Full Circle

Photo: Charles McColm

PROGRAM IN FREEPASCAL

```
 File   Edit   Search   Run   Compile   Debug   T
─[■                                          Com
program test;                    Main file: test1.
                                 Done.
uses                             Target: Linux for
      Crt;                       Line number:
                                 Used memory:      1
var                              Total errors:
      i : integer;
                                            Com
begin
      Clrscr;
      for i := 1 to 10 do begin
            Writeln(i, ' -> '
```

NEW

BOOK REVIEW

LINUX
SERVER SECURITY
HACK AND DEFEND

CHRIS BINNIE

WILEY

REFURBISH AN OLD COMPUTER
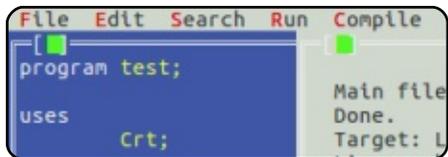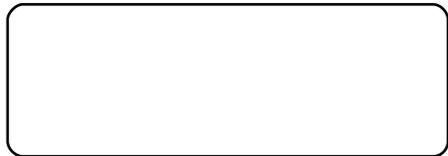## THE STEPS TAKEN TO CHECK, WIPE AND RESTORE A PC
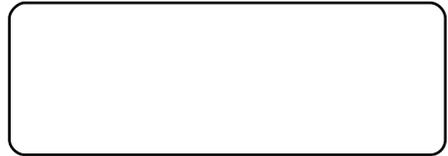
Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

EDITORIAL

WELCOME TO ANOTHER ISSUE OF FULL CIRCLE.

This month we have the start of a new short series. Let Alan Ward teach you the delights of the Pascal programming language using FreePascal. OK, so it's not the most popular programming language in the world, but it's still widely used and respected. We also have a short piece on the various methods used to install software in Debian/Ubuntu distributions. If you thought that apt (formerly known as apt-get) was the only way to install software then prepare to be amazed. To round off the How-To treble this month we have, of course, Greg's Python article.

Our cover feature is Charles' Linux Labs article on how he, and his colleagues, prepare a used PC for recycling. And by recycling I mean reuse. I have to be honest. I thought most PC reuse places would simply format the hard drive, install something-or-other, and put it up for adoption. Charles and his people certainly go the extra mile. Everything from clean-up to port testing, HDD wiping to Xubuntu install, are all meticulously done, logged, and put through a rigorous QA system. Highly impressive. I just hope other reuse places use such care with formerly used hard drives. But I suspect not.

As you can see from the contents page, things are pretty sparse. I really need more How-To articles, stories, reviews, anything! Please write about something. If you regularly use a piece of software then write one/two How-To articles on it. Write a review on a piece of software you use. If there's something you think we're not writing about then write an article on that missing subject. As long as it features Linux of some kind, we can use it. '*I don't have the time*' is **NOT** a valid excuse. The current regular writers are all busy people, and we shouldn't reply on them to fill the magazine each and every month. In short: **write!**

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

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

IT RUNS ON THE CLOUD, AND THE CLOUD RUNS ON LINUX. ANY QUESTIONS?

Like it or lump it, the cloud is taking over IT. We've seen the rise of the cloud over in-house IT for years now. And, what powers the cloud? Linux.

A recent survey by the Uptime Institute of 1,000 IT executives found that 50% of senior enterprise IT executives expect the majority of IT workloads to reside off-premise in cloud or colocation sites in the future. Of those surveyed, 23% expect the shift to happen next year, and 70% expect that shift to occur within the next four years.

This comes as no surprise. Many of us love our servers and racks, but it often doesn't make financial sense to run your own data center.

Source:
http://www.zdnet.com/article/it-runs-on-the-cloud-and-the-cloud-runs-on-linux-any-questions/

openSUSE USERS ARE GROWING BY THE MONTH

The number of users of openSUSE, the community GNU/Linux distribution supported by the German SUSE Linux company, has grown, with an average of 400,000 DVD images being downloaded each month.

openSUSE sees 1600 new installations and 500.000 package installations per month, according to Alberto Planas, a member of the openSUSE development since 2012 and an employee of SUSE.

He mentioned the figures during a presentation at the annual openSUSE conference held in Nuremberg from 22 June to 26 June.

openSUSE has three streams of development: Tumbleweed is a rolling distribution that has all the latest stable versions of software.

Source:
http://www.itwire.com/business-it-news/open-source/73513-opensuse-users-are-growing-by-the-month.html

CANONICAL DEMONSTRATES HOW EASY IT IS TO CREATE A VENDOR-INDEPENDENT SNAP STORE

The Snappy vs. Flatpak story continues, and Canonical is now demonstrating how easy it is to roll out a vendor-independent Snap store on the recently released Fedora 24 Linux operating system.

A couple of days ago, Canonical and Ubuntu founder Mark Shuttleworth finally answered one of the big questions many members of the GNU/Linux community had been asking since the unveiling of Snaps as universal binary formats for major Linux kernel-based operating systems.

Now that we know Snap stores

are simple HTTP web servers, Canonical's Dustin Kirkland has taken the time to show us how easy it is to create a Snap store. For this simple example, he chose to use an AWS instance of Fedora 24, but you can do the same on any other GNU/Linux operating system that supports Snappy.

Source:
http://news.softpedia.com/news/canonical-demonstrates-how-easy-is-to-create-a-vendor-independent-snap-store-505664.shtml

FEDORA 24 PUSHES LINUX BOUNDARIES

Fedora Linux is the community version of Red Hat Enterprise Linux, or RHEL. Fedora 24 is comprised of a set of base packages that form the foundation of three distinct editions: Fedora 24 Cloud, Fedora 24 Server and Fedora 24 Workstation.

Delayed four times during its development cycle, Fedora 24 includes glibc 2.23 for better performance, and improvements to POSIX compliance and GNU Compiler Collection 6. All base packages have been rebuilt with GCC 6, providing better code optimization across all Fedora 24 editions, and improving the overall stability of each addition.

Fedora 24 is the result of Fedora's drive to provide the latest powerful open source tools and components to a variety of end users, according to Matthew Miller, Fedora project leader. Those users range from developers to system administrators.

Source:
http://www.linuxinsider.com/story/83642.html

KaOS 2016.06 MOVES THE DISTRO TO LINUX KERNEL 4.6, ADDS FULL-DISK ENCRYPTION

First and foremost, the devs have decided to move the distribution from the long-term supported Linux 4.4 kernel series to Linux kernel 4.6, which makes it possible to fully automate the early microcode update. Furthermore, the default desktop environment has been migrated to the Beta of the upcoming KDE Plasma 5.7.

Another goodie that has made it into today's KaOS 2016.06 ISO image is the brand new Qt 5.7.0 GUI toolkit with its amazing technologies, which are perfect for a distribution running on top of the latest KDE software. Talking about KDE, the devs have also managed to add the latest KDE Application 16.04.2 software suite and the KDE Frameworks 5.23.0 collection of add-ons for Qt5.

However, probably something that most of KaOS users have expected for such a long time is full-disk encryption, which allows them to install the GNU/Linux operating system on an encrypted file system. Of course, the full-disk encryption support has been implemented in the installer via LUKS, for both the automated install and manual methods.

Source:
http://news.softpedia.com/news/kaos-2016-06-moves-the-distro-to-linux-kernel-4-6-adds-full-disk-encryption-505656.shtml

AUTOMOTIVE GRADE LINUX WANTS TO HELP OPEN SOURCE YOUR NEXT CAR

Recently I rented a Jeep Cherokee Limited edition, that included a touch-screen console with what was supposed to have all the bells and whistles. That touch screen wound up to be less-than user-friendly, not even remotely yielding to what I wanted it to do, and served little purpose other than to navigate my wife and I through Miami, Florida, listen to music, and view the rear-facing camera for backing up. The in-console display had serious issues connecting to any smartphone we had, so music was limited to satellite.

Needless to say, I wasn't impressed.

In the coming few years, that will all change, thanks to our friends at the Linux Foundation. The foundation started Automotive Grade Linux (AGL) to create open source software solutions for automotive applications. Their initial focus is on In-Vehicle-Infotainment (IVI) and their long-term goals include

the addition of instrument clusters and telematics systems. Already AGL has the likes of Ford, Jaguar, Land Rover, Mazda, Mitsubishi Motors, Nissan, Subaru and Toyota on board and that list will only continue to grow.

Source: http://www.techrepublic.com/article/automotive-grade-linux-wants-to-help-open-source-your-next-car/

Subuser uses Docker containers to deliver desktop apps for Linux

Distributing desktop applications for Linux has long been a headache, in large part because apps have to be repackaged for each Linux distribution. And while an app-containerization technology like Docker makes it easier to bundle and distribute apps, it wasn't really designed for distributing desktop applications.

Subuser is a new application-packaging system that allows Dockerized desktop apps to be run as if they were regular Linux applications. It provides just enough permissions to allow the Dockerized app to interact with the local system -- for instance, to work with the X11 display server -- while still keeping it locked down.

Most of the focus with Docker has been on building distributed applications or enabling consistent workflow for an app through its entire lifecycle. There's been relatively little discussion of how Docker can also be used to deliver and manage desktop applications. Subuser hints at a lot of untapped potential.

Source: http://www.infoworld.com/article/3088574/desktop-apps/subuser-uses-docker-containers-to-deliver-desktop-apps-for-linux.html

Microsoft releases cross-platform .NET Core 1.0 at Linux event

Microsoft has announced the release of .NET Core 1.0 and ASP.NET Core 1.0, the open source, cross-platform fork of the .NET Framework, letting people know at the Red Hat DevNation summit in San Francisco.

"This makes Red Hat the only commercial Linux distribution to feature full, enterprise-grade support for .NET," said Red Hat's blog post on the subject. The company has even registered the domain redhatloves.net, which redirects to a site with developer resources for .NET on Red Hat Enterprise Linux.

Why is Red Hat so keen? This line in the release says it all: "New .NET Core workloads can now be easily moved from a Windows Server environment to Red Hat Enterprise Linux, even if development was primarily done via Windows."

The company has its eye on Windows developers who now have a familiar route to deploying on Linux rather than Windows Server.

On Microsoft's side, the company is betting that any loss in Windows Server licences will be made up by increased use of its cloud services. Linux can be deployed on its Azure cloud, and there is strong .NET support for both Office 365 and Azure. Another factor is the forthcoming SQL Server for Linux, which will integrate nicely with .NET Core. Microsoft can profit from SQL Server licences even if you are not using Windows Server.

Source: http://www.theregister.co.uk/2016/06/28/microsoft_releases_crossplatform_net_core_10_at_linux_event/

Android phones can now read books, signs, business cards via Google's Mobile Vision

Google has introduced a new Text API for its Mobile Vision framework that allows Android developers to integrate optical-character recognition (OCR) into their apps.

The new Text API appears in the recently-updated Google Play Services version 9.2, which restores Mobile Vision, Google's system to make it easy for developers to add facial detection and barcode-reading functionality to Android apps.

The Text OCR technology currently can recognize text in any Latin-based language, covering most European languages, including English, German, and French, as well as Turkish.

The Mobile Vision API arrived last year with just the Barcode and Face APIs. Besides facial recognition, apps could also identify features on a face such as the eyes, nose, and mouth, as well as tell whether faces are smiling or if the individual's eyes are open.

Source:
http://www.zdnet.com/article/android-phones-can-now-read-books-signs-business-cards-via-googles-mobile-vision/

CANONICAL RELEASES NEW KERNEL UPDATE FOR UBUNTU 12.04 LTS AND UBUNTU 14.04 LTS

Ubuntu 15.10 will be supported for only one month or less, Canonical has released multiple Linux kernel updates for the other LTS (Long Term Supported) Ubuntu releases, Ubuntu 14.04 LTS (Trusty Tahr) and Ubuntu 12.04 LTS (Precise Pangolin).

According to Ubuntu Security Notice USN-3018-1 and Ubuntu Security Notice USN-3021-1, a total of seven kernel vulnerabilities affected both Ubuntu 14.04 LTS and Ubuntu 12.04 LTS operating systems, including the latest point release of Trusty Tahr, which has received new hardware enablement kernel from Vivid Vervet (Ubuntu 15.04) and Utopic Unicorn (Ubuntu 14.10).

Among the security issues fixed in Ubuntu 12.04 LTS and Ubuntu 14.04 LTS, we can mention information leaks in the core USB, ALSA, X.25 Call Request handling, and Rock Ridge implementations, a use-after-free issue in the generic PPP layer, and bugs in the CDC Network Control Model USB driver and InfiniBand interfaces.

Source:
http://news.softpedia.com/news/canonical-releases-new-kernel-update-for-ubuntu-12-04-lts-and-ubuntu-14-04-lts-505735.shtml

UBUNTU AMBIANCE THEME PORTED TO GTK 3.20, NAUTILUS 3.20 IS READY FOR YAKKETY YAK

Debian and Ubuntu developer Iain Lane writes today on one of the mailing lists of Ubuntu about the fact that he's been working for the past few weeks on bringing GTK+ 3.20 support to Ubuntu.

According to Iain Lane, the latest GTK+ 3.20 GUI toolkit, which is usually shipped by default with the GNOME 3.20 desktop environment, is ready for upload on the Ubuntu software repositories, for Ubuntu Desktop. Still, he doesn't mention for which Ubuntu version, so we can only guess that he's talking about Ubuntu 16.10.

However, he provides us with a PPA (Personal Package Archive) repository that contains the latest GTK+ 3.20 packages, the Adwaita icon theme, a GTK3 port of Mozilla Firefox 47.0, as well as Baobab and Nautilus 3.20 apps. The PPA is only for Ubuntu 16.10, so do not attempt to install it on Ubuntu 16.04 LTS (Xenial Xerus) because it won't work.

Source:
http://news.softpedia.com/news/ubuntu-ambiance-theme-successfully-ported-to-gtk-3-20-radiance-to-follow-soon-505765.shtml

SNAPPY MOVES IN THE MAIN ARCH LINUX REPO, SNAPD 2.0.10 RELEASED TO FEDORA COPR

Snapd 2.0.10 comes two weeks after the release of the 2.0.9 version, which introduced full Snap confinement on the elementary OS 0.4 "Loki" operating system, among several other goodies. However, the most interesting part of today's announcement for snapd 2.0.10 is that it also landed for Fedora 24 users in the COPR repository.

So if you don't want to use Flatpak, you can now install the latest Snappy implementation on Fedora too. "Fedora users can now get snapd 2.0.10 from the COPR repository. There are many bug

fixes and new features in this release," said Zygmunt Krynicki, Technical Lead in UES Commercial Engineering at Canonical.

As for the changes implemented by the Snappy development team in the snapd 2.0.10 release, we invite our tech-savvy readers and developers who would like to package their applications as Snaps for various GNU/Linux distribution to take a look at the changelog attached at the end of the article. There are a total of 82 changes in snapd 2.0.10.

Source:
http://news.softpedia.com/news/snappy-moves-in-the-main-arch-linux-repo-snapd-2-0-10-released-to-fedora-copr-505964.shtml

CERTIFIED UBUNTU IMAGES AVAILABLE IN SOFTLAYER

SoftLayer, an IBM Company and a world-leading IaaS provider, is now an Ubuntu Certified Public Cloud partner for Ubuntu guest images. For users, this means you can now harness the value of the best Ubuntu user experience, optimized for SoftLayer bare metal

and virtual servers.

This latest announcement nicely complements the recent work and broad partnership with IBM across platforms such as LinuxOne, Power and Z Systems. Canonical builds, continually maintains, tests, and updates certified Ubuntu images, making the latest versions available through Softlayer within minutes of their official release by Canonical. This means that you will always have the latest and most secure version of Certified Ubuntu images.

Source:
https://insights.ubuntu.com/2016/07/01/certified-ubuntu-images-available-in-softlayer/

AMD HAS YOU COVERED: ITS NEW LINUX GRAPHICS DRIVER ALREADY SUPPORTS THE RADEON RX 480

AMD's efforts to catch up with Nvidia on Linux are bearing fruit. AMD has historically lagged on Linux support for new graphics hardware, but its new AMDGPU-PRO 16.30 driver offers day-one

support for the impressive AMD Radeon RX 480.

This driver is currently available for download from AMD's website. At the moment, it's only officially supported on 64-bit versions of Ubuntu 16.04 LTS. It's very similar to the earlier beta release and AMD still calls it a beta, but it's reportedly very stable. AMD's website offers installation instructions.

The beta version of SteamOS 2.83 also includes the new AMDGPU-PRO driver, so AMD graphics may finally appear in future Steam Machines, too. Importantly for the future of SteamOS, AMDGPU-PRO offers support for Vulkan graphics.

Source:
http://www.pcworld.com/article/3090213/linux/amd-catches-up-its-new-linux-graphics-driver-already-supports-the-radeon-rx-480.html

LUMENERA CORPORATION ANNOUNCES THE RELEASE OF THEIR 'LUMENERA LINUX SDK 2.1'

Lumenera, a leading manufacturer and developer of high performance digital cameras and custom imaging solutions, is pleased to announce the launch of Lumenera Linux SDK 2.1, their new software development kit designed for embedded platforms using Linux ARM based processors.

Combining powerful features with reliable performance, the Lumenera Linux SDK 2.1 software maximizes the high-speed, high-resolution features of Lumenera's USB 3.0 cameras. Lumenera chose to formally support Ubuntu Linux with their cameras as it is used by many development teams around the world because of its strong presence globally, versatility, reliability, and extensive developer libraries. Ubuntu has regular Long Term Support (LTS) releases that are supported for five years, making it an ideal choice for developers.

Source:
http://www.benzinga.com/pressreleases/16/07/p8180882/lumenera-corporation-announces-the-release-of-their-lumenera-linux-sdk-

Pivotal Adds Ubuntu to Cloud Foundry

The steady shift to cloud native infrastructure continues with a partnership between enterprise software vendor Pivotal and Linux specialist Canonical that will provide secure images from Canonical's Linux distribution Ubuntu on the Pivotal Cloud Foundry.

The partners said they also would continue working to "harden" the Cloud Foundry distribution of the Ubuntu operating system to comply with federal benchmarks. Those include a U.S. military standard overseen by the Defense Information Systems Agency called the Security Technical Implementation Guide along with the Center for Internet Security benchmark.

The hardening initiative reflects growing security demands among government and highly regulated customers, the companies noted.

Meanwhile, Pivotal and Canonical said they are collaborating on a proposed industry standard set of security certifications for running Ubuntu on cloud native platforms. They said more details would be released later this year.

Source:
http://www.enterprisetech.com/2016/07/06/pivotal-adds-ubuntu-cloud-foundry/

Star Cloud PCG03U is a compact Ubuntu PC for $90

Chinese device maker Star Cloud has been offering tiny Windows and Android computers for a few years, but the company first came to my attention back in 2012 when I learned that the Android-powered Mele A1000 TV box was also able to run Linux.

This year the company started selling some products with Ubuntu Linux pre-installed, and the latest is the PCG03U, a compact computer/TV box with 2GB of RAM, 64GB of storage, an Intel Atom Bay Trail processor and Ubuntu 14.04 Linux.

It's available from AliExpress for $90.

Source:
http://liliputing.com/2016/07/star-cloud-pcg03u-compact-ubuntu-pc-90.html

Artificial Intelligence Is Setting Up the Internet for a Huge Clash With Europe

Neural networks are changing the Internet. Inspired by the networks of neurons inside the human brain, these deep mathematical models can learn discrete tasks by analyzing enormous amounts of data. They've learned to recognize faces in photos, identify spoken commands, and translate text from one language to another. And that's just a start. They're also moving into the heart of tech giants like Google and Facebook. They're helping to choose what you see when you query the Google search engine or visit your Facebook News Feed.

All this is sharpening the behavior of online services. But it also means the Internet is poised for an ideological confrontation with the European Union, the world's single largest online market.

Source:
http://www.wired.com/2016/07/artificial-intelligence-setting-internet-huge-clash-europe/

32-Bit Ubuntu Alternatives

Some folks may find the idea of using a 32-bit distribution of Linux to be downright silly. After all, we live in a 64-bit world these days, right? Well, that depends on who you ask. The fact of the matter is there are still a lot of fully functional PCs out there that run 32-bit Linux. Up until recently, this was all well and good. Then the news came down that Ubuntu would no longer be supporting 32-bit systems come the next Ubuntu release. Clearly not everyone is thrilled about his news.

Rather than throw in the towel and recycle these PCs, I think it's important to realize there is a world beyond Ubuntu. Yes, many other distros have also stopped support 32-bit distros. However for

the time being, there are still options to choose from.

The idea that 32-bit computers are going to suddenly stop working tomorrow is silly. Even if you choose to stick with Ubuntu, the 12.04, 14.04 and 16.04 LTS releases are all still supported. Ubuntu 16.04 is supported with updates until the year 2021. By then, your old 32-bit PC will likely be dead and you'll be commuting to work with a jetpack. That said, it's good to know that there are still oodles of great desktop Linux distros designed for those who are ready to use non-Ubuntu based distributions.

Source:
http://www.datamation.com/open-source/32-bit-ubuntu-alternatives.html

AT&T OPEN SOURCES ECOMP TO LINUX FOUNDATION, HOPES TO MAKE IT INDUSTRY'S STANDARD FOR SDN

AT&T today announced it will release its Enhanced Control,

Orchestration, Management and Policy (ECOMP) platform to the wider telecom industry as an open source offering managed by the Linux Foundation. The goal, the company said, is to make ECOMP the telecom industry's standard automation platform for managing virtual network functions and other software-centric network capabilities.

Donovan explained that ECOMP is a key element of AT&T's efforts to virtualize its network operations as the carrier moves into a software-defined network (SDN) architecture. According to AT&T, ECOMP enables the automation of service delivery, service assurance, performance management, fault management and SDN tasks. What's more, ECOMP is designed to work with OpenStack but can be extended to other cloud and compute environments.

Source:
http://www.fiercetelecom.com/story/att-open-sources-ecomp-linux-foundation-hopes-make-it-industrys-standard-sd/2016-07-13

BODHI LINUX 4.0 TO BE BASED ON UBUNTU 16.04.1 LTS, ENLIGHTENMENT'S EFL 1.18

Jeff Hoogland, developer and creator of the Ubuntu-based Bodhi Linux operating system, informs the community about a few important facts related to the upcoming Bodhi 4.0.0 release.

It looks like a pre-release Alpha version should happen soon, around the date of July 18, after EFL (Enlightenment Foundation Libraries) 1.18 hits Alpha stage of development.

Bodhi 4.0.0 Alpha release will be available next week even if Enlightenment delays its EFL 1.18 Alpha build, with the current 1.17.x versions of the EFL and Elementary packages for the Enlightenment desktop environment, on which Bodhi's Moksha interface is based. The even greater news is that Bodhi 4.0.0 will be based on Canonical's upcoming first point release for Xenial Xerus, Ubuntu 16.04.1 LTS.

Source:
http://news.softpedia.com/news/bodhi-linux-4-0-to-be-based-on-ubuntu-16-04-1-lts-enlightenment-s-efl-1-18-506257.shtml

LINUX CLOUDY TIE UPS: SUSE AND MICROSOFT, CANONICAL AND PIVOTAL

It's a big week for Linux cloud tie-ups, with SUSE and Microsoft expanding their partnership, and Canonical becoming Pivotal's preferred operating system in Cloud Foundry.

In the SUSE/Microsoft deal, the Linux outfit joins two Redmond programs: the Microsoft Enterprise Cloud Alliance, and the its test drive program.

The latter will be interesting to supercomputing geeks, because Azure customers will be able to try out SUSE's HPC Edition on Redmond's cloud.

The HPC offering includes Intel's Message Passing Interface (MPI) packages, and access to an Infiniband back end via Remote Direct Memory Access (RDMA).

The test drive uses SUSE Studio to create the images, and SUSE Manager, which is integrated with Microsoft Systems Center to manage the Linux Vms.

VMWare spinout Pivotal and Ubuntu master Canonical have worked together for years, so their tie-up is really about formalising their relationship.

They will be providing secure, certified Ubuntu images in Cloud Foundry, with upgrades and security patches automated so users don't have to worry about them. Unbuntu will also be providing support for its users on Cloud Foundry.

Source:
http://www.theregister.co.uk/2016/07/07/linux_cloudy_tie_ups_suse_and_microsoft_canonical_and_pivotal/

IBM Launches Blockchain Cloud Services On Linux Server

IBM has introduced a cloud-based blockchain service for business-to-business networks that allows companies to test performance, interoperability and privacy of blockchain ecosystems. The company noted in a press release that the service is suited to organizations in regulated industries.

IBM's blockchain cloud is supported by IBM LinuxOne, which IBM considers the industry's most secure Linux-only server. LinuxONE addresses the security requirements of the healthcare, government and financial sectors.

IBM's cloud services are designed to run blockchain in a production environment to enable clients to easily access a secure, portioned blockchain network they can deploy, test and operate. The IBM blockchain cloud protects data and entry points.

The service is currently in limited beta.

Source:
https://www.cryptocoinsnews.com/ibm-launches-blockchain-cloud-services-linux-server/

ChaletOS: A Linux that Provides an Uncanny Resemblance to Windows 7

The Linux platform has seen a surge of new users, who are usually migrating from Windows or at least they are trying Linux for the first time. But often, but they are afraid the interface will be too alien. Some developers think that it's a good idea to give users something familiar, so that their first experience on the open source platform won't be all that strange.

For many, the computer interface has always followed the same formula: Start Menu - Panel - System tray - Desktop icons

With those simple elements, people have happily interfaced with their hardware and done their jobs for a very long time. For those working within the world of Windows, the best take on that formula was Windows 7. So, it should come as no surprise that some Linux distributions have adopted that formula to create a desktop with which users would feel a sense of immediate connection.

Linux users do have a large selection of distros to choose from and there are plenty of friendly options out there, but some of the developers want to offer more than a simple implementation of a desktop environment.

Source:
http://www.techworm.net/2016/07/chaletos-linux-provides-uncanny-resemblance-windows-7.html

Intel's SGX tiptoes towards Linux

Intel has fulfilled a promise made in April to open-source a Linux driver for its SGX technology.

SGX – Software Guard Extensions – first landed in 2013, and allows programmers to lock up code and data inside containers enforced by the CPU. The idea is to create an environment to assure people "clouding" their enterprise systems that not even admins in the data centre can spy on what's going on.

Back in April, Chipzilla promised an SGX SDK for Linux, and a few weeks ago – with so little fuss we overlooked it – it made good over at GitHub.

The current implementation is very Alpha-looking, with just one distribution anointed to run SGX – Ubuntu 14.04-LTS 64bits. The hardware requirement is a Skylake system configured with SGX enabled.

Its Linux SGX implementation includes driver, SDK, and platform software. Intel notes that the driver isn't yet incorporated into the Linux main tree.

Source:
http://www.theregister.co.uk/2016/07/18/intels_sgx_tiptoes_towards_linux/

SUSE EXPANDS LINUX RESELLER PROGRAMME TO MSPS

SUSE has introduced a new programme that allows resellers of SUSE subscriptions to host SUSE software products on behalf of their end-user customers, providing more flexibility to both resellers and end customers beyond customers' on-site data centres.

A recent report from 451 Research found managed services at cloud service providers are growing at a compound annual growth rate of 26 percent and are projected to account for 36 percent of total cloud market revenue in 2018.

SUSE Reseller Hosting simplifies the move to managed services, allowing SUSE resellers to provide hosting services around SUSE products within their existing SUSE relationship. Resellers now have a low-touch way to realise new recurring revenue streams by extending their suite of offerings through managed hosting, said the vendor.

SUSE Reseller Hosting complements the existing SUSE reseller programme which enables SUSE partners to resell SUSE subscriptions for use in customers' on-premise data centres.

Source:
http://www.channelbiz.co.uk/2016/07/19/suse-expands-linux-reseller-programme-to-msps/

HOW TO INSTALL THE PRIVACYIDEA AUTHENTICATION SYSTEM ON UBUNTU

PrivacyIDEA is a modular authentication system that can manage authentication on your network. It's powerful, it's flexible, and it can be set up for free on an existing Ubuntu server.

Unlike a lot of authentication systems, privacyIDEA isn't that hard to install and set up; in fact, you can have your next authentication system up and running in minutes. I'll use the Ubuntu 16.04 platform to show how to set it up, which might seem like a problem because there isn't a privacyIDEA release for anything later than 14.04—fortunately, there's a very easy workaround.

You need a Ubuntu server that's up and running, and that server will need to have a full LAMP stack.

You can install privacyIDEA with NGINX, but I'm going to stick with what I know best: Apache.

Source:
http://www.techrepublic.com/article/how-to-install-the-privacyidea-authentication-system-on-ubuntu/

The
ubuntu®

STICKERS COLLECTION

~

Get the new exclusive Ubuntu stickers with a **15% discount** using the coupon

UBUNTU15

www.unixstickers.com/ubuntu

As a guitarist who is also a geek at heart, I'm always looking for a way to organise my various tab books alongside my PDF tabs. In preparation for my summer, I've therefore started collecting some of the tabs I like in an electronic format. At first, it was via scanning, but the quality of that was often less than overwhelming. Instead, I've begun writing the musical notation and tabs using LilyPond.

What is LilyPond?

For anyone who knows what LaTeX is, LilyPond is (in very basic terms) a typesetting language for music. In the music world, it's not called typesetting, but engraving. LilyPond is a language where you describe the music, and it then takes the information (keys, notes, etc), and generates staffs and/or tabs. Quality-wise, I've found the PDFs it produces to be better quality (from font-size, to spacing in general) than things like TuxGuitar or Tab Pro. Best of all, it is cross-platform, and editable from anywhere (as the source files are just text).

How does it Work?

LilyPond uses *.ly files as source code. Once the code has been written, it is then compiled into a PDF using the lilypond CLI tool. The way the source code works is by defining a Staff - this is the normal musical notation, or TabStaff (tablature), or a variety of other options, none of which I have tried thus far. Each Staff can have multiple voices (for splitting up upper and lower registers, for example).

How can I get started?

The website is http://lilypond.org, although the odds are the CLI tools are already in your local repository. The Manuals section of the site (http://lilypond.org/manuals.html) offers a good place to start. Once you've gotten comfortable with the syntax (or have run into a use case that doesn't seem officially supported), then the user repository of snippets is an excellent spot to look: http://lsr.di.unimi.it/LSR/Search. For anyone who wants some actual examples, jump to the section called "The Basics", as I will walk you through some examples.

It seems complicated - why can't I just drag and drop?

If you want to edit visually, you'll need to look for other software to work with. Instead, I recommend LilyPond for anyone who prefers having total control (and finds a keyboard to be faster and more efficient than a mouse). I also highly recommend this approach to anyone who wants to learn to read musical notation better (as you'll learn the note names, and therefore the location on the staff), or for anyone who wants to learn the locations of the notes on their instrument (specifically, guitar). This is because in thinking about the music in terms of notes (instead of string or fret numbers), you'll learn which note is where - especially as you move on to more complex songs.

The Basics

The template for the file can be found here: http://pastebin.com/pyJS56zj. Copy and save the template as FCM-example.ly (or any other name you may want).

A single Voice, on a normal musical staff (no Tab):
• Add the following between \pointAndClickOff and \score

```
scale= \relative c {
  a ais b c cis d dis e f fis g gis
  \bar "|."
}
```

• As you can see, you're simply defining the notes that are to appear (and terminating with a closing bar). In LilyPond **is** stands for sharp (#), and **es** is flat (b). Keep in mind that a# is also bb (hence why no flats are present in the example above).
• Before you can compile it, you'll need to add a Staff to the score. Inserting a Staff into the score is

done by placing the following between the << and >> inside the \score{}

```
\new Staff = "guitar" \with {
   } <<
      \time 4/4
      \context Voice = "guitar"
{ \clef "G_8" \voiceOne
\scale }
   >>
```

• The \new Staff should be self-explanatory. The "guitar" simply assigns a name (in case you have multiple instruments on one page, for example). The \with{} is used for settings, and I leave it in my template so that I can more easily change settings for a Staff. If you prefer to leave it out, you can do so (but leave the <<). The \time is the time signature. Then comes the important step - defining a Voice. I've defined it as "guitar", with a G_8 clef (the clef is the symbol right at the beginning of the Staff), and defined VoiceOne using the variable scale we created earlier (step 1).

• Compilation time! Be sure to replace FCM-example.ly with whatever you called your file.
```
lilypond FCM-example.ly
```
• This should result in a PDF, showing a scale of the notes. If this is the case, move on. If not, double-check the steps and the error messages.

• What if we want to add in another octave? To do so, you can use ' and , to increase, and lower, the octave of a note. See the example below.
```
scale= \relative c {
      a ais b c cis d dis e f
fis g gis
      a' ais b c cis d dis e f
fis g gis
      \bar "|."
}
```

• If you look at this example, you'll see that I've inserted only one apostrophe. This is because it is a toggle, and not a one-time modifier. Meaning if you use two, your second note will be two octaves higher. Also, most people will notice that this actually skipped an octave. This is due to the fact that LilyPond is smart enough to realise that, after an increasing scale like above, that a new a should probably be an octave higher. You can disable this feature, but I find it quite useful. So to get the first 3 octaves, you can do this instead:
```
scale= \relative c {
      a ais b c cis d dis e f
fis g gis
      a ais b c cis d dis e f
fis g gis
      a ais b c cis d dis e f
fis g gis
      \bar "|."
}
```

• This should result in a single scale, covering three octaves. LilyPond does offer a \repeat unfold option, to avoid having to type the same line multiple times. However, it would not work in this case, as there is an implied difference between each line (their octave). The below will repeat the same scale 3 times, in the same octave. This is useful to keep in mind.
```
scale= \relative c {
      \repeat unfold 3 {
         a ais b c cis d dis e
f fis g gis
      }
      \bar "|."
}
```

• Prettying up the source code is useful, especially when you write whole songs. A bar division is indicated with "|", and LilyPond will track the number of notes in a bar (to ensure you keep the time signature). It will throw warnings, but still compile, if this is wrong. See below:
```
scale= \relative c {
      a ais b c |
      cis d dis2 |
      e f |
      fis |
      g gis
      \bar "|."
}
```

• This example will result in the warning "barcheck failed at 1/2". This is because I changed the d# into a half note (2 beats), and this change is automatically carried through to all following notes. By placing the f# in a bar by itself (and therefore being only 2/4), it results in the warning. In the PDF, there is no noticeable issue, but this will cause problems if using multiple voices (or if the song were longer). Instead, we need either a rest, or to adjust the duration of the notes. Correct the bar to read:
```
fis r2 |
```
• This adds a 2 beat rest into the bar, and fixes the timing issue.

• Lastly, let us add a TabStaff, in order to create actual tablature. To do so, insert the following after the first >> (following the closing bracket of \context Voice).
```
   \new TabStaff = "guitar" <<
      \time 4/4
      \context TabVoice =
"guitar" { \clef "G_8"
\voiceOne \scale }
   >>
```

• As you can see, it's a similar approach. We're even reusing the same variable with the notes. Compiling at this time will yield a typical guitar tablature

There are many other things you can do - chords, multiple voices, arpeggios, repeats, etc. But for a basic example, this will have to do. If you want to see what my file looked like at the end of this tutorial, it can be found here:
http://pastebin.com/eqBwrkX9

If you have any issues, or questions, feel free to contact me. For some more examples, I highly recommend the snippets and documentation page. Or just simply creating a tab you already know. As always, I can be reached at lswest34+fcm@gmail.com.

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.

ubuntu uk podcast

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

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

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

podcast.ubuntu-uk.org

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!**

Last month, we worked with the DS18B20 Temperature Sensor. This month we will start to interface a 16x2 LCD display to show our temperatures. Don't tear down your setup, but make sure you have enough room to mount the display on your breadboard. You'll need about 32 pinholes for the length of the device and 16 for the pins to connect to. You will have only three pinholes left if you mount the display at the bottom of the vertical holes, so you will need to use some jumpers to connect the bottom verticals to the top verticals.

Of course, the 16x2 display has 16 characters on two rows. The backlight comes in many colours. I chose a blue one. We can address each of the 32 character positions individually, or print pretty much like we do to the regular monitor.

We will be making 8 connections to the RPi as well as the three that we used for the temperature sensor last month. You will need the following additional items for this month:

• 10K Potentiometer
• 16x2 LCD Display
• Many breadboard jumpers, Male to Male and 8 Male to Female

By the time you are done, the wiring diagram (and the resulting board) will look like a bit of a rat's nest, but go slowly – make sure you have the wiring correct.

As you can see in the graphic above, it's pretty crazy, so I'll lay out all the wiring for you in text.

First, you will need to put a jumper between the two horizontal busses on both the top and bottom. That way, you'll have power and ground on both busses. I chose to do it on the left side, but you can put it anywhere that is convenient for you. The next thing to do is to wire in the potentiometer. One side (it doesn't matter which) needs to go to ground and the other side to our 5 volt supply. The center contact (the wiper) will wire to pin 3 of the LCD display. This controls the contrast, so you can control how bright the characters appear. You
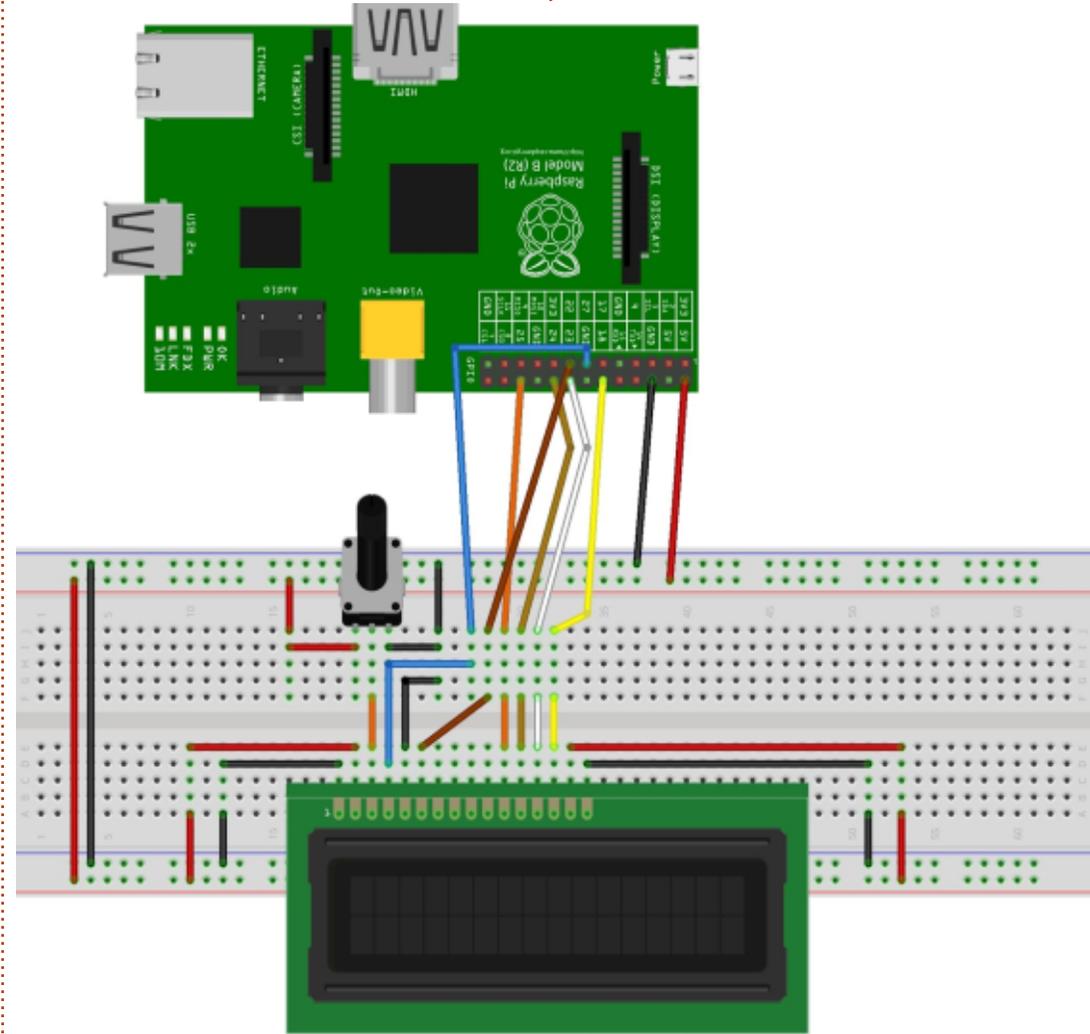
should already have 5 volts to the board, as well as ground, from last month.

On the display, connect pin 1 to

ground and pin 2 to the +5 volt buss. That makes three connections out of the twelve we need. Pin 6 of the display goes to pin 22 of the RPi. This is the Enable

pin. Pin 5 on the display goes to ground, and pin 4 to pin 27 on the RPi. We are up to 6 connections so far. That makes us halfway there. Because we have to use pin 4 for our sensor, we can't control the backlight.

Now we will work backwards from pin 16. Pin 16 goes to ground, and pin 15 to +5v. Pin 15 is actually the backlight voltage on mine. If you find the display too bright, you could put the wiper of another potentiometer connected between +5v and ground and control the display backlight.

Now for the data lines. There are actually 8 data lines, but thankfully, we will be using only 4. Pins 11 to 14 are D4, D5, D6 and D7 (counting from 0). Here is the connection list.

Display Pin	Raspberry Pi Pin
11	25
12	24
13	23
14	18

Now everything is hooked up, so we will continue with some sample code to test the display. But we need to get the Adafruit python library for LCDs. In a

```python
#!/usr/bin/python
# Example using a character LCD connected to a Raspberry Pi or BeagleBone Black.
import time
import Adafruit_CharLCD as LCD
# Raspberry Pi pin configuration:
lcd_rs      = 27   # Note this might need to be changed to 21 for older revision Pi's.
lcd_en      = 22
lcd_d4      = 25
lcd_d5      = 24
lcd_d6      = 23
lcd_d7      = 18
lcd_backlight = 4
# Define LCD column and row size for 16x2 LCD.
lcd_columns = 16
lcd_rows    = 2
# Alternatively specify a 20x4 LCD.
# lcd_columns = 20
# lcd_rows    = 4
# Initialize the LCD using the pins above.
lcd = LCD.Adafruit_CharLCD(lcd_rs, lcd_en, lcd_d4, lcd_d5, lcd_d6, lcd_d7,
                           lcd_columns, lcd_rows, lcd_backlight)
# Print a two line message
lcd.message('Hello\nworld!')
# Wait 5 seconds
time.sleep(5.0)
# Demo showing the cursor.
lcd.clear()
lcd.show_cursor(True)
lcd.message('Show cursor')
time.sleep(5.0)
# Demo showing the blinking cursor.
lcd.clear()
lcd.blink(True)
lcd.message('Blink cursor')
time.sleep(5.0)
# Stop blinking and showing cursor.
lcd.show_cursor(False)
lcd.blink(False)
# Demo scrolling message right/left.
lcd.clear()
message = 'Scroll'
lcd.message(message)
for i in range(lcd_columns-len(message)):
    time.sleep(0.5)
    lcd.move_right()
for i in range(lcd_columns-len(message)):
    time.sleep(0.5)
    lcd.move_left()
# Demo turning backlight off and on.
lcd.clear()
lcd.message('Flash backlight\nin 5 seconds...')
time.sleep(5.0)
# Turn backlight off.
lcd.set_backlight(0)
time.sleep(2.0)
# Change message.
lcd.clear()
lcd.message('Goodbye!')
# Turn backlight on.
lcd.set_backlight(1)
```

terminal window, type the following...

```
git clone
https://github.com/adafruit/A
dafruit_Python_CharLCD

cd Adafruit_Python_CharLCD

sudo python setup.py install

cd examples
```

Now load char_lcd.py into your favorite editor. Or, you could type it in from the previous page.

Ignore the backlight messages, but you should see...

```
Hello World!
Show Cursor_
Blink Cursor_
Scroll (right and left)
Flash backlight in 5
seconds...
Goodbye!
```

If everything worked, we are ready to proceed. If not, go back and check your wiring.

Here is the modified program from last month that includes snippets from this example (top right) from Adafruit. (New code is in bold.)

That's about it for this month. Next month we will look at

```
from w1thermsensor import W1ThermSensor
from time import sleep
import Adafruit_CharLCD as LCD
# Raspberry Pi pin configuration:
lcd_rs        = 27
lcd_en        = 22
lcd_d4        = 25
lcd_d5        = 24
lcd_d6        = 23
lcd_d7        = 18
lcd_backlight = 4
lcd_columns = 16
lcd_rows    = 2
# Initialize the LCD using the pins above.
lcd = LCD.Adafruit_CharLCD(lcd_rs, lcd_en, lcd_d4, lcd_d5, lcd_d6, lcd_d7,
                           lcd_columns, lcd_rows, lcd_backlight)

sensor = W1ThermSensor()
while 1:
    # temp_in_celsius = sensor.get_temperature()
    temp_in_fahrenheit = sensor.get_temperature(W1ThermSensor.DEGREES_F)
    print temp_in_fahrenheit
    lcd.clear()
    lcd.message(str(temp_in_fahrenheit))
    # print temp_in_celsius
    sleep(3)
```

changing out our regular 16x2 display for a 16x2 I2C display (which uses only 2 lines for data and all control, and 2 lines for power.) We will also discuss the different ways of using serial communication for interfacing displays and other devices. Until then, have fun!

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

If you are new to Debian/Ubuntu, then you'll need help installing and managing apps on it. Sure, you can do it via a GUI installer/package manager if a specific Linux distribution already has one pre-installed. But where's the fun in that? If you are managing a virtual server, you probably would not have a GUI. Installing and removing software through a command line is not hard to do, and people have fun with it. You can even do it by using different methods which we will explain in the article below. For the purposes of this tutorial, we will use our Debian 8 VPS, but the same instructions apply to any other Debian-based distribution, including Ubuntu, both for desktop and server versions.

MANAGE SOFTWARE USING DPKG

Ever seen a .deb file and wondered what it is? Well, that's the Linux version of the Windows ".exe". You can't click on it to install it, but you can use dpkg, which is a package manager for Debian-based distributions. Note that dpkg can't automatically download and install software for you, you'll have to download a .deb file first, and then install it using dpkg.

So, to install a locally-available .deb file using dpkg, use the following command syntax:

```
sudo dpkg -i app-file-name.deb
```

Of course, replace app-file-name.deb with the actual filename of the app you are trying to install.

And that's it. It's that simple. If you want to uninstall a package, you can do so using the following command:

```
sudo dpkg -r packagename
```

Although, removing packages with dpkg is not recommended, because it will remove the main package, but it will retain other dependencies that the package has.

There's a lot more you can do with dpkg, like listing all installed packages using:

```
dpkg -l
```

Beware, this may give you a huge output, depending on the number of packages you have installed on your system, so you may want to use the less command, for example:

```
dpkg -l | less
```

Refer to the manual for more information and help, by typing the following command:

```
man dpkg
```

USE APT-GET TO DOWNLOAD, INSTALL AND REMOVE APPS

Using apt-get is a lot easier way of installing packages or applications. The big upside of using apt-get instead of dpkg is that you can do pretty much everything with apt-get. You don't have to download a package before installing it with apt-get, you can do it instantly by using the following command:

```
sudo apt-get install
package_name
```

That's it. The package is installed. You can update a single package using:

```
sudo apt-get --only-upgrade
install package_name
```

Note that this will not install any new packages. This just upgrades it. Before installing or updating a package or application, it is a good idea to download the package lists from the repositories and update them to get information on the newest versions of packages, so do it with this command:

```
sudo apt-get update
```

This will update your package lists and the versions you have installed on your system. You should always run the update command first, and then run:

```
sudo apt-get upgrade
```

HOWTO - INSTALL SOFTWARE ON DEBIAN/UBUNTU

which will update all of your software to their latest versions. The difference between apt-get update and apt-get upgrade is that the former updates your package information and the latter command actually upgrades your packages.

To uninstall (remove) an app from Debian/Ubuntu, use the following commands:

```
sudo apt-get remove
package_name
```

This will remove all the files from the package, but it would not remove any configuration files or dependencies. To remove (just about) everything, including configuration files, use:

```
sudo apt-get purge
package_name
```

You can remove the dependencies of a package by using this command:

```
apt-get autoremove
```

That's just about it. There's a lot more stuff you can do with apt, just refer to the manual page by using:

```
man apt
```

INSTALLING SOFTWARE FROM SOURCE CODE ON LINUX

Although not the easiest way of installing software, especially compared to apt-get, it's still fun to compile and install software from source code on Linux. Sometimes it's the only way of doing it. This method should work on almost every Linux distribution. Always refer to the software's official documentation/manual to get the right information. This method can be very different depending on the software you are trying to build.

To compile and install an application from source code, you must first unpack it.

If the archive is in .tar.gz format, use:

```
tar xvzf package_name.tar.gz
```

If the archive is in tar.bz2 format, use:

```
tar xvjf package_name.tar.bz2
```

If the archive is in .tar format, use:

```
tar -xvf package_name.tar
```

If the archive is in .zip format, use:

```
unzip package_name.zip
```

Then, you need to configure the software using the ./configure command. First, navigate to the directory where you've unpacked the software:

```
cd package_name
```

Then enter the following command to use the default configuration for the software:

```
./configure
```

This is the recommended way of doing it, unless the official software distributor recommends doing some other changes. You can see a list of all the options you can use with:

```
./configure --help
```

The next step is to compile the software. To do so, use the following command:

```
make
```

The compiling process can take from a few minutes to a couple of hours, even days. It all depends on the software you are installing and your machine's hardware specs. For most apps it usually takes a few minutes or so. Watch out for some errors in the output. The most common ones are missing dependencies. You'll need to install them if you notice such errors where there's an output with "not found", "unable to locate" or something similar. Refer to the official documentation of the app you are compiling to get help with any errors.

After the compiling is finished, you can finally install the app using the following command:

```
make install
```

Uninstalling software you've previously manually compiled and installed can be tricky. If the developer has included the option, you can run the following command:

```
make uninstall
```

and the software should be removed, but in some cases that option may not be available to you,

so you can run:

```
make clean
```

to remove any files the installation process has outputted, but this will remove only the files in your source/build tree. It won't affect the rest of the filesystem. Again, you should refer to the official documentation of the app you are trying to install/remove. They'll most probably have an in-depth tutorial of doing everything you need.

We've covered most methods of managing software on Debian-based distros. What's your favorite way of installing and managing software on Linux? Which method do you prefer?

RoseHosting.com is a privately held company founded in 2001 and based in St. Louis, Missouri - RoseHosting offers quality managed Linux VPS hosting with 24/7 support.

Many PC programmers started with one of Borland's products, back in the day. Both Turbo Pascal and Turbo C ran on the MS-DOS operating system, and had excellent integrated development environments comprised not only of the compiler itself, but also a nice editor. To the point, they also contained a system of text-based widgets (buttons, radios, editing areas, menus…) called Turbo Vision that was not only used to build the IDE itself, but was also at the disposition of programmers for their own applications.

For simplicity and for elegance, it has not yet been surpassed as a text-oriented human interface. This is not only my own very honest opinion, but also that of some others - as seen on the Internet. So, it must be true. It was with interest that I discovered the existence of Free Pascal (www.freepascal.org), a project led by Florian Paul Klämpfl, that aims to bring back a similar Pascal compiler, but open-source and adaptable to many different architectures. It also contains the interface environment, but adapted to modern times. This has been named Free Vision. As an aside, Lazarus (http://wiki.freepascal.org/Lazarus_Faq) is a companion project that continues along the same lines as the Borland Delphi graphical development interface; Vince DuBeau gave us a brief description of that project some time back in **FCM#77** (http://fullcirclemagazine.org/issue-77/).

In this series of articles, I will be building a text-based application with Free Pascal, using its text-based interface for user interaction. This will be combined with other, more modern, technologies such as database access using SQL, and Web access with HTTP. The final aim of the project is to demonstrate how Pascal can be used to build a modern application, while avoiding the overhead associated with a graphical interface (using a widget set such as GTK or Qt). This first part will describe installing the tools and writing some simple programs.

Why?

Before delving into the technical intricacies of setting up the Free Pascal and Free Vision combination, perhaps it would be judicious to answer this basic question: why is it worthwhile for me to spend time on this technology? After all, there is a plethora of other languages for the programmer to choose from, such as Go, Swift, Haskell… just to name a few. The obvious answer is: "speed!" Just like that, with an exclamation mark. This stuff is seriously lightweight, and thus it also runs seriously fast. It gets to the point where a Raspberry Pi running an interface written in Free Vision can beat a proper modern computer with a multicore CPU and an SSD running a similar application, but interfaced through the Web. Which also makes sense, in a way, since the RPi has more processing horsepower than the 286 or 386 on which most Turbo Pascal programming was done.

Meanwhile, the program using Free Vision is a compiled, executable binary that does not need to contend with all the quirks of a graphical interface or HTTP server. These are available, but the programmers do not need to waste CPU or RAM on them unless they are actually useful to what they are doing.

A second, equally valid, answer may be: "for old times' sake." Pascal is also a nice, clear and structured programing language that can help us get into good habits - habits that JavaScript is not known for nurturing, just to give a quick example.

Installation

Free Pascal is installable using standard tools (apt, synaptic…) from Ubuntu 16.04's repositories, where it is identified as package fp-ide. Package fp-units-base will also be required to make standard units available, as well as fp-units-fv for Free Vision, fp-units-db to access databases, etc.

HOWTO - PROGRAMMING WITH FREEPASCAL

Just to see if there were differences with the original developer's version, I downloaded FreePascal version 3.0.0 for the Intel 86-64 (64-bit) architecture, from the project's web page at http://www.freepascal.org/download.var . The download is a .tar file, that in turn contains a further three compressed files (in the .tar and .tar.gz formats) that contain the compiler itself in binary or executable format, examples of programs, and the documentation. An install script is also provided.

Running the install script as root will allow us to install the lot, choosing between either /usr or /usr/local as a starting point for our files. The decompression process is fast and hassle-free. Once done, we have the appropriate PATH variable set in our system, and can proceed with creating a directory for our source files, and opening it within a terminal. Something along the lines:

```
mkdir Desktop/pascal
cd Desktop/pascal
fp
```

There were no detectable differences with the version from the repositories. The final command launches the integrated Pascal IDE (Integrated Development Environment), which will bring back memories to users of Turbo Pascal. Some of us have spent hours staring at this screen.

A configuration file is created for the environment within this directory. At this point, we can immediately begin programming. However, if we need to use units such as Crt, we will need to direct the IDE to use the appropriate unit directory. This is a step the installer does not seem to take care of. Simply choose Options > Directories in the main menu, and add the following two lines. If you have installed in /usr instead of in /usr/local, do the appropriate modifications:
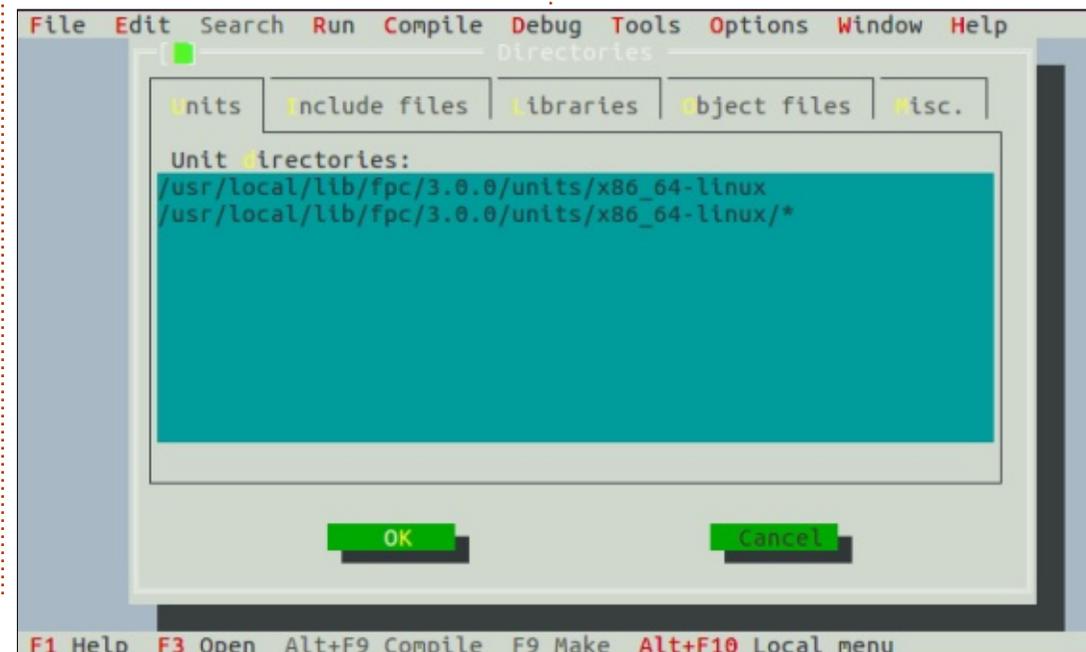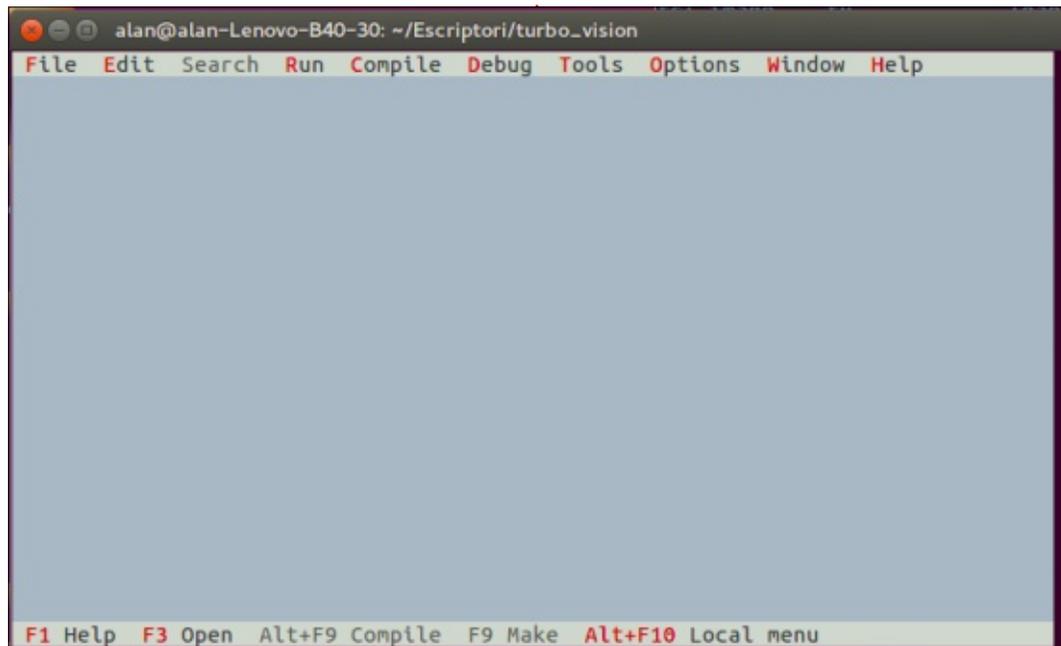
```
/usr/local/lib/fpc/3.0.0/units/x86_64-linux
/usr/local/lib/fpc/3.0.0/units/x86_64-linux/*
```

The first line is to include the unit directory itself, while the second also includes any subdirectories within the search path.

COMPILING A CONSOLE PROGRAM

To start, let us write a very simple program. This will serve just

HOWTO - PROGRAMMING WITH FREEPASCAL

to test the compilation process, and see if units are in fact correctly detected. Choose File > New in the menu, and type in a test program (such as http://pastebin.com/wJhgry5A). Keyboard shortcut command F2 saves the file (with extension .pas), and F9 compiles it (top right).

The program can be both compiled and executed in one step using command Ctrl+F9. The result is shown in the terminal, and execution goes back to the IDE whenever a key is pressed.

A FIRST FREE VISION APPLICATION

Let us get on to using the modern equivalent of Borland's Turbo Vision text-based user interface, now called Free Vision. The corresponding units are grouped in subdirectory:

```
/usr/local/lib/fpc/3.0.0/unit
s/x86_64-linux/fv
```

The first one to use is simply the App unit itself. With the ten-liner in http://pastebin.com/13UfNBNS, we can set up a sample text-based application.

The corresponding code is quite minimalistic. We have simply created a new application type, TMyApp, based on the pre-existing TApplication. This application is then instantiated in a variable. The application is initialized in method Init. Method Run is the main loop, where user keyboard and mouse events are handled. Finally, we clean up in method Done. That's it - but it is already a working application, written in less lines of code than are needed to describe the process.

```
var MyApp : TMyApp;

begin
     MyApp.Init;
     MyApp.Run;
     MyApp.Done;
end.
```
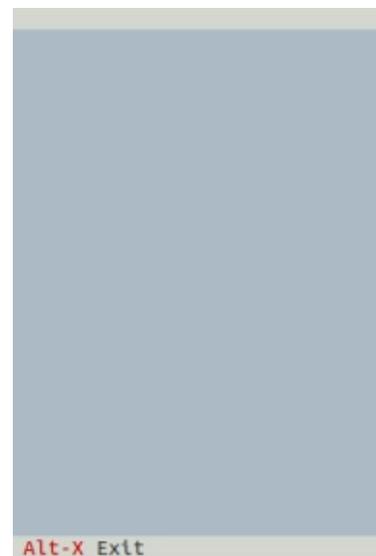
One point that can be noticed is that the resulting executable files are rather small, even if using FreeVision. The first one - using merely Crt - takes up about 400 kBytes, and the second - with fv - takes less than 800 kBytes. This is quite small by modern standards.

```
$ ls -lh test? test?.pas
-rwxrwxr-x 1 alan alan 416K
jun  7 09:05 test1
-rw-rw-r-- 1 alan alan  175
jun  7 09:04 test1.pas
-rwxrwxr-x 1 alan alan 781K
jun  7 09:06 test2
-rw-rw-r-- 1 alan alan  162
jun  7 08:45 test2.pas
```

This will be all for this first part in our series on Free Pascal - and Free Vision. In the next part, we will make the user interface responsive, and create some simple dialog windows.

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 systems administration at the Open University of Catalunya (UOC).

UbuCon Europe 2016

18.-20. November 2016 | Unperfekthaus, Essen, Germany, Europe

ubucon.eu

www.ubuntu.com

ubuntu

It appears that last month's announcement about forums was premature. Suffice to say that politics and personalities have been at loggerheads in the world of Inkscape support forums, but things have since calmed down.

So I'll stick to pure facts: Both the inkscapeforum.com and inkscapecommunity.com forums continue to operate, each with a different subset of users (and some degree of overlap). Support requests posted to either will generally elicit a response, and there's no need for normal Inkscape users to concern themselves with the behind-the-scenes shenanigans. None of this has any impact on the development of Inkscape itself. If and when there is any more concrete information about an official forum, I'll write about it, but until then I'll be keeping forum politics well out of this column!

Now, where were we...? Oh yes, filters. Let's return to the single-colored drop shadows of Part 49 to show you a simpler way to create the same effect. Previously, I introduced you to the Color Matrix primitive as a means of converting one color to another, but when all you need to do is to introduce a specific fixed color into your filter chain, it's usually easier to use the Flood primitive.

As you might have guessed from the name, the Flood primitive floods an area with color. You may now be thinking along the lines of the bucket tool in Inkscape or other graphics programs, which typically floods an area by working outwards until it hits a differently colored boundary line. But there's no such finesse here; the flood primitive simply fills the whole of the "filter effects region" with a flat color. The filter effects region is the rectangle defined by the Filter General Settings tab (see part 48), and is typically larger than the bounding box of your selected objects.

Starting again with some simple text, create a filter and add the Flood primitive. With the primitive selected, use the controls at the bottom of the Filter Effects dialog to choose a color and opacity, and you should get a result something like this (note that it doesn't matter what the input of the Flood primitive is connected to, as it has no effect on the output):
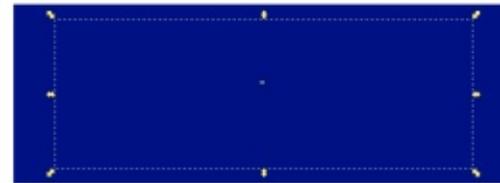
Not terribly inspiring, is it? So the question now is how to turn this big blue rectangle into a softly shaped drop shadow. If you followed last month's tutorial, you'll know that the Composite filter (used in "In" mode) can be used to crop the blue rectangle into the shape of our text.

It's not very common that you'll want the output from your filter to be strictly rectangular, rather than following the shapes and curves of your selected objects. So whenever you see a Flood primitive in a filter chain, there's a good chance that there will be a Composite filter following along shortly afterwards to trim it to shape.

Now that we've got a colored version of the text, it's a straightforward matter to offset and blur it, before merging it with the Source Graphic – you should be adept at those steps by now, so I'll spare you a detailed description and instead present the result, a graph version of the chain, and a screenshot from Inkscape.

Whilst Flood provides you with a rectangle of a single color, the Turbulence primitive gives you a rectangle filled with a chaotic mix of colors. It's not strictly random, in the mathematical sense, as the output is well defined and repeatable (meaning that your
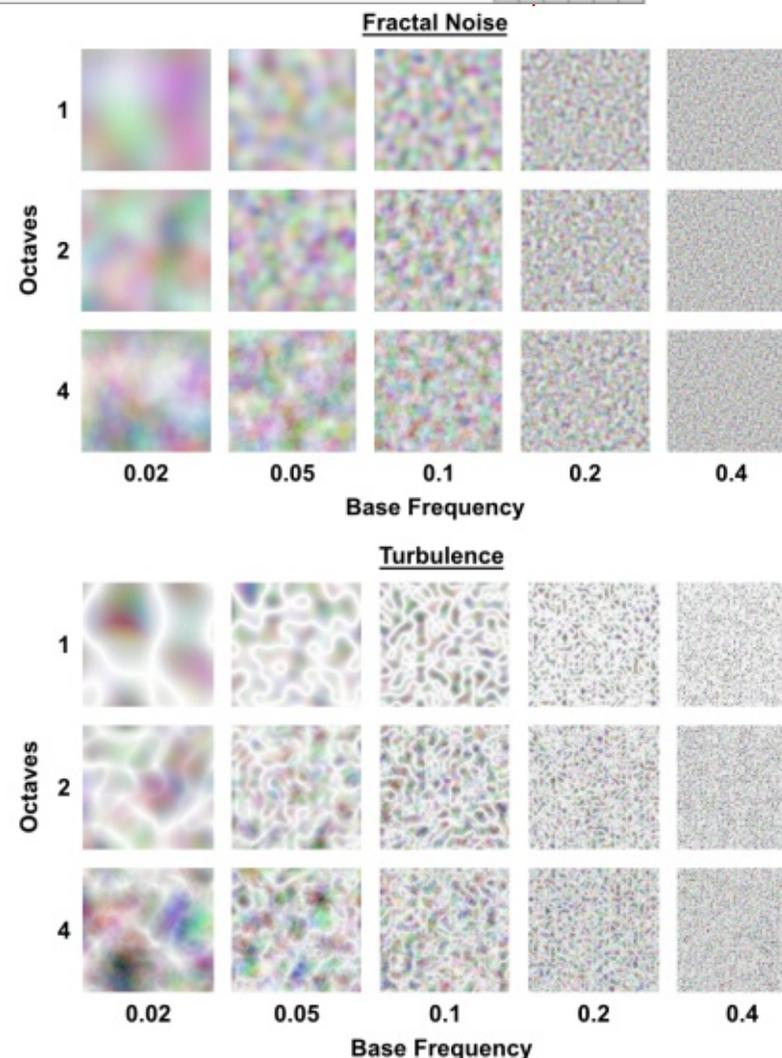
Fractal Noise

Turbulence
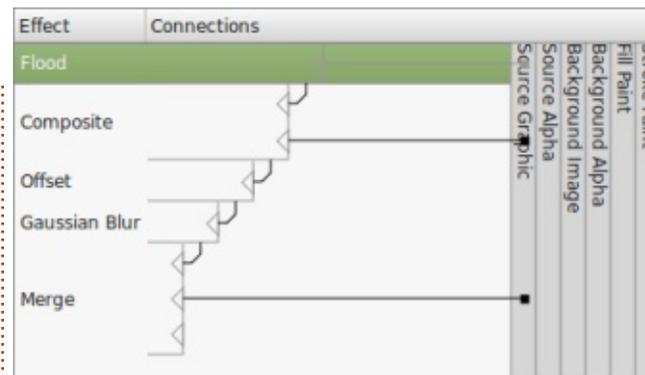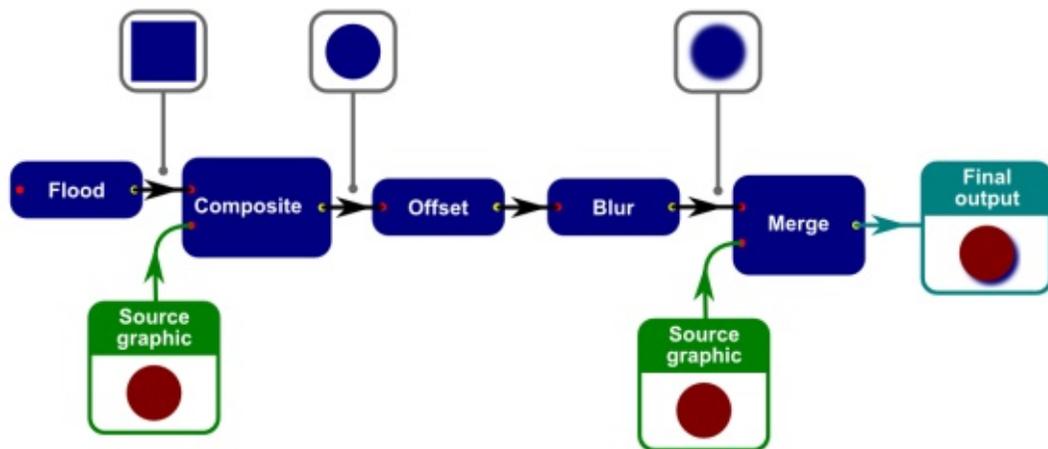
filters should look the same in any renderer – although in reality that may not be the case), but, in the colloquial sense, it's this primitive that you should head for if you want to add a degree of randomness or noise to your image. It has two modes: Fractal Noise and Turbulence. The difference between them is that the latter has more "troughs" in the output, where the background shows through, giving the appearance of joined up lines running throughout the output, whilst the former has more of a cloudy appearance.

Whichever mode you choose, the rest of the controls remain the same. The Base Frequency sliders control how "dense" the noise appears – low values give slow, smooth transitions, whereas higher values result in transitions that change more rapidly, making the output more reminiscent of "snow" on an old un-tuned TV set. The horizontal and vertical frequencies are usually the same, but can be changed independently by toggling the Link button to the right. The Octaves slider controls how detailed or complex the noise appears; taking this much beyond about 4 is rarely worthwhile as the increased detail is too small to see, and it imposes an extra load on the processor. Finally, the Seed value can be used to prime the pseudo-random number generator at the core of the filter to give you a slightly different output pattern without changing the other parameters.

The following images show the effect of varying the Base Frequency and the Octaves sliders, for both the Fractal Noise and Turbulence modes.

You'll notice that the images are fairly pastel in tone. This is because all four channels (R, G, B, A) are calculated independently – each pixel actually consists of a combination of four pseudo-random numbers. The value of the Alpha channel will override all the others, so even if you happen to have a strong color from the RGB components, a low Alpha can knock it back to a translucent shadow of its former self.
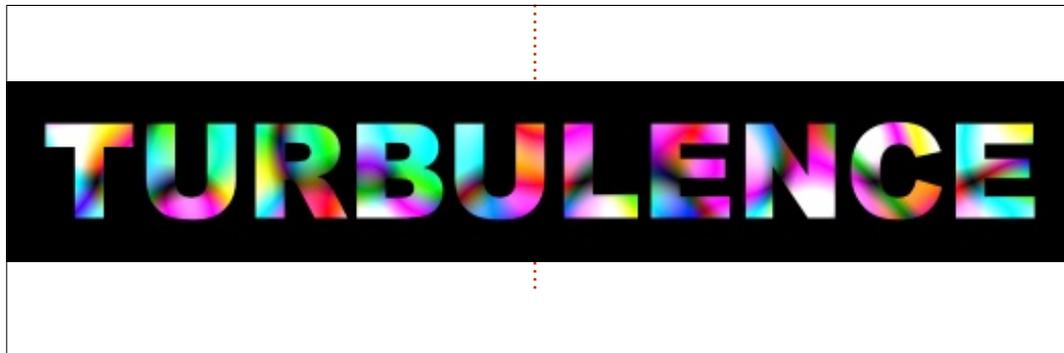

You can use a Color Matrix to extract a single channel, or to stretch the output to make it more vibrant. In this example, I've done the latter, as well as wiping out the Alpha channel entirely and replacing it with a fixed value of 1 (fully opaque). The cyan color of the original text doesn't show through at all in this case (any cyan in the result comes purely from the Turbulence filter), but I have used a black background to make the colors stand out even more.

To extract a single channel from the output, zero everything in your Color Matrix, and then populate just one of the first four columns, depending on what you want to get out. For example, setting a value of 1.00 in every field in the third column would take the 0-255 values in your Blue channel and map them to RGBA in the output. A Blue value of 63 on the input would produce (63, 63, 63, 63) – a

translucent gray color – as an output. You might want to take the Alpha channel out of the equation by setting a value of 1.00 in the bottom right corner (the Fixed Value column for the Alpha output). In this example, I've used the Green channel to set only the Alpha of the output, and stretched the values a bit by using 3.00 rather than 1.00. This gives an image that runs from opaque black to transparent black, so, by Compositing it (to clip it to shape) and then Merging it with the cyan Source Graphic, it's easy to create an "electric" or "plasma" effect.
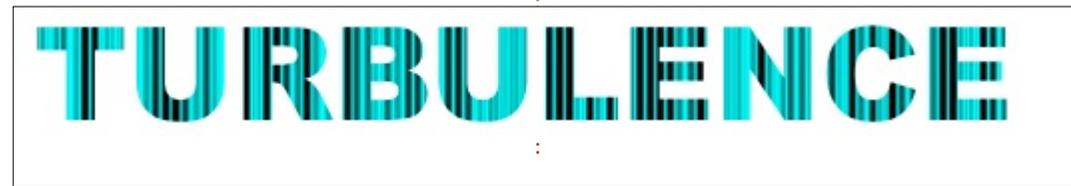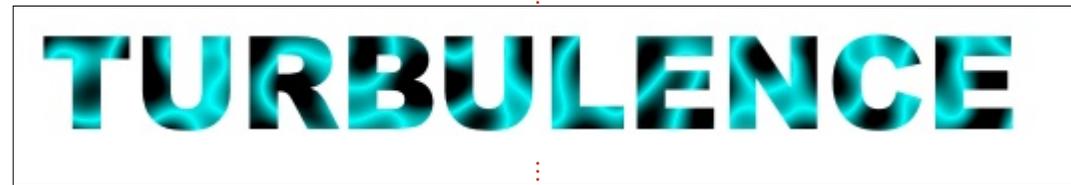
You could also bring a Flood primitive into a chain like that to ensure that your result has the right color in the filter, regardless of the color of the object it's applied to.

Don't forget that the Base Frequency controls can be un-linked. By keeping the values close to each other, you can introduce a slight stretching or bias into the patterns, whilst separating them by some distance can result in almost horizontal or vertical lines appearing. Here's the previous filter, but with the horizontal Base Frequency set quite high, and the vertical Base Frequency at zero – the result isn't what you might usually think of as "turbulent", but can be a useful addition to your filter arsenal nonetheless.

By now you should be starting

to appreciate the power and flexibility of filters. By combining a few primitives in various ways, you can quickly create complex results. Throw in a little pseudo-random chaos and you're well on your way to everything from clouds to marble, whilst Flood primitives can ensure that the important colors in your filter are independent of the objects they're applied to.

One common theme between Flood and Turbulence is that they fill the filter region entirely, usually requiring a Composite operation to trim them to shape. Next month, we'll look at the last of these "fill" primitives – uncovering limitations in Inkscape along the way – then progress on to other ways to change their shape.

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

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.

"

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

Richard Trotter
Geoteric

A big weakness in the Chromebook is the 16 GB SSD and the use of cloud storage. I intended to write an article discussing Network Attached Storage (NAS) and using SSH. I picked up a cheap Pogoplug and reprogramed it into a NAS while using a Chromebook for access. There is plenty of online documentation that outlines this process: Pogoplug as NAS (http://projects.doozan.com/debian/). You simply replace the native Arch Linux on the Pogoplug to Debian by following those instructions. This device will allow a "home private cloud" to be established for personal use.

```
Welcome to crosh, the Chrome OS developer shell.

If you got here by mistake, don't panic!  Just close this tab and carry on.

Type 'help' for a list of commands.

crosh> ssh
The 'ssh' command has been removed.  Please install the official SSH extension:
https://chrome.google.com/webstore/detail/pnhechapfaindjhompbnflcldabbghjo
crosh>
```

However, as I started reviewing the Pogoplug OS replacement steps, I realized I never reviewed the Chrome OS command line terminal which is called Crosh. To access Crosh, you press Ctrl+Alt+T, and the terminal opens up in the Google browser.

So, exactly what are the available Crosh Commands? Ping, SSH, Battery test, Bluetooth, Debugging, Memory Test, and some other commands. A partial list of possible commands can be found here:
• How to Geek (http://www.howtogeek.com/170648/10-commands-included-in-chrome-oss-hidden-crosh-shell/)
• Samsung Chromebook Support (http://goo.gl/HLQ7ff)
• Crosh Commands at krypted.com (http://krypted.com/unix/helpful-chrome-os-shell-crosh-commands/)

When I searched the Chromium OS projects for further information about Crosh, only 15 items populated. These fifteen items were off-topic and merely mentioned Crosh. These commands are useful; however Crosh is not intended to be a major tool in the Chrome OS. For a power user, Crosh will be a benefit in debugging and SSH. For most users, Crosh is just a digital vestigial organ.

Yet I wonder why Crosh is so minimalist? After careful consideration, Crosh is limited since the Chrome OS is tied to the cloud. Since there are very few apps that reside on the SSD, it makes sense that Crosh is limited. Perhaps as the Google Playstore becomes tied into the Chrome OS, the Crosh command options will expand. I foresee more locally controlled apps versus cloud usage in the Chrome OS future. Consequently, the Crosh will become more user friendly. And on a sidenote, I can SSH into my Pogoplug using my Chromebook.

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.

The impact of the amperstamp. This is the @, that is found in all email. The person that created and sent the first "Internet" email is Ray Tomlinson. Tom used the @ in his code. He recently passed away.

Ray was born and raised in New York State, USA. He graduated from Rensselaer Polytechnic Institute in 1963 studying electrical engineering. He was then accepted into Massachusetts Institute of Technology graduate program to continue his electrical engineering.

He eventually was hired by Bolt Beranek and Newman, also known as BBN. At BBN, he helped develop the Tenex OS that worked with Arpanet. He wrote a program called CYPNET that allowed for file transfer to the various sites that were connected to Arpanet. He added a few lines of code that enabled users to send messages between different computers. Previously only computers on the same site network could leave messages for each other. Ray introduced the @ and some code that enabled the different Arpanet sites to send email. The first email was sent to himself as test. He believed the email contained QWERTYUIOP.

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.

Have you ever wondered what happens to a computer once you've donated it to a community computer refurbisher? Different refurbishers have different processes – so I can't speak for every refurbisher – but I imagine many refurbishers have similar processes to The Working Centre's Computer Recycling project. Our project volunteers work off a checklist of steps.

The first step is an initial assessment. We see enough equipment come through our project that some computers are worth more to us recycled than they are reused. Ideally, we would refurbish everything that came through our doors, but the reality is there is very little demand from people for Pentium III. It's sometimes hard to tell a donor that their several thousand dollar computer (they bought 10 years ago) has more value as scrap than as a refurbished computer. Ideally, we'd love to refurbish every machine, but not everyone loves TWM, Emacs, and CLI. (We have refurbished the odd lower spec machine for clients who really want one).

Our project is at the point that even Pentium 4 computers are getting too old for us to refurbish. Pentium 4 computers tend to choke on high definition videos, and since a lot of YouTube videos now default to high definition, we aimed a bit higher for our basic build. We could use Firefox plug-ins to manage the default resolution for YouTube, but then we'd have to account for other video sites. In the end, we decided it was just better to set our base refurbished model specifications a little bit higher, at least a Core 2 Duo.

When we first look at a computer, we examine the outside of the computer and check for a number of problems:
• Is the computer missing parts we don't have, such as a specific kind of side panel?
• Is the computer too yellow from sunlight?
• Is the computer full of tar from heavy smoking near the computer?
• Does the computer look colossally old?
• Are there signs of insect/animal droppings on the back of the computer (this step surprises many, but people have donated computers where we've had to bring in pest control).

The next step is an internal examination:
• Again we look for insects/critters first. Sometimes we have to give a computer a little bit of a shake to see movement. If there are critters, we bag the computer immediately and call our on-call pest control expert.
• If a motherboard has bad capacitors we'll examine the board to determine if it's worth fixing or not. We might replace a capacitor or two on a system capable of dual core with SATA and PCIe –

provided there are not too many bad capacitors.
• If there are missing parts (hard drive, RAM), we make a note of it.
• If the computer is dirty we use a Data Vac to suction out large chunks of dirt. Our Data Vac has two ends, one end vacuums, the other blows compressed air. We find the blower end much stronger than the suction end, but we generally vacuum first to avoid spreading dust and germs. Volunteers wear protective eye gear and masks for the process.

External cleaning:
• Remove any asset tags that identify the donor. Often we'll get computers donated with the hostname, or a person's name, on the computer. We remove these stickers/tags. Educational institutions and government bodies often put asset tags on their equipment. We have putty knives and razors to help lift the tags.
• Next we give the computer case a good rub down with a diluted mixture of isopropyl alcohol and water. For stickers, we use a product called Goo Gone, or a competitor called Goo Off. There are lots of sticker removal products on the market, but you
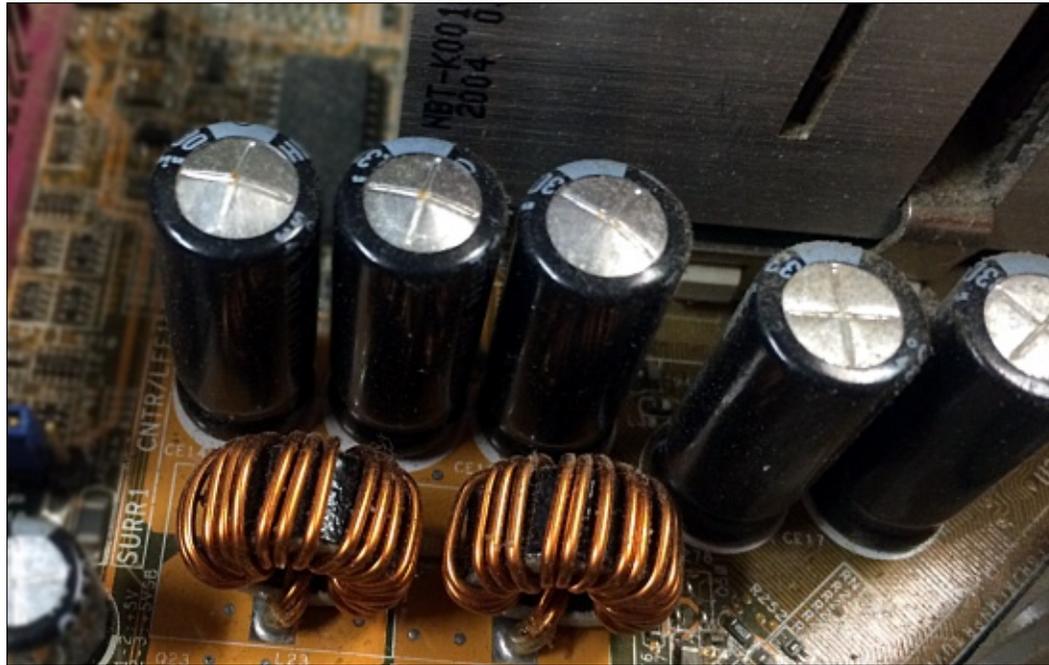
want to make sure you don't buy something too abrasive.

Booting/Building a machine:
• We have almost all our software tools and installers on a PXE-enabled server. To access the server, we enable PXE network booting in the BIOS of the computer. If a computer doesn't have PXE booting (this was another reason we decided to set our minimum build at dual core or better), we boot off a gPXE CD. gPXE, or Etherboot, is a bootable CD full of network drivers that lets you boot from the network by first booting from a CD on machines

that don't support network booting in the BIOS.
• Once we have a hard drive in the computer, we ensure that the BIOS SMART feature is enabled if it's available. SMART just lets us see if a drive is starting to go bad.
• We then boot a Debian live environment that has several tools we use for different steps. The first of which is phoronix-test-suite – a benchmarking tool that can be used to also gather system specifications. We use Phoronix and the computer BIOS itself to gather the following information about a PC:
o Original Equipment

Manufacturer Serial Number
o CPU Manufacturer (Intel/AMD)
o CPU Model
o CPU Speed
o Hard Drive Manufacturer
o Hard Drive Model
o Hard Drive Size
o RAM Type: DDR, DDR2, DDR3, OTHER
o RAM Size
o Motherboard Manufacturer
o Motherboard Model
o Optical Drive
o Type of Network card: 10/100/1000
o Whether there's a sound card or not (this is almost a useless field because, with the exception of servers, we have one on every machine)
o Video Card Manufacturer
o Video Card Model
o Video Card RAM

We also have a spot for notes, and, at the top, logos for Windows or Xubuntu, along with a spot just below for the estimated price.

If the specifications are good enough, we'll take the machine to our back build area. Our general rule for a minimum Xubuntu Linux build is a Core 2 Duo with 2GB of RAM and an 80GB SATA Hard Drive. We break this rule sometimes if we

don't have extra RAM or are low on SATA Hard Drives. If a computer has sticks of RAM larger than 2GB, we normally remove and replace them with smaller, tested sticks of RAM. It's very rare for us to get a computer with 6 or 8GB or RAM, but it does happen once in awhile. Only recently have we switched from 32-bit to 64-bit installs of Xubuntu.

Unless the machine in question is a quad core, we usually remove any hard drives that are 500GB or larger. We keep the 500GB hard drives for the occasional quad core machine that comes in.

Our old process was to wipe (dban) old hard drives during the build process, but this took a lot of time, so we switched to replacing any drives in a system with a drive that has already been wiped. Darik's Boot And Nuke (dban) is a great program that can overwrite a drive with garbage data as many times as you tell it. Typically, we use the DOD-Short standard (3 passes) except when asked by specific clients to do more. Putting drives in our dban machine is an easy job for new volunteers to tackle. There is a little bit of troubleshooting sometimes when

there's a really bad drive, or when our wiping machine doesn't network boot, but we have other volunteers around to help troubleshoot.

Running DBAN on a hard drive only wipes the drive. If we have a really bad drive, DBAN fails. Unfortunately, DBAN doesn't report all drive errors so we need to use a third party tool to do basic and (sometimes) extended tests on the drive. To test hard drives, we use Gsmartcontrol, a Linux tool that reads the drive's SMART data. If a drive is showing any pink/red, we recycle it and use a different

drive. I touched on Gsmartcontrol before in FCM issue 108.

CD-ROM drives have been dead to us for about two years. We generally remove any CD-R/CD-RW drives and replace them with a DVD-RW. We started doing this because our optical media shelf was getting so full, we were finding ourselves recycling drives every couple of weeks. Only using DVD-RW drives also means we don't get questions about why a movie doesn't work in a CD drive.

The computers we get don't always come in complete.

Sometimes they're missing front panel pieces or there are empty holes at the back of the computer where someone has removed an expansion card. Internally we check to make sure that all cables are connected correctly and that there aren't extra cables hanging loose in the case. Good cable management allows air to flow through the computer.

Our next step is to test the RAM. Unlike the hard drives, we normally test RAM in the computer. There are a couple of reasons why we test the RAM in the system:
• To ensure the RAM slots are not bad (if we suspect a slot is bad, we'll test the RAM at a workstation where we know the RAM slots are good, and we retest in the build computer with other RAM)
• Testing RAM usually takes less than 30 minutes, whereas wiping a drive can take 5 hours – or more for large drives.

Next comes the OS install. I mentioned earlier that all our tools run off a PXE boot server. Our Xubuntu installer is scripted so no user interaction is required and we get a complete environment including some special
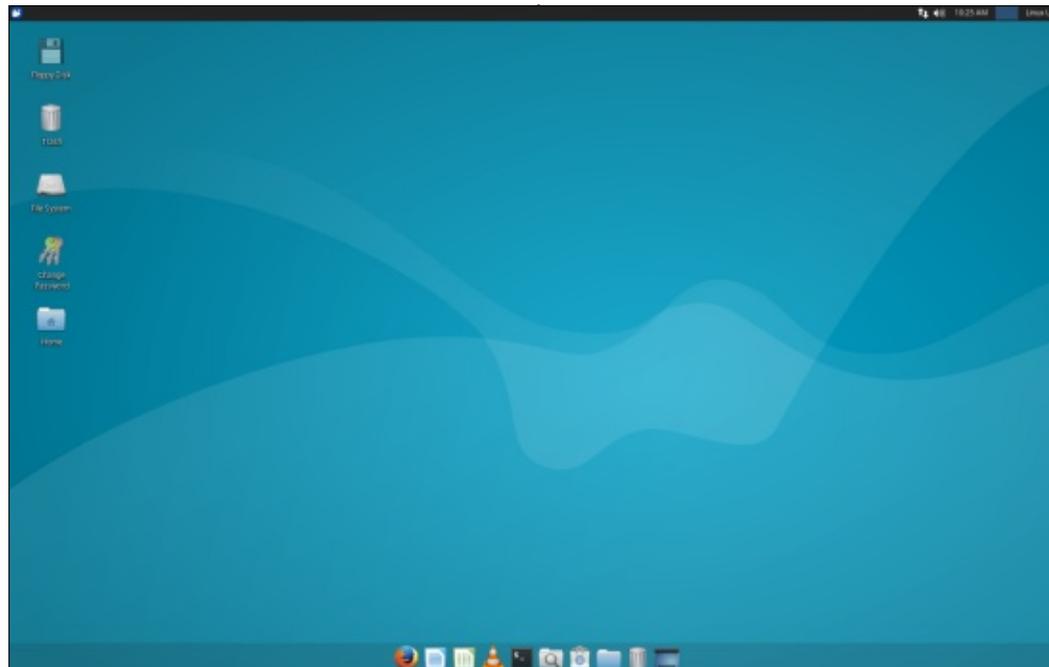
scripts/tools that help us do remote support. Our scripts are shell scripts (vnc) that provide two-way authentication, so both our project and the client have to provide passwords for remote access to occur.

Once the install is complete, we review the machine to ensure autologin is working, and web browsers and flash are working. I normally test video by going to Youtube because we can also test sound at the same time (we probably should change this since much of youtube is HTML5). We also test any floppy drives and all the USB ports.

Our installer is pretty great; we also get recent updates thanks to apt-dater. Once a machine is finished installing, it generally is right up-to-date. The additional step we need to do sometimes is enable any proprietary video drivers for Nvidia or AMD video cards. For this, we just use *buntu's built-in Additional Drivers tool (this may change for us in the near future since 16.04 has some changes and doesn't appear to work the same as 14.04).

At this point, the build is done

and the machine gets set on a shelf with a ready for QA (Quality Assurance) sticker. We have 2 volunteers check the build to make sure all the build steps were followed. QA includes looking for things like bad capacitors, ensuring all the ports are working, and checking to see if the system is clean and up-to-date. Generally, we have a more senior volunteer look over a build because they'll look for things a less experienced volunteer might ignore (missing screws, cabling issues, sounds that the computer shouldn't be making).

Finally, if the computer passes both QAs, it gets a price/specifications sticker, and we enter it into our point-of-sale inventory. Our point-of-sale is a custom PHP/HTML/Javascript front end we had written on top of the eCommerce suite OSCommerce. Our original intention was to mirror our inventory on our web site, but we removed OSCommerce from our web site years ago and replaced it with a Drupal + commerce solution. Any volunteer can enter a computer into the point-of-sale, though it's only our most experienced volunteers who

tend to do so.

RESOURCES:

• gPXE/Etherboot:
http://etherboot.org/wiki/index.php
• Phoronix Test Suite
http://www.phoronix-test-suite.com/
• DBAN
http://www.dban.org/
• GSmartcontrol
http://gsmartcontrol.sourceforge.net/home/
• Apt-dater
https://www.ibh.de/apt-dater/
• OSCommerce
https://www.oscommerce.com/Products
• Xubuntu
http://xubuntu.org/

Charles is the author of Instant XBMC, and the project manager of a not-for-profit computer reuse project. When not building PCs, removing malware, and encouraging people to use GNU/Linux, Charles works on reinventing his blog at http://www.charlesmccolm.com/.

This list is only a highlight of the few changes available in this update. Please check the detailed changes for all the changes included in this OTA.

IMPORTANT FEATURES

• Complete MPRIS support for playlists
• Fingerprint reader support for the Meizu PRO 5 (turbo)
• Convergence
    • New libertine scope for converged devices
    • Animated mouse cursors
    • Windows vertical/horizontal maximization
    • On-screen-keyboard support for X apps
• Location permission prompt on scopes first use
• Video consumption from scopes
• Message forwarding in messaging-app
• Color emoji support in the ubuntu-keyboard
• Additional useful toggles in various indicators
• Oxide 1.15
• Updated core-apps
• Webbrowser
    • Touch-selection improvements
    • Hovered hyperlink URL displayed in an overlay
    • Only loading the current tab when the session is restored
    • Optimized new tab view loading time
    • Zoom support
    • Consistent page headers across the UI

BUGFIXES

• Fixed screen dimming when indicators are open
• WebRTC fixes for selecting cameras - fix video rotation issues Connectivity fixes
• Various stability improvements
• Fixing AltGr not working for external keyboards for accented characters
And many many more...!

DETAILED CHANGES

Commitlog:
http://people.canonical.com/~lzemczak/landing-team/ota/ota-12.commitlog

Milestone:
https://launchpad.net/canonical-devices-system-image/+milestone/12

KNOWN ISSUES

The following is a list of issues that notable enough for explicit mention in the release notes.

LP: #1600582
(https://bugs.launchpad.net/canonical-devices-system-image/+bug/1600582) -
Installation of "Desktop Applications" breaks BQ Phones - please do not install puritine on non-convergent devices
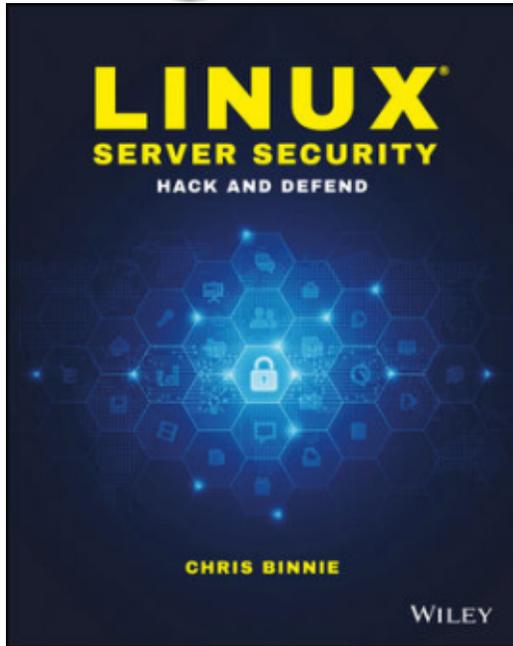
Linux Server Security - Hack and Defend
Chris Binnie
126 pages
Wiley : ISBN: 978-1-119-27765-1

When I got this book, I looked at the fact that it is so thin and wondered if there would be enough covered in sufficient detail to make this a good book. Well, as they say, good things come in small packages and it's true with this book.

Chris Binnie does a wonderful job taking the complex subject of Linux security into a comprehensive narrative. He states, right up front, that this book is not for everyone. In fact, he uses the phrase "mid-level admins, software hackers, and other IT professionals". This however, doesn't mean that a "normal" level person couldn't do some research on a particular subject of interest, and then jump to Chris's book for some higher level information.

Well laid out, the book provides information on both system AND network level hacks that you might encounter, and provides information on how to mitigate the potential damage as well as some software packages to help out that you might have not heard of.

Subjects covered include things like making your server invisible without disrupting services, how to turn the hacker's favorite tools against them, learning how hackers identify your system's weak points, how to defend against malware and DDoS attacks, and more...all packed into 126 pages.

I learned a tremendous amount of tips and tricks from this book and am happy to give it 5 out of 5 stars.

⭐⭐⭐⭐⭐

Table of Contents

LETTERS

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

REASONS I LIKE UBUNTU

I had an old HP notebook with Windows XP with the same problem - took minutes to boot and same to get online, and sometimes just would not load an app.

I was about to junk it when I tried Xubuntu 12.04 -- tried 14.04 but no PAE. I found that Xubuntu 12.04 or earlier didn't require it. Xubuntu 12.04 works like a charm. I'm sure it's even less than 40 sec to load. I use Xubuntu 14.04 on an old desktop, 2 GB RAM and Ubuntu 14.04 on the same computer as a dual boot.

Tried Ubuntu 15.10 but it kept crashing. I use this computer just for testing, although, because of age, it's not too useful for distros.

On another old desktop – not sure of age as it was one that someone was going to discard – 2 GB RAM single core, I put 12.04 on it. My husband uses it for Internet and a bit of other stuff. I think it's very slow, but he is happy with it.

I use Mint 17 on my main computer and am very happy with it. I am 80+ and using Inkscape tutorials from Full Circle which I find fun. They are great. Just did lesson #49.

Trudy Thorgeirson

PYTHON PROBLEM

I have learned to code with HTML 5 and CSS 3. I have made two websites and uploaded them. Now reading through Full Circle, I thought it is time to learn Python. I opened "Program in Python, Part 1". I saved the first bit to hello.py. Very nice.

Then trouble! It reads: *Before we can run the program, we need to set it to be executable. Do this by typing chmod +x hello.py in the folder where you saved your python file. Now let's run the program.*

greg@earth:~/python_examples$ ./hello.py

My question: how do I type in a folder. This is not possible. It says

run the program. How?

Pieter

Greg says: *You should use a terminal window. In that window change to the folder (in my case it was "/home/greg/python_examples/"), but yours would be different. Like this...*

```
cd ~/python_examples
```

but whereever you saved your program code. You want to change the permissions to be able to execute it as a program, not just a text file. You would then type in the terminal...

```
chmod +x hello.py
```

Then to run the program, you would type into the terminal...

```
./hello.py
```

or you could force Linux to use python like this...

```
python ./hello.py
```

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&A
Compiled by Gord Campbell

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

Q How might I get VLC Media Player to play all the music tracks of a CD in sequence?

A For this task, Audacious is much easier to use.

Q I'm considering making some changes to my system setup. Is there a way I can create a "restore point" in case it doesn't work out?

A Google: **systemback**

The web site is in Launchpad, and you can install it from a PPA.

Q I get an error message when I try to update Google Chrome.

A Enter this command:

```
wget -q -O -
https://dl.google.com/linux/l
inux_signing_key.pub | sudo
apt-key add -
```

You might think it did nothing, because it immediately responds, "OK". But the error message goes away.

Q I have an ASUS router with a shared folder which I am able to connect to from Windows 10, and a few Mac clients, without any problems. When I attempt to connect from Ubuntu by clicking 'Connect to Server' and using the IP address of the router, similarly as I've done from Windows, I'm prompted for a username and password, and the Domain is filled in as 'WORKGROUP'. I have the correct username and password, but every time, I get no connection, and the dialog box for username/password keeps popping up.

The command: smbtree -d3 produces output which includes:

```
Server does not support
EXTENDED_SECURITY but 'client
use spnego = yes and 'client
ntlmv2 auth = yes'
```

A (Thanks to **bab1** in the Ubuntu Forums.) The client needs to degrade security as we can't update the ASUS server. Enter this command:

```
sudo gedit
/etc/samba/smb.conf
```

In the [global] section, add these lines:

```
client use spnego = no

client ntlmv2 auth = no
```

Then run your smbtree command again.

Q I would like to encode with mencoder stream from my /dev/video0 camera, and send it directly to mplayer for immediate watching.

A If you just want to watch video from the camera, use camorama or cheese.

Q Latest Kubuntu 16.04. A bit disappointed to find that Raw thumbnails in Dolphin are not working.

A (Thanks to **Rog131** in the Ubuntu Forums) Bug #1574049 kdegraphics-

thumbnailers 4:15.12.3-0ubuntu1 uses the wrong path for plugin libraries. The solution posted in Launchpad was to link where the plugins are, to where Dolphin expects them:

```
sudo ln -s /usr/lib/x86_64-
linux-gnu/plugins/*
/usr/lib/x86_64-linux-
gnu/qt5/plugins/
```

TOP QUESTIONS AT ASKUBUNTU

* How to block input to webcam?
http://goo.gl/txVRzh

* CPU temperature embedded in Bash command prompt
http://goo.gl/jY0JdQ

* How do I cite Ubuntu in an academic thesis?
http://goo.gl/eFuX2M

* Transplanting Drive
http://goo.gl/2yq8ot

Q&A

* Command to display an arbitrary message if a particular file exists
http://goo.gl/BWq4bq

* How to paste commands into a TTY?
http://goo.gl/tMzJD3

* How do I retrieve only the needed line from terminal output?
http://goo.gl/GI1FUZ

* Generate white noise to calm a baby
http://goo.gl/tNFOWW

* How do you protect Ubuntu login?
http://goo.gl/loml18

TIPS AND TECHNIQUES

SOLVING MY OWN PROBLEM

In 2009, I built a fairly high-performance PC, using components which were well reviewed on Newegg. It worked OK, not perfectly, with the Ubuntu of the day. That's common when you use the latest and greatest hardware with Linux.

It performed perfectly with Ubuntu 10.10 through 12.04. By that time, I was using Linux Mint 13, based on Ubuntu 12.04, dual-booting with various test versions of Linux.

After 12.04, things went poorly. Everything would work fine for a while, then my computer would completely lock up. It might happen after a day, or after a week. It really had nothing to do with Ubuntu or Mint, it was the Linux kernel. Eventually, I shifted my attention to Xubuntu, and used Xubuntu 15.10 in a dual-boot with Mint 13. Xubuntu would lock up two or three times a month.

When I started testing 16.04, it was much worse. It was rare that I could run for 24 hours without a lock up.

I had previously made half-hearted efforts to solve the problem, and found some things which didn't work. It became more urgent when Xubuntu 15.10 was approaching the end of support, and Mint 13 would also soon be unsupported.

I even tried using Android for PC, and it was no better. Same Linux kernel, same problem. I was starting to wonder if I would need to give up on Linux until I could build a new PC.

Then I found the magic Google search. My motherboard is a Gigabyte MA770T-UDP3. I Googled: MA770T-UDP3 linux

A couple of the search results took me to pages which talked about how to specify "acpi=off", and that was the solution.

After a fresh install of Xubuntu 16.04, I rebooted and paused when the GRUB menu appeared. I edited the Linux command line, to add "acpi=off".

When the system was running, I used sudo to edit /etc/default/grub, adding "acpi=off" to the line which contained "quiet splash". Then I ran: sudo update-grub

It worked. The only downside that I have seen is that the CPU always runs at 3.1 GHz; apparently CPU frequency scaling depends on acpi. The good news is, I can continue to use Linux on my high-performance PC.

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

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

Few game franchises have been as successful as Tomb Raider. The first Tomb Raider game was released back in 1996 for Microsoft Windows PCs, the original PlayStation, and the Sega Saturn. Ever since then, there have been at least 10 game titles in the main series with additional titles coming from spin-offs and alternate time-lines. In addition, there have been two box-office hit movies released, as well as an animated series, with a third film adaption in the works. The third film in the series will be based on the 2013 Tomb Raider reboot which focuses on Lara Croft's first adventure as a tomb raider.

Tomb Raider was developed by Crystal Dynamics, published by Square Enix, and released for Microsoft Windows, PlayStation 3 and Xbox 360 in 2013. In 2014, Tomb Raider was released for OS X, and, more recently, in April 2016 it was released for Linux. The Linux game was published by Feral Interactive, and, judging by the first 40 hours I've spent playing it, they've done a marvelous job at bringing it over to Linux. You can buy it direct from the Feral store, from Humble Bundle Store, or from Steam, for around $19.99 at the time of publishing.

The 2013 Tomb Raider reboot takes us back, way back, to Lara's very first expedition, before she was considered a tomb raider. It tell us the story of how Lara matures into the heroine we've always known her to be. The game begins with Lara and other explorers in a ship demolished by a storm near the Dragon's Triangle somewhere near Japan. As the ship is torn apart by the storm, Lara manages to cling to her life by barely making it to the island, but she's immediately captured by some strange man. When Lara awakes, she manages to escape, but not without first watching her
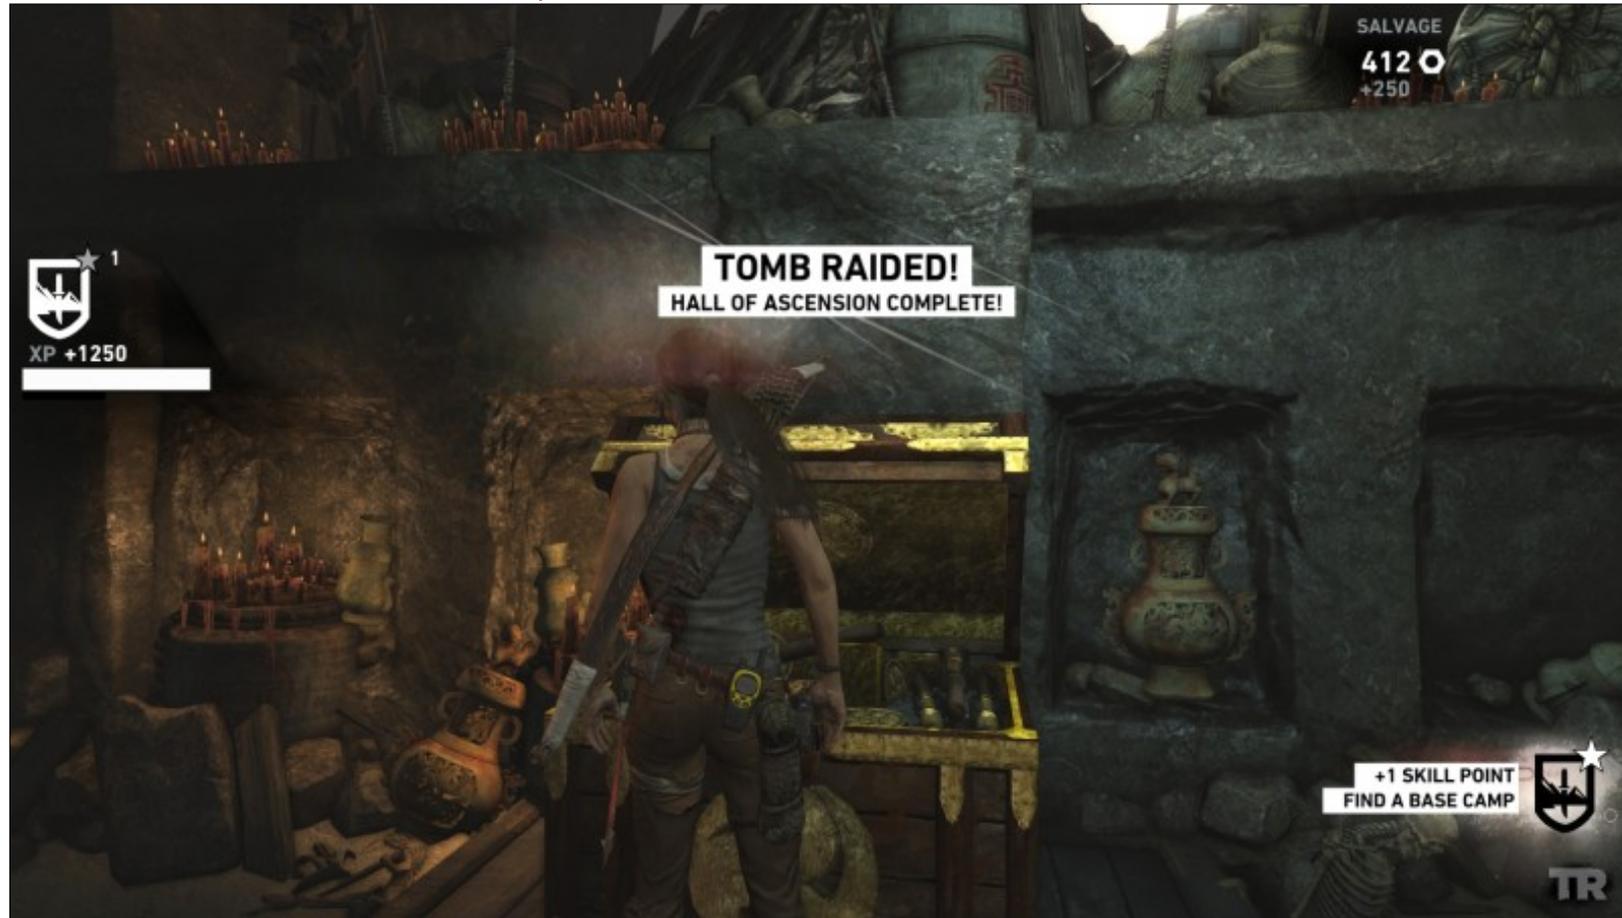
captor be killed by the crumbling cave where she was being held captive. From here on, Lara undergoes many transformations on her way to find out the truth about the fictitious Yamatai island, while, in the process, discovering many valuable archaeological artifacts, raiding her fair share of tombs and learning many valuable skills along the way.

One of the most memorable moments in the game is when Lara has to kill someone for her first time in order to survive. Having to push the button that pulls the trigger was somewhat of an emotional sequence even for me. This is the kind of game that Tomb Raider is; it sinks its teeth into your psyche while keeping your reflexes sharp, engaging you in some frightful experiences and torturing your brain into solving countless puzzles, all while keeping you at the edge of your seat.

Tomb Raider could be considered a third-person, action-adventure, exploration, platformer, puzzle and survival video game, with elements borrowed from other genres as well. So it's basically a well-crafted amalgam of a wide range of genres

– all rolled into one engaging & entertaining work of art that seems more like an action movie in which the main character is you, the player. Some sequences go from having Lara climb a mountain to a third-person shootout against the bad guys, then into a cinematic scene that culminates with Lara sliding down a mountain while a crashing airplane is hot on your tail devouring everything that lies on its path, yourself included – unless you manage to dodge it and all of the debris that the plane has dragged along with it on its furious path of destruction.

The game can be played with a mouse/keyboard or with a game

controller. I tried both and they both work well. My one and only complaint about Tomb Raider is the options menu placement of the default gamepad controller button-mapping layout. After playing the game for a while with a mouse/keyboard combination, I decided to try it out with a game controller and was able to adapt to it instantaneously. Having played some of the earlier franchise titles on the PlayStation family of consoles, it felt more natural for me to play it with a controller than with a mouse/keyboard. However, the problem came when I didn't know how to sprint (turns out you can't), and I had to look up the default gamepad button layout.

When you pause the game (by pressing the "Start" button), one of the choices on the menu screen is "Button Mapping," yet, when you go to it, it refers only to a mouse/keyboard layout. The default gamepad button layout is buried in the "Gameplay Options" under the "Gamepad" tab. It makes absolutely no sense why a game would have a button-mapping section, but include the actual button layout in a completely different section of the options menu. Once I found it, I knew where it was, and, from time to time, when I had to refer to it, I knew where to go. But, even then, it takes two steps too many to get to something that could have been located in the actual "Button Mapping" section along with the mouse/keyboard button layout. Other than this inconvenience, which is rather minor in the big scheme of things, there was nothing else crucial enough for me to give the game a lower grade. The game shines in all other aspects.

As mentioned earlier, an integral part of the game is its seamless blending of intense game-play sequences with cinematic cut-scenes which renders

UBUNTU GAMES

some gorgeous graphics. From the lush scenic environment to the seemingly infinite variations of light, to the immaculate detail in Lara's hair, the graphics are some of the best in any current generation game available on Linux. The sound is superb, especially its one-of-a-kind voice acting – which is on par with some of today's top movies. In fact,

special care, research and intense casting was required to find the right girl to play Lara in this game. Camilla Luddington is the actress who got the role for Tomb Raider as well as its sequel Rise of the Tomb Raider.

   All in all, Tomb Raider is perhaps the best addition to the Linux game ecosystem in the first half of

2016. Feral Interactive has done a great job at making the game feel right at home on Linux boxes. In my opinion, Feral Interactive are among the best Linux publishers since they began publishing games on Linux two years ago – Tomb Raider is proof of that.

⭐⭐⭐⭐✬

MINIMUM SYSTEM REQUIREMENTS:
OS: Ubuntu 14.04 or Steam OS 2.0 (64-bit)
CPU: Intel i3 or AMD FX-6300
Memory: 4GB RAM
HDD: 15 GB HD space
GPU: Nvidia GeForce 640 with 1GB of Video Memory, AMD R7 260X

MY GAMING BOX:
AMD FX-6100 3.3GHz CPU (over-clocked to 3.5GHz)
Gigabyte Windforce GeForce GTX 960 graphics card with Nvidia's 361 driver
16GB of Kingston Hyper X RAM
Ubuntu 14.04.4 LTS with Unity desktop

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

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
Chris Giltnane
Joao Cantinho Lopes

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
Max Catterwell
Frank Dinger
Paul Weed
Jaideep Tibrewala

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#112
Deadline:
Sunday 07th Aug 2016.
Release:
Friday 26th Aug 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.