what they used to call that weirdo hanging out downtown.

Of course, users can add what they like outside of the deepin Software Center by merely adding Synaptic, and the adventurous can always use apt-get commands via the terminal (excuse me, deepin Terminal).

As for general desktop speed and usage, I found no problems. Again, I'm pushing grossly outdated hardware, so anytime an OS works smoothly, I'm happy.

And there's no arguing that the overall design is top-notch. I can think of no other OS that has such a nice, functional interface. Windows and OS X could take a note or two here.

However, there is a kink in other aspects of deepin OS that can make it dicey, and that relates with program downloads.

In short, they range from blazing fast to dead, desiccated slug-on-the-sidewalk slow. On several occasions, server speeds went as low as 5 KB/s. Notice that
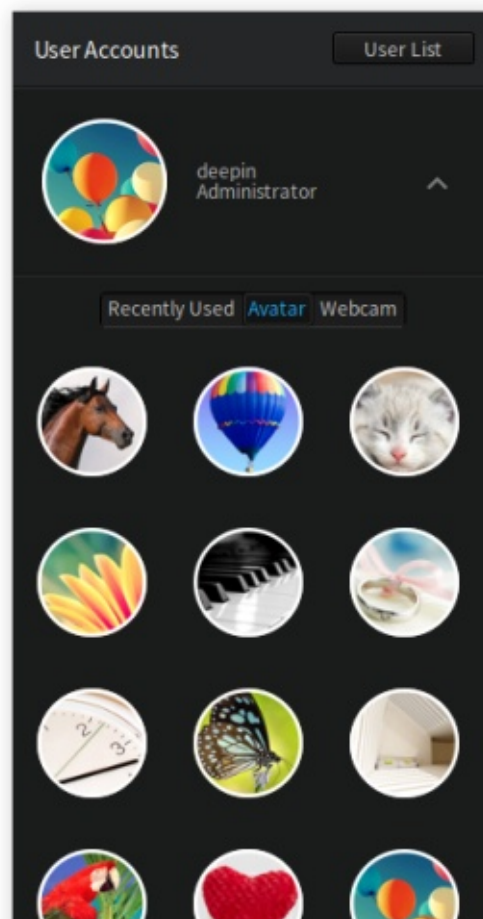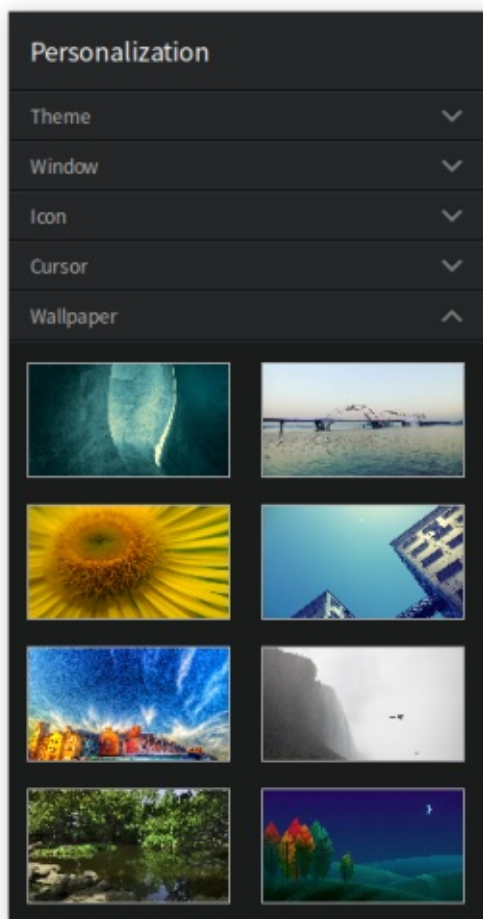
figure is not MB, and several sub-50MB downloads took an hour or more.

And before you mention it, I did change servers -- several times. There is just no telling whether or not a program will download quickly.

So, is it worth giving deepin OS a try?

All depends upon your level of patience. You do get a creative desktop that works quite well (it never did lock up or fail me), but you may have to live with occasional download speeds that come from the dial-up modem days.

Once you do get it all set up and ready to go, it is a force to be reckoned with, though.

LETTERS

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.

Q How do I Play DVD Movies in Ubuntu 14.04 LTS.

A See this: https://help.ubuntu.com/community/RestrictedFormats/PlayingDVDs

Q I'm doing a Something Else installation on an old netbook. The root, swap and /home partitions are all sorted. But where does the bootloader go? This is what's available:
```
/dev/sda
/dev/sda1 Windows Recovery
Environment (loader) (NTFS)
(which XP sees as PQSERVICE)
/dev/sda2 WinXP (NTFS)
/dev/sda6 / root (ext4)
```

A /dev/sda

Q I have ipcop on a flash drive, how do I install it? I have just installed Xubuntu 15.10 on my desktop.

A (Thanks to *pauljw* in the Ubuntu Forums) DO NOT attempt to install IPCop on your Xubuntu machine. IPCop is a standalone firewall/gateway server and will wipe a hdd clean prior to install.

Q I am trying to install Ubuntu Server 14.04.3 LTS on a Lenovo ThinkServer RD650. I have already configured the server RAID controller to RAID5 and all eight disks are within this. When I boot from the Ubuntu disc, I see dev/sda and dev/sdb. Do I need to use "Software Array configuration"?

A (Thanks to *darkod* in the Ubuntu Forums) If you use the server raid controller card and configure one or more RAID5 arrays, then they are presented to the OS as simple disks. You do NOT RAID them further with software RAID.

What you have is two RAID arrays of four drives each; if you want a single array, you need to set that up from the controller card. With eight drives, you should consider RAID6, to protect against a second drive failing before you can replace a single failed drive.

Q I need to do a project on a Raspberry Pi. Can I use my laptop as a monitor for the Pi?

A No. To set up the Pi, connect a keyboard, mouse and monitor, and set it up to host a remote desktop session. You can test the setup from your laptop. Once that is done, you can remove all the peripherals.

Q Would anyone know, does the Raspberry Pi 3 support Ubuntu 14.04 or perhaps 16.04 when it becomes available?

A (Thanks to *QIII* in the Ubuntu Forums) Unless someone creates an image, I don't think it's available. It's not in NOOBS, and I don't see it in the third party downloads on raspberrypi.org.

I think that Unity would be beyond the capability of the graphics module anyway.

Q (Asked in the Ubuntu Forums) What are some other good Ubuntu forums that offer assistance? My question seems to be out of the scope of those here, so I was curious of others I might be able to post on and get some feedback/help.

A You can try askubuntu, but note that the protocol is different there: it is expected that you have checked that your question has not been asked before. As well, you might try refining your Google searching skills. I find that I can often get a useful result by specifying the best four words. (And once, just once, I did a Google search which provided exactly one response, and it was the answer I wanted!).

TOP QUESTIONS AT ASKUBUNTU

* What does this "printf" command

do?
http://goo.gl/yqbhjZ

* Is my system able to run 64-bit OS?
http://goo.gl/mKiVt6

* Is there a way to create multiple directories at once with mkdir?
http://goo.gl/WQ6Jk3

* Is swap area required? Can we install Ubuntu without a swap area?
http://goo.gl/IQ5qHo

* Ubuntu settings after adding additional RAM
http://goo.gl/x5A4sv

* How to find ONLY duplicate files that have different names?
http://goo.gl/2pnSeG

* Is it possible to recover data stored on hard drive with windows with Ubuntu live?
http://goo.gl/qTlM1D

* My USB flash drive became completely bricked
http://goo.gl/52kKUG

* Why do Windows & Ubuntu have similar shell commands?
http://goo.gl/Q7MlEa

TIPS AND TECHNIQUES

Six years in the saddle

This issue marks six years that I've been preparing the Q&A column, so I am pondering what has changed and what is the same. I'm a bit surprised by how little has changed.

I'm still using the same primary computer. I upgraded the memory from 4 GB to 16 GB because I wanted to become familiar with Windows Server for work; it has made no difference in Linux. A blu-ray player fell into my lap, so I installed it, but I have never had a blu-ray disc. The biggest and most expensive upgrade was a 24-inch Dell Ultrasharp monitor, well worth it. (Other brands also deserve consideration, most notably HP.)

I also added a single-port USB 3 card, which now works nicely. (See last month's column.)

The AMD Phenom II dual-core CPU at 3.1 GHz, and the 640 GB Western Digital Black hard drive, still meet my needs. If I spent two grand, I could probably build a system which would be three or four times as fast, but that wouldn't make much difference to me. When I power up, I usually go and make myself a coffee, so the faster boot time would be wasted. I do a little video editing, so the speed would help there, but it's not crucial. Perhaps my lack of desire for a new system illustrates the sad state of the PC industry.

On the peripheral front, I've gone from a "competent" feature phone to a Moto G smartphone, a huge upgrade.

Software has progressed slowly. When I set up the computer, it was dual-booting the trial version of Windows 7 and Ubuntu 9.04. Now, I have Linux Mint 13 (based on Ubuntu 12.04) and Xubuntu 15.10. In between, there were a lot of versions which were not stable on my hardware, including 14.04. I hope that I will be able to upgrade to 16.04 and stay there for a while. I know a lot of work has gone into the kernel and the distros, but the only improvement which has helped me is being able to plug in my phone and see it as a flash drive.

When Ubuntu adopted the Unity interface, I decided it was not for me. I've used Mint or Xubuntu to get stuff done since then. (I still test an average of one distro/version a month.) I have heard Xubuntu described as "lightweight," but I haven't noticed any missing functionality. It takes a bit more effort to set up network shares, but I have also managed to do some pretty complicated stuff on that front.

In the near future, my most likely upgrade will be a faster Internet connection. For computer hardware and software, I can't think of a single thing which I would describe as, "I want to buy..."

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.

PRO 5
ubuntu® edition
`32GB` $ 369.99

16:04
Thursday, April 21, 2016

33 minutes talk time

MEIZU

14:10
lunes, 2 de marzo de 2015

1 video grabado hoy

BQ AQUARIS E4.5 & E5HD
Life at your fingertips

Ubuntu reinvents the way you interact with your smartphone.
Everything you need in your day is now at your fingertips.

AVAILABLE WORLDWIDE

About a year ago, I signed a petition requesting that both Saints Row the Third and Saints Row IV be ported to Linux. At the time of this writing, it seems like work is currently underway to port Saints Row the Third over to Linux, but for now we can all rejoice and enjoy Saints Row IV which has recently been ported to Linux. Originally developed by Volition and published by Deep Silver back in August 2013 for Microsoft Windows, PlayStation 3 & Xbox 360, it was eventually ported to other systems such as PlayStation 4, Xbox One, as well as Linux. Virtual Programming is responsible for porting the game to Linux in December of 2015, and after playing it for the last two months without a single crash, I can honestly say they have done a superb job.

Saints Row IV is an open-world, action, comedy, co-op game in which almost anything is possible. Currently the game is available for the low price of $14.99 through Steam & Humble Bundle Store. There are also a number of DLC add-ons available, most of which are included in the Saints Row IV: Game of the Century Edition currently selling for $19.99 which is a good price since it includes Saints Row IV (the game itself), Saints Row IV – Enter the Dominatrix, Anime Pack, Bling, Bling Pack, College Daze Pack, GAT V Pack, SR4: Volition Comics Pack and much more. These add-ons are also available individually for about $1.99 each.

For anyone who's ever played any prior games from the Saints Row series, or perhaps any of the games from the Grand Theft Auto series, Saints Row IV will feel right at home. Saints Row IV begins with you (the player) saving the world from a possible nuclear holocaust which then leads to you being elected as president – while the rest of your Saints Row crew become your close advisors. However, things go south almost immediately when, on your way to give a press conference, the world is invaded by aliens from the Zin Empire led by their leader Zinyak. After a very brief battle against

the Zin invaders, you wake up in an absurd alternate reality in a 1950's sitcom (complete with title sequence and recorded audience laughter). After cruising around town for a bit, you're contacted by Kinzie (a Saints computer expert) who informs you that you're trapped in a sort of computer simulation created by the Zin aliens. It now becomes evident that you must try to break free from this simulation. From this point on, the game takes place inside the simulation of fictional Steelport, but you eventually break free from this simulation with help from Kinzie and vice president Keith David by stealing a Zin spaceship which allows you to enter/exit the simulation as needed. From this point on, you go back and forth, in & out of the simulation, as you try to free the rest of the Saints who are trapped in their own simulations (a product of their nightmares), while also trying to weaken Zinyak's simulation so you can fight him back in the real physical world.

The game doesn't just take a left turn into wonderland; it leaps, time-warps, explodes, and even implodes into places never before seen in a video game. Its dark humor and flat-out absurd game-play elements make it a most entertaining and unique experience. While in previous Saints Row games you were expected to hijack cars in order to get around the city (much like GTA), in SR4 all of that goes out the window as you gain some awesome superpowers that help you sprint faster than any car available, jump higher than humanly possible, and, while in the air, gliding makes you feel like you are flying, not to mention the super-punches that you're able to dispense on your enemies. In addition, you're allowed to use most alien weapons that the aliens bring with them, some of which zap, abduct, or, in the case of the Rectifier Probe, do unimaginable things to aliens and humans alike.

As an open-world game, you're allowed to explore Steelport as much as you want. You can walk, run, sprint, jump, fly, drive a vehicle, fly an alien ship, the possibilities are almost endless. Your character can be fully customized, so much so that if you manage to log on the community site (which has been problematic for all users regardless of platform), you'll be able to find characters that look just like movie stars, movie characters, historical figures, other video game characters, etc. Some examples of characters I've seen created on the community website are Kobe Bryant, Lara Croft, Elvis Presley and more. At the beginning of the game you get to customize your character but if you want a different character later on, you don't need to worry as that can be arranged. Modifications can be made not only by buying different clothes and getting different hairstyles, but also via plastic surgery, tattoos, piercings and you can even get a surgery to change your gender. In order to do all of these things, you have to first hack the various fashion stores, gun stores, auto mechanic shops and surgical doctors, so you can then have access to them and buy your upgrades as needed.

One feature I really like about this game is that you can get the game to auto-detect your hardware and adjust graphics settings according to what the game thinks will give you the best performance, while still managing to look as good as possible. You can also choose Low, Medium & High graphics settings as well as Custom – where you can manually tweak individual items like shadows, anti-aliasing, and other video options. I opted to go for the automatic-detect but I later changed a couple of the individual settings to get slightly better graphics without taking a performance hit. So far, I've only had a few color glitches on some characters but these glitches were very brief and didn't affect the actual performance of my game. The input game-controls are very intuitive, especially since I'm familiar with the Saints Row game series as well as the Grand Theft Auto game series.

While playing Saints Row IV, I mostly used an Xbox 360 game controller but I've also played it a few times with keyboard/mouse; both input methods have been flawless. The graphics throughout the game have been pretty impressive, especially because it switches from an open-world third-person perspective to a driving game and has some cinematic cut scenes throughout the main story. The voice acting is funny and very well done; my only complaint is that some of the NPCs do tend to repeat a few of the same lines

often – depending on what neighborhood you happen to be in. The music is another plus because you have the option to have different radio stations playing while driving cars – ranging from classical, rock, hip-hop, talk radio, reggae and more. You can change the radio stations on-the-fly while driving a car, which is a great feature, and it makes it so that you can literally play for many hours before having a song repeat itself. I was pleasantly surprised to find out that the Linux version of Saints Row IV supports game mods.

There is a strong and thriving modding community for this and other games in the Saints Row series so for those interested in modding, all you've got to do is find one of the many Saints Row modding sites on the net, and download the scripts required. I tried adding a couple of mods, and, with a little help from my friend, Google, I was able to successfully get them to run in my game. There is no performance downside to adding most of the mods available. Unfortunately, logging in to the online community has been problematic for Linux gamers and everyone else regardless of platform. They are supposedly

working to correct the problem but for now it makes me take a half star off the full five-star rating.

I most strongly recommend getting Saints Row IV as it is a very fun game to play. You can very easily get hooked on it or you can just as easily make it a casual game that you play only from time to time. As for me, I've embraced it and I'm on my way to complete it with 100% progress. The game IS rated M for Mature and very deserving of its title due to its language, violence and sexual content, so this is something to take into account before buying it. But, other than that, it's a must-get title that you should add to your

game library.

★★★★½

Minimum Requirements:
Ubuntu 14.04 or newer
CPU: Intel Core i3-2100 @3.1 GHz / AMD Athlon II X4 645 AM3 @3.1 GHz
Memory: 4 GB RAM DDR3
Hard Drive: 10 GB free space
GPU: Nvidia GeForce GT 440 / AMD Radeon HD 6670 (1GB VRAM)
Additional Notes: Nvidia Graphics drivers 352, 355 or 358 series / AMD Catalyst (fglrx) 15.9 or better. MESA drivers & Intel graphics are not currently supported.

Well for instance, in ETD - that's short for Enter the Dominatrix - the Zin arrived with much fanfare, assaulting the city of Steelport... but, you know, I don't want to spoil it for you, see the mission and you'll get what I wanted to change.

Oscar graduated with a music degree from CSUN, is a Music Director/Teacher, software/hardware beta tester, Wikipedia editor, and active member of the Ubuntu community. You can email him at: 7bluehand@gmail.com

The Kindred is a Steam Early Access game. This review pertains to version 0.3.35, which is the latest release as I write (Feb. 2016).

Intro

I can describe The Kindred as only a mashup between Minecraft (the blocky look, mining and crafting), and something like Settlers (looking after your townsfolk). The idea behind it is that you take your people (Kin) and keep them fed, watered, and working.

Tutorial

The game has a brief tutorial mode that walks you through the basics of mining, crafting and building. After that you're on your own to look after your own group of Kin.

Playing

You start the game with about half a dozen Kin and have to start by chopping down some trees, building some beds, till the soil for planting crops, and various other crafting things.

Along the bottom of the screen is your menu with their relevant keyboard shortcuts. For example: clicking 'Mine' lets you choose blocks to dig up. An icon showing what that block is is shown at the bottom right of the screen. You can click 'cancel' (in the mine menu) to choose blocks to ignore and not dig. Clicking Back (or using right-click on the mouse) will take you back one menu item. Till is the same idea, but you remove the grass to expose dirt for planting seeds.

Clicking 'Housing' will grant you access to beds, planks, wood, windows, and the like, but these have to be crafted first. The one downside of the crafting is that it doesn't seem to show what you need to make an item. Another downside is that you have to go through the crafting table menus to find the item to craft, then through the normal menu to find

the same item now available (since you just crafted it). In other areas though, it is helpful. For example: if you try and milk an animal (with no bucket available), it will say it can't be done. Similar thing for crops that are not ready for harvesting.

Crafting also includes electricity. You have to craft up a wind turbine to generate electricity to have lighting. You don't get stick and coal torches in The Kindred.

Moving the camera is done using WASD while rotating, tilting and zooming is done using YU,GH, and VB.

While doing all that mining and crafting, you still have to look after your Kin and provide them with beds and food. At the top left of the screen is a list of your nutritional value requirements (per day) and available in storage. Also the Kin alive (six to start with), and the amount of electricity you're generating, storing and require. Clicking a Kin will give you info on

the person, such as name, age, strength, status, gender, also fatigue and hunger percentage.

For such a simple looking game, there's a lot to juggle. Which is good, as there's always something to do.

LOOK

The look of the game is very like Minecraft but with what look like Lego minifigs. There is a day and night cycle, but no enemies are in The Kindred. This is more about building a little town, looking after your people, and managing resources.

The developer is also very active, posting updates almost daily, and answering questions in the Steam discussion forum.

Even for early access, it's very polished with regards to graphics and the look and feel of the menus. A bit more documentation (or recipes) wouldn't go amiss, but this is definitely a great start.

Steam link:
http://store.steampowered.com/app/373410/

Minimum requirements:
Dual Core 2GHz
2GB RAM
512MB VRAM and GPU with passmark of 2000
Resolution of 1280x720 (or better)
400MB storage space
Sound card

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.

contents ^

PATRONS

PATRONS

Bill Berninghausen
Jack McMahon
Linda P
Remke Schuurmans
Norman Phillips
Tom Rausner
Charles Battersby
Tom Bell
Oscar Rivera
Alex Crabtree
Ray Spain
Richard Underwood
Charles Anderson
Ricardo Coalla
Chris Giltnane
William von Hagen
Mark Shuttleworth
Juan Ortiz
Joe Gulizia
Kevin Raulins
Doug Bruce
Pekka Niemi
Rob Fitzgerald
Brian M Murray
Roy Milner
Brian Bogdan
Scott Mack
Dennis Mack
John Helmers
JT

Elizabeth K. Joseph
Vincent Jobard

DONATIONS

John Niendorf
Daniel Witzel
Douglas Brown
Donald Altman
Patrick Scango
Tony Wood
Paul Miller
Colin McCubbin
Randy Brinson
John Fromm
Graham Driver
Chris Burmajster
Steven McKee
Manuel Rey Garcia
Alejandro Carmona Ligeon
siniša vidović
Glenn Heaton
Louis W Adams Jr
Raul Thomas
Pascal Lemaitre
PONG Wai Hing
Denis Millar
Elio Crivello
Rene Hogan
Kevin Potter
Marcos Alvarez Costales
Raymond Mccarthy

CHA CHA CHA CHANGES

Our admin went AWOL for months, and I had no idea if/when the site would/wouldn't get paid. Initially the plan was to move the site and domain name to my hosting, but eventually I managed to track him down and get the FCM domain name, and site hosting transferred to me.

The new site is now up. **HUGE** thanks to **Lucas Westermann** (Mr. Command & Conquer) for taking on the job of completely rebuilding the site, and scripts, from scratch, in his own time.

The Patreon page that I've set up is to help me pay the domain and hosting fees. The yearly target was quickly reached thanks to those listed on this page. FCM is not going away. Don't worry about that.

Several people have asked for a PayPal (single donation) option, so I've added a button to the side of the site

A big thank you to all those who've used Patreon and the PayPal button. It's a big help.

https://www.patreon.com/fullcirclemagazine

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 **site** via: fullcirclemagazine.org

FCM#108
Deadline:
Sunday 10th Apr. 2016.
Release:
Friday 29th Apr. 2016.

For the Full Circle Weekly News:

You can keep up to date with the Weekly News using the RSS feed: http://fullcirclemagazine.org/feed/podcast

Or, if your out and about, you can get the Weekly News via Stitcher Radio (Android/iOS/web): http://www.stitcher.com/s?fid=85347&refid=stpr

and via TuneIn at: http://tunein.com/radio/Full-Circle-Weekly-News-p855064/

Getting Full Circle Magazine:

EPUB Format - Most editions have a link to the epub file on that issues download page. If you have any problems with the epub file, email: 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.

Magzster - You can also read Full Circle online via Magzster: http://www.magzter.com/publishers/Full-Circle. Please share and rate FCM as it helps to spread the word about FCM and Ubuntu Linux.