Full Circle

Hard Drive

CPU

RAM

UPGRADING AN OLD LAPTOP

IS IT EVEN WORTH IT THESE DAYS?
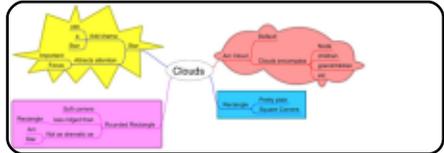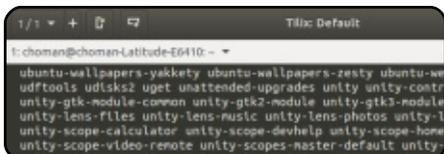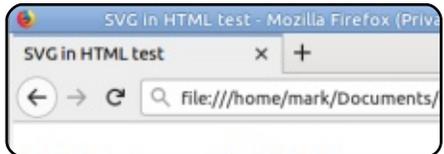
Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

EDITORIAL

Welcome to the latest issue of Full Circle.

We have more Python, Freeplane, and Inkscape for you this month. Due to personal reasons Miguel and his Ubuntu Touch series has had to miss this month, but hopefully he'll be back next month. Also no Great Cow Basic this month, but it has been replaced with an interesting article about upgrading with aptfast. It's a bit unconventional, but it is topical what with 18.04 (Bionic Beaver) being released last month.

This month's *My Opinion* gives cause to think: is it really worth it to upgrade a laptop these days? Not just software, but with hardware too. What are your thoughts on this? Feel free to email them to letters@fullcirclemagazine.org.

Elsewhere, we have a game review (Deus Ex) and a distro review (Neon). Neon uses the (KDE) Plasma 5 desktop.

Last but by no means least, don't forget that the FCM 2018 survey is still up: http://bit.ly/fcm2018.

All the best, and keep in touch!
Ronnie
ronnie@fullcirclemagazine.org
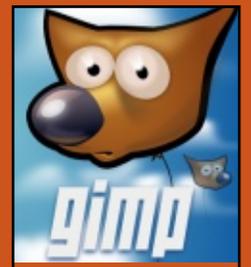
Find Full Circle on:

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/f
eed/podcast

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

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

FULL CIRCLE 2018 SURVEY

It's that time of the year again where we ask what you think of FCM, Ubuntu, and Linux.

Some questions are a requirement, some you can skip over if not applicable.

Your answers will help shape Full Circle, so please use your constructive criticism. If you don't tell us what you think, or what we're doing wrong, then we won't know.

Survey URL:
http://bit.ly/fcm2018

LINUX LITE 4.0 OS ENTERS BETA WITH NEW LOOK AND FEEL, BASED ON UBUNTU 18.04 LTS

Dubbed "Diamond," based on Ubuntu 18.04 LTS, and powered by the Linux 4.15 kernel, the Linux Lite 4.0 operating system enters beta stages of development today to give us a first glimpse of the upcoming release, which was slated for worldwide release on June 1, 2018.

According to the developer, Linux Lite 4.0's biggest changes are both internal and visual as the operating system comes with a brand new icon and system theme, namely Papirus and Adapta, Timeshift app by default for system backups, and new, in-house built Lite applications.

Among the new Lite apps, we can mention the Lite Desktop, which manages application icons and other objects on the desktop, and Lite Sounds, a tool designed to help users manage system-wide sounds. Also, Linux Lite 4.0 ships with the MenuLibre tool to help you easily edit application menu entries and Shotwell for basic image management.

Another major change in Linux Lite 4.0 is that 32-bit installations are no longer supported. The release would be available only for 64-bit computers, though the developer promises to offer support for existing Linux Lite 3.x installations until April 2021, when Canonical would end support for its Ubuntu 16.04 LTS (Xenial Xerus) operating system series.

Source:
https://news.softpedia.com/news/linux-lite-4-0-os-enters-beta-with-new-look-and-feel-based-on-ubuntu-18-04-lts-520892.shtml

KALI LINUX 2018.2 ETHICAL HACKING DISTRO IS HERE

Earlier this year in February, Offensive Security pushed the first release snapshot of Kali Linux for the year 2018. The Kali developers have already shifted their release model to rolling but they keep releasing these snapshots from time to time to provide a fresh ISOs to new users.

Following that tradition, Kali Linux 2018.2 has been pushed as the second release. 2018.2 is the first Kali ISO to be based on Linux kernel 4.15. It brings along the expected fixes for the Meltdown and Spectre flaws. So rest assured about these CPU flaws.

On the hardware support front, now the graphics cards of AMD are better supported. Moreover, AMD users can also take advantage of Secure Encryption Virtualization for better security and virtual memory encryption.

It goes without saying that Kali 2018.2 features lots of package updates to give you the latest of everything. The updated packages include Bloodhound, Reaver, PixieWPS, Burp Suite, Hashcat, etc.

A notable change comes in

Metasploit that makes it easier for you to access the scripts. Kali devs have included the links to all scripts in the PATH and they start with msf-.

Source: https://fossbytes.com/kali-linux-2018-2-download-iso-torrent-features/

POP!_OS 18.04 RELEASED — GET SYSTEM76'S BEAUTIFUL UBUNTU-BASED LINUX DISTRO

Out of the large crop of PC manufactures out there, System76 is one of the very few ones with a focus on Linux-based operating systems. Last year wh±en Ubuntu decided to ditch Unity development and switch back to GNOME, System76 chose to fork Ubuntu and create something new and beautiful.

Pop!_OS is the result of that decision, and its development is gaining pace. Last year, in an interview with Fossbytes, System76's CEO Carl Richell told that they created their products

with a focus on "software development, scientific computing, DevOps, and the maker movement – essentially STEM." Well, Pop!_OS isn't any different.

As Pop!_OS is based on Ubuntu, the release of Ubuntu 18.04 LTS was bound to be followed by Pop!_OS 18.04. System76 continues to refine the operating system and keep adding its own set of tweaks for a better experience. For this release, the installer has been refined and a good amount of effort has been put. The overall visual design and drive selection panel have been updated. It also supports custom partitions and LVM partitions. System76's own firmware notifications can now natively appear on operating system's notification system–both on Ubuntu and Pop!_OS.

The battery indicator has also been improved and now one can see the power status more clearly. A power profile picker has also been added to manage the settings in a better way. Now one can choose between High Performance, Balance, and Battery Life.

Source: https://fossbytes.com/pop_os-18-04-release-download-features/

FEDORA 28 "CUTTING EDGE" LINUX DISTRO RELEASED WITH NEW FEATURES

Red Hat-sponsored Fedora operating system is known for its bleeding edge features and leading the way for other Linux distros. Very often Fedora makes way for technologies that are often adopted by other distro developers. It's also Red Hat's testing ground as it acts as the upstream source of RHEL.

GNOME 3.28 is the default desktop environment shipping with this release. As a result, you get a wide variety of improvements and fixes. The most notable GNOME 3.28 features are updates to Files, Contacts, Clocks, Calendar, and on-screen keyboard apps.

Fedora 28 comes included with some popular 3rd-party software repos like Chrome, PyCharm,

Steam, etc. You're asked to enable them when you launch GNOME Software 3.28 for the first time.

On the hardware front, Thunderbolt 3 support has been added. This Fedora 28 feature will allow GNOME shell to automatically enroll Thunderbolt devices when they're connected.

This little change is going to make the lives of many Linux enthusiasts more comfortable–especially for those who love to hop distros and try out things on virtual machines. So, from Fedora 28, any installation of Fedora will have Virtual Box Guest Additions working out of the box. This will bring capabilities like auto desktop resizing, seamless mode, copy/paste b/w host and guest etc.

Source: https://fossbytes.com/fedora-28-download-features/

MAJOR DEBIAN LINUX KERNEL PATCH FIXES 8-YEAR-OLD PRIVILEGE ESCALATION FLAW

NEWS

The Debian Project released new major Linux kernel patches for the Debian GNU/Linux 8 "Jessie" and Debian GNU/Linux 9 "Stretch" operating system series to address a total of 27 security vulnerabilities, including an 8-year-old privilege escalation flaw.

First and foremost, the security update again patches Debian GNU/Linux's kernel against both variants of the Spectre vulnerability (CVE-2017-5715 and CVE-2017-5753). These could allow an attacker that has control over an unprivileged process to read memory from arbitrary addresses, including kernel memory.

While Spectre Variant 2 was mitigated for the x86 architecture (amd64 and i386) via the retpoline compiler feature, Spectre Variant 1 was mitigated by first identifying the vulnerable code sections and then replacing the array access with the speculation-safe array_index_nospec() function.

Another important bug (CVE-2018-8781) patched with these new kernel updates for Debian GNU/Linux is a recently unearthed privilege escalation flaw that was introduced in the Linux kernel no less than eight years ago. It affected udl (DisplayLink) driver's mmap operation, allowing a local attacker with access to a udl framebuffer device to gain root access by overwriting kernel memory.

Source:
https://news.softpedia.com/news/major-debian-linux-kernel-patch-fixes-8-year-old-privilege-escalation-flaw-520942.shtml

Linux Apps Are Coming to Chromebooks and You Can Try Them Right Now

Rumored to come to a Chromebook near you, support for Linux apps just entered beta testing in the latest Chrome OS Dev channel, as confirmed by Kevin Tofel of About Chromebooks. The feature appeared in the Settings and needed to be turned on if you want to use Linux tools, editors, and IDEs on your Chromebook.

The first sign of Linux app support in Chrome OS appeared two months ago when a Reddit user discovered a Chromium Gerrit commit explaining a new device policy designed to allow containerized Linux apps to Chromebooks. Then, someone discovered a Terminal app, whichfailed to install, but suggested upcoming support for Linux apps.

Now that the feature is finally accessible to the masses, you can try Linux apps on your Chromebook by switching to the Chrome OS Dev channel. To do that, access the Chrome OS Settings, scroll down on the left side panel list and click on "About Chrome OS," then click on "Detailed build information," click "Change channel" next to "Channel," and choose the Developer channel.

Click "Change Channel" to change the channel and apply the settings. Chrome OS will then automatically download an update to your Chromebook and ask you to restart. Once you're back on the desktop, open the Chrome OS Settings again, go to "About Chrome OS" and make sure you're running at least Chrome OS 68.0.3416.0. Go back and turn on Linux (Beta) feature under the Linux Apps section.

Source:
https://news.softpedia.com/news/linux-apps-are-coming-to-chromebooks-and-you-can-try-them-now-in-chrome-os-dev-520966.shtml

Ubuntu MATE And Ubuntu Budgie Dropping 32-bit Hardware Support

Many leading Linux distros like Ubuntu, Arch Linux, Manjaro, etc., have already ditched the support for 32-bit architecture and decided to focus on 64-bit machines. Now, following their footsteps, Ubuntu Budgie and Ubuntu MATE have also joined the league.

The main reason behind these moves is the decreasing number of users actually running these operating systems on 32-bit machines. That's why it doesn't make a lot of sense to keep allocating resources and time to the hardware that people don't use anymore.

Ubuntu's flagship edition

decided to make this change from 17.10; Ubuntu Budgie and Ubuntu MATE will be doing the same from the next 18.10 release.

MATE team has said that this change will let save time and offer better support to ARM devices like Raspberry Pi. The development teams also have limited access to 32-bit hardware, which doesn't ensure a rigorous testing.

Don't worry. If you're still stuck with your older hardware and looking to get some work done, you're all set for the next 3 years. As this change will take place from 18.10, the 32-bit 18.04 images will continue to get support for the next 3 years, i.e., April 2021.

Source:
https://fossbytes.com/ubuntu-mate-ubuntu-budgie-drop-32-bit-support/

KAOS KDE-FOCUSED ROLLING LINUX DISTRO CELEBRATES 5TH BIRTHDAY WITH UPDATED ISO

When a person celebrates a birthday, they often receive gifts, eat cake, and spend time with friends and family. After all, it is intended to be a happy occasion as the person marks another year of life.

But what the heck does a Linux distribution do to celebrate a birthday? In the case of the excellent KaOS operating system, the answer is simple -- release an updated ISO. Yes, as a way to mark the 5th birthday of the KDE-focused distro, version 2018.04 is now available for download. You get Plasma 5.12.4, Linux kernel 4.15.7, Libreoffice 6.0.3, and more.

"It is five years ago this month that KaOS started, a nice way to commemorate is with releasing 2018.04. This ISO has a complete redesign of the Midna theme for 2018. Some 2,500 new icons in use, rewritten sddm login theme and a KaOS community selected new wallpaper (created by Jomada)," says The KaOS Team.

The team further says, "Also new is KaOS' creation Croeso (Welsh for welcome) for helping with configuring a new install. It

will run on the newly installed system and offers to adjust some 15 commonly used settings and replaces the formerly used, PyQt based first run wizard Kaptan. It also includes a custom Wallpaper selector, distribution info, and news. It is written in QML and fits well with the Welcome application used in the Live system. The latter now includes a fully rewritten (also in QML) Installation Guide."

Source:
https://betanews.com/2018/05/04/kaos-kde-linux-5th-birthday-iso/

UBUNTU 18.10 OPERATING SYSTEM DUBBED "COSMIC CUTTLEFISH" BY MARK SHUTTLEWORTH

We already knew for about a week or so, when the first daily build ISO images popped-up on the main servers, that Ubuntu 18.10's codename will start with Cosmic, but we did not new the animal mascot it will be accompanied by. So today Mark Shuttleworth announced the full codename of Ubuntu 18.10 as the "Cosmic Cuttlefish".

According to Mark Shuttleworth, Ubuntu 18.10 will focus more on the overall security of the Linux-based operating system. Not that Ubuntu or any other GNU/Linux distribution out there is not already very secure by design, thanks to the powerful Linux kernel, but Shuttleworth wants to raise the bar even higher on security for the Ubuntu 18.10 release.

Apart from the on-going security enhancements, Ubuntu 18.10 will also come with various other novelties, including the upcoming GNOME 3.30 desktop environment by default, a new system theme by default, called communitheme, and all the latest GNU/Linux technologies like GCC, Linux kernel, systemd, Wayland, X.Org Server, Mesa Graphics Stack, etc.

However, Ubuntu 18.10 will be a testbed for future LTS (Long Term Support) versions on the operating system, which means that it will be supported with security and software updates for only nine months.

Source:

https://news.softpedia.com/news/ubuntu-18-10-operating-system-dubbed-cosmic-cuttlefish-by-mark-shuttleworth-521029.shtml

MALICIOUS PACKAGE FOUND ON THE UBUNTU SNAP STORE

An attentive Ubuntu user has spotted today a cryptocurrency miner hidden in the source code of an Ubuntu snap package hosted on the official Ubuntu Snap Store.

The app's name is 2048buntu, a clone of the popular 2024 game, packaged as an Ubuntu snap —a relatively new app format for Ubuntu OS.

According to a GitHub user named Tarwirdur, the app contained a cryptocurrency mining application disguised as the "systemd" daemon, along with an init script that provided boot persistence.

The code mined the Bytecoin (BCN) cryptocurrency for a user account with the email address "myfirstferrari@protonmail.com."

Following the report, the Ubuntu Snap Store team removed the app, along with the app developer's other snap packages "pending further investigations."

The Ubuntu Snap Store does not provide an install count, so the number of affected users is unknown.

Anyone can create and submit a snap package to the Ubuntu Snap Store. Submitted snaps do not go through a security check, similar to the apps submitted to the iOS App Store and Google Play Store.

With this incident, the Ubuntu Snap Store takes its place in the pantheon of app stores that have suffered malware infections, right next to the Chrome Web Store, the Google Play Store, the Apple App Store, and the Windows Store.

Source:
https://www.bleepingcomputer.com/news/linux/malicious-package-found-on-the-ubuntu-snap-store/

FEDORA 28 IS OUT AND DEVELOPERS ARE IN LOVE WITH IT

Fedora, the love of many Linux enthusiasts just got an upgrade and brings a lot of new features. Fedora is not something a novice Linux user uses, if you are one, try out Ubuntu or Linux Mint as they are simpler and have a lot of resources to help you if you are stuck somewhere. But, if you are a serious Linux user, you would have tried Fedora at some point.

Fedora is a Linux distribution developed by the Fedora Project which is community-supported and sponsored by the Red Hat. The latest version of Fedora uses Linux Kernel 4.16.3. The latest Fedora28 comes in three different editions for a different use- cases: Fedora 26 Workstation, Fedora 28 Server, and Fedora 28 Atomic Host.

All the mentioned versions are basically built from a common set of base packages and then extended to serve different purposes. The new packages come with a lot of bug fixes and performance tweaks along with some new features. The latest

Fedora28 base package comes with all the latest compilers including GNU Compiler Collection(GCC) 8, Ruby 2.5 and Golang 1.10.

The latest Fedora version tries to solve a very common problem most programmers face. It brings the modular repository feature which lets the developer use alternative software versions rather than the default package. This means, if you are a Ruby developer, you can use different versions of Ruby in different projects and maintain a good workflow without needing to change much.

Fedora is making this possible with their new optional repository Modular or the "Application Stream" or AppStream in short. This software ships the software versions on independent life cycles. This helps in choosing the right version of the repository and still keep the operating system up-to-date.

This is done via a new optional repository Modular, aka the "Application Stream" or AppStream for short.

CentOS 7 1804 Linux Distro Available For Download

While making a list of free operating systems that can be used both as a daily use system as well as a server, CentOS gets an early mention. Based on RHEL base, CentOS is known for being a stable and manageable platform. Just recently, the developers have shipped the sixth CentOS-7 release.

It makes CentOS 7 1804 the most current release, which based on Red Hat Enterprise Linux 7.5 released on April 4th, 2018.

This release comes with tons of changes to cryptographic abilities of different packages. As a result, many outdated and insecure algorithms and protocols are deprecated.

Critical PGP Flaws Can Expose Encrypted Emails In "Plaintext" — Disable It Right Now

Email users who use PGP (based on OpenPGP) and S/MIME to encrypt and decrypt their communications are at "immediate risk." The reason is that a team of European researchers has found critical flaws in the encryption standards and currently there are no fixes available.

The vulnerabilities dubbed EFAIL are harmful as they can reveal the contents of messages in plain text, even for the messages from the past.

If an attacker gains access to a victim's encrypted emails through methods like eavesdropping or compromising email accounts, EFAIL can be used to "abuse active content of HTML emails, for example externally loaded images or styles, to exfiltrate plaintext through URLs," reads the website detailing the vulnerabilities.

A modified encrypted email sent by the attacker to the victim is decrypted by their email client. While doing so, the client loads any external content, thus, exfiltrating the plaintext to the attacker.

The PGP encryption is mostly used by political activists, journalists, and whistleblowers as an extra layer of encryption. On the other hand, S/MIME is used mainly in enterprise infrastructure.

Why this should be taken seriously is because the Electronic Frontier Foundation (EFF) is also spreading the word. The Foundation which has been in communication with the researchers has advised users to "temporarily stop sending and especially reading PGP-encrypted email".

Users should immediately disable or remove any tools that automatically decrypt PGP-encrypted emails until the flaws are understood and fixed, EFF said. They have published guides for Thunderbird, Apple Mail, and Outlook.

Opera Touch is a dream Android browser for users who are always on the go

If your life is hectic, you probably have become very familiar with using your Android smartphone with one hand. You're power-walking through the company campus, or down the sidewalk on your way to work or a meeting and you only have one hand to use to navigate with your mobile device. For some, that may not be a problem. For most, however, using the phone with one hand can be an exercise in clumsiness. So when Opera announced the release of their Opera Touch browser, one that could easily be used with one hand, my curiosity was certainly piqued. After using the browser for a while, I've discovered that Opera has hit that particular nail square on its head. Opera Touch makes one-handed usage not only possible for all, but easy.

FULL CIRCLE 2018 SURVEY

It's that time of the year again where we ask what you think of FCM, Ubuntu, and Linux.

Some questions are a requirement, some you can skip over if not applicable.

Your answers will help shape Full Circle, so please use your constructive criticism. If you don't tell us what you think, or what we're doing wrong, then we won't know.

Survey URL:
http://bit.ly/fcm2018

For those who frequently find themselves using their mobile devices with one hand, Opera Touch might well be the best laid out mobile browser for this purpose. Combine that with the Flow feature and Opera is seriously onto something. Give Opera Touch a go and see if it doesn't wind up your default Android browser.

Source:
https://www.techrepublic.com/article/opera-touch-is-a-dream-android-browser-for-users-who-are-always-on-the-go/

UBUNTU 18.10 COSMIC CUTTLEFISH MAY INCLUDE OUT-OF-THE-BOX ANDROID INTEGRATION

If you're reading this article, there are great chances that you might be owning a computer as well as a smartphone. In such case, very often we feel the need for a seamless integration between both the devices. While Apple is known for how well its devices work with each other, Microsoft has also been making similar attempts.

In an encouraging development for Linux enthusiasts, Canonical has expressed interest in shipping Ubuntu 18.10 with out-of-the-box Android integration, thanks to GS Connect—a GNOME Shell implementation of KDE Connect without KDE/Qt dependencies.

Just like KDE Connect, you can use GS Connect to get Android notifications on your Ubuntu computer, share data, get device information, browse the file manager, etc.

Canonical has also invited the developers from the community to contribute in the process.

The development work for this integration is expected to take place the way Ubuntu Dock and Dash-to-Dock were added to the Ubuntu desktop. It means that the developers won't be forking the project; they'll be branching the upstream. Ubuntu's Will Cooke has also hoped that someone would work on an iOS port of the KDE/GS Connect mobile app.

Meanwhile, before this feature lands in Ubuntu 18.10 Cosmic Cuttlefish, you can use the GS Connect app in Bionic Cuttlefish and control your Android phone from Ubuntu desktop.

Source:
https://fossbytes.com/ubuntu-18-10-cosmic-cuttlefish-android-integration-gs-connect/

CANONICAL SAYS THERE'S NO RULES AGAINST MINING CRYPTOCURRENCIES THROUGH SNAPS

Last week, users discovered that two of the snap packages uploaded by user Nicolas Tomb in the Snap Store, namely 2048buntu and Hextris, mined cryptocurrency in the background while the applications were running without user's knowledge. Canonical immediately removed the apps from its Snap Store.

Now, the company behind the popular Ubuntu Linux operating system is addressing the issue saying it has no rules against mining cryptocurrencies through snap apps if the developer informs users about this. As Nicolas Tomb didn't inform users that his apps are mining for cryptocurrencies,

the apps were removed.

Canonical also said that mining cryptocurrency is not illegal, nor unethical, so the only thing the publisher did wrong is not to inform users about his snap apps mining for cryptocurrency in the background. Nicolas Tomb informed Canonical that his goal was to "monetize software published under licenses that allow it."

In the lengthy blog post, Canonical explains that it doesn't have the manpower to review hundreds of thousands of incoming source code lines from snap packages published in its Snap Store every single day. Therefore, it urges users to install apps only from trusted sources and developers.

With that in mind, the company promises to enforce the security of its Snap Store by implementing the ability to flag specific publishers as verified, helping users with their decision to install a certain snap from the Snap Store, which currently contains more than 3,000 packages for open source and closed source apps.

Source:
https://news.softpedia.com/news/canonical-says-there-s-no-rules-against-mining-cryptocurrencies-through-snaps-521132.shtml

ENDLESS OS 3.4 RELEASED WITH NEW FEATURES, LINUX 4.15, AND PHONE COMPANION APP

Founded in 2011, Endless Mobile, Inc. develops Linux-based Endless OS and hardware running the same. The firm has recently shipped Endless OS 3.4, the latest and major release of the operating system.

As it's a big feature release, it goes without saying that Endless OS 3.4 comes with many changes and new features. Thanks to an updated open source core, 3.4 also ships with stability and performance improvements.

That being said, let's tell you about the major new features shipping with Endless OS 3.4.

This release has gained the functionality of detecting what kind of internet connection one is running; it'll let you make sure that your system doesn't try to suck tons of data on a limited data connection. In case Endless isn't able to detect the type of connection, there's a toggle to mark it manually.

A special new feature for Android users is also coming to Endless OS. Called Endless Companion App, it'll let you browse the content of your phone with ease and get notifications.

In other improvements, there's GNOME 3.26, Linux 4.15 kernel, polished applications, latest LibreOffice, etc.

Source:
https://fossbytes.com/endless-os-3-4-released-new-features-download/

CRITICAL LINUX FLAW OPENS THE DOOR TO FULL ROOT ACCESS

Red Hat has patched a vulnerability affecting the DHCP client packages that shipped with Red Hat Enterprise Linux 6 and 7. A successful exploit could give an attacker root access and full control over enterprise endpoints.

According to an alert issued Wednesday from US-CERT, the critical-rated flaw, first reported by Google researcher Felix Wilhelm, would "allow attackers to use malicious DHCP server responses to execute arbitrary commands on target systems over the local network," if those systems use NetworkManager and are configured to obtain dynamic IP addresses.

An attack would take advantage of the way the DHCP protocol is used to dynamically assign IP addresses to computers; i.e., the fact that the endpoint clients essentially broadcast out a query over the local network to obtain an address from a DHCP server.

An attacker may be able to compromise the legitimate DHCP server itself in order to be able to send out the payload in the first place, Ayer added; or, he or she could set up a fake node on the network to masquerade as a legitimate DHCP server, sending out malicious, spoofed responses

to normal network addressing queries. In both cases, the bad actor would need to be attached to the same local-area network as the targeted systems in order to exploit the flaw.

Source: https://threatpost.com/critical-linux-flaw-opens-the-door-to-full-root-access/132034/

BODHI LINUX 5.0 ENTERS DEVELOPMENT BASED ON UBUNTU 18.04 LTS, FIRST ALPHA IS OUT

Now that Canonical released Ubuntu 18.04 LTS (Bionic Beaver), more and more Ubuntu-based GNU/Linux distributions would want to upgrade to it for their next major releases, including Bodhi Linux with the upcoming 5.0 series. The first Alpha is here today to give us a glimpse of what to expect from the final release.

Besides being based on Ubuntu 18.04 LTS, the Bodhi Linux 5.0 operating system will be shipping with the forthcoming Moksha 0.3.0 desktop environment based on the

Enlightenment window manager/desktop environment, and it's powered by the Linux 4.9 kernel series. Also, it supports 32-bit PAE and non-PAE systems.

During the development cycle of Bodhi Linux 5.0, the development team will focus mainly on integrating the in-house built Moksha 0.3.0 desktop environment with the Ubuntu 18.04 LTS base. That's why they don't recommend users to upgrade to Bodhi Linux 5.0 when its out later this summer, especially that Bodhi Linux 4.5.0 is supported until April 2021.

Source: https://news.softpedia.com/news/bodhi-linux-5-0-enters-development-based-on-ubuntu-18-04-lts-first-alpha-is-out-521219.shtml

HANDS-ON WITH FIRST LUBUNTU 18.10 BUILD FEATURING THE LXQT DESKTOP BY DEFAULT

Lubuntu developer Simon Quigley approached us earlier

today to inform that the latest Lubuntu 18.10 daily build is quite usable as he and his team did a lot of work in the past week to accommodate the LXQt desktop environment by default instead of the LXDE desktop environment.

The main difference between LXDE and LXQt is that the former is written with the GTK+ 2 technologies, which will eventually be phased out in favor of the more advanced GTK+ 3, and the latter is built using the Qt framework. However, it doesn't look like there are any plans for LXDE to move to GTK+ 3.

That's why LXQt was created in the first place, as a fork of LXDE that offers computer users a classic desktop environment with a modern look and feel, and supported for a long time thanks to the Qt technologies, on which the popular KDE Plasma desktop environment is built.

Long story short, the first Lubuntu release with the LXQt desktop environment by default will be Lubuntu 18.10, due for release on October 18, 2018, as part of the Ubuntu 18.10 (Cosmic Cuttlefish) series. And the first

daily build is now available for testing on both 64-bit and 32-bit architectures.

Source: https://news.softpedia.com/news/hands-on-with-first-lubuntu-18-10-build-featuring-the-lxqt-desktop-by-default-521200.shtml

500,000 ROUTERS IN 54 COUNTRIES HACKED TO CREATE MASSIVE BOTNET ARMY

The Cisco security research team has unearthed a hacking attack that targets more than 500,000 consumer Wi-Fi routers. These infected routers can be used to create a massive botnet army and launch a heavyweight cyberattack.

As per the findings, this attack seems to be a work of some state-sponsored actor. The malware used to infect the devices, in this case, has been termed VPNFilter. As the malware shares lots of code with the BlackEnergy malware used to deploy large-scale attacks on Ukraine, it could be tied to the

Russian government.

It's worth noting that the Cisco report didn't directly mention Russia but the code overlap surely hints the same. In a related development, the FBI has seized an important server used in the attack. As per the agents, the server, ToKnowAll.com, was being used by Russian hackers to spread a second stage malware attack.

Coming back to currently infected routers, the devices belong to major companies, including TP-Link, NETGEAR, Linksys, and MikroTik.

The VPNFilter malware responsible for the attack is particularly concerning as it contains code to steal website credentials and make the infected router unusable. Moreover, it has "the potential of cutting off internet access for hundreds of thousands of victims worldwide."

This multistage, modular platform malware persists through a reboot in its initial stage. In the second stage, it uses different command and control mechanisms to find the IP address stage 2 deployment server and proceeds with the intelligence-collection process; it also has a self-destruct capability. The stage 3 modules further acts as plugins for stage 2 malware.

Source: https://fossbytes.com/routers-hacked-botnet-army-vpnkill-malware/

PARROT 4.0 ETHICAL HACKING LINUX DISTRO RELEASED

It goes without saying that Parrot 4.0 ships with all the fixes and updated packages released since 3.11 apart from the new features. This release also took a lot more time as compared to previous versions.

Starting with 4.0, the developers have decided to ship netinstall images. It'll enable the interested security researchers to create their customized working environment with only the bare core and the software they need.

This version ships with Linux kernel 4.16, which includes tons of fixes, optimizations, and better hardware support.

The official Parrot Docker templates have also been shipped. With Parrot on Docker, one can access the system on top of any work environment that supports Docker. Apart from Docker on Parrot, different images with bare system, complete environment with tools, and Metasploit container has also been made available.

In other changes, the sandboxed applications have been made more stable for better security; nginx is the new default web server daemon; md raid support has been added by default; LibreOffice 6, Firefox 60, and MATE 1.20 have been added.

Source: https://fossbytes.com/parrot-4-0-release-download-features-hacking-linux/

LIBREOFFICE 6.1 BETA ARRIVES NEXT WEEK FOR SECOND BUG HUNTING SESSION ON MAY 28

Now that the first bug hunting session, which took place last month on April 27 for the alpha milestone, was a success leading to 91 bugs (8 of them marked as critical and 4 already fixed) being reported by those who attended the event, it's time for a second bug hunting session at the end of May to discover and squash more of those pesky bugs and issues that may block the release of LibreOffice 6.1.

As mentioned before, the tests that will be performed during the second bug hunting session to find, report and triage bugs will be carried against the beta version of LibreOffice 6.1, which should be available for download on all supported platforms, including GNU/Linux, macOS, and Microsoft Windows in the coming days. Users will be able to run LibreOffice 6.1 beta alongside LibreOffice 6.0.

Once the beta is out, the Document Foundation could release a second beta throughout June if more bugs are discovered on May 28, and plan a third bug hunting session, which could be announced at a later date. According to the release schedule,

at least three RC (Release Candidate) builds are planned throughout July, and the final release of the LibreOffice 6.1 office suite is expected in early August.

Source:
https://news.softpedia.com/news/libreoffice-6-1-beta-arrives-next-week-for-second-bug-hunting-session-on-may-28-521256.shtml

Ubuntu 18.04 LTS Gets First Kernel Update with Patch for Spectre Variant 4 Flaw

Canonical released the first kernel security update for its Ubuntu 18.04 LTS (Bionic Beaver) operating system to fix a security issue that affects this release of Ubuntu and its derivatives.

As you can imagine, the kernel security update patches the Ubuntu 18.04 LTS (Bionic Beaver) operating system against the recently disclosed Speculative Store Buffer Bypass (SSBB) side-channel vulnerability, also known as Spectre Variant 4 or CVE-2018-3639, which could let a local attacker expose sensitive information in vulnerable systems.

Canonical also notes the fact that to fully mitigate Spectre Variant 4, users must also update the processor microcode firmware. They've been working with Intel to offer future microcode updates that fully address Spectre Variant 4, but AMD users will need to contact their vendor for firmware updates. To enable the kernel mitigations for Spectre Variant 4, the SSBD (Speculative Store Bypass Disable) feature is required on i386 (32-bit) and amd64 (64-bit) machines, which will be implemented by BIOS vendors with future updates. Therefore, users must also ensure they're using the latest BIOS version on their computers.

Canonical is urging all Ubuntu 18.04 LTS (Bionic Beaver) users to update their systems as soon as possible and install the new kernel version is made available for 32-bit and 64-bit architectures, Amazon Web Services (AWS), Microsoft Azure Cloud, and Google Cloud Platform (GCP) systems, OEM processors, and cloud environments.

Source:
https://news.softpedia.com/news/ubuntu-18-04-lts-gets-first-kernel-update-with-patch-for-spectre-variant-4-flaw-521248.shtml

FULL CIRCLE 2018 SURVEY

It's that time of the year again where we ask what you think of FCM, Ubuntu, and Linux.

Some questions are a requirement, some you can skip over if not applicable.

Your answers will help shape Full Circle, so please use your constructive criticism. If you don't tell us what you think, or what we're doing wrong, then we won't know.

Survey URL:
http://bit.ly/fcm2018

I've recently begun reading up on machine learning, due in large part to recent advances and news articles. Specifically, recent reports on Leela Zero (Go AI), combined with the AlphaGo documentary I recently watched on Netflix. So, for anyone who may be interested in this topic, I wanted to cover a few basic points. If you're not sure what machine learning is - feel free to also read on!

WHAT IS IT?

Machine learning is a term that is generally associated with artificial intelligence. While machine learning is used in creating artificial intelligence, it is not the same thing. Machine learning is the approach to teaching a computer how to accomplish a task without explicitly programming it. Personally, I think of it as teaching a computer to recognize patterns and extrapolate conclusions, which works fine for a basic understanding.

Here's an example - say you want to create an application capable of recognizing ingredients for a recipe from a single image. Programmatically, you'd have to define all the visual variables and somehow write a program to classify the key points properly and reach a conclusion. However, using machine learning, you would instead teach it by showing it images - various pictures of flour, eggs, butter, etc. By having control over your sample size (and correctly identifying them in advance), you can train your application to recognize what the images represent. Once the program understands the individual parts, you could train it to then recognize them in larger images.

IS THIS AI?

It depends on your definition of AI. In the Sci-Fi world, AI is often portrayed as (essentially) a super intelligent digital being (who often concludes that the human race is the cause of any number of problems). However, machine learning at its core only develops so-called "dumb" AI. This is essentially a piece of software that is extremely good at doing one task - in the example above, recognizing ingredients in photos. So far, any kind of AI we have or use falls into this category. And in the case of machine learning, I would make the case that not all applications are automatically AI - for me, AI implies a level of intelligence beyond "If A do B", but more reasoning along the lines "If A is true, then perhaps Y or Z would be better than B".

WHO NEEDS AI/MACHINE LEARNING?

Machine learning is generally applied in fields of study where explicit programming would fail due to the sheer amount of information and analysis the programmer would have to account for. Specifically, things like facial recognition (such as on Facebook or Google Photos), or recommendation engines (such as Amazon's recommended products).

These both deal with massive amounts of data that would need to be continually updated (for every new face, for example). Instead, we use machine learning to allow a program to recognize the same face (or pet), or to find similar products in a growing assortment.

ISN'T THIS A NEW TECHNOLOGY?

The term "machine learning" was actually coined in 1959. And the concept, as such, has been around just as long. Until recently, however, the average computer was not up to the task and so research and availability of machine learning was severely limited. With the increase in computing power and the fall in costs, this has allowed more people to begin working with these systems.

HOW DOES IT WORK?

There are new techniques and approaches that we have started

to use. The most popular/common one seems to currently be the idea of a neural network, and is also what all of my examples typically describe. A neural network is very loosely inspired by biological neural networks (ie, the human brain), and is an interconnected group of neurons that have a specific function. Take the example of recognizing specifics of pets - the machine would learn aspects from a set of manually labelled images. And each neuron may be in charge of a defining trait (ie, does it have fur?). When an unknown image is input, the system then "fires" each of these neurons, and therefore eliminates possibilities until it has a reasonable chance at guessing correctly.

There are other approaches - decision trees, association rule learning, etc. But in most cases, the basic approach is the same - give the system a set of prepared data to train it, and then begin testing unlabeled data.

What about AlphaGo and Leela Zero?

Both of these AIs use something called deep learning -

this is a step further from machine learning. While machine learning is dedicated to teaching systems to recognize specific patterns or accomplish specific tasks, deep learning is more abstract. Take AlphaGo for example - it was taught the rules of Go and was fed basic data from a few Go records. However, it then began playing against itself and learning from those encounters. As such, it didn't have a specific goal in mind (beyond "play Go"). This application of pitting two versions of the system against each other helped AlphaGo learn the uncommon strategies it employed against Go professionals.

It also approaches the game with a fundamentally different method than human players. It is very certain in its ability to count the points on the board. As such, while human players may use moves to gain a substantial lead in points due to their uncertainty in counting, AlphaGo will instead focus on the percentage chance a move has to win the game - regardless of whether it wins by one point or one hundred. So while a human player may attempt a desperate move to increase their lead, it may have a lower chance of succeeding, while a safer move that keeps its lead the same will be preferred by AlphaGo.

Can we learn from AI?

Yes! AlphaGo released a teaching tool a while back, that essentially shows the decisions for each move it makes (represented as percentages). In doing so, people have discovered that AlphaGo believes the komi (the points white receives for going second in a game of Go which has steadily increased over the years) currently puts black at a disadvantage from the beginning of the game (the opening moves are all under 50% chance of winning). This has led professionals to re-examine some of the assumptions about the game. It has also popularized a few moves that were previously deemed "bad" and led to new styles. It has also massively increased the popularity of Go around the world (after AlphaGo played against Lee Sedol there was a worldwide shortage of Go boards).

How can I get into machine learning?

This depends on your current knowledge. If you're comfortable in Python, the Google Developers YouTube channel has a good

introduction to machine learning. A great deal of languages have machine learning packages, and so you should look for a tutorial in your language of choice. To get started, you shouldn't need any hardware outside of a relatively modern computer.

DISCLAIMER

I feel that (in this type of article) I should clearly state that I am not an expert, and all details found here are my best understanding of the available facts and information I've found. As such, there may be mistakes.

I hope this article has inspired you (or informed you). If anyone has noticed any issues or corrections (or has a great resource for learning), feel free to email them to me. Similarly, if you have requests for articles or feedback for this one, you can reach me at lswest34+fcm@gmail.com.

FURTHER INFORMATION

https://youtu.be/lvzekeBQsSo - Part 2 of a series. The video creates a recommendation engine for color schemes in the web.

https://youtu.be/cKxRvEZd3Mw - Part 1 of the Google Developer series (Python).

https://www.tensorflow.org/ - An open source machine learning framework

https://www.alphagomovie.com/ - the AlphaGo documentary (available on Netflix).

https://github.com/gcp/leela-zero - Leela Zero

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.

L ast month, we took a brief look at MQTT.

This month, we'll continue our discussion of MQTT, especially how MQTT can support quality of service and then we'll modify our original program from Full Circle issue #109 to support MQTT.

QoS

QoS stands for Quality of Service. MQTT supports three levels, numbered 0, 1 and 2. If you don't have to be sure that those who are subscribed to your topic messages, then you simply use a QoS of 0. Here is a breakdown of the levels:
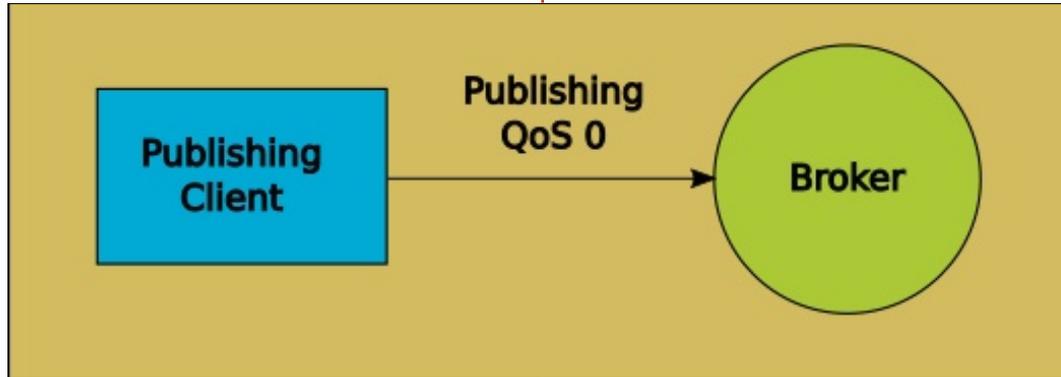
0 - At most once
1 - At Least once
2 - Exactly once

When the broker transfers a message to a subscribing client, it uses the QoS of the subscription made by the receiving client

LEVEL 0

Level 0 guarantees a best effort delivery. The message won't be acknowledged by the receiver (broker or client), nor will it be stored and re-delivered by the sender.

LEVEL 1

Using Level 1 is guaranteed that a message will be delivered at least once to the receiver (publisher to broker or broker to subscriber). The sender stores the message until it receives a PUBACK message from the receiver. Note that the message can be sent multiple times to the receiver until the PUBACK is sent back to the sender.

LEVEL 2

As you can see from the diagram, the client sends a publish message to the broker and the broker sends back a PUBREC. The sender then stores a reference to the packet identifier, which is then stored and then sends out a PUBREL message. If the receiver responds with a PUBCOMP message, the sender then knows it can safely discard the message and reference. If a message gets lost somewhere along the way, the sender (either the publishing client or the broker) is responsible for re-

sending the last message.
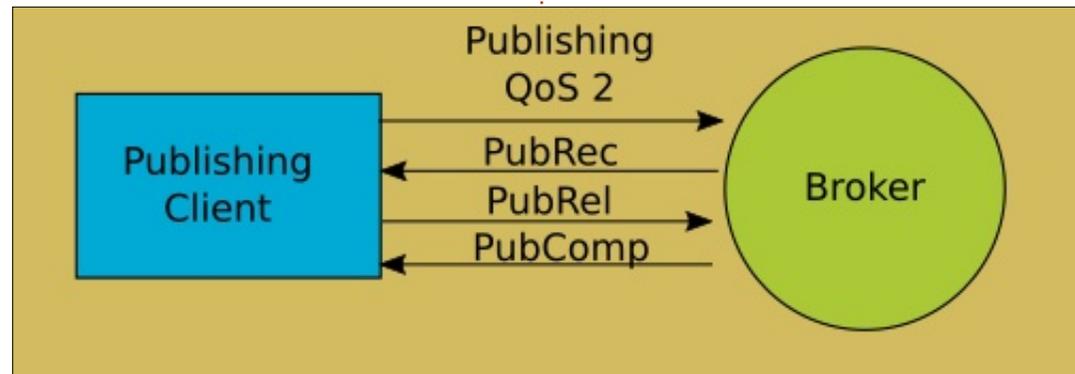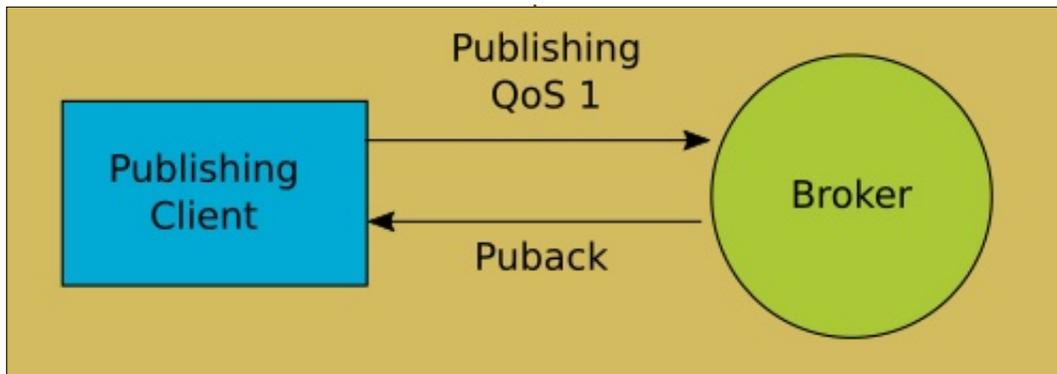
Of course, QoS Level 1 is faster than Level 2, since there are only two messages to deal with where Level 2 has four and QoS Level 0 is faster than Level 1.

When the broker sends a message to a subscribing client, the QoS used is what the client used when subscribing to the topic. Therefore, the QoS can be downgraded from what the publishing client intends.

When should you use one level over another? Basically if your publishing client doesn't need to worry if there might be a missed transmission or two, and your connection to the broker is reliable 24/7, you could pretty much feel safe to use Qos Level 0. If you need to have a guarantee that the published message will get to (or from) the broker with somewhere near complete assurance that the messages will make it through, then use Level 1. If a message absolutely positively must get through, then use Level 2, but remember, nothing in life is 100% perfect, especially when relying on connecting to a broker (on either end) via an external connection to

10K Ohm Resistor Between Pin 1 and Pin 2   DHT11 or DHT22

Pin 1 to RPi Pin 2 (5VDC)
Pin 2 to RPi Pin 7
Pin 4 to RPi Pin 6
Pin 3 - NO CONNECT

10K Ohm
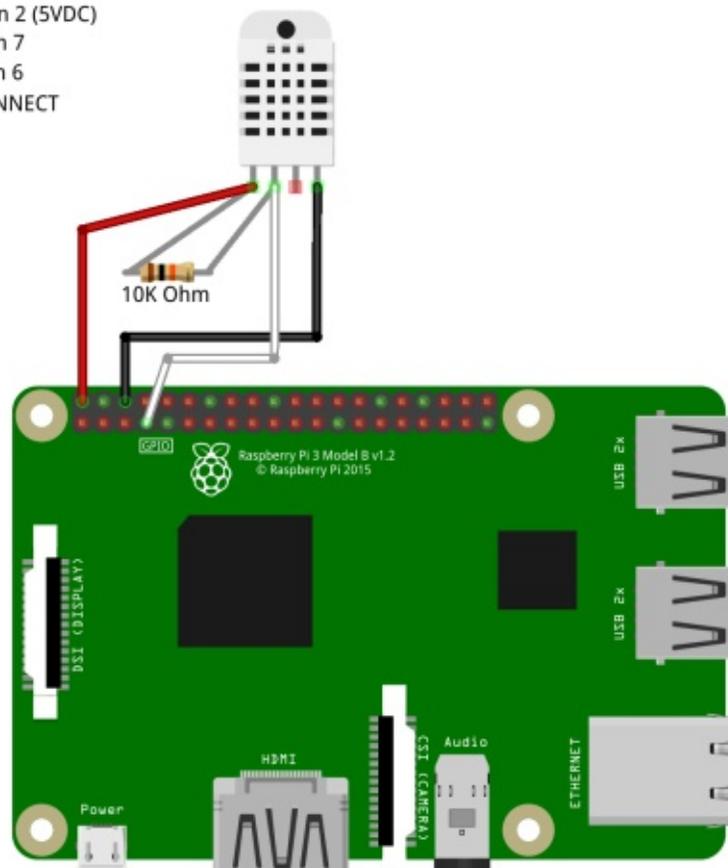
the internet.

We will be using a Qos Level 1 for our communications in our demo, just for practice.

THE HARDWARE

If you remember the original project we did back in FCM 109, we used a DHT11 connected to our Raspberry Pi. We used the Adafruit DHT22 python library and slightly modified their sample code.

I've recreated the Fritzing diagram showing the hardware connections...

Once you have the hardware all set up, now we have to download

and install the library. If you don't currently have the library, you can find it on GitHub at https://github.com/adafruit/Adafruit_Python_DHT.git . Remember, you have to install the library on the Raspberry Pi. Here are the steps to clone the library from github.

```
cd /home/pi
```

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

```
cd Adafruit_Python_DHT
```

```
sudo apt-get update
```

```
sudo apt-get install build-
essential python-dev python-
openssl
```

```
sudo python setup.py install
```

The sample code is in the /examples folder, called simpletest.py. Be sure to change the line that defines which pin your sensor is on to match your setup. The original code assumes your are putting the data pin of the sensor on GPIO 23 (line 35). In the setup I present, use GPIO 4. Also be sure to comment out line 31 if you are using a Raspberry Pi instead of a Beaglebone board.

Here is our updated code (right, and continuing on to the next page). You can try modifying the original code, or just start over. It's not that long.

In those last lines (next page), we check to make sure that we got data from the sensor, and if so, we create two messages, one containing the humidity value and one containing the temperature value. In the next two lines, we actually publish the messages. Notice that we are using a QoS value of 1 with a retain value of True. This should ensure delivery if the broker is there and if there are any clients listening for our data. Finally, we sleep for 5 seconds and repeat the loop. Feel free to change the sleep value to whatever you wish over 5. The DHT11/22 doesn't like to be accessed more than about once every 5 seconds. A realistic value would be somewhere between 30 and 60 seconds.

```python
client.publish(MQTT_PATH1,hms
g,qos=1,retain=True)

client.publish(MQTT_PATH2,tms
g,qos=1,retain=True)

else:
```

```python
#!/usr/bin/python
#  -*- coding: utf-8 -*-
# ==============================================================
# simpletest.py
# --------------------------------------------------------------
# Code modified from Adafruit sample code simpletest.py from the
# DHT library
# --------------------------------------------------------------
# Modifications by G.D. Walters for Full Circle Magazine #132
# ==============================================================
```

Here are our imports, the Adafruit DHT library, the paho library and the time library.

```python
import Adafruit_DHT
# import the paho library
# also import the time library for the sleep function
import paho.mqtt.client as mqtt
from time import sleep
```

MQTT_SERVER is the IP address of my RPi where the broker is running. Yours would probably be different, so put your IP in here. As to the path definitions, you can keep mine, or use your own.

```python
# Set some global definitions
MQTT_SERVER = '192.168.1.224'
MQTT_PATH1 = 'greghouse/dht22/humidity'
MQTT_PATH2 = 'greghouse/dht22/temperature'
```

I'm using a DHT22 in my system. Simply change the 'DHT.DHT22' to 'DHT.DHT11' if you are using the other DHT sensor.

```python
sensor = Adafruit_DHT.DHT22

pin = 4    # GPIO Input pin for data from the sensor Physical pin 7
sleep(3)   # Wait 3 seconds before starting up for things to stabilize…
```

If you remember from last month, the function on_connect() is a callback that runs whenever the program connects to the broker. We make the initial topic subscribe calls here.

```python
def on_connect(client,userdata,flags,  rc):
    print('Connected with result code ' + str(rc))
    client.subscribe(MQTT_PATH1)
    client.subscribe(MQTT_PATH2)
```

[CONTINUED ON NEXT PAGE]

The on_message() callback function is, in this case, only really to verify that messages are being published to the broker.

```python
# The callback for when a PUBLISH message is received from the broker.
def on_message(client, userdata, msg):
    print(msg.topic+' '+str(msg.payload))
    # more callbacks, etc.
```

Here we create our MQTT client and define the callback routines and connect to the broker.

```python
client = mqtt.Client()
client.on_connect = on_connect
client.on_message = on_message

client.connect(MQTT_SERVER, 1883, 60)
```

We enter a while loop to keep processing the sensor. This will continue until the program is halted by a keyboard interrupt <Ctrl-C>.

```python
while 1:
    humidity, temperature = Adafruit_DHT.read_retry(sensor, pin)
    # Temp in centigrade
    temperaturec = temperature
    # Temp in fahrenheit
    temperaturef = temperature * 9/5.0 + 32
```

In the above lines, we get the humidity and raw temperature data. We also convert the temperature from centigrade to fahrenheit. Sorry, but I have a REAL hard time thinking in metric.

```python
    if humidity is not None and temperature is not None:
        # If using Celsius, change the '*F' to '*C'
        # and use temperaturec
        print('Temp={0:0.1f}*F  Humidity={1:0.1f}%'.format(
                                        temperaturef,
                                        humidity))

        hmsg = '{0:0.1f}'.format(humidity)
        # Change the next line to temperaturec if you want to
        # publish in centigrade
        tmsg = '{0:0.1f}'.format(temperaturef)
```

```python
print('Failed to get reading.
Try again!')

    sleep(5)
```

Of course, I'm assuming that your broker is running on the same Raspberry Pi. If not, be sure the

change the IP address that MQTT_SERVER pointing to.

Now you can run your program

(be sure the broker is running BEFORE you start the program). Next, run mosquitto_sub on the RPi subscribing to one or both of the topics we set up in our program. You should see a message each time the loop runs.

I've put our new program on Pastebin at https://pastebin.com/gqPLNsiw, just in case you don't want to do all the typing.

Next month, our project will be a GUI subscriber client to monitor our sensors. This client can be either on the RPi or on another computer on your network. In fact, if you use a public broker, you can monitor it from anywhere in the world.

Until then, have a great month.

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

When you learn or create, associations speed up the process. By associating something known with something new, you fix the idea in your mind. You can take the new idea and associate it with another idea, creating a chain of ideas. Many of the great thinkers of history used association in their work. One most people are familiar with is Leonardo da Vinci. In his work, he created pictures, codes, and lines. These "doodles" created associations between ideas.

Mind maps let you develop associations through their connections between nodes. But connections are not always direct. An idea in one branch could associate with an idea in another branch. Or a portion of a branch is more important than the rest. You could find a group of sibling ideas that lead to another idea by their association. In Freeplane, you can create these connections through clouds, connectors, and summary nodes.

CLOUDS
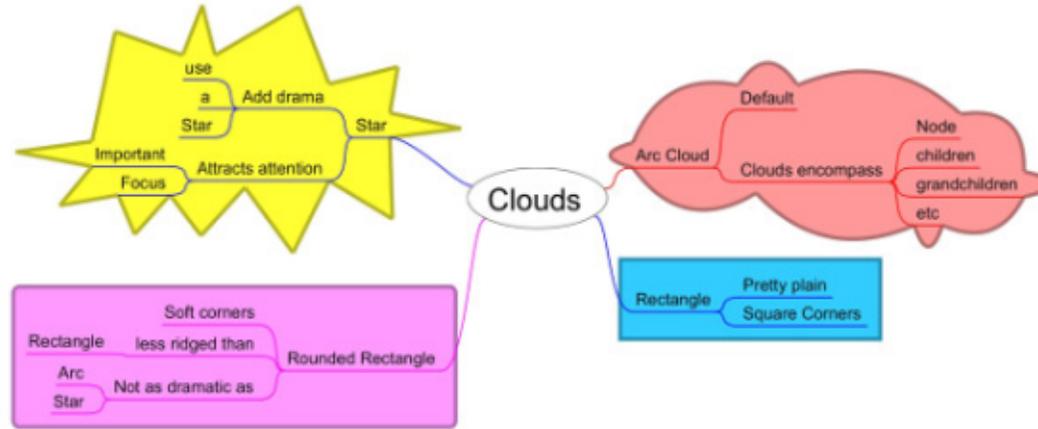
Clouds group a node and its children, grandchildren, etc. You can use the cloud to emphasize a certain section of a branch by creating a border around the nodes. The cloud surrounds the node and those attached to it down the branch until it reaches the leaves. The cloud will enclose any nodes added to the cloud's nodes.

To create a cloud, you select the node where you want the cloud to begin. On the keyboard, you use the keyboard combination CTRL + SHIFT + B. The cloud will surround all the nodes connected to the selected node. You can add a cloud through the menus using Edit >

Add/Remove Cloud. The menu item is also available in the right-click context menu. The main toolbar has a Cloud button, too. Repeating the action on the node will remove the cloud. If you select a node inside a cloud, you can create a cloud inside a cloud.

By default the shape of the cloud is the Arc style, creating the classical idea of a cloud. You have other options for the cloud shape. You can select from Arc, Star, Rectangle, and Rounded rectangle. In your code, the different shapes can have different meanings. Star for something exciting, or rectangle for something serious. To change the shape, use the drop-down list in the Cloud section of

the Tool Panel. From the menu, you select Format > Cloud shape.

Color is another way to code your maps, and Freeplane lets you change the color of the clouds. The cloud color changes the background and outline color of the cloud. As you have seen in previous articles, you select colors from a color dialog. To open the cloud color dialog, you can click on the color swatch in the Cloud section of the Tool Panel. You can access the cloud color through the menus, too, at Format > Cloud shape > Cloud color. Finally, the main toolbar has an icon for the cloud color. From the color picker dialog, you select the color you want for your cloud background. You click the OK button to accept the color.

CONNECTORS

Connectors associate nodes with relationships other than parent-child. Connectors are good for connections between siblings or nodes in different branches. As an example, a group of siblings
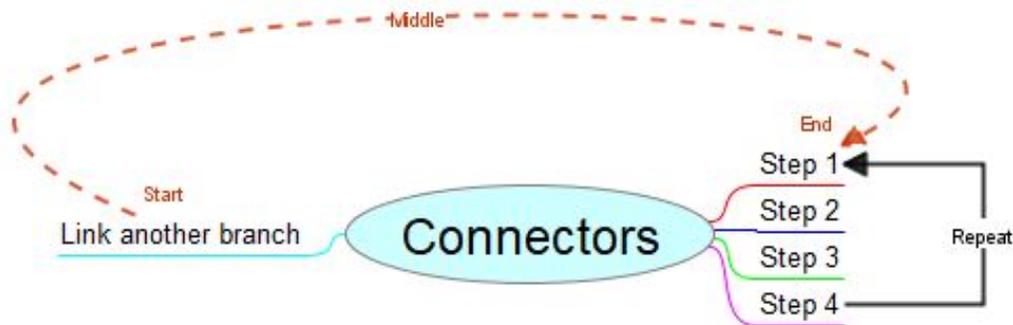
defining a repeatable process. A connector would allow you to loop the last process back to the first one. The connector can contain text to explain the relationship. Arrows can show the direction of the connection. You can edit the line style, color, and thickness to match your coding system.

The easiest way to create a node is to right-click on the start node and drag to the end node. By default, an arrow points toward the end node. You can change this, but it's worth thinking about before you create the connector. You can also select the start node and hold down SHIFT + CTRL while you drag with the left mouse button to the end node. Finally, you can select the start and end nodes, right-click, and select Connect from the menu. With two nodes selected, you can use the CTRL + L combination to connect

the nodes.
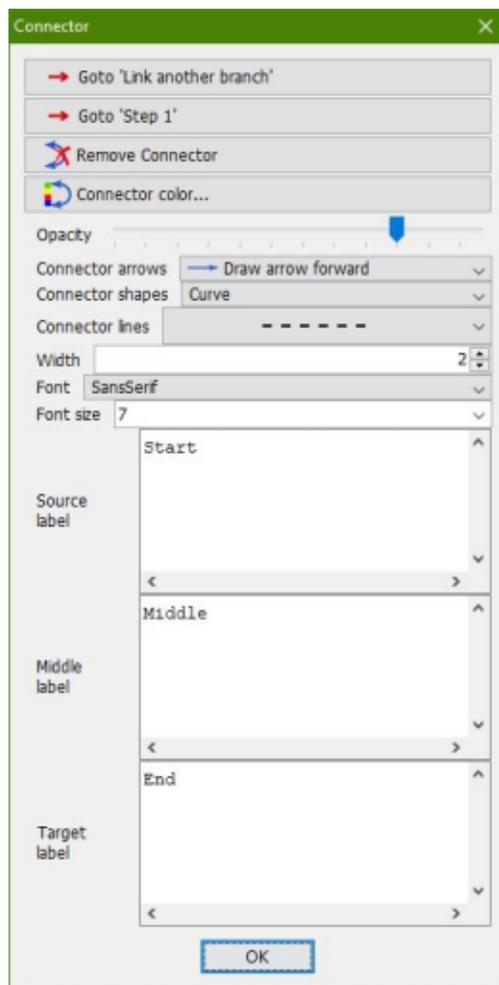
You can connect many nodes at once. Select the nodes you want to connect. The last node selected becomes the focus node. All the other nodes will point to the focus node. With the nodes selected, you right-click and select Connect from the menu. You can use the CTRL + L combination, too. All the selected nodes will connect to the focus node (last selected).

You can edit a connector by right-clicking it. A Connector dialog will open. The dialog lets you change the style, color, shape, thickness, and opacity of the connector. You can add text to the connect to define the reason for the connection.

The two Goto buttons let you select which node is active. You can switch between the two nodes at the connectors endpoints.

The Remove Connector button removes the connector and closes the connector dialog.

Use the Opacity slider to increase or decrease the visibility of the connector. Move the slider to the left to decrease the visibility, and move it to the right

to increase the visibility. By default, the slider is 75% to the right.

The Connector arrow drop-down list lets you select the arrows on the connector. You select None to create a relationship without direction. Draw arrow forward points the arrow at the endpoint. To point the arrow at the start point, you select Draw arrow back. You show a back and forth relationship by selecting Draw arrow forward and back. One of the ways I have used the arrows is when creating a repeating process. The process is outlined as a series of sibling nodes, then the last sibling points back to the first sibling. Using this method, I create a looping or repeating process.

By default, the connector is drawn as a looping curve. The Connector shape drop-down list changes the shape. You select line to create a straight point-to-point line. Linear path draws a line with sharp corners. Finally, you can Simulate edge. This choice will create a line that simulates the edge of the endpoint.

The Connector line drop-down list sets the type of line. As always,

contents ^

you have a choice between solid, long dash, short dash, dot, and dot-dash. The line types can become a part of your code. A dotted line can show two nodes that are similar but have no direct relationship. But a solid line can show a direct relationship, such as a node that is a child of two different nodes. The Width spinner box sets the width of the line. The width is a scale value rather than a unit measurement.

The font setting controls the font and size of the text used on the connector. The Font drop-down list selects the font from the fonts installed on your PC. The Font size spinner box sets the size in points.

With the font set, the connector has three label positions. The top text box is for the Source label. The source label is the start node of the connector. The Middle label places text in the middle of the connector. To place a label at the endpoint node, enter your text in the Target label text box. If you use the source and target label, you need to set the connector to show in front of the nodes. The setting is in the preferences, Tools > Preferences.

Click on the Appearance tab. In the Connectors section, you uncheck the box for Paint connectors behind nodes.

Click the OK button to close the Connector dialog.

When you left-click the connector, you see two adjustment handles. You drag the handles to adjust the shape and location of the connector. Moving the handles moves the location of the arrows and labels. This is another method to prevent the label and node's core from blocking each other.

SUMMARY NODES

What do you do when you need a node that pertains to several sibling nodes? For example, say you have a list of fruits and vegetables as siblings to a node. You would like to label which ones are fruits and which ones are vegetables. You could create a

child node for each of the nodes stating its food category. But this seems redundant and time-consuming.

Enter the summary node. A summary node groups together a series of sibling nodes using a curly bracket. The summary node allows you to group your fruits and vegetables using one node for each group.

To create a summary node, select the series of sibling nodes. With the nodes selected, press the keyboard combination SHIFT + ALT + INSERT. The summary node created looks like any node except its edge is a curly bracket. Enter the text of the node and press ENTER. You can press ENTER again to create a sibling of the node, or you can press INSERT to create a child node. From the summary node, your map can continue as you would with any other node.

There are a couple of alternate methods for creating summary nodes. After selecting the node series, you can right-click and choose New summary node. From the menus, Edit > New Node > New summary node creates a summary node of the selected sibling nodes.

When you create and learn, associations allow you to build on previous knowledge. While the parent-child relationship handles most associations, sometimes you need to create others. Freeplane gives you the tools you need to create those associations. With the cloud, you can encompass a group of nodes, creating a section set off from the others. You can use connectors to associate two nodes that are not parent-child. Finally, summary nodes allow you to group sibling nodes together.

Summary of summaries
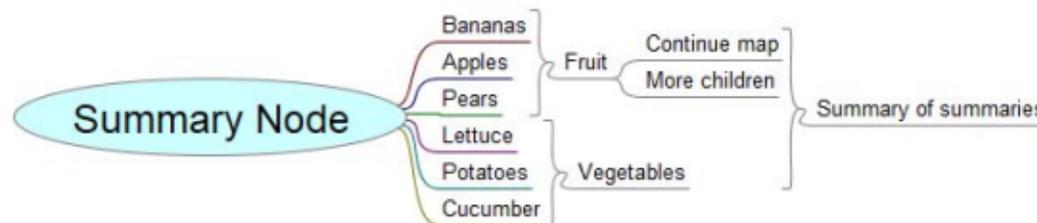
Elmer Perry is a technical support rep for an international keyless access company. He enjoys writing, woodworking, and technology. He lives in Leicester, NC with his wife.

Ah, that time of year is here again. A new Ubuntu release. There are typically two thoughts on updating; fresh and upgrade. This article will focus on the upgrade. The rolling upgrade can be time consuming though, downloading all those new packages. What if I told you that we can speed up our upgrade to 18.04.

Well we can; using apt-fast we can make your updates go about 30% faster. While I am thinking about it, here is a link to the github page: https://github.com/ilikenwf/apt-fast. But for those just reading on, what is apt-fast?

In short, apt-fast (ppa:apt-fast/stable) is a wrapper around apt-get or apt, and that uses aria2c to perform a multithreaded download of the needed updates. But what if we could use apt-fast to assist in a release upgrade - to upgrade to 18.04. Many user's might be upgrading from 16.04 or 17.10, heck there might even be some upgrading from 14.04. Either way, guess what, you can use apt-fast to speed up your release upgrades.

While it's possible to use this even with the upgrade GUI. I'll focus on the command line tool; do-release-upgrade. So what do we need to do to perform such a miracle? Multiple terminals, that's what. Or some tmux/screen knowledge**.

LET'S GET STARTED

First off, you'll need two terminals. But may I suggest three. Consider it a safety thing. I'll explain as we go.

Step 1: In Terminal 1, execute the do-release-upgrade command as in figure 1. One may notice I am passing a "-d" switch. This is to

```
Hit http://ppa.launchpad.net/webupd8team/terminix/ubuntu artful InRelease
Err http://ppa.launchpad.net/tmsu/ppa/ubuntu artful Release
  404  Not Found
Fetched 0 B in 0s (0 B/s)
Reading package lists... Done
Building dependency tree
Reading state information... Done

Updating repository information

Third party sources disabled

Some third party entries in your sources.list were disabled. You can
re-enable them after the upgrade with the 'software-properties' tool
or your package manager.

To continue please press [ENTER]
```

check the beta channels for an update. Use this switch at your own risk. Otherwise you'll need to wait until after the final release.

Step 2: After you launch the command, your repository information will be updated to pull the new release. If you use PPAs on your system, you will see something like the image shown above, alerting you that they will be disabled. If you do not use PPAs, you may move on to Step 3.

When you see the prompt, acknowledge it by pressing the "enter" key. The update will now refresh to the new release repos.

Step 3: This is very important if you are going to speed up the release upgrade. When you see what is shown in the image below, **DO NOTHING**.

When you see this prompt, it is time to move to a new terminal.

```
choman@choman-Latitude-E6410:~$
choman@choman-Latitude-E6410:~$
choman@choman-Latitude-E6410:~$ sudo do-release-upgrade -d
```

```
27 packages are going to be removed. 241 new packages are going to be
installed. 2192 packages are going to be upgraded.

You have to download a total of 2,548 M. This download will take
about 12 minutes with your connection.

Installing the upgrade can take several hours. Once the download has
finished, the process cannot be canceled.

Continue [yN]  Details [d]
```

```
choman@choman-Latitude-E6410:~$ sudo apt-fast dist-upgrade
```

```
xorg-docs-core xpra xserver-common xserver-xephyr xserver-xorg xserver-xorg-core
xserver-xorg-input-all xserver-xorg-input-libinput xserver-xorg-input-wacom
xserver-xorg-legacy xserver-xorg-video-all xserver-xorg-video-amdgpu
xserver-xorg-video-ati xserver-xorg-video-intel xserver-xorg-video-openchrome
xserver-xorg-video-radeon xterm xtightvncviewer xvkbd xwayland xxd yelp-xsl zeal
zeitgeist zeitgeist-core zeitgeist-datahub zenity zenity-common
2192 upgraded, 241 newly installed, 27 to remove and 0 not upgraded.
Need to get 0 B/2,548 MB of archives.
After this operation, 674 MB of additional disk space will be used.
Do you want to continue? [Y/n]
```

Proceed to Step 4.

Step 4: In a new terminal window. Enter the dist-upgrade command as shown above.

NOTE: Do <u>not</u> add the "-y" flag to this command.

When you hit enter, all the required packages for the release upgrade will be downloaded using the aria2c multithreaded download manager. Your terminal will look like the image below.

Step 5: This is a good time to open a third terminal and keep it in focus. It's really easy to mess this up, and this helps to prevent other issues. Off jog the window a bit so you can see when the dist-upgrade command is done running. When it looks like the image top right, proceed to the next step.

Remember DO NOTHING...

Step 6: When the dist-upgrade command is finished downloading, you may close your safety terminal. In the terminal representative of Figure 7, enter "n". That's right, we do not want to continue. Why?

Other things happen in the do-release-upgrade command; such as removing obsolete packages. We want to take advantage of this. So enter "n" and close the terminal. Now we come "full circle" to terminal 1.

That's right, we are now ready to proceed with the upgrade. Enter "y". Here's what will happen next. The system will attempt to download the required packages and notice that they are already downloaded. So then it will commence with the upgrade and start installing packages.

This is where your system will dictate when you will finish upgrading as CPU, RAM, drive speed, all come into play. At the end of the package installation, you will be prompted to remove obsolete packages. Then you will be prompted to reboot the system to finish the upgrade.

Congratulations, not only have you now learned how to speed-up the release upgrade. But you have also made it through my first article. So thanks for reading. My hope is that you like the extra bump in speed, and add this process to your 6-month or 3-year update plans. A couple of notes before I let you go.

NOTES

• Try this in a VM before a physical system, remember snapshots. Get familiar with the flow.
• Remember your PPAs are disabled.
• This works with the upgrade GUI. Follow the prompts to know when to perform the dist-upgrade.

Thanks for reading, and enjoy the following flow as a cheat-sheet.

CHEAT SHEET

TERMINAL 1:
• Execute the do-release-upgrade command:
    • sudo do-release-upgrade
    • If you're bold, you may also do this with the do-release-upgrade -d, but I suggest you try this first in a VM to get comfortable with the process.

```
04/10 17:12:17 [NOTICE] Download complete: /var/cache/apt/archives/apt-fast/rename_0.20-
7_all.deb

04/10 17:12:17 [NOTICE] Verification finished successfully. file=/var/cache/apt/archives
/apt-fast/perl_5.26.1-5_amd64.deb

04/10 17:12:17 [NOTICE] Download complete: /var/cache/apt/archives/apt-fast/perl_5.26.1-
5_amd64.deb
```

HOWTO - UPGRADE WITH APT-FAST

• If you do not use PPAs, continue to the next bullet. If you do use PPAs, you'll eventually see a prompt that they will be disabled. Hit yes and let the process continue.

• **ALERT**. Had to get your attention. You'll eventually see a large prompt saying:
    • X number of packages are obsolete
    • Y number of packages are new
    • Z number of packages will be updated.
    • STOP STOP, just read.
    • Then a prompt saying "press y to continue".

• DON'T do anything, Move on to the next section.

TERMINAL 2:

• Execute the apt-fast command:
    • sudo apt-fast dist-upgrade
    • DO NOT use the "-y". But go ahead and hit enter once.

• This is where I suggest a third terminal and put it in focus. It's far too easy to mess things up. And I'm here to say this is safe.

TERMINAL 3:

• When you open the third terminal, move it a bit out of the way so that you can SEE terminal 2.

• Just wait, since terminal 2 is downloading multithread style, it should only take a few minutes.

• When you notice Terminal 2 finished, it will be prompting you to continue. The ANSWER is no.

TERMINAL 2:

• Return to terminal 2 and enter 'n' to leave the apt-fast command. That was easy. Let's go back to terminal 1.

TERMINAL 1:

• Now we're ready to continue with the release upgrade.

• Enter 'y' to continue.

• Here's what will happen:
    • The next step for do-release-upgrade is to download all required packages. But we just did that using apt-fast.
    • So, when do-release-upgrade notices the cache is up-to-date, it will casually move on to the next step; Installing the packages...
    • This is where apt-fast may not seem like a winner. But it is. While installing the updates, this part of the process is plagued by your hardware: RAM, CPU, drive speed.
    • At the end of this, you;ll be prompted to remove obsolete software, go ahead and enter "y".
    • At the end of this, you'll be prompted to reboot the system, go ahead and enter "y"


```
1/1 ▾  +  ⊡  ⬚                    Tilix: Default            Q  ≡  ⊟⊡⊗
1: choman@choman-Latitude-E6410: ~ ▾                               □  ✕
Get:30 http://us.archive.ubuntu.com/ubuntu bionic/multiverse DEP-11 64x64 Icons [216 kB]
Fetched 47.6 MB in 6s (3,524 kB/s)

Checking package manager
Reading package lists... Done
Building dependency tree
Reading state information... Done

Calculating the changes

Calculating the changes

Do you want to start the upgrade?


49 installed packages are no longer supported by Canonical. You can
still get support from the community.

27 packages are going to be removed. 241 new packages are going to be
installed. 2192 packages are going to be upgraded.

You have to download a total of 2,548 M. This download will take
about 12 minutes with your connection.

Installing the upgrade can take several hours. Once the download has
finished, the process cannot be canceled.

Continue [yN]  Details [d]y
```

O ver the past six years, this series has detailed just about every part of Inkscape, from basic shapes through to more complex features such as filters and path effects. With nothing more to cover, it's time for this column to take a sharp turn from its previous tack: I'm going to look at how you can use graphics you've created in Inkscape on the world wide web.

On the surface, this might seem like a simple task – worthy of an instalment or two at most. But SVG is unlike any other graphics format; its XML heritage lends it to the same kind of manipulation with CSS and JavaScript as HTML, allowing for graphical tricks that go way beyond simply displaying a static image on a page. At the same time, you still have the option to treat SVG as a dumb image format, if that's a better fit for your requirements.

Broadly speaking, there are four ways to include an Inkscape-created image in a web page. Each has its own pros and cons, which are outlined in the table shown

Method	Scalable	CSS animations	JavaScript	Difficulty
Convert to raster	✘	✘	✘	Low
Link to SVG	✔	✔	✘	Low
Object element	✔	✔	✔	Medium
Inline SVG	✔	✔	✔	High

above.

The first method is to simply avoid the whole SVG question entirely and export your image to a PNG file – possibly even converting it to a JPEG afterwards to reduce the file size, if the image doesn't require transparency. The resultant raster image can be included in a web page via an <img> tag, used on social media sites or blogs, and generally be treated the same way as a photo from a phone or digital camera. Because you're effectively creating a snapshot of the image as it appears in Inkscape, you also don't have to worry about missing fonts, flowed text, or mismatches in browsers' rendering engines. The downside is that the exported

raster image is no longer infinitely scalable in the same way as a vector image. You also lose any concept of the individual objects that made up the image, so it's not possible to manipulate them individually using CSS animations or JavaScript.

The second approach relies on the fact that most modern browsers do a pretty decent job of rendering SVG – so why not just use the SVG file directly, rather than converting to a raster format first? This is as simple as referencing an SVG file in your <img> tag, rather than a PNG or JPEG. Because the internal structure of your SVG is preserved, the browser knows about individual objects, so can animate

them using CSS. In theory it could also allow JavaScript, but that's prevented when using images like this, for security reasons.

What does security have to do with an SVG image, you may wonder. Consider that many social media platforms, forums, and other websites, will let you upload your own SVG file, then expose it to other people on the site by using an <img> tag. This would effectively give a malicious user a means to run their own JavaScript under the guise of the hosting website. The code could capture user's keypresses, spoof a login dialog to record passwords, or redirect the user to another site entirely. Given that being able to upload an SVG file to a site and

have it rendered as an image is generally a good thing, but that unfettered permission to run JavaScript in such a file could easily be abused, browsers take the sensible approach of allowing CSS animations to run (giving SVG an edge over raster images), but disallowing any JavaScript.

Unfortunately this method does come with one big drawback – though it's a problem that diminishes every year. Because you are relying on your browser's rendering engine to draw each object in your image, you no longer have complete control over the appearance of your image. For simple drawings, this isn't generally a problem. But add in more advanced features, such as filters, and the output of the browsers begins to diverge.

Text rendering is a particularly significant problem. If the person viewing the image doesn't have the same font installed on their machine that you used to create the document in the first place, it will be replaced with a different font. In HTML, this can be annoying, but usually the browser adjusts the page layout a little and no real harm is done. In an image format, however, it can be much more of a problem. Consider something like a comic strip: if the text is rendered at the wrong size it might not fit the speech bubbles; if it's rendered in a different font it could change the feel of the narrative entirely. For this reason, I post all my comic strips as PNG or JPEG images, and link to the SVG files separately where I want to make the source files available for people to download.

The third approach is essentially another way of linking to an SVG file, but this time it uses the <object> element rather than an <img>. Whereas forums and social media sites are happy to expose a user's file via an <img> link, they tend not to do the same via <object> – so, if you want to take this approach, you need to be able to edit the underlying HTML directly. This usually implies that you are creating content for your own domain, rather than just uploading something to a third-party site. What you gain from this extra requirement is an enhanced level of trust from the browser, allowing JavaScript in your image to execute. With sufficient knowledge and experience, it's even possible to create two-way interactions between code on the page and in your image. But, again, you're also at the whim of the browser in terms of the final rendering of your file.

The last approach is a relatively recent option. For years, SVG was neglected by the web developer community, largely because Microsoft refused to support it in earlier versions of Internet Explorer, at a time when it commanded the majority share of the market. But with later releases, Microsoft softened their stance, adding SVG support and paving the way for SVG to be promoted to a first-class citizen in the world of HTML5. Now, you can put a chunk of SVG code directly into your HTML and it will just work. And if you're hand-coding your SVG, you no longer have to concern yourself with the complexities of namespaces in XML – just use the SVG tags in the same way as you would HTML.

With this last method, the SVG is a direct part of your web page, so writing code to manipulate the image is much simpler than with the <object> approach. But it does mean that your SVG code is right there in the middle of your HTML.

If all you're doing is drawing a couple of simple objects, that might not be so bad, but including all the code from a complex drawing can quickly overwhelm the rest of your page, making it harder to work with the surrounding HTML. And, once more, your browser's rendering engine might make poor work of your image.

Enough theory, let's put it into practice. Next month, I'll describe how to use <object> and inline SVG, but, this time, we'll deal with the <img> methods. Begin by loading a simple image into Inkscape, or creating one from scratch if necessary. I'm going to work with this basic circle (the square around it is the Inkscape canvas):
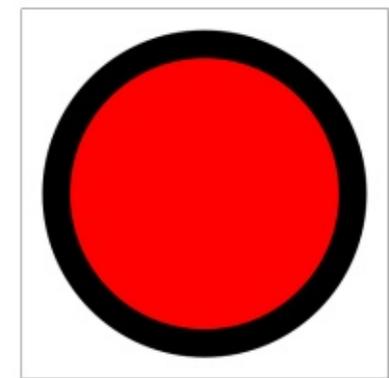
Let's export this to a PNG image via the File > Export PNG Image…

menu item. This opens a dialog in which to set the details for the export:
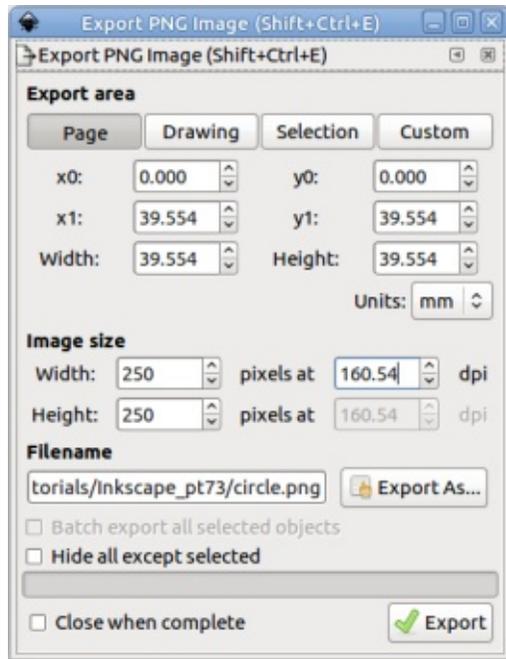
I covered this dialog way back in part 12 of this series, but, for the purposes of creating a PNG to use online, you can ignore many of the options. You do need to decide what to export using the buttons at the top: usually this will be the entire page, the entirety of your drawing (which could be larger or smaller than the page), or a rectangle that encompasses the currently selected objects. I've chosen to export the whole page, as I want to include a little bit of blank space around my circle.

Now comes the most important decision: how big do you want your image to be? This is set in the Image Size section, by adjusting the image dimensions or dpi (dots per inch). For web use, you can generally ignore the dpi option and go straight for setting the width and height in pixels. If you're creating an image for a profile picture or banner ad, you'll usually be told the dimensions your graphic must be. Note, however, that you can't set these fields independently, so you'll either have to create your design to the right proportions, or export a larger area that you can then crop in a bitmap editor such as The GIMP or Krita.

The last setting required is the filename to export to. Click the Export As… button and navigate to your destination folder in the file selector. Enter a filename in the selector, making sure it ends in ".png", and then use the file selector's Save button to return your selected location and name to the export dialog. Despite the name of the button, this won't have saved a PNG file anywhere; it's just saved the export location. In order to actually save the file, you have to click the Export button.

You should now have a PNG file on disk that can be uploaded to just about any site that accepts user images. If you want to use it in your own HTML page, that's easy too. If you're already familiar with writing HTML then you almost certainly know how to link to your image. But if you've never written a line of HTML in your life, why not give it a try now – it's not as tricky as you might think. In a text editor type the code shown below.

Modify the "circle.png" string to suit the filename of your own PNG, and save the document into the same directory as the exported image, making sure to give the filename an extension of ".html". Open the web browser of your choice, press CTRL-O to open the file selector, and load your HTML file. You should see a text heading with your PNG image displayed below it.

That wasn't too tricky, was it? Linking directly to an SVG file is just as simple. If your original Inkscape file isn't in the same directory as the HTML file, use File > Save a Copy… to put a copy of it there. When you look in the directory you should now see your HTML file, your exported PNG, and an SVG file. To get the latter to appear in your web page, just add the following lines below the existing image link:

```
<h1>IMG tag (SVG)</h1>

<img src="circle.svg"></img>
```

Save the file, switch to your

```
<!DOCTYPE html>

<html>

<head>
   <title>SVG in HTML</title>
</head>

<body>
   <h1>IMG tag (PNG)</h1>
   <img src="circle.png"></img>
</body>

</html>
```

browser, and press F5 to reload the page. You should see both your PNG image and the SVG version. There's a good chance that the images aren't the same size, but SVG files are scalable so we can fix that easily. Given that my PNG was exported at 250px in width, I can modify the <img> tag for the SVG file to set it to the same:

```
<img src="circle.svg" width="250"></img>
```

Another press of F5 in the
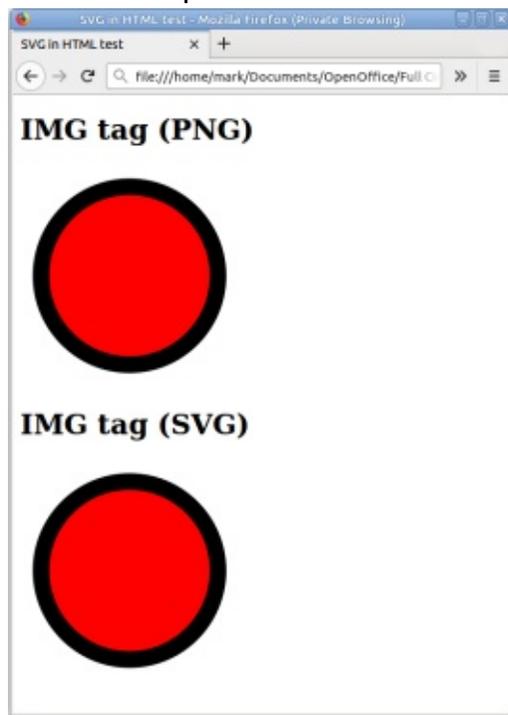
browser and you should find that both images are the same size. But what if we wanted them both to be larger? Try setting a width attribute on both of them, with a value of 1000 or more. Notice how blocky the PNG becomes, whilst the SVG is still rendered as smoothly as possible. That's the main advantage of using an SVG file in a web page instead of a PNG.

Linking to an image via an <img> tag is a common way to display user-submitted images on the web. But, when coding your own pages, images are often used in other ways as well, typically via CSS. You can use an SVG file wherever you would usually use a PNG in these cases. For example, let's fill the background of our web page with copies of the SVG file using this chunk of CSS, inserted between the </head> and <body> sections of the HTML code:

```
<style>

body {

   background:
url("circle.svg");

   background-size: 50px;

}

</style>
```

Reload the page to see the result. As a rule, modern browsers will let you use an SVG image wherever a raster is allowed. But although the browser might be happy to do that, many websites still only let you upload pure raster formats such as PNG, GIF and JPEG. If you write your own HTML you have full control over what formats to use, but if you're trying to upload an image to a third party site you may find your options limited. If you can use an SVG, though, I encourage you to do so. Only by ever more SVG content making its way online will browser vendors have an incentive to provide better support for the format, and the SVG working group have a chance to introduce more of the new features that would make this useful format even better.

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

Recently, Microsoft announced that Linux will be a subsystem operation in Windows 10. My research institutes employs Linux in various forms. Unfortunately, as an end user I am in the crux between our software developers and our ITS. The developers write programs using Linux, but it is not supported by our ITS. The developers are offsite, so our natural support group, ITS, shuts down and refuses to service the Linux ecosystem.

So, I am left wondering would a linux subsystem suffice for our needs and force ITS to support the laptop? Can we drift away from dual booting? These are just a few of my thoughts. Perhaps it would be beneficial to review the Linux subsystem in Windows 10. So how does one install the subsystem? It appears it follows the same credo as Crouton; boot from developer mode and install it from an app store. The two choices for Linux are openSUSE and Ubuntu. These are not GUI intensive programs, but merely a command line terminal. It is the correct step for

Microsoft to acknowledge Linux's kernel in the world of computers.

However I must diverge from my main article's intention. I gleefully used Korora as my main distro for my HP laptop at work and home. Yet I recently found out that the team are taking a sabbatical. I understand the need for breaks for work-life balance, but, in my work, I need a consistently reliable OS. I would like to give Korora a large thanks

for being my distro for the last 2 years.

So I returned back to using Fedora LXDE and XFCE on my laptops. I choose the upstream product for one simple reason: solidarity. I extensively know the command line under RPM based distro. My Perl, R, and PHP batch files are a fit with Fedora. I would not have to revamp my work style nor my understanding. I love rooting for the underdog, however

I do not see the corporate Linux companies going out.

I do not use my Fedora laptop as my daily computer nowadays. Instead, it is the perfect compliment to my office needs. Many times I use my batch files to automate my tasks or I use my work laptop for automation. I rather use the free and open software options to complete a job task. I dislike SaaS options that are now frequent under Windows OS compliant software. Next month I will review: R and Gnuplot.

I frequently have to develop basic stats for my work.
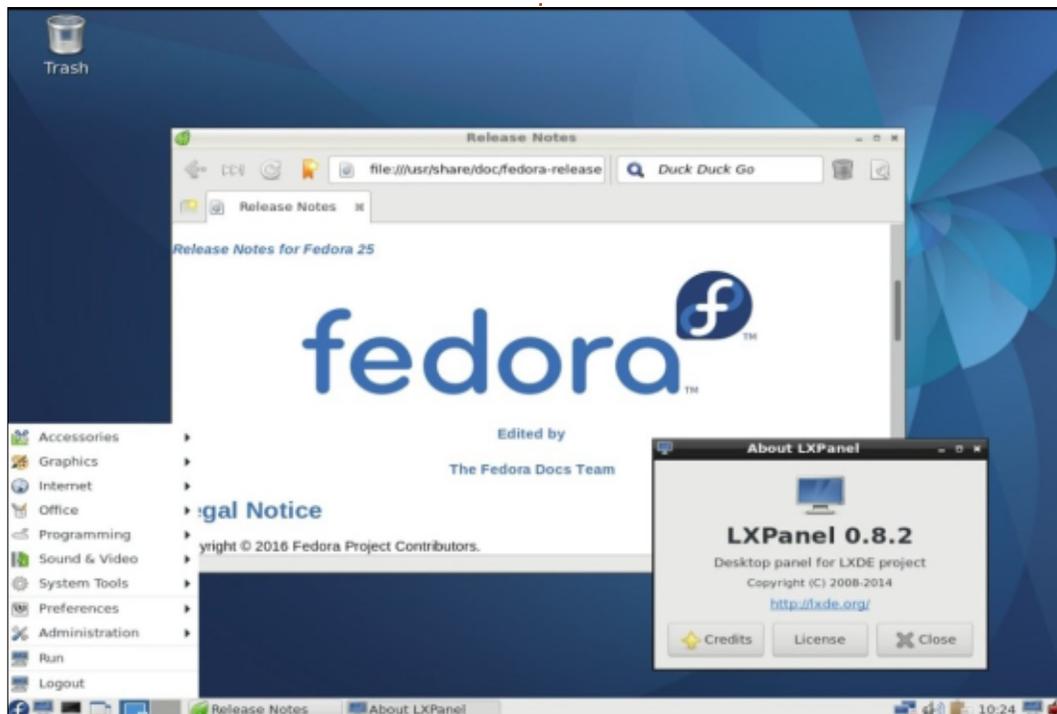
SJ Webb is a researcher coordinator. When he is not working, he enjoys time with his wife and kids. He thanks Mike Ferarri for his mentorship.

As promised last month, we're going to now look at basics of Ubuntu's System Settings.

Every operating system I've ever used, including Linux, has default system settings that just make me scratch my head in puzzled bewilderment, but it's fair to observe that it probably all comes down to personal preference.

To get to the System Settings in Unity, click the 'Gear' icon on the upper right:
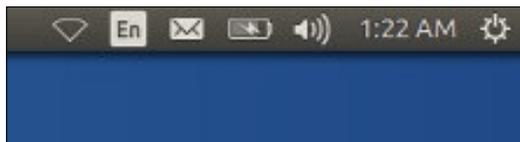
This will bring up a control menu that includes 'System Settings'. Click it and you'll get the screen shown right.

We'll go through the commonly used ones in order. First: Appearance.

The Appearance dialogue looks like the screen shown top right.
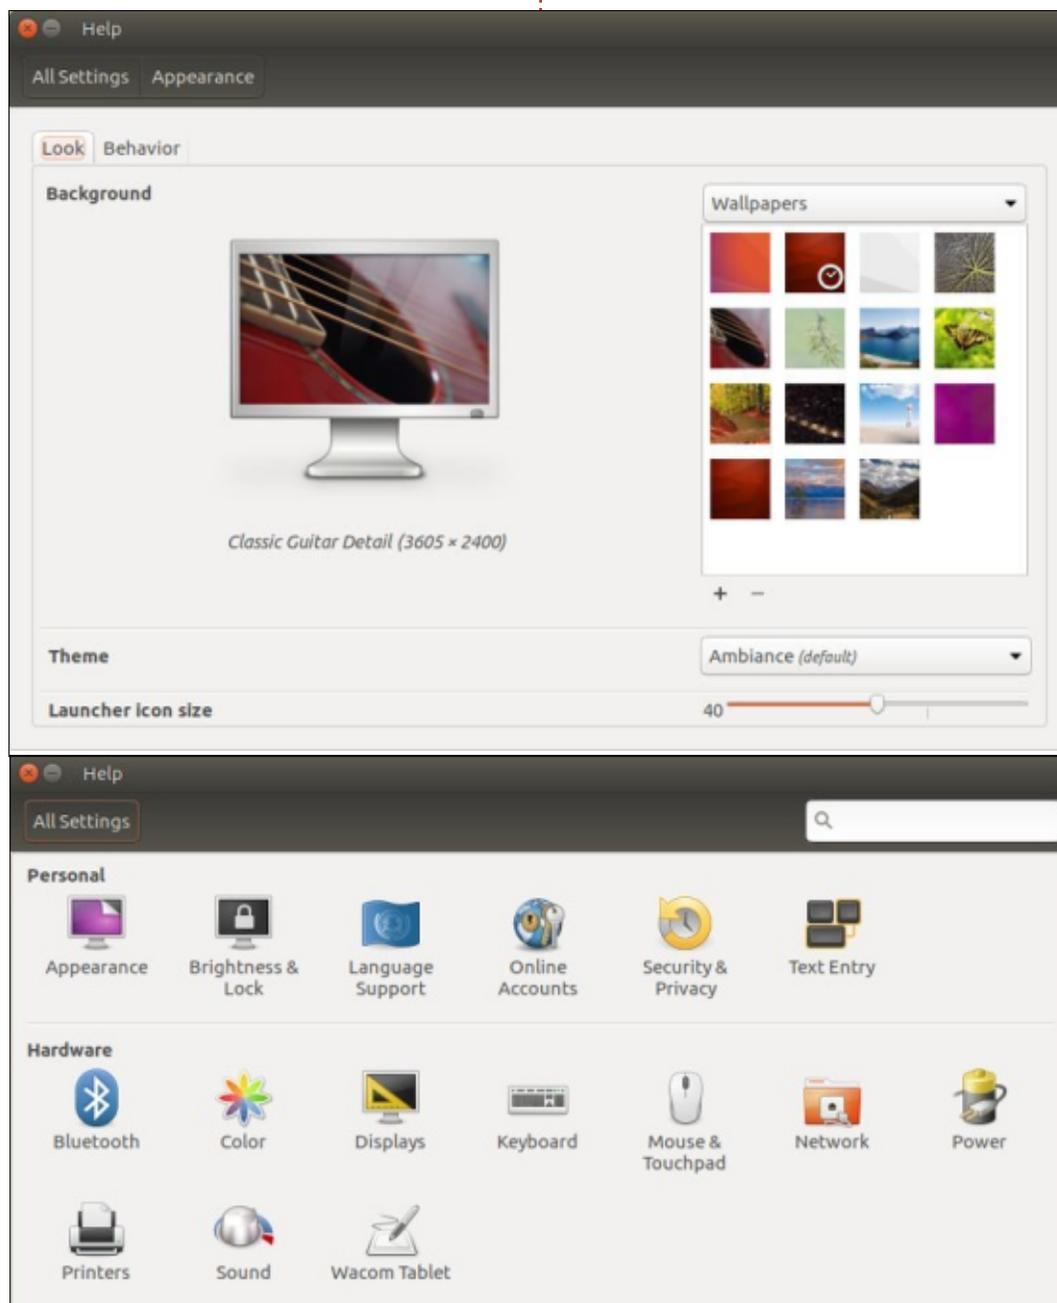
There are a number of attractive wallpapers built into Ubuntu, and I am always personally inclined to use any system's default ones, a habit that dates from the days when hard-drive space was at a real premium, and disk space used on wallpapers could be better used for other things. There are a lot of attractive wallpapers online available to be downloaded, so you should certainly feel free to go find ones you consider appealing. You can click the down-arrow at the top to change from Wallpapers to Pictures or Colors and Gradients. To add other graphics to use as wallpapers, click the + button below and browse to the location where you have stored the desired picture.

Below the wallpaper display, you can pick a theme. Adding themes is surprisingly difficult in Unity, but can be done either by downloading and using the Unity Tweak tool, or by finding themes with self-installing routines, sometimes requiring the setup of

specific repositories (check the previous Everyday Ubuntu column in issue #130 for more info on repositories). I think most users will find that the results are not really worth the effort, so we'll leave it aside for brevity's sake.

Below the Theme selection area, there is an option to change the icon size for icons on the Launcher. It will default to 40; I like to drop mine to around 34 so I can get more icons on the Launcher, which I use a good bit. The trade-off is that the larger icons are easier to see, but that becomes less of a problem as the icons

become more familiar over time. The name of the application appears if you hover the mouse, so that also makes it easier when using the smaller icons.

At the top of the dialogue box are two tabs; the one we have been looking at so far, the default, is Look. Switch to the Behavior tab shown below.

Here you can decide if you want the Launcher to auto-hide, you can add a Show Desktop icon to the Launcher, and you can change how menus work. The Show Desktop amounts to minimizing all open

windows at once and can be quite handy.

Show menus for a window defaults to 'in the menu bar', which is more like Mac OS, whereas 'in the window's title bar' is more like Windows behavior. I prefer the second option because I find it more intuitive, and it's more efficient to not always have to move the mouse pointer to the top of the screen rather than the top of the active window. If you come to Linux from Mac, you may prefer the default.

I'm personally not really sure why you wouldn't want your menus always visible, but you can opt to

have them hidden until your mouse hovers in the title bar. I choose 'Always displayed'.

Click All Settings at the top to go back to the main System Settings, then try Brightness and Lock:

I usually go here first when setting up the system to turn off the screensaver (select Never) and turn screen lock off, at least until initial setup is complete. Having the screen lock or dim while you're actively trying to set the system up is very annoying, in my opinion. You can always go back here and tweak things to your liking after setup has been completed.
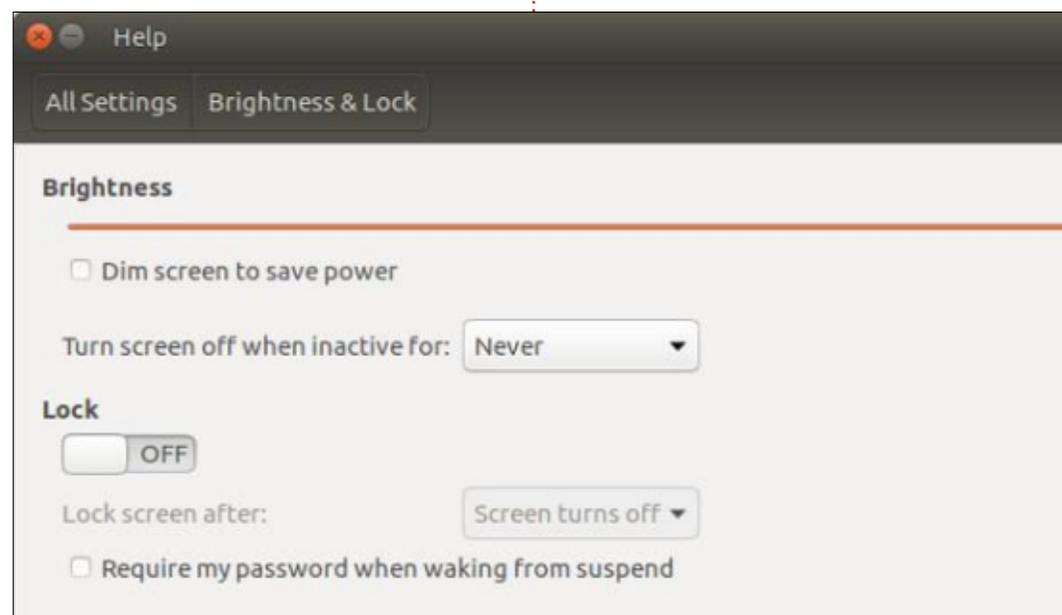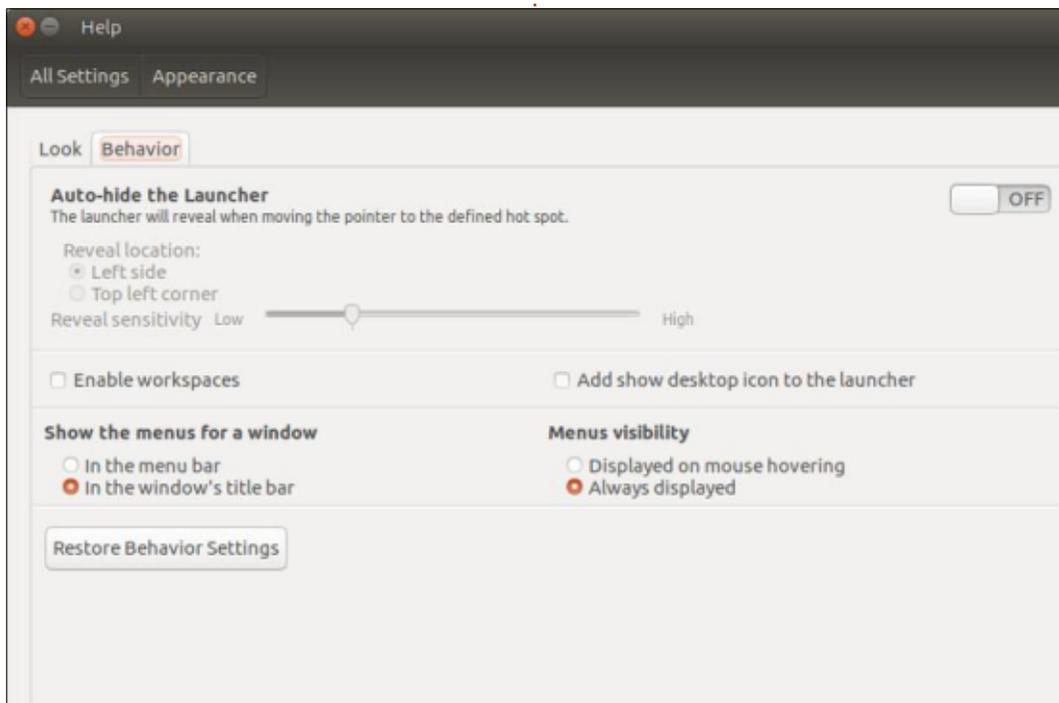
Back in the System Settings, we'll take a quick look at Online Accounts (shown below).

You can add Facebook, Flickr, or Google accounts here.

Your average user may not often use Security and Privacy settings, but there are still a few things to look at here:

I personally turn off 'Require my password when waking from suspend' because I'm not using my machine in an open environment where someone (with ill intent) could gain physical access. This is another of the first things I change when doing setup, since repeatedly inputting your password when you're sitting at your own desk at home is pointless
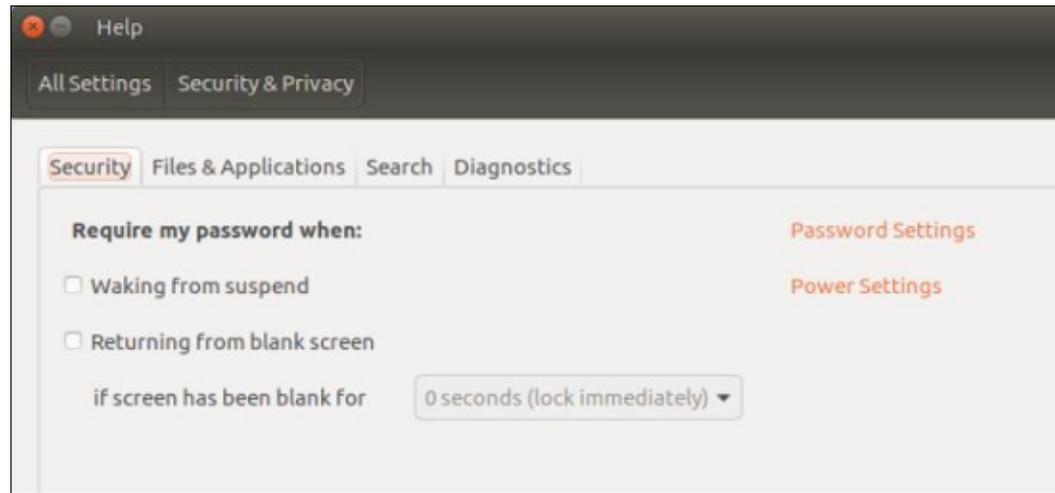
and inefficient. Most of the other settings here won't affect home users.

Under the Hardware subsection, you'll find Bluetooth. Here, you can turn Bluetooth off and on, and manage Bluetooth devices.

The Color management part of

System Settings will usually only be important to professional graphics artists and the like, so we'll skip that and look at Display (shown bottom right).

Here you can manage your displays, resolution, screen rotation, size of menu and title bars, the placement of the

Launcher on all displays or a specific one, Sticky Edges, and scaling. Sticky Edges controls auto-sizing of windows when dragged to screen edges and is probably best left at default in most cases.

I invite feedback on easier/better ways to do things. Any such submissions in response to articles or content will be considered the property of Full Circle Magazine for publication purposes, without remuneration, unless the writer/commenter specifies otherwise. That said, commentary and feedback are heartily encouraged and appreciated, at acer11kubuntu@gmail.com.
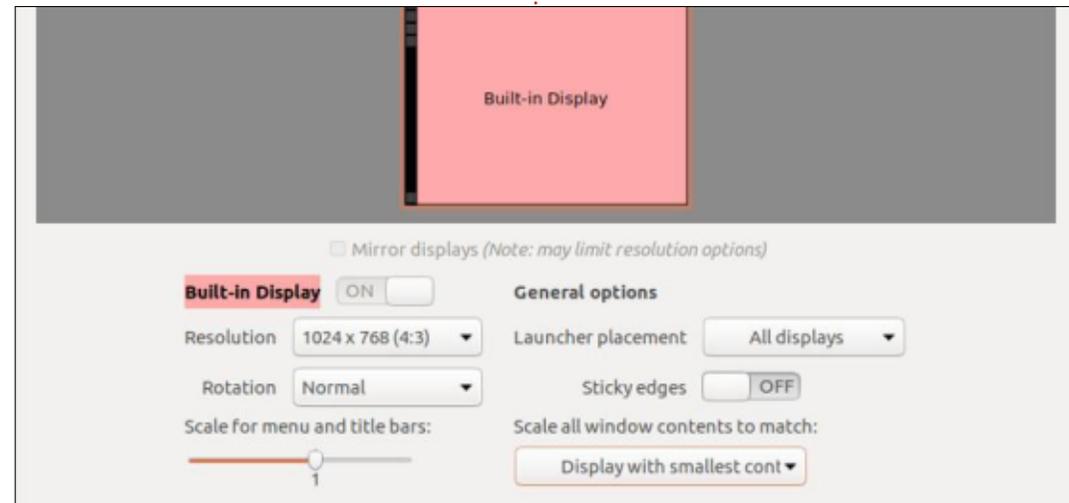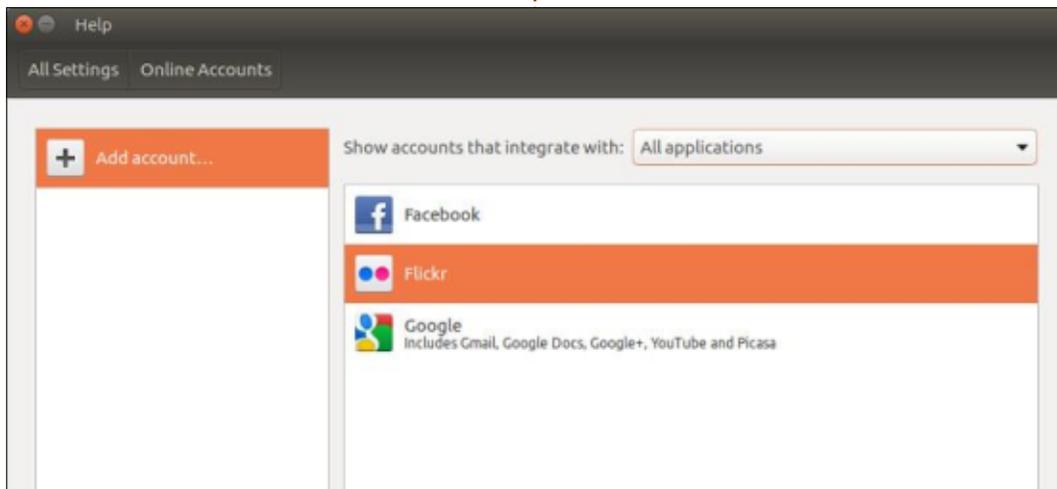
A quick recap of last month's introduction. I wasn't happy with how Rhythmbox handles song sorting information and decided to brush up my programming skills to write a little fix/improvement. In the process, I reinforced my praise for open source software and picked up a new language thanks to Full Circle. Let's start with the former point, and keep in mind that most of what is valid for the Gnome music player can be easily applied to other open source software.

My first idea was to see if I could patch the source code. This is one of the ways you can contribute to the OSS community. See if there's a bug or room for improvement, and propose your solution. Before I could do that I had to understand if it was a feasible option since I quit programming long ago and didn't have too much free time. I searched the Rhythmbox project page and landed on a Gnome wiki packed with information. Scrolling down on the page, you can reach a Development Resources section (see Illustration) where you are presented with several options. At the time I am writing this, the Getting Started link goes to a non-existing page, I am sure it will be fixed by the time you read it, in any case I believe the right page should be this one:
https://wiki.gnome.org/Rhythmbox/FAQ#Development_information

At the top, there is a link to a Newcomers Guide, which explains the process of how to get involved. On that page you find this sentence: "There are no requirements to start development. It can be an advantage if you know a bit of object oriented programming and git.". I do know OOP (Object Oriented Programming), but only a faint idea of how git works. Nevertheless, I dug into the list of recommended reading, and was fascinated with the whole Gnome project, and I encourage you to do so. Although I eventually went down a different road, I would like to point out how all the tasks I mentioned are easy and achievable in the open source world where projects are shared, source code is available, and there is a community of people who are working to spread the knowledge and encourage others to contribute. With proprietary software, you normally find Application Programming Interface (API) documentation, where available, and that's it!

Anyway, going back to my browsing the wiki, I realised that I had to digest some information before I could start typing some code, but along the way I came to understand that sorting information is stored in the Rhythmbox DB, and that is stored in an XML file on the computer. I believe the location of this file has changed over time, but now you should be able to find it in ~/.local/share/rhythmbox, and you should find a file there with the name rhythmdb.xml. If you use nautilus to browse the file system, you have to enable the view of hidden files (shortcut Ctrl+H) to see the .local folder.

I took a peek at the file itself, you can do it too by opening the file in a text editor or in a browser. XML files are human readable by definition, and that simplified my task a lot. I could easily identify the structure of information and understand that every song's information was packed in between two tags: <entry type="song"> and </entry>. Inside those two markers is stored all the information for each track; you can see an example in Text. I immediately realised where sorting information was stored... towards the end there are two fields: <mb-artistsortname> and <album-sortname>. Note that these are not part of the standard MP3 tag structure, so Rhythmbox can rely only on the one inside its database. Browsing the file, I discovered that some tracks had no sorting information despite that I entered them in the program itself. I also realised that if I entered them manually in the file they would stick, and that would partially solve my problem. If not, that it is a tedious job, you have to enter them for each and every track, and if you do something

wrong, you could possibly destroy the whole database.

Therefore I took the decision of writing some code that would do the job for me. I knew several languages with C and C++ my favourites. I never wrote anything in Python, but I had the chance of following the lessons from Greg D. Walters on this same publication. It sounded like a good compromise for me: I didn't go for the hard job of learning the whole Gnome/Rhythmbox development structure, but at least I would learn something new. Plus Greg and I both enjoy cooking!

As most people do when picking up a programming language, I went back to the first chapters of Greg's lessons and tinkered with the "Hello World" programs to understand a little how Python works. It was very easy (see FCM#27 onwards), and soon I felt ready to take on my challenge.

My objective was plain: given any album, I wanted to be able to enter quickly the sorting information. As an example, for Hunky Dory by David Bowie, I wanted all the tracks to have "Bowie, David" in the field mb-

artistsortname, and 1971 in the album-sortname field (the year it was released so that all the records would be chronologically sorted). To achieve that, I had to create something to manage the XML structure, look for track entries that match the given one (David Bowie/Hunky Dory), check if the sorting information is already there, if not create it, and then populate it. Now XML files are text files, but directly manipulating them seemed not a good option. Could I avoid the task of interpreting XML? A quick search in Google provided the answer: "python xml library" returned at the top of the page: https://docs.python.org/3/library/xml.html.

Now I had an idea, plenty of documentation, a programming language, and already a library to tackle the issue. Next month we'll see the code I wrote and how it works.

One of the best things about GNU/Linux is that it does not always require the latest and greatest hardware in order to work – it performs quite well at times when other operating systems would very simply require a new computer to run their latest versions, or the latest versions of user applications.

So there is a tendency in the community to hold on to hardware for longer periods of time, especially when the owner base may be slightly more computer literate than other sectors of the population. However, as with any technological product, it is to be expected that at some point the law of diminishing returns kicks in, and it becomes simply impractical to run any kind of modern GNU/Linux distribution on a very old computer. Using a computer based on an Intel i386 or i486 to do any real work should probably be considered highly unlikely at this point in time, but, in practice, where should one draw the line between usable and no longer usable hardware? In this piece, I

would like to consider a couple of real-world cases that may give the reader some points of reference.

As a first example, let us consider my Acer Extensa 5220 laptop that came out in 2008 and that I acquired new that year. This

computer sufficed quite well for my needs at the time, with one GB of RAM, an 80 GB hard drive - with the then-new SATA interface - and Windows Vista. That sets one back a step in time, I suppose. Although I upgraded the RAM to two GBytes at some point, it is clear that some

of the technical specifications cannot be belied: RAM must be of the slow DDR2 type since this is all the motherboard will support, the USB 2 ports must remain what they are, and I would be hard pressed to find a suitable replacement for the 802.11g WiFi Card. The processor is a single-threaded Intel Celeron 32-bit CPU running at all of 1700 MHz, and this is not upgradable. However, it still has a fairly nice keyboard, and the screen is fine for my eyes, so it stayed running off and on for the last decade. Even the battery retains about 10 minutes of charge, enough to set off the effects of a sudden power drop.

This computer is running the latest and greatest Linux Mint at the time of writing, version 18.3 Sylvia XFCE (i386) with all applications updated to their latest versions. Aside from Dropbox - which does seem to bog things down a bit with slower hard drives - I have every reason to be satisfied with the computer's speed. The boot process, from GRUB menu to desktop, takes 126 seconds,

MY OPINION

Firefox starts up in 10 seconds, and LibreOffice in 7. All this means that if I make some provision for booting (ie, go get my coffee while it does its thing), I can actually use it for some productive work in the fashion it was meant to be used. Even the Firefox and Chrome browsers seem to work relatively well in JavaScript hell, though I do admit to using mostly LibreOffice for text editing on this particular machine.

So, is it still worth my money to upgrade further the existing hardware? At this point, just about the only part that could be swapped for something new would be the hard drive. A rather cheap SSD drive is an obvious candidate. In internal 2.5" form, with a SATA connector, these days they can cost anywhere from about 40€ (35 GPB, 50 USD) new - if one is not too demanding on quality and capacity - to about zero if there is a spare drive lying around in a drawer somewhere. Some of the earlier drives with capacities of 32 or 64 GBytes could very well be doing just that, having been replaced with larger sizes in a production computer somewhere along the line. This is the path I took, with a 32 GByte SSD drive or,

more precisely, a 32 GByte m-SATA drive, in card format, placed inside a 2.5" format adapter. I connected the new drive to the computer with an external USB-to-SATA interface, formatted it using btrfs and transferred both system and user files over using the btrfs

send/receive command. I then reinstalled GRUB on the new drive with the grub-install command using the --boot-directory parameter, and unplugged the usb drive. I then opened up the laptop, and physically swapped the old rotational drive for the new one.

One nice thing about working with this class of computers is that one is not overly concerned about messing things up. The financial implications would be rather low, even if the worst happened. A second point is that access to the computer's internals is actually

Windows Vista OEM sticker (a relic of the past)

Battery receptacle

Disc drive

WiFi

RAM

CPU

quite good for this particular model, as for other laptops of the same time-period. A single cover needs to be removed (held on with 8 Phillips 00 screws), and then RAM, hard drive and WiFi card can all be accessed together. Perhaps this could be a good way of introducing youngsters into the arcane arts of actually working on their devices. But I digress…

Once the hard drive had been swapped, I timed the same actions. Boot from GRUB to desktop was now 103 seconds. So, about a quarter of a minute had been shaved off, but both the limitations of the CPU and of the hard drive's SATA I bus (just 150 MBytes/s, compared to ATA's 133 MBytes/s) mean that the speedup is certainly there, but may be seen as marginal as regards actual usability. As for applications, Firefox now started up in 7 seconds, while LibreOffice took a measly 4 seconds. So, while the new values are better than the old, it is clear that the difference is probably not something that could radically alter our current use of this computer.

For this class of machine, I would probably conclude that upgrading some of the hardware would certainly have an educational purpose to it, but that the benefit for practical use of the laptop would at best be marginal.

As a second example, let us take my Acer Aspire One (AO-722) netbook. This one dates from 2011, has an AMD C60 processor with one core but two threads - still at 1600 MHz total frequency, but each thread can take 800 MHz separately - and it is actually a 64-bit machine. The original complement was a 320 GByte rotational hard drive, two GBytes of the slightly faster DDR-3 RAM, and Windows 7 Home Edition. The screen is a 11.7" unit that has its quirks, but still works well. Unlike the Extensa, this netbook does not have an optical unit. Battery life is pretty low after this time, but, with a bit of care, one can eeke out about one hour of power. It currently runs Linux Mint 18.3 Sylvia with the Cinnamon desktop (64-bit). On the positive side for upgraders, access is really easy, with a bottom panel that comes out by unfastening a single Phillips 0 screw.

With the original two Gbytes of RAM and rotational hard drive, boot-up time from GRUB to the (cinnamon) desktop was 65 seconds, launching Firefox took 10 seconds, and LibreOffice 13. This result is interesting in itself. This computer could be expected to be quite a bit faster than the Extensa, by virtue of its superior CPU, faster hard drive connection (SATA-II instead of SATA-I), and faster RAM. While general system boot-up times are coherent with this line of reasoning, Firefox seems to stick to its 10-second times, and LibreOffice is now noticeably slower than before. I am loathe to give a precise explanation for this behavior, but perhaps some of the blame can be placed on using 64-bit versions of this software, which are known for taking up slightly more disk space and RAM than the 32-bit versions used on the Extensa. Other causes related to system libraries may also come into play.

Upgrading the RAM installed to four GBytes gives no speed advantage in the least: 68 seconds to boot the desktop, 10 seconds to launch Firefox and 14 to launch LibreOffice. Within the margin of error of the previous measurements, it is clear that putting more RAM into a computer that is not actually using it (e.g. under heavy use) is not much use at all.

As before, I ended up swapping the rotational hard drive for an SSD, this time a NGFF 32 GByte unit also installed inside an adaptor case to fit into the Aspire's 2.5" SATA bay. As before, I mounted the new hard drive using an exterior USB cable, formatted it using btrfs, and transferred the system and data snapshots over to the new drive using the btrfs send/receive commands. I then installed GRUB on the new drive, and installed it physically in the computer.

With the new SSD drive, boot time went down to 45 seconds, while Firefox and LibreOffice both launched in 8 seconds. So a noticeable improvement has been seen here, making the drive swap perhaps a worthwhile investment for this class of computer. This bump in speed is actually even more visible when the computer has a more powerful CPU. A Lenovo from a similar time frame (late 2011) boots up in less than 20 seconds, helped no doubt by its then-powerful Intel Core i5 dual core (four-threaded) processor.

Summing up all these factors, my final recommendations would be as follows. In the first place, when working with older 32-bit hardware, a maximum of two GBytes of RAM should be plenty to run most applications. It is also probably not worth the time and hassle to change the hard drive for something faster, since the motherboard will likely not be able to use the supplementary speed. It would be best, to my mind, to simply use lighter distributions such as Xubuntu or Linux Mint XFCE editions, with Linux Mint MATE or Lubuntu as strong alternatives.

On the other hand, once we are speaking of hardware that was formerly of a slightly superior class, such as a 7-to-8 year-old Core i5 or other 64-bit processor, in that case upgrading to a 64-bit version of our favourite distribution will mean that RAM requirements should probably begin at about two GBytes, and move on up from there. Most of these upper tier user-grade or entry-level business-class computers will usually have come with about four GBytes of RAM, which is perfect for most modern GNU/Linux distributions.

Naturally, owners with special use-cases, such as video editing, may wish to invest in more RAM - and yes, I have successfully edited quite long videos on my 2011 Lenovo with no difficulties.

In any case, this class of computer will benefit more from a faster hard drive, so upgrading to an SSD model can be seen as a smart move. Besides the clear advantage in speed, it may also be noted that these drives with no moving parts tend to support better use and abuse when on the move, beside requiring less power to work. So changing the drive will usually impact battery life favourably, and make the computer less susceptible to loosing data from hits and falls.

A final reflection I would like to make is that most GNU/Linux developers are definitely in the power-user category. As such, they tend to favour high-end hardware. For this reason, when on the lookout for an older computer to repurpose with any of the variants of Ubuntu or most other mainstream distributions, it may be worthwhile to try and get a

more powerful piece of equipment than is really needed. Business-level computers are more likely to be completely supported than some entry-level consumer products. If in the market for a second-hand unit, price differences are no longer quite as noticeable as when buying new, so why not? These computers may also turn out to offer more upgrade paths than lower-end units. An example of this would be my decidedly lower-end Aspire, where one of the RAM slots has never even been solderdered into the motherboard. This would never have been omitted in a power laptop. As usual, though, some precautions need to be taken, such as bringing your bootable USB along when choosing the computer.
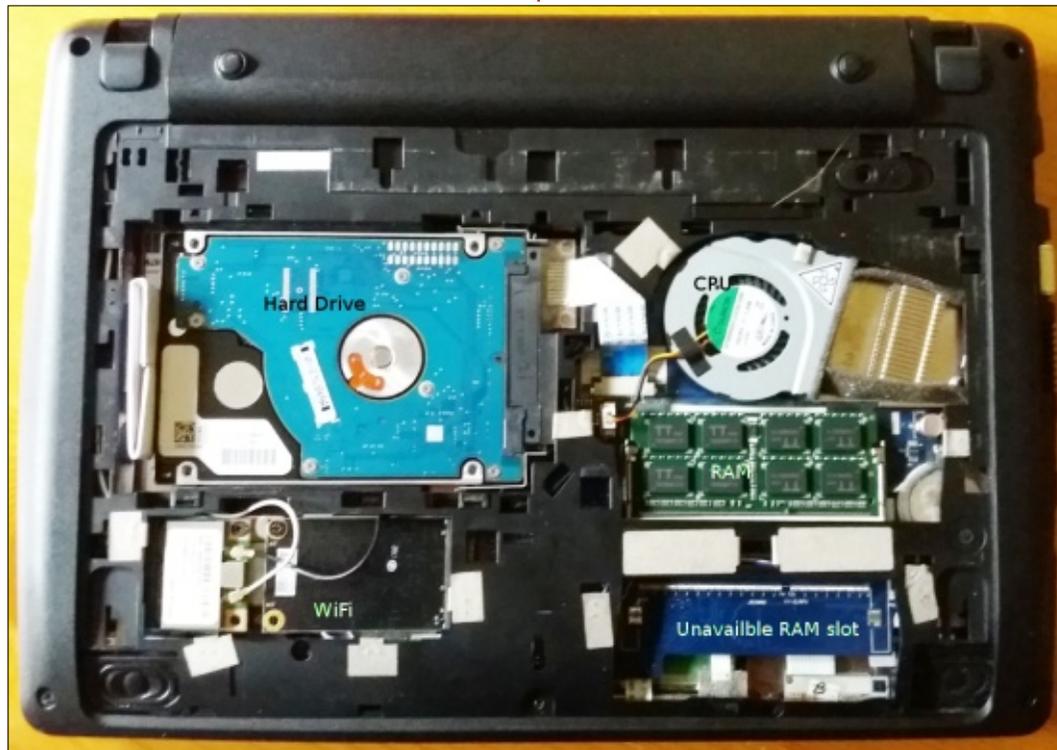
Alan holds a PhD in Information and the Knowledge Society. He teaches computer science and engineering at Escola Andorrana de Batxillerat (high-school). He has previously given GNU/Linux courses at the University of Andorra and taught GNU/Linux systems administration at the Open University of Catalunya (UOC).

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.

As some readers may remember, I am still on a quest to get some real use out of my HP 255 G3 laptop. This rather economical computer - less than 250€ (US $300) when new - came with a rather underwhelming AMD E1-2100 system-on-a-chip, 4 GB of RAM, a 1 TB rotational hard drive: not bad for the price, but certainly more comparable to early Chromebooks than a professional laptop. Both the original Windows 8.1, and the Windows 10 upgrade, struggled to perform in any meaningful way, and - much to my chagrin - so did most Ubuntu derivative distributions back in early 2016. The result was that I have been using this system mostly in text mode since then.

Two years later, however, I have noticed that many GNU/Linux distribution developers seem to have been busy making their offerings more agreeable to lower-powered computers, specifically since Ubuntu 17.10 came out. Earlier versions also have been progressively made slightly less resource-greedy or, at least, that has been my impression concerning Ubuntu 16.04 LTS. This optimization means that even the most modern distributions can often be used on earlier hardware. I wanted to find out if this progress could actually be translated into making my laptop useful again as an actual laptop, with a full graphical environment. To do so, I kept the hardware at precisely the original specifications, though, if things go well, an SSD hard drive upgrade may be in the works.

The Neon distribution (https://neon.kde.org/) is a distribution centered on the KDE Plasma-5 desktop, and based on Ubuntu 16.04. Though it does share quite a few similarities with Kubuntu, the aims of this distribution are rather different. Neon focuses on the desktop itself, which it showcases with the most recent versions of the software. The fact that it is based on Ubuntu seems to be less of a concern for its developers, who basically just needed a very stable operating system as a platform to build their desktop on. Ubuntu 16.04 seems to have fit the bill quite well.

The original version of the Plasma desktop in Kubuntu 16.04 was version 5.5.5, going up to 5.10.5 in Kubuntu 17.10. Newer versions are available in the backport PPA repository, currently with version 5.12.0 of the software. In Neon, the project repositories are added out-of-the-box, with newer versions of the Plasma desktop and auxiliary software:

```
deb
http://archive.neon.kde.org/user xenial main

deb-src
http://archive.neon.kde.org/user xenial main
```

In my opinion, recent versions of the Plasma desktop have had time to evolve into a nice, mature environment. This can be seen with seemingly simple operations, such as using the Alt + Tab key combination to switch between applications. A very clear list of open windows is presented to the left of the screen, and the user can see precisely which application is under focus. Application windows are also seen overlayed in
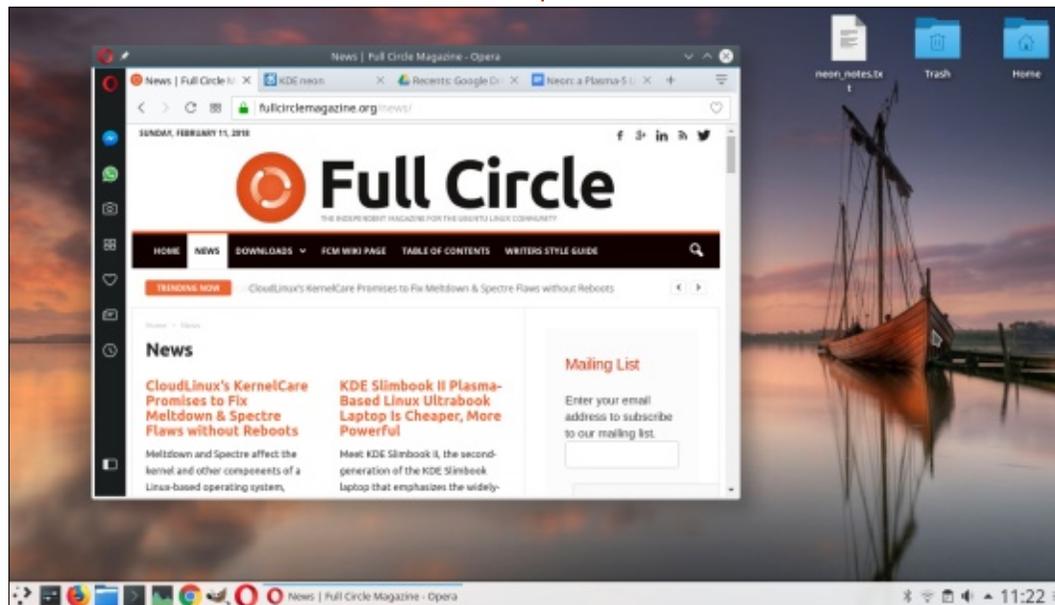
transparency in the main desktop area, with the topmost completely opaque and easy to see. This is of help when choosing between various windows of the same application - think of several Konsole terminals open at once, for instance.

In essence, one could very easily transform a Kubuntu installation into a Neon, in much the same way an Ubuntu 17.10 can be transformed into a Xubuntu or Ubuntu Mate, for example. However, installing Neon instead of another *Buntu distribution has two specific advantages, that may or may not be pertinent for specific users. In the first place, Neon has the most recent Plasma software, with no need to fiddle around with the repositories. This makes sense if one wishes to stay as current as possible with the KDE desktop. The second point is that Neon has no supplementary software, which makes for a smaller download file (1.2 GB, instead of up to 2 GB or more for some recent distributions) and a quicker installation process. Applications such as LibreOffice are not included by default, though they can be added with ease using the regular Ubuntu

software management techniques. Depending on which use the system will be given, it may be more handy to install just those packages that are needed by the user.

Speed is quite good. System startup and shutdown are performed in a reasonable time for the hardware. There is, in fact, not much difference between Neon and Xubuntu 17.10. For some reason, Neon seems to be more responsive than a comparable version of Kubuntu or Linux Mint 18 KDE version. To my mind, this is a symptom of good systems integration between the underlying operating system, the SDDM display manager, and the desktop itself. Applications also come up nicely, though with a bit of a delay between clicking on the icon and seeing a window, especially when the desktop is still settling down after initialization. This slight lag will probably not be a nuisance for calm users. Perhaps power users would need to consider running Neon on more powerful hardware, which they would probably already use in any case. As usual with KDE, the system configuration application has many options that can be used

to activate graphical effects as required. I found that the default set works well and is not too intrusive. However, users using just the Intel processor internal graphics circuit may wish to quit some of the animations to get better performance.

The sore point will, as usual, be the web browser. Though truly necessary to make actual use of a computer in this day and age, browsers are also known to hog processor power and computer memory, moreso when that bane of the Internet known as JavaScript is in use. I will admit that it does have some advantages when trying to make a web page more pretty and perhaps make browser-to-server communication more responsive. However, it must also be said that some abuse of this technology can be seen out there – for nefarious purposes or not, much in the same way Adobe Flash had previously become a vector of both the best and the worst in programming techniques.

At the time of writing, a fully upgraded version of Neon comes with Firefox "Quantum" version 58.0.1. It works quite well even with Google Drive and suchlike,

Information About This System

KDE neon 5.12
http://neon.kde.org/

Software
KDE Plasma Version: 5.12.0
KDE Frameworks Version: 5.42.0
Qt Version: 5.9.3
Kernel Version: 4.13.0-32-generic
OS Type: 64-bit

Hardware
Processors: 2 × AMD E1-2100 APU with Radeon™ HD Graphics
Memory: 3,3 GiB of RAM
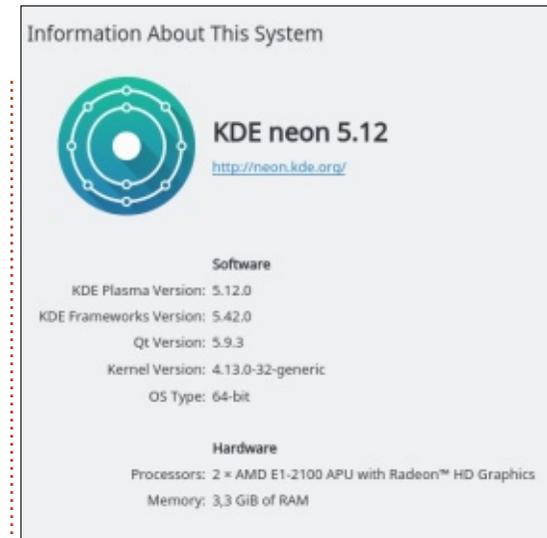
much quicker in fact than previous versions of this browser. This concurs with my impression that developers are in fact listening to users; fair play to them for that. I actually quite enjoyed using Firefox for browsing and for some online editing, to the point that I saw no need to install a more lightweight browser such as Midori or Epiphany. Both of these are excellent choices for simple browsing; however, many of the heavier Web platforms - such as Google Drive or Apple's iCloud - tend to whinge a bit about "your browser not being compatible".

Google's Chrome and the Opera browser are two modern alternatives that can very easily be installed in Neon. For Chrome, just

download the .DEB package from https://www.google.com/chrome/, then install using the following commands (please note some extra packages need to be pulled down using the aptitude or apt commands):

```
sudo bash

cd Downloads

dpkg -i google-chrome-stable_current_amd64.deb

aptitude install fonts-liberation libappindicator1
```

In the case of Opera, the same procedure applies after downloading the .DEB package from http://www.opera.com :

```
dpkg -i opera-stable_51.0.2830.26_amd64.deb

aptitude install libcurl3
```

All of these three browsers work well, within hardware limitations. In most workflows, RAM usage can be kept underneath 2 GB, which, on a computer with a mere 4 GB, still leaves us some free space. However, the processor does tend to struggle a bit on these low-spec, cheap laptops. With my AMD E1-2100, both cores are pegged at
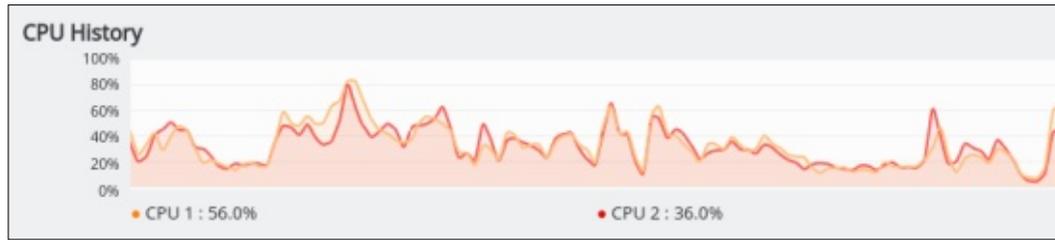
about the 50% mark even while doing relatively light tasks. In this screen capture, I was editing a single Google Document (this article) in Google Chrome. This gives you an idea of what to expect under real-world conditions: as long as you do one thing at a time, and keep unnecessary windows closed, the computer and operating system should hold up quite well. However, multitasking and editing lengthy documents would probably be better left to a more capable piece of kit.

The takeaway on Neon is that this modern distribution offers a very nice desktop environment, with all the bells and whistles of the most recent version of Plasma 5. Even so, it can be used with good results on a rather low-powered computer. Some caveats do apply, as usual. The main point I should stress is that the user should definitely not expect much from a system with 2 GB of RAM (or less). As usual with KDE and Plasma, more processor power, or

a better graphics card, can be put to good use, when available. On a system such as this one, swaping the rotational hard drive for an SSD unit will surely be of help to speed things up a bit. However, the limiting factor is clearly the processor, and putting in a faster disk would no doubt result in less of an increase in performance than with a more capable CPU. In this sense, the Neon distribution would surely marry itself quite well to a computer that is a bit older, but equipped with an Intel Core i5 or i7 CPU and modernized with an SSD drive. So, if the reader has one lying around and is in need of a project… Why not?

Alan holds a PhD in Information and the Knowledge Society. He teaches computer science and engineering 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).

LETTERS

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.

BACK NEXT MONTH

Deus Ex: Mankind Divided

Probably one of the best games available for Linux at the moment is Deus Ex: Mankind Divided. Mankind Divided is the sixth and latest game in the Deus Ex video game series. The first game, titled Deus Ex, was released back in 2000 for Microsoft Windows & Mac OS X with an eventual release for PlayStation 2 in 2002. Deus Ex: Mankind Divided was released in 2016 for Microsoft Windows, PlayStation 4, Xbox One & Linux, with a follow-up release for Mac OS X in 2017. Mankind Divided is the first game in the series to be released for Linux. Mankind Divided was developed by Eidos Montreal & released by Square Enix with the Linux port provided by none other than Feral Interactive.

Deus Ex: Mankind Divided is currently selling for $29.99 through Steam, Humble Bundle, or you can get it straight from Feral Interactive. This game can morph into many different genres – which is one cool thing that makes it very interesting. I've seen the game described as a stealth, sci-fi, cyberpunk, action, single-player, FPS, RPG game.

The year is 2029, and, revolutionary advances in technology and medicine have brought us what are known as augmented humans. When someone needs a transplant, or another type of life saving surgery, they can get robotic transplants – which turn that person into some sort of bionic human, or, as the game calls them, augmented humans. In Deus Ex: Mankind Divided, you play the role of Adam Jensen, an augmented officer working for Task Force 29, a special Interpol unit. Adam Jensen has had his share of accidents and medical emergencies, so, as a result, he's had a number of body parts replaced and augmented to the point of being a very efficient augmented officer.

What sets the Deus Ex series

apart from many other stealth games is that, in addition to its stealth elements, you also have to deal with the augmentation aspect of the game. For example, every now and then, when Adam needs surgery, he gets an augmentation added to his list of abilities. One of the first augmented abilities that Adam gets is the ability to see through walls. You have to push a key on your keyboard (or the Left Shoulder Button on a gamepad controller) to activate (or deactivate) this visual augmentation. Some of the other augmentations you get are a the Glass-Shield Cloaking system which basically lets you become invisible, the NanoShield which makes you invincible, and, later, you get a fist-mounted electro-shock system which basically replaces the stun gun that you have in your inventory during the early levels.

However, using any of these augmentations drains your Energy level. Each augmentation has a minimum level at which it can be used – which also means that while you're using an augmentation, and your energy level is dropping, eventually your energy drops so low that your augmentation can no longer work. This energy system being tied in with the augmentations makes it so that you can't just become invisible at the beginning of a level and walk through the entire level without being noticed. Speaking of trying to complete a level without being noticed, apparently it's possible to beat the entire game without being noticed, and without killing anyone, but I'm assuming you need lots of patience and practice to do so.

There are many ways to play the game – which makes the RPG concept really shine in this regard. You can go the route of the full stealth player that moves about unnoticed, or you can try to play as a mean ruthless killing machine. As
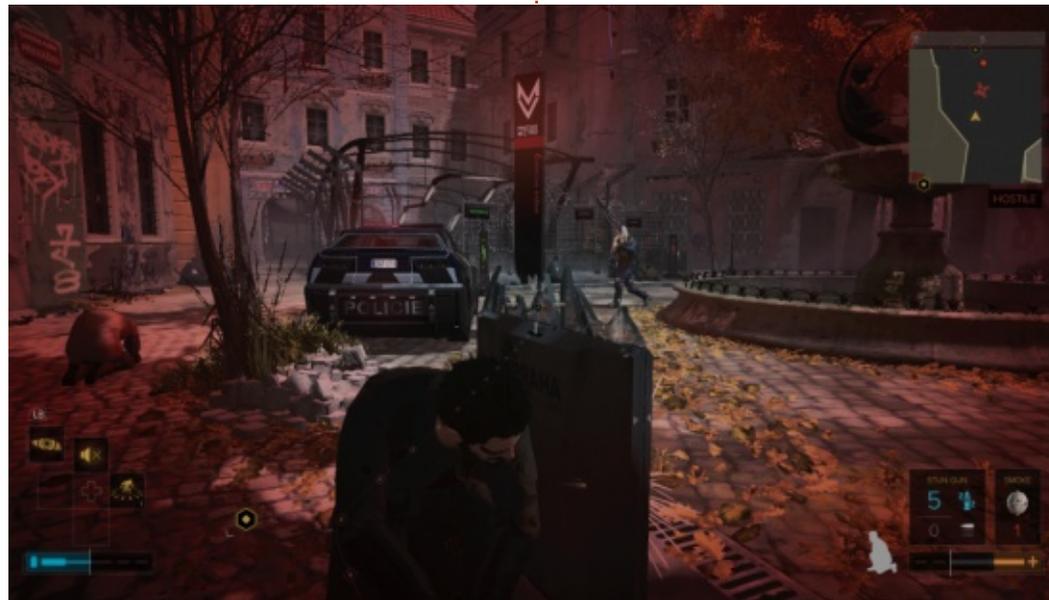
for me, I play the game in a very sloppy yet efficient manner.

The game is a true joy to play. Being that it's first and foremost an action, RPG, stealth game, I decided to play it with a gamepad controller, but you can also use a mouse/keyboard if you so desire. I had to make a couple of adjustments to match my prior experience with stealth games, but, once I had everything set-up the way I like it, I was able to play the game in a very natural manner. You can have a running tutorial, which teaches you new moves as you get new augmentations or learn new tricks. Even with the tutorial, there is still some extra stuff you'll have to learn on your own, such as hacking terminal servers which deactivate security measures in the buildings which you're infiltrating.

With the visually stunning scenes, and the superb voice-acting, Deus Ex: Mankind Divided feels like a complex cinematic gaming experience. The sound effects are also very impressive. Being that this game has some pretty high minimum system requirements, the better hardware you have, the better the game will look and feel. Yet, I was pleased to discover that even with my slowly aging Linux box, I was able to fully enjoy this game. The trick is to let the game decide what settings are optimal (which it will do on its own anyway), or choose one of the lowest default pre-sets so as to get better performance. If your PC can handle it, then you're in for a cool ride as you'll be enjoying some really cool graphics. Also, before you fire up the game, make sure that there are no other programs running that might try to steal some resources from the game. A few times I made the mistake of having my web browser running in the background, which would sometimes cause the game to lag a little bit. With the advice I've given,

and the recommended hardware, the game should play pretty smoothly.

Only twice has the game frozen on me and both times were right after I tried taking a screenshot with Steam's built-in Screenshot recorder. For the screenshots I've included in this article, I relied on the default Linux PrtScn via the button on my keyboard.

Other than the problems I've described – which are related to the demanding system requirements – I have only one other complaint about this game. The very first time I played the game, it took about 20 minutes for the game to load. I was kind of concerned that it would always take that long to load, but, from the second time I played it until the present, it's taken about 3 minutes total to load up, a huge improvement from the initial 20 minutes. However, after a gpu driver update, expect another long wait lasting about 10 minutes. A minor nuisance, but then again, lots of AAA titles will have a high demand on hardware, take up a lot of space in your hard drive, and take a long time to load up.

FYI - On my computer, the game took up around 45 GB hard drive space. Also, even though my CPU is just shy of the minimum requirements (AMD FX6300 vs the required AMD FX8350), the game ran just fine as long as I didn't have any other resource-hungry app running in the background.

Overall, I strongly recommend this game. It's got a great story-line, awesome graphics, great sound-effects, amazing voice acting, and it plays like a charm. It's a great game if your system can handle it. One star deduction for the long load-up time & the demanding minimum requirements.
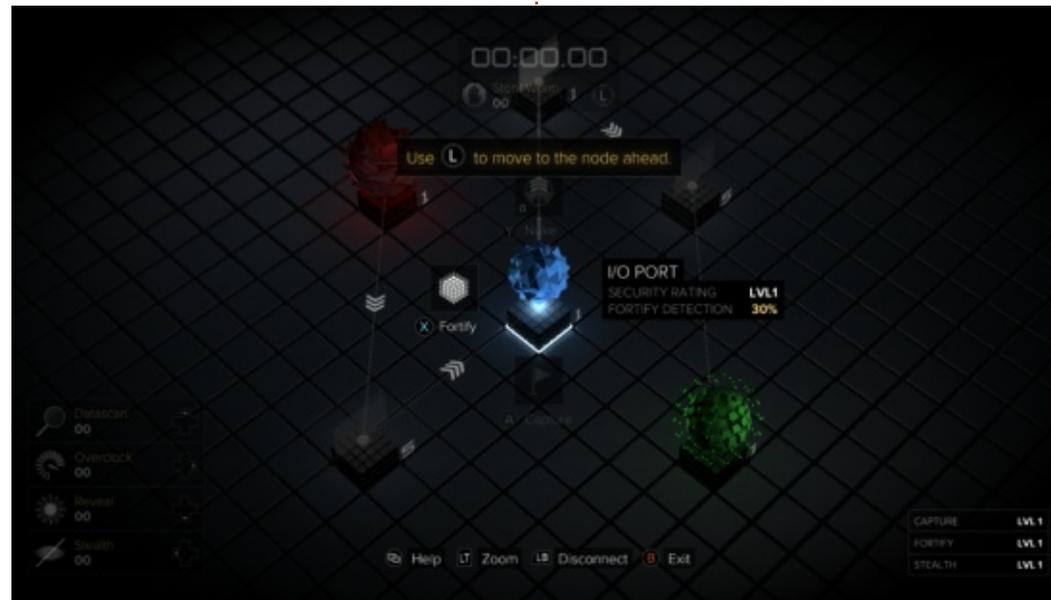
MINIMUM REQUIREMENTS

OS: Ubuntu 16.04 64-bit, SteamOS 2.0

CPU: Intel Core i3-4130, AMD FX8350

RAM: 8 GB RAM

GPU: 2GB Nvidia 680 (driver version 367.57 or higher), AMD GPUs not supported

Hard Drive: 67 GB available space

MY GAMING BOX

AMD FX-6100 3.3GHz CPU (over-clocked to 3.5GHz)

Nvidia GeForce GTX 960 graphics card with Nvidia 396.24 driver

16GB of Kingston Hyper X RAM

Ubuntu 16.04 LTS (64-bit) with Unity desktop

Oscar graduated from CSUN, is a musician, game enthusiast and has been working with Bitcoin and other alt-coins. You can follow him at: https://twitter.com/resonant7hand or email him at: 7bluehand@gmail.com

PATRONS

Monthly Patrons

2016 - Present:
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
Joao Cantinho Lopes
John Andrews

2017 - Present:
Matt Hopper
Jay Pee
Brian Kelly
J.J. van Kampen

2018 - Present:
John Helmers
Kevin O'Brien
Kevin Raulins
Carl Andersen
Charles Stewart
Dave Nelson
Brian Bogdan
Dennis Shimer

Single Donations

2017:
Linda Prinsen
Shashank Sharma
Glenn Heaton
Frank Dinger
Randy E. Brinson
Kevin Dwyer
Douglas Brown

Daniel Truchon
John Helmers
Ronald Eike
Dennis Shimer
Iain Mckeand
Jaideep Tibrewala
Kevin Dwyer

2018:
Yvo Geens
Graig Pearen
Carlo Puglisi
James A Carnrite
John Holman
P G Schmitt
Robert Cannon
Thomas A Lawell
Ronald Le Blanc
Luis Eduardo Herman
Glenn Heaton
Peter Swentzel
Alain Mallette
Christophe Caron
Linda Prinsen
Ronald Eike
Anthony Cooper
Louis W Adams Jr
Joseph Tong
Robert G. Wells
Robert Kaspar

The current site was created thanks to **Lucas Westermann** (Mr. Command & Conquer) who took on the task of completely rebuilding the site, and scripts, from scratch, in his own time.

The Patreon page is to help pay the domain and hosting fees. The yearly target was quickly reached thanks to those listed on this page. The money also helps with the new mailing list that I set up.

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

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

https://www.patreon.com/ fullcirclemagazine

https://paypal.me/ronnietucker

https://donorbox.org/recurring -monthly-donation

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#134
Deadline:
Sunday 10th June 2018.
Release:
Friday 29th June 2018.

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 you're 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 issue's 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.