



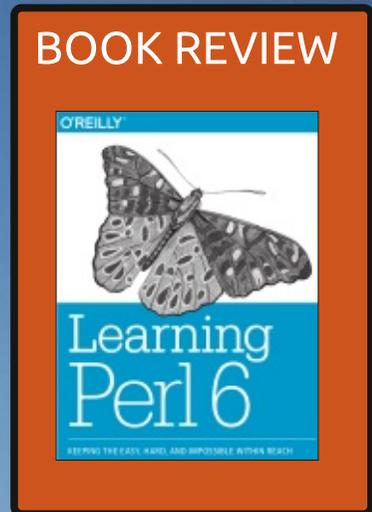
Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

ISSUE #145 - May 2019



Photo: Martin Svedén
(Flickr.com)



GRAMPS

A QUICK LOOK AT CREATING A FAMILY TREE

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



HowTo



Python p.23



GRAMPS p.26



Using Netcat Pt2 p.28



p.XX



Inkscape p.29

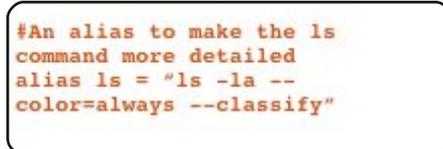


Graphics

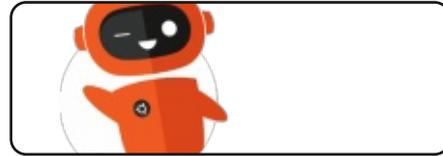


Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY



Command & Conquer p.20



Ubuntu Devices p.40



Book Review p.44



Q&A p.50



Linux Loopback p.33



The Daily Waddle p.41



Letters p.XX



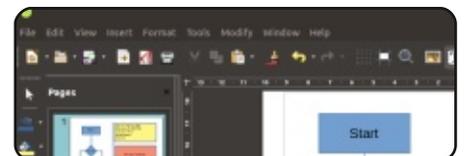
Linux Certified p.53



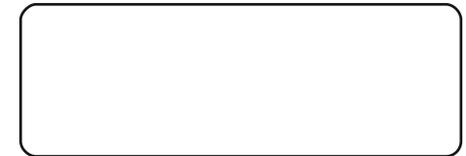
Linux News p.04



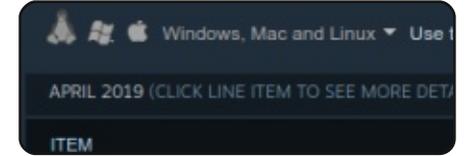
Everyday Ubuntu p.37



My Opinion p.42



Interview p.34



Ubuntu Games p.56



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

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



WELCOME TO ANOTHER ISSUE OF FULL CIRCLE!

Joining Python and Inkscape we have a second article from Erik on netcat. With no Freeplane for you this month, we have a quick look at GRAMPS. If you've ever thought about doing a family tree then GRAMPS is an essential piece of software to keep all that info in order. I've used it myself in the past. Incredibly powerful for keeping sources in order and creating those essential family reports.

SJ has another BSD article for you this month and an interview with GhostBSD developer, Eric. Next month he'll be looking at Project Trident. Elsewhere, Richard explains where to look if you need help with Ubuntu, Erik gives his opinion on LibreOffice Draw, answers your questions, talks more about getting Linux Certified, and tells you what he thinks about the book Learning Perl6. Oscar marks his return to Ubuntu Games with a look at the annual Steam Hardware Survey.

This month also marks the release of Ubports Touch OTA-9. More bug fixes for everyone's favourite mobile OS.

All the best, keep in touch, and heres to another 12 years!

Ronnie

ronnie@fullcirclemagazine.org



This magazine was created using :



Find Full Circle on:



goo.gl/FRTMl



facebook.com/fullcirclemagazine



twitter.com/#!/fullcirclemag



<http://issuu.com/fullcirclemagazine>



<https://mastodon.social/@fullcirclemagazine>



Weekly News:



<http://fullcirclemagazine.org/feed/podcast>



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



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

FCM PATREON : <https://www.patreon.com/fullcirclemagazine>



DEEPIN 15.10 IS HERE -- DOWNLOAD THE MOST BEAUTIFUL LINUX DISTRIBUTION NOW!

Windows 10 is a functional operating system, but goodness, it is not at all attractive. When I use Microsoft's OS, I don't get feelings of joy or happiness. Windows 10's design is very bland and seemingly uninspired. By comparison, Apple's macOS makes me very happy. And yes, appearance matters -- a good user interface and design can motivate the user and help them to be more creative.

As great as macOS is, there is a Linux distribution that rivals its beauty. Called deepin, this operating system is now based on Debian Stable (previously Debian Unstable), and provides an absolutely drop-dead gorgeous user interface. In fact, I am confident to say it is the most beautiful Linux distro. The newest version, deepin 15.10, is now available for download. It is chock

full of bug fixes, and also, many new features. One of the most significant changes is dde-kwin is now the default window manager. Users should experience improved performance as a result.

Speaking on new features, the team explains, "When the 'Auto merge' option is checked in desktop context menu, files on desktop will be automatically grouped into different folders named by Videos, Music, Pictures, Documents, Applications, and Others, keeping your files on desktop in order. When the 'Wallpaper Slideshow' option is checked in wallpaper settings, users can set the slideshow interval to have a new wallpaper constantly! The settings for sound effects (such as shut down, log out, wake up etc.) in Control Center are separated, so that users can turn on or off a certain system sound through the switches, and have a better control of desktop environment."

Source:

<https://betanews.com/2019/04/28/deepin-1510-linux-debian/>

DEBIAN GNU/LINUX 9.9 RELEASED WITH OVER 120 BUG FIXES AND SECURITY UPDATES

Debian GNU/Linux 9.9 is here two and a half months after the Debian GNU/Linux 9.8 point release as yet another up-to-date installation media containing all the latest security updates and bug fixes released on the main archives. It can be used for fresh installations without downloading all updates after the installation.

The Debian GNU/Linux 9.9 "Stretch" maintenance update includes a total of 122 changes, consisting of miscellaneous bug fixes that add important corrections to 70 packages, as well as 52 security updates addressing various of the latest security vulnerabilities and other related

issues.

Therefore, if you're already running the Debian GNU/Linux 9 "Stretch" operating system series on your computers, all you have to do is update your installations by installing all the updates from the official software repositories. To do that, you can run the "sudo apt-get update && sudo apt-get dist-upgrade" command in a terminal emulator.

Debian GNU/Linux 9.9 "Stretch" install and live mediums are not yet available to download at the moment of writing this article, but advanced users can use the network installation method. We'll let you know when Debian GNU/Linux 9.9 ISOs are released so you can download them if you want to reinstall or deploy Debian Stretch on new computers.

Source:

<https://news.softpedia.com/news/debian-gnu-linux-9-9-released-with-over-120-bug-fixes-and-security-updates-525803.shtml>

DEBIAN GNU / LINUX 9.9 RELEASED

About two months after the release of Debian 9.8, Debian 9.9 followed this weekend. The ninth update of the current stable release of Debian has an average size and fixes major bugs in 70 packages. As always, all such changes are limited to what is absolutely necessary to avoid regressions. In addition, security issues have been resolved from 51 Debian Security Advisories (DSA). Five packages were removed from the distribution due to incompatibility with current versions of Firefox ESR and Thunderbird.

The announcement points to something special when updating to Debian 9.9. Users who still use the outdated apt-get tool should work with the dist-upgrade command while apt or aptitude users should use the upgrade command.

Problems were fixed in the Debian kernel, in the Debian installer as well as in Flatpak, LibreOffice and Systemd. The five

removed packages are add-ons to Firefox ESR and Thunderbird, which are no longer compatible with their current versions. Among the packages that received security fixes, packages from the web space are particularly prevalent. These include Drupal, Firefox ESR and Thunderbird, which each have closed several gaps. Systemd also resolved two security issues.

Source: <https://www.pro-linux.de/news/1/27010/debian-gnulinix-99-freigegeben.html>

NVIDIA CREATES FREE VIRTUAL LINK DRIVER FOR LINUX

The current Turing generation of graphics cards from Nvidia offers a so-called Virtual Link, which is designed as a commercially available USB-C socket and should support future VR headsets via a single cable. Another part of the necessary free Linux drivers is now included in the USB branch of kernel developer Greg Kroah-Hartman. The driver should therefore appear with the upcoming version 5.2 of the Linux kernel. Previously, Nvidia had

already used firmware and a driver for the USB-C controller.

The now available different contributions to support the technology are mainly used to implement the Virtual Link as Alternate Mode for USB-C. As can be seen from the patches, the Virtual Link is not officially standardized by USB Implementers Forum, so the manufacturer has to use its own ID for this mode.

In addition, the Virtual Link extends the Displayport Alternate Mode to use the USB 2 pins, which is not actually provided in the Alternate Mode. This is probably why Nvidia has to rely on its own manufacturer ID for the technology. Since the communication to Displayport is not different, the individual devices need no driver of their own, as stated in the patch. Nvidia instead relies on a review of the devices on the feature.

In addition, an Intel developer was involved in the implementation, which is responsible for the Displayport Alternate Mode in the kernel. These patches prepare for the use of Alternate Mode via UCSI. It is a

special interface (PDF) used to control the controller responsible for USB-C.

Source: <https://www.golem.de/news/vr-nvidia-erstellt-freien-virtual-link-treiber-fuer-linux-1904-140930.html>

LINUX-POWERED ATOMIC PI IS A BITE-SIZED PC WITH INTEL CPU

Atomic-Pi is an ultra-small Raspberry-pi alternative made to utilize the power of Linux. With a \$35 price tag, the Atomic Pi features more bang per buck and an Intel CPU to boot.

The Atomic Pi comes pre-installed with Linux and has several connector pins, external storage capabilities, faster memory, and a pretty huge heat sink. It is suitable for people who have always wanted to run x86-based apps on their miniature computers.

Atomic Pi uses an x86-based Quad-Core Atom CPU whose performance is similar to that of Snapdragon 801. The CPU model

number is x5-Z8500 and it is 80MHz faster than the older Z8500. The Intel Atom CPU uses Turbo Boost very efficiently due to its updated Airmont architecture. Intel Atom CPU has a base clock of 1.44 GHz and a boost clock of 1.92GHz on all cores.

The CPU features Intel Cherrytail integrated graphics. It is quite old and based on Intel 8th Gen architecture. However, the GPU does support 4K/H.26 video acceleration and Direct X 11.2 support. The graphics processor is capable of running at 500MHz and has 12 Execution Units. Its performance is slightly lower than a Qualcomm Adreno 330.

The Linux-powered Atomic Pi has 2GB of DDR3 RAM, which is plenty enough for light computing. The small PC features MediaTek RT5572 broadband Wifi capable of connecting to the 2.4GHz or 5GHz network. For wired connectivity, a Realtek RTL8111G-CG ethernet controller device is also present.

Atomic Pi features Bluetooth 4.0 connectivity, HDMI output, 16GB of eMMC flash memory and 256GB expandable storage slots. The small developer PC also

features 26 GPIO pins to connect sensors, actuators, etc.

The size of the board with all its components is 130x100x50 mm. Atomic Pi has a large heat sink almost 30 mm in size and it features a physical clock and a battery for the increased power requirement of the X86 processor.

The Linux-powered Atomic Pi also supports Windows OS. The base price of the device is \$35 but you'll need an additional power supply that delivers up to 15W to utilize its full potential.

Source:
<https://fossbytes.com/linux-atomic-pi-intel-cpu/>

PARROT 4.6 LINUX DISTRO FOR ETHICAL HACKING RELEASED WITH NEW KDE DESKTOP OPTION

When we talk about Kali Linux alternatives, options like Parrot Linux and BlackArch often turn out to be the top contenders. There are Windows-based options like Commando VM as well, but

Linux-based ethical hacking distros are the go-to options for the security researchers.

The Parrot Linux team recently announced the release of the latest Parrot Linux 4.6. It's a result of three months of a rigorous development cycle. It's a big milestone for the team as well as they've now shifted everything to their own infrastructure and this is the first release utilizing the same.

The Debian "Testing" branch-based Parrot has added a new ISO image that's based on the KDE Plasma desktop. This new flavor is available both in Security and Home editions. Prior to this change, Parrot shipped with the default MATE desktop.

On the design front, new boot-splash animation and desktop background have been added. While the icons and theme remain the same, the desktop-base and wallpapers have been updated to match the new appearance.

Another major change comes in the form of APT's enforcement of HTTPS. Now Parrot 4.6 is configured to redirect to HTTPS mirrors when possible. If HTTPS

mirrors aren't available, the signatures are still verified.

Moving on to better hardware support, Parrot 4.6 ships with Linux 4.19 Kernel. While it might not be the latest 5.0, 4.19 is a long term support kernel. The Nvidia drivers are also updated to the new 410 version to ensure better performance.

Source:
<https://fossbytes.com/parrot-4-6-linux-distro-hacking-download/>

OPENBSD INTRODUCES SYSUPGRADE

The Unix derivative OpenBSD originated in 1995 from NetBSD with a strong focus on security. The project is known for its uncompromising attitude towards correct code and software licenses. Even the software is subject to a BSD license. Another focus of the product is reliability, which always occupies an outstanding position in the project. Among other things, regular auditing should find bugs and eliminate them as much as possible. For this reason, the

project also introduces new technologies rather hesitantly and does not allow, for example, some self-evident practices directly. One such matter of course is an automated update of the complete system.

However, as the team now announced, the upcoming version of the operating system will introduce `sysupgrade`, a tool that allows upgrading OpenBSD to a higher version or snapshot. The tool first loads all data relevant for an upgrade into the `/home/_sysupgrade` directory, verifies the data with `signify` and automatically starts the update with `/bsd.upgrade`.

Source: <https://www.pro-linux.de/news/1/27020/openbsd-f%C3%83%C2%BChrt-sysupgrade-ein.html>

FEDORA 30 RELEASED WITH GNOME 3.32 AND LINUX KERNEL 5.0, HERE'S WHAT'S NEW

After more than six months in development, the Fedora

Linux 30 operating system is finally here to give fans access to some of the latest and greatest GNU/Linux technologies and Open Source software. Besides up-to-date components, Fedora 30 comes with many new features, optimizations, and several other improvements for a richer Fedora Linux experience.

There are some great additions in Fedora 30, such as the ability to install the Deepin and Pantheon desktop environments alongside existing and renowned flavours like GNOME, KDE Plasma, Xfce, LXQt, MATE, Cinnamon, and others. Of course, Fedora 30 ships with the latest GNOME 3.32 and KDE Plasma 5.15 desktop environments, and it's powered by Linux kernel 5.0, GCC 9, Bash 5.0, and PHP 7.3.

Under the hood, Fedora 30 ships with several improvements to the DNF package management system, which powers everything you install, update or remove on your Fedora computers. The DNF repository metadata is now compressed with the `zchunk` format in addition to the `gzip` and `xz` ones to speed up deltas, which makes updating your Fedora 30

operating system a breeze.

Support for ARM devices has been improved as well in the Fedora 30 release, which can now run on both Raspberry Pi 2 and Raspberry Pi 3 single-board computers.

Source: <https://news.softpedia.com/news/fedora-30-released-with-gnome-3-32-and-linux-kernel-5-0-here-s-what-s-new-525820.shtml>

LINUS TORVALDS RELEASES LINUX 5.1 WITH LOTS OF NEW FEATURES

Linux boss Linus Torvalds has released the stable version of Linux kernel 5.1 after seven release candidates. Kernel releases generally achieve the stable status after seven or eight release candidates, so it's great for Linux 5.1 to take shape a week early. Also, the "Shy Crocodile" codename of the latest release remains the same as the past ones.

Compared to other releases, Linux 5.1 arrived a bit late in the day due to some last-minute pull

requests. As per Phoronix, Linux 5.1 comprises of about 17.8 million lines of code and 3.3 million lines of comments; it comes loaded with many new features.

The new high-performance I/O interface is the biggest highlight of Linux 5.1. The new `io_uring` interface is expected to bring fast and scalable asynchronous I/O to Linux. Moreover, it also adds a user space library that lets apps set up an `io_uring` instance without needing to know the ins and outs of the `io_uring`.

This release also improves the `fanotify` interface that helps one monitor Linux file system for changes and adds "super block root watch" feature as a scalable way to keep track of changes.

Another significant change in Linux 5.1 is the support for the usage of persistent memory as RAM. It has already been made clear in the past that doing so could bring along some performance-related compromises, but it's good to see Linux giving a choice to let people use NVDIMMs as additional RAM.

With 5.1, Linux kernel also

continues to work on the year 2038 problem. Other major features include the new hardware support, R-Pi 3 Model A+ mainline kernel support, Intel 22260 WiFi support, etc.

Source:

<https://fossbytes.com/linux-5-1-kernel-features-download-linux/>

MOZILLA ISSUES NEW FIREFOX UPDATE TO FIX ADD-ONS PROBLEMS AND WARNS USERS NOT TO TRY DUBIOUS WORKAROUNDS

Mozilla has pushed out Firefox 66.0.4, properly addressing a problem that prevented add-ons from working in the web browser.

On Friday, an expired security certificate caused frustration and confusion for Firefox users as extensions were disabled and rendered unusable. Having delivered a patch through its Studies system (which did not work for everyone), Mozilla has now issued a browser update which it says will fix the problem for more people, although it warns that

"there are remaining issues that we are actively working to resolve".

In an update to an earlier blog post, Mozilla product manager Kevin Needham writes: "A Firefox release has been pushed -- version 66.0.4 on Desktop and Android, and version 60.6.2 for ESR. This release repairs the certificate chain to re-enable web extensions, themes, search engines, and language packs that had been disabled (Bug 1549061). There are remaining issues that we are actively working to resolve, but we wanted to get this fix out before Monday to lessen the impact of disabled add-ons before the start of the week".

In the release notes for Firefox 66.0.4, Mozilla says that with the build it has "repaired certificate chain to re-enable web extensions that had been disabled". It also notes that "a small number of add-ons may be listed as unsupported or may not appear in about:addons. Their data is not lost; users should be able to re-install the add-ons and recover the data".

Source:
<https://betanews.com/2019/05/06/firefox-update-add-ons-fix/>

FREESPIRE 4.8 OFFICIALLY RELEASED, BASED ON UBUNTU 18.04.2 LTS

Launched last year on August, the Freespire 4.x operating system series continues to be updated with new point releases, Freespire 4.8 being the latest in the series, packed with some of the latest software updates and security fixes. Freespire 4.8 is the latest version and it's recommended for new installations.

Highlights of the Freespire 4.8 release include the long-term supported KDE Plasma 5.12.7 desktop environment, which is accompanied by the KDE Frameworks 5.44.0 software suite built against the Qt 5.9.5 LTS libraries. Freespire 4.8 is powered by the Linux kernel 4.18 from the Ubuntu 18.04.2 LTS (Bionic Beaver) operating system.

Freespire 4.8 also comes with the latest versions of pre-installed

apps, among which we can mention the Chromium web browser, Calligra office suite, Geary email client, VLC media player, Amarok music player, Shotwell image viewer and organizer, Synaptic package manager, and DOSBox x86 emulator with DOS.

Ice SSB is included as well in Freespire 4.8 to let you install the web browser of choice in case you don't want to use Chromium, and many of the standard KDE apps are present as well, including Karbon vector drawing application, KolorPaint paint program, KPatience card sorting game, as well as DreamChess game of chess.

Source:
<https://news.softpedia.com/news/freespire-4-8-officially-released-based-on-ubuntu-18-04-2-lts-525902.shtml>

UBUNTU 14.04 (TRUSTY TAHR) REACHED END OF LIFE, UPGRADE TO UBUNTU 18.04 LTS Now

Released on April 17th, 2014, the Ubuntu 14.04 (Trusty Tahr) was an LTS (Long Term Support) version supported with security and software updates, as well as regular maintenance releases (the last one being Ubuntu 14.04.6, released on March 5th, 2019) for a total of five years, until April 25th, 2019, when Canonical announced the availability of the extended maintenance support.

Now that Ubuntu 14.04 (Trusty Tahr) will no longer receive security patches, nor software updates, users can choose to either upgrade to a supported release, such as the Ubuntu 16.04 LTS (Xenial Xerus) or Ubuntu 18.04 LTS (Bionic Beaver), or keep their Ubuntu 14.04 installations updated by purchasing the Extended Security Maintenance (ESM) package from Canonical.

We encourage all Ubuntu 14.04 users to upgrade their systems to the Ubuntu 18.04 LTS release, which will be supported until April 2023. However, the upgrade is done in stages via the Ubuntu 16.04 LTS release, which will be supported until April 2021.

Source:

<https://news.softpedia.com/news/ubuntu-14-04-trusty-tahr-reached-end-of-life-upgrade-to-ubuntu-18-04-lts-525899.shtml>

GNU LINUX-LIBRE 5.1 RELEASED

Already in 2002, Richard Stallman, founder and president of the FSF, criticized the fact that the Linux orientation makes the systems unfree. According to Stallman, the kernel uses a lot of code that is not GPL-compliant and challenges kernel redistribution. Not a small number of device drivers contain series of numbers that represent the firmware of the chips and are not available in the source code. In the opinion of the FSF president, this part of the Linux kernel violates the GPL and thus makes the core unfree.

Therefore, a solution has been promising for some time a modification of the Free Software Foundation Latin America (FSFLA) kernel, directed by Alexandre Oliva. The developer maintains an alternative version of the kernel,

which dispenses with all components not available in the source code and completely devotes itself to the idea of free availability. To do this, Oliva adjusts the sources of the kernel, eliminating dependencies and straightening calls. Meanwhile, it works quite smoothly, so that the developer can publish a clean version of Linux almost at the same time as the shares of the regular kernel.

Linux-libre 5.1 came just hours after the release of Linux 5.1 and is based on the official release of Linus Torvalds' Kernel 5.1. In addition to the regular changes, the current "libre" kernel cleans a bunch of drivers and eliminates loading of non-free components in mt7603 and goya. The "deblobbing" in wilc1000 has been improved and the handling optimized in iwlwifi, soc-acpi-intel sound, brcmfmac, mwifiex, btmrvl, btmtk and touchscreen_dmi. In addition, the cleanup of components that no longer exist was switched off.

The clean version of the kernel can be downloaded from the source code page of the FSFLA. In addition, the project also provides

xdelta packages.

Source: <https://www.pro-linux.de/news/1/27036/gnu-linux-libre-5-1-freigegeben.html>

THERE ARE NOW 2.5 BILLION ACTIVE ANDROID DEVICES

Ten years after its first launch, Android is still setting device records. Today at the I/O developer conference, Google announced that there are currently 2.5 billion active Android devices. It's a staggering number for Android, and a sign of just how successful Android's modular approach has been in reaching new users and hardware partners.

"We get to celebrate a milestone together," said Android senior director Stephanie Cuthbertson onstage at the event.

Since the number is based on Google's Play Store statistics, it doesn't include non-Play Store forks like Amazon's Fire OS or most of China's Android devices. Google made the announcement as part of the launch of Android Q beta 3.

The number of active devices is growing quickly. Google publicly reached 2 billion active devices in 2017, announced at that year's I/O conference.

Those numbers also underscore the scale of the fragmentation challenge, as Google looks to apply basic updates and security standards to all Android devices across different versions, regions, and manufacturers. According to the Google distribution dashboard's October report, just under half of Android devices are running Oreo or Nougat, the two most recent versions of Android.

Source:

<https://www.theverge.com/2019/5/7/18528297/google-io-2019-android-devices-play-store-total-number-statistic-keynote>

UBUNTU TOUCH OTA-9 RELEASED FOR UBUNTU PHONES WITH REFRESHED LOOK, IMPROVEMENTS

Ubuntu Touch OTA-9 comes two months after the OTA-8 update with a refreshed look

consisting of new and updated Suru symbols and folder icons to give users a better Ubuntu Phone experience, improvements for the Nexus 5 camera so users can now record videos again, better detection of the system-wide dark theme, as well as a new "Busy" indicator.

Also included in this release is support for the OpenStore V3 API in the update handler of System Settings, the ability to save images using the previously used compression settings, improvements to the characters counter for messages, support for searching the Web with Lilo, simplified transitions for the Stack View, and a new "Paste and Go" option in the browser.

The Ubuntu Touch OTA-9 software update is now rolling out to all supported Ubuntu Phone devices, including Fairphone 2, Nexus 5, Nexus 4, OnePlus One, BQ Aquaris M10 FHD, BQ Aquaris M10 HD, Meizu MX 4, Meizu PRO 5, BQ Aquaris E4.5, BQ Aquaris E5, and Nexus 7. Users will be able to install the OTA-9 update from the Software Updates panel in System Settings.

UBports said that the rollout of the Ubuntu Touch OTA-9 update should complete on Sunday, May 12th, 2019. By that time all users should have received the update on their devices, so make sure you install it as soon as it's available if you want to have a more stable and reliable Ubuntu Phone experience.

Source:

<https://news.softpedia.com/news/ubuntu-touch-ota-9-released-for-ubuntu-phones-with-refreshed-look-improvements-525949.shtml>

URBAN COMPUTING FOUNDATION FOUNDED

As the Linux Foundation (LF) announced now, the Urban Computing Foundation was founded under its umbrella. It aims to create a community of developers who can help steer urban development through open source. Contributors to the Urban Computing Foundation include developers from Uber, Facebook, Google, HERE Technologies, IBM, Interline Technologies, Senseable City Labs, StreetCred Labs, and the University of California San Diego

(UCSD).

The website of the Urban Computing Foundation outlines the objectives of the new foundation: "As cities and transport networks evolve into ever more complex systems, urban computing is becoming an important area to bridge the gap between development, visualization and traditional transportation system analysis , However, these advances depend on the compatibility of many technologies in different public and private companies. The Urban Computing Foundation will provide a neutral forum for this critical work, including the adaptation of geospatial and temporal machine learning techniques and urban environments, as well as simulation methods for modeling and predicting urban phenomena. "

This will enable developers, data scientists, visualization specialists and engineers to improve the urban environment, people's quality of life and city operating systems, and to build a networked urban infrastructure. This will be done through an open governance model that promotes the participation and technical

contribution of all stakeholders, and will provide a framework for long-term care by companies and individuals investing in the success of Open Urban Computing.

As an aid to unobtrusive and ubiquitous sensor technologies, advanced data management and analysis models, new visualization methods and scale simulation of urban systems are used. The first project hosted at the LF is Kepler, a geospatial analysis tool developed by Uber to create large data sets. Kepler was released in 2018 to facilitate the creation of meaningful visualizations of non-coding location data.

Source: <https://www.pro-linux.de/news/1/27039/urban-computing-foundation-gegr%C3%83%C2%BCndet.html>

BUG IN ALPINE LINUX DOCKER IMAGE LEAVES ROOT ACCOUNT UNLOCKED

A security vulnerability in the Official Docker images based on the Alpine Linux distribution allowed for more than three years logging into the root account using

a blank password.

Tracked as CVE-2019-5021, the vulnerability has a critical severity score of 9.8. It was initially reported in build 3.2 of Alpine Linux Docker image and patched in November 2015, with regression tests added to prevent it from occurring in the future.

"This lead to logic that may have caught this regression being simplified, causing these tests to be incorrectly 'satisfied' if the root password was once again removed," Cisco Talos says in a report today.

A subsequent commit removed the "disable root by default" flag from the 'edge' build properties file, allowing the bug to regress in the next builds of the image, starting v3.3 to 3.9.

The result was a blank sp_pwdp field in /etc/shadow - the configuration file user account management where passwords are saved in encrypted form, allowing logging as root without typing in any password.

eter Adkins of Cisco Umbrella found the problem again earlier

this year and put it into the limelight. The official Alpine Linux Docker image has over 10 million downloads.

The vulnerability was fixed and closed on March 8, 2019, but it could have been solved sooner as it was rediscovered and reported on August 5. It slipped through because it was not flagged as a security problem.

All supported builds have been updated and are "now only generated from upstream minirootfs tarballs," shows a commit from Natanael Copa, the creator of Alpine Linux. Release and update scripts have been refactored and moved to the official Alpine Linux image repository on the Docker portal.

To mitigate the issue on systems that still run vulnerable builds of the Alpine Linux container, Cisco Talos recommends disabling the root account.

"The likelihood of exploitation of this vulnerability is environment-dependent, as successful exploitation requires that an exposed service or application utilise Linux PAM

[Pluggable Authentication Modules], or some other mechanism which uses the system shadow file as an authentication database," Cisco Talos says.

Source:

<https://www.bleepingcomputer.com/news/security/bug-in-alpine-linux-docker-image-leaves-root-account-unlocked/>

"RED HAT WILL REMAIN INDEPENDENT; I'M NOT BUYING THEM TO DESTROY THEM," SAYS IBM CEO

Ever since IBM announced that it's going to acquire the open source giant Red Hat, Linux and open source enthusiasts have been voicing their concerns regarding the future of Red Hat? Is it going to become just another IBM subsidiary or is Red Hat going to adopt IBM's corporate culture?

At the Red Hat Summit in Boston, executives of both companies shared their thoughts and reiterated their commitments to move forward and drive more innovation. When asked if the

Linux giant would remain independent, IBM CEO, chairman, and president Ginni Rometty said: "I don't have a death wish for \$34 billion."

"I'm not buying them to destroy them. It's a win-win for our clients. It's a way to drive more innovation," she added.

These statements further affirm that both companies are working together to ensure that the \$34 billion deal, which is scheduled to be completed in the second half of 2019, stays on track. In a related development, the U.S. Department of Justice has also approved the deal.

During the summit, Red Hat CEO and president Jim Whitehurst said (Via: SiliconANGLE) they are all about making open source the default choice in the enterprise (and world). He added that both the companies have recognized that working together is the best way possible to serve their customers.

Talking about the possible culture clash, Whitehurst said that it's possible to "celebrate the strengths of each other's cultures."

On the same matter, Rometty said that both the companies agree that their mission is to scale open source — so, preserving each other's values is important.

Source: <https://www.pro-linux.de/news/1/27039/urban-computing-foundation-gegr%C3%83%C2%BCndet.html>

LINUX KERNEL PRIOR TO 5.0.8 VULNERABLE TO REMOTE CODE EXECUTION

Linux machines running distributions powered by kernels prior to 5.0.8 are affected by a race condition vulnerability leading to a use after free, related to net namespace cleanup, exposing vulnerable systems to remote attacks.

Potential attackers could exploit the security flaw found in Linux kernel's `rds_tcp_kill_sock` TCP/IP implementation in `net/rds/tcp.c` to trigger denial-of-service (DoS) states and to execute code remotely on vulnerable Linux machines.

The attacks can be launched

with the help of specially crafted TCP packets sent to vulnerable Linux boxes which can trigger use-after-free errors and enable the attackers to execute arbitrary code on the target system.

The remotely exploitable vulnerability has been assigned a 8.1 high severity base score by NIST's NVD, it is being tracked as CVE-2019-11815 (Red Hat, Ubuntu, SUSE, and Debian) and it could be abused by unauthenticated attackers without interaction from the user.

Luckily, because the attack complexity is high, the vulnerability received an exploitability score of 2.2 while the impact score is limited to 5.9.

According to CVSS 3.0 the impact metrics, the CVE-2019-11815 flaw comes with high confidentiality, integrity, and availability impact which makes it possible for would-be attackers to gain access to all resources, modify any files, and deny access to resources after successfully exploiting the vulnerability.

Source:

<https://www.bleepingcomputer.com/news/security/linux-kernel-prior-to-508-vulnerable-to-remote-code-execution/>

GOOGLE SAYS ALL NEW CHROMEBOOKS WILL BE LINUX-READY

One of the biggest announcements from this year's Google I/O related to the Pixel 3a and Pixel 3a XL, but this was far from being all there was to get excited about. Microsoft may be increasingly embracing Linux in Windows 10, and Google is doing the same with Linux on Chromebooks.

Support for Linux apps on Chromebook is nothing new, but Google has now announced that all Chromebooks that launch from this point forward will be fully functional Linux laptops, regardless of whether they are ARM or Intel devices.

At Google I/O, the company said: "all devices launched this year will be Linux-ready right out of the box". This is great news for anyone

looking for a cheap Linux laptop.

There's no suggestion that Chromebooks will come pre-installed with anything other than (the admittedly Linux-based) ChromeOS, but being Linux-ready should mean that it is easier than ever to run a different operating system.

Source:

<https://betanews.com/2019/05/12/linux-ready-chromebooks/>

WHATSAPP VULNERABILITY ALLOWS HACKERS TO INFECT IPHONES, ANDROID PHONES

A report from The Financial Times reveals that the security flaw made it possible for malicious actors to inject Israeli spyware on mobile device using nothing more than a typical WhatsApp call.

By the looks of things, users didn't even have to answer the calls, as the malicious code was executed in the background. The missed call was most often

removed from the call history, which means that users who didn't notice it in the first place had no idea they were hacked.

The malicious code is linked to Israeli company NSO Group, according to the same report, albeit the company claims it "would, or could not, use its technology in its own right to target any person or organization." NSO Group builds cyber tools that are then sold to governments across the world for a wide variety of purposes, many of them related to terrorism.

For example, such code can be used to extract data from locked devices that are involved in criminal investigations.

WhatsApp, on the other hand, says it has already resolved the bug, and users are now recommended to install the latest version of the messaging app as soon as possible. WhatsApp has more than 1.5 billion users on Android and iOS, but parent company Facebook can't share any details as to how many of them might have been targeted in such attacks.

Source:

<https://news.softpedia.com/news/whatsapp-vulnerability-allows-hackers-to-infect-iphones-android-iphones-526019.shtml>

MICROSOFT SELLS UNAUTHORIZED ARCH LINUX

For a long time, Microsoft's redmond manufacturer has been using a dedicated Linux subsystem to run Linux applications on Windows. After initially speculating on the meaning of the development, later all indications indicated that Microsoft wants to natively support Ubuntu and integrate it as an integral part in Windows 10. Over time, the manufacturer expanded the offering and added more distributions to the Windows Store.

Ubuntu, Debian and OpenSuse can be installed directly from the company's store. Now, anyone interested in Arch Linux will also find another popular distribution included with the offer. WSL Arch Linux allows users to use the terminal and use the standard tools on the Bash command line.

Additional applications can also be installed with the regular Arch Linux tools. Graphical operation is not installed by default.

In contrast to the distributions already offered, WSL Arch Linux is neither an official product of Arch Linux nor an authorized modification of the operating system. While the other distributions offered in the Microsoft Store are managed by the official distribution teams, WSL Arch Linux is rather a private hobby project of a Chinese developer. To make matters worse, the system in the standard version includes servers that are not officially maintained by the Arch Linux team.

The use of the distribution without the consent of Arch Linux could therefore be quite problematic. As Robin Broda of Arch Linux writes on Twitter, the company violated in his opinion not only against the license terms of the distribution, but also against their own terms of use.

Source: <https://www.pro-linux.de/news/1/27059/microsoft-vertreibt-unautorisiertes-arch-linux.html>

LINUX-BASED OS IS SAVING \$430 MILLION IN INDIAN STATE OF KERALA

Using Linux-based operating systems have tons of advantages like better security and freedom to customize the open source software. Another major advantage that attracts different organizations and schools is the cost saving that comes along the way. In the past, we have reported various European cities going the Linux way to cut down the costs and save the public's money.

As per a recent report published in Financial Express, schools in the Indian state of Kerala are saving about Rs 3,000 crore by moving to a Linux-based operating system. This news

follows a previous report from 2017 that mentioned that Kerala is saving Rs. 300 crore each year. If the report is to be believed, it seems that the South Indian state is making great progress in making open source software available in schools.

As per K Anvar Sadath, the executive director of KITE (Kerala Infrastructure and Technology for Education), more than 200,000 computers will be running Linux-based operating system next year and more than 150,000 primary teachers will be trained for the same.

Sadath adds that if the state had chosen proprietary software like Windows, Microsoft Office, etc., each PC would have added about Rs 1.5 lakh in licensing. If you calculate the cost-saving for the

200,000 computers, the total saving turns out to be \$430 Million (Rs 3,000 crores).

Source:

<https://fossbytes.com/linux-based-os-is-saving-430-million-in-indian-state-of-kerala/>

ELIVE ELEVATES LINUX WITH ENLIGHTENMENT

The Elive distro's integration of the Debian Linux base and the Enlightenment desktop is a powerful combination. Together, they offer a unique computing platform that is powerful and flexible.

Elive is not like most Linux distributions today. It does not have a team of workers supporting

multiple desktop offerings cranking out frequent upgrades each year. It also does not have a thriving community.

In fact, Elive is one of a few Linux distros that aggressively asks for donations in order to download the installation ISO file. You can get the download without donating, but the process requires you to verify your email address and wait for the download link.

Elive first appeared in January 2005. The second stable version came in 2010. Eight years later the third stable version arrived, version 3.04.

Developer Samuel F. Baggen announced the release of version 3.05 on April 29. It is based on Debian 7 "Wheezy," with a customized Enlightenment 17 desktop.

The customization is key to what gives Elive the edge over the few other distros running the latest version of Enlightenment, which is E22. However, this latest Elive version is likely the last update in the Elive 3 series.

The developer is focused on the



DistroWatch.com

Put the fun back into computing. Use Linux, BSD.



next release, which will be based on Debian 10 "Buster." That release could be well in the future, though, because donations from users have not been sufficient to support the developer's continued efforts so far.

The silver lining is that this latest Elive release is updated with some of the internal improvements Baggen developed for the next version of Elive. So this latest release provides an early look at what may be coming next.

Elive is a fast and very configurable Linux OS that has an unusually pleasing appearance. It is designed to run fast on older computers with more modest hardware specs. It is blazingly fast on newer computers with more memory and better graphics circuits.

Source:

<https://www.linuxinsider.com/story/Elive-Elevates-Linux-With-Enlightenment-86009.html>

CLEAR LINUX FROM INTEL BRINGS BEST PERFORMANCE ON INTEL CPUs

The latest release ships with a new installer for the easy setup process and a Clear Linux Store for all bundles, apps, and container images for the distribution. The tools on the store are categorized in Developer Tools, Education, Games, Security, Productivity, Programming, etc.

The Developer Edition is focused towards Linux developers and Intel will ship one image daily. To help developers who are looking to push optimized code regularly, Clear Linux ships with GCC9 and Intel plans to upgrade to GCC10 as soon as it's released. It also comes loaded with basic programming bundles to provide an out-of-the-box experience to developers.

On the security front, Clear's rolling release model lets developers get the latest update as soon as possible. This removes different roadblocks and ensures a secure platform for developing software. Clear Linux is also helping out the Linux kernel in

general by helping the kernel become more efficient on Intel hardware.

When it comes to hardware requirements, Clear Linux can run on hardware as low as a single core CPU, 600MB storage, and 128MB storage. However, different applications obviously mean different configurations. Generally, Intel recommends a 64-bit processor that supports UEFI and SSE v4.1 streaming SIMD instructions.

Source:

<https://fossbytes.com/clear-linux-intel-best-developer-performance/>

"JOHN THE RIPPER" 1.9.0 RELEASED

John the Ripper (JtR) is a free tool by Alexander Peslyak to test passwords and authentication facilities. The software makes it possible to use brute force or dictionary attacks to guess encrypted passwords and thus identify insecure systems. To do this, JtR compares the hashes generated by the application using regular expressions with encrypted

strings. If both are equal, then either the password or a hash collision was found. With the help of Distributed John (djohn), the calculation actions can also be divided into several computers.

With "John the Ripper 1.9.0-jumbo-1" is four and a half years after the last version, a functionally upgraded version of the software ready. The software essentially builds on the "John the Ripper 1.9.0 core" released in April, which included improvements in the speed of processing long passwords and various processor optimizations. In addition, the benchmark functionality has been improved.

Jumbo-1 now adds extra functionality to the core, and in the latest version has, among other things, a broader support for external hardware. CUDA has been abandoned in favor of OpenCL and the number of supported CPU optimizations has been further increased. Also new is the improved support for archive formats and improvements in the use of GPU resources. In addition, the methods for password recognition have been improved and the "single crack" mode has

been improved.

Source: <https://www.pro-linux.de/news/1/27068/john-the-ripper-190-freigegeben.html>

GCC: SWITCHING TO GIT IS APPROACHING

It is undeniable that GCC is one of the most complex free projects. Launched more than 30 years ago by GNU founder Richard Stallman, the »GNU Compiler Collection« has become standard on many operating systems. GCC supports more architectures, programming languages and operating systems than any other compiler. The compiler collection is portable, mature, standards compliant, and generates optimized code that is only surpassed by other compilers in individual cases. GCC provides the tools that make up all Linux distributions and countless embedded systems.

However, a disadvantage of the system is the long development time and the tools used for the creation. Currently, the maintenance of the sources takes place in an SVN repository, which

has, however, been reaching its limits for quite some time. Efforts to convert the entire development to Git are therefore neither new nor revolutionary. However, they almost always failed due to the complexity of the project. Eric S. Raymond, for example, attempted a conversion using "Reposurgeon" in the summer of last year, but did not get any useful results either.

A new impetus could be the work now through a patch series by Maxim Kuvyrkov. The three scripts convert the entire repository of the GCC project and the associated branches as well as the numerous Git version notes, taking into account the various pitfalls and inconsistencies. These include, for example, deleted or renamed branches as well as orphaned entries.

Source: <https://www.pro-linux.de/news/1/27071/gcc-umstieg-auf-git-kommt-n%C3%83%C2%A4her.html>

SECURITY RESEARCHERS DISCOVER LINUX VERSION OF WINNTI MALWARE

For the first time, security researchers have uncovered and analyzed a Linux variant of Winnti, one of the favorite hacking tools used by Beijing hackers over the past decade.

Discovered by security researchers from Chronicle, Alphabet's cyber-security division, the Linux version of the Winnti malware works as a backdoor on infected hosts, granting attackers access to compromised systems.

Chronicle says it discovered this Linux variant after news broke last month that Bayer, one of the world's largest pharmaceutical companies, had been hit by Chinese hackers, and the Winnti malware was discovered on its systems.

During subsequent scans for Winnti malware on its VirusTotal platform, Chronicle said it spotted what appeared to be a Linux variant of Winnti, dating back to 2015 when it was used in the hack of a Vietnamese gaming company.

Chronicle says the malware they discovered was made up of two parts. A rootkit component to hide

the malware on infected hosts, and the actual backdoor trojan.

Further analysis revealed code similarities between the Linux version and the Winnti 2.0 Windows version, as described in reports by Kaspersky Lab and Novetta.

Other connections with the Windows version also included the similar way in which the Linux variant handled outbound communications with its command-and-control (C&C) server -- which was a mixture of multiple protocols (ICMP, HTTP, and custom TCP and UDP protocols).

Last but not least, the Linux version also possessed another feature that was distinctive to the Windows version, which was the ability for Chinese hackers to initiate connections to infected hosts without going through the C&C servers.

Source: <https://www.zdnet.com/article/security-researchers-discover-linux-version-of-winnti-malware/>

VMWARE TAKES OVER BITNAMI

The previously privately owned company Bitnami is the driving force behind the eponymous application technology »Bitnami«. The project provides ready-to-install, easy-to-install application packages for various open source projects, which include the necessary infrastructure to run them when shipped. The packages called "stack" are encapsulated, leaving the user's enclosing operating system untouched.

Vmware plans to expand its portfolio through the acquisition of Bitnami and enhance the offered cloud services through the technology. "Bitnami enables our customers to easily deploy application packages in any cloud - public or hybrid - and in the optimal format - as a virtual machine, container, or Kubernetes helmet control chart," the statement said. "In addition, Bitnami will be able to strengthen our existing efforts to provide VMware customers with a marketplace that offers a variety of applications and development environments in addition to

infrastructure software." The aim of the manufacturer is therefore to offer users simplified ways to use open source software applications and frameworks.

For users of Bitnami nothing should change. As the founders write in the in-house blog, the catalog of applications to be developed and expanded. "There are a lot of great software, many of them open source, which is unattainable for many developers and system administrators because it's too complex to set up and maintain. Our goal is to make software available to as many users and developers as possible," the company said. The purchase price is unknown.

Source: <https://www.pro-linux.de/news/1/27075/vmware-%C3%83%C2%BCbernimmt-bitnami.html>

SOUTH KOREA WANTS TO SWITCH FROM WINDOWS 7 TO LINUX

The government of the country South Korea wants to operate

its administration in the future on the free Linux, instead of as before on Windows. This is reported by the English-language daily The Korea Herald in its online edition. According to the report, the Ministry of Homeland Security announced this decision last week.

It also states that the ministry would like to first run a test run with Linux systems on its computers as well as a security check. In addition, an attempt should be made to establish compatibility with previously used Windows programs and web applications on the new system or to test out this part of the migration. Should these tests succeed, the Linux systems should be rolled out throughout the government administration.

As the main reason for the change from Windows 7 to Linux, the report calls above all the expected very high costs for the safe continued operation of the outdated system from Microsoft. Microsoft's regular Windows 7 support officially ends in mid-January 2020, but the manufacturer also offers extended support over a three-year period, with license fees rising annually.

For the purchase of new computers and migration to Linux systems, the government estimates about 780 billion won (about 585 million euros). In addition to saving costs through the use of open source components, the government also wants to make the step more independent of a single vendor for the operating system used.

Source:

<https://www.golem.de/news/verwaltungs-suedkorea-will-von-windows-7-auf-linux-wechseln-1905-141406.html>

ANTERGOS LINUX PROJECT IS DEAD

A group of developers started Antergos as a hobby project in 2012 and pushed the first release in July 2012 under the name Cinnarch. Later, it got its name Antergos and gained popularity as an Arch Linux-based distribution created for all the users.

In a development that will surprise the Antergos enthusiasts, the developers of the operating

system have announced their plans to end the project. In other words, the Antergos 19.4 ISO refresh was the last release and the developers don't wish to continue the project any further.

In their announcement post, the developers have mentioned that they don't have enough free time to properly maintain Antergos and the regular updates weren't being pushed. As they rightly mentioned, it would have been a "huge disservice to the community" if they had chosen to continue the Antergos project in a similar manner.

For those who are interested in numbers, the Antergos image has been downloaded about 1 million times since 2014, which is a pretty good number for a Linux distro that was being developed by volunteer developers.

Source: <https://fossbytes.com/antergos-linux-dead-alternatives/>

TAILS 3.14 RELEASED

Tails stands for "The Amnesic Incognito Live System" and is used to anonymize the Tor network, through whose node computer the network traffic is routed. It is designed as a live system for use on USB sticks or DVDs and specializes in anonymity and maintaining the privacy of its users. The basis is Debian »Stretch« 9.9.

After the retracted version Tails 3.13.2, which fixed a critical hole in the Tor browser, Tails 3.14 has now been released in the schedule. Tails 3.13.2 had become necessary in response to Mozilla's blunder, where an expired certificate disabled all extensions, exposing the Firefox-based browsers to potential attacks.

Tails 3.14 also closes some gaps, among others in the kernel, the Tor Browser, in Bind, Cups and Samba. In addition, a new Intel microcode has been integrated, which is directed against the ZombieLoad, RIDL and Fallout baptized newly discovered gaps in Intel processors. In this context, Simultaneous Multithreading (SMT) was disabled

for the affected processors. The kernel was raised to 4.19.37 and many firmware packages were updated. Tor Browser received a major update to version 8.5.

Removing less-used applications reduced the size of the image by about 40 MB. In addition to around 10 tools for the command line, the graphic apps Gobby, Pitivi and Traverso were also removed. All removed packages can be reinstalled using the Additional Software feature. An issue where Tails will start after installation but not then will continue to be investigated. Affected users are requested to contact the developers.

Because of the closed gaps, users are encouraged to update to Tails 3.14 in a timely manner. Automatic updates are available from 3.12, 3.12.1, 3.13, 3.13.1, and 3.13.2. The release of Tails 3.15 is scheduled for July 9th. For the foreseeable future, the project's roadmap includes support for Secure Boot, the move to Wayland, the improvement of the technical infrastructure, and Tails 4 following the release of Debian 10 Buster.

Source: <https://www.pro-linux.de/news/1/27088/tails-314-ver%C3%B6ffentlicht.html>

PEPPERMINT 10 OPERATING SYSTEM OFFICIALLY RELEASED, BASED ON UBUNTU 18.04 LTS

As Peppermint 10 is based on the Ubuntu 18.04 LTS repositories as of May 14th, 2019, it means it's in fact based on Ubuntu 18.04.2 LTS, which ships with updated kernel and graphics stacks from the Ubuntu 18.10 (Cosmic Cuttlefish) operating system. As such, Peppermint 10 is powered by Linux kernel 4.18.0-18.

Highlights of the Peppermint 10 release include support for automatically install the Nvidia proprietary graphics drivers, including support for Nvidia Optimus setups, Ice 6.0.2 with support for isolated profiles for Chromium, Google Chrome, and Vivaldi web browsers, and a new utility for setting font DPIs.

Among the updated components included in the

Peppermint 10 operating system, we can mention Nemo 4.0.6 file manager, mintinstall 7.9.7, mintstick 1.39, Neofetch 6.0.1, Xed 2.0.2, XPlayer 2.0.2, and XViewer 2.0.2. Additionally, XReader replaces Evince as default document viewer.

The light-locker and light-locker-settings packages now replace i3lock, the network-manager-pptp-gnome and network-manager-openvpn-gnome plugins for NetworkManager are now installed by default, and Peppermint 10 also ships with a new xfce-panel-switch profile for resetting the panel.

Source:

<https://news.softpedia.com/news/peppermint-10-operating-system-officially-released-based-on-ubuntu-18-04-lts-526146.shtml>

KALI LINUX 2019.2 RELEASED WITH NETHUNTER 2019.2 AND NEW KERNEL

Offensive Security, the makers of Kali Linux, have shipped

their second release in 2019. The new Kali Linux 2019.2 distribution is now available for ethical hackers and security researchers. This release brings along many bug fixes and updated packages that are surely worth upgrading.

Before you move ahead to explore the new changes in Kali Linux 2019.2, let me tell you about our new list of best Kali tools for hacking and pen-testing. These tools are highly recommended if you are willing to kickstart a journey in the field of ethical hacking.

Coming back to the latest Kali 2019.2. Offensive Security adopted a rolling release model a few years back and it continuously keeps updating the existing Kali installations. But what if a new user needs to perform a clean installation? To address this issue, the developers keep releasing fresh Kali builds from time to time and ensure that new downloads contain bug fixes, new Linux kernel, and other updates.

The biggest feature that ships with Kali Linux 2019.2 is the freshly baked Nethunter 2019.2 that now supports more than 50 mobile

devices, running Android versions from 4.4 through 8.0. The announcement post specifically mentions 13 new images for new Android versions for popular Nethunter devices like Nexus 6, Nexus 6, OnePlus 2, and Galaxy Tab S4

Source:

<https://fossbytes.com/kali-linux-2019-2-released-nethunter-download/>



One of the programs I've written that I use pretty much every day is something I call "media-tracker". It's a Ruby on Rails app that allows me to add/track movie (cinema and DVD) and game releases. Lately, I've also started tracking current episodes of shows from streaming services (Netflix, Crunchroll, Amazon Prime Video) due to the frequency with which I start a show, watch a chunk of it, and then notice it's been removed from the streaming service. This way I can pick up where I've left off if the show pops up somewhere else or I buy a DVD of it. Long story short - I start this app in Tmux manually every time I log in. I've finally gotten tired of this, and instead created a small systemd service to run the script on boot.

NOTES

This article will focus on getting a Rails app up and running through systemd. If you're using something different, the broad strokes will be the same, but you may need to

adjust more environment variables.

All commands have been run on an ArchLinux system. If your distribution has a different format (ie, `systemctl <service> <command>`), then stick to your distribution's format. If you're unsure, then use the commands I list here and see what errors occur.

GETTING SET UP

First, you want to make note of where your program's files are, and any paths you need to know (such as the `$GEM_HOME` and the relevant section of your `$PATH`). You can find these by simply running the following in your terminal of choice:

```
echo $PATH
echo $GEM_HOME
```

This is important, as you'll need to tell Systemd what to set these variables to, otherwise commands will fail or simply not be found.

DECIDE - USER OR SYSTEM

SERVICE?

Most modern systems have systemd set up to be run as both system, and a user-specific version. They are separate (and have separate folders they watch). The approach is similar for either one, but here are the main differences:

1. System-wide systemd services, when enabled, run on boot. User-based services will run on login.
2. User-based services also can't be configured to run as another user (for example if you want to run Apache as an 'html' user or similar).

Then you'll need to use the system-wide version.

Think about it and make your decision. Then skip ahead to the relevant section.

SYSTEM-WIDE SERVICE

The media-tracker.service file I created looks like the code shown below.

As you can see, there are a few important aspects to this file. First, I set Type to "simple", which means the ExecStart command is

```
[Unit]
Description=Media-Tracker

[Service]
Type=simple
User=lswest
Environment="PATH=$PATH:/home/lswest/.gem/ruby/2.6.0/bin"
Environment="GEM_HOME=/home/lswest/.gem/"
WorkingDirectory=/home/lswest/Repositories/rails-development/media-tracker
ExecStart=/bin/bash -lc 'bundle exec rails server'
TimeoutSec=30
RestartSec=15s
Restart=always

[Install]
WantedBy=multi-user.target
```

expected to run just while the service is alive. There are other types (for example if a command runs and then quits, like a bash script). For most cases, I imagine “simple” will be sufficient, though. I then set the user to my username, to ensure the service has access to my app files and the installed Gems.

Speaking of Gems, I needed to configure the environment variables to point to my .gem directory paths, otherwise the service failed to find the correct commands. Most articles I saw on this topic claimed that `bash -lc` would load the user’s shell profile (and therefore the variables), but that didn’t seem to be the case in ArchLinux. If you want to test to see if these lines are required on your machine, just delete them and check the output of your service through `journalctl`.

Lastly, I set the `WorkingDirectory` (path to my app files), and then `ExecStart` to run `bundle exec rails server`.

The other options are relatively self-explanatory, or shouldn’t need to be changed.

Copy/place the file in `/etc/systemd/system/` and call it `<something>.service` (the `<something>` part can be anything you want). Once the file is there, you can start/enable it with the following commands:

```
sudo systemctl start media-tracker.service
```

```
sudo systemctl enable media-tracker.service
```

To stop/disable the service, just replace `start` or `enable` with the word `stop` or `disable`. Similarly, you can run `status` to get the exit code and current status of the service.

If you make changes to the service file, you may get a warning that the files need to be reloaded. To do that, run:

```
sudo systemctl daemon-reload
```

To debug issues, you can use the following command:

```
journalctl -u media-tracker
```

Replace `media-tracker` with the file name you chose.

USER-SPECIFIC SERVICE FILE

```
[Unit]
Description=Media-Tracker

[Service]
Type=simple
Environment="PATH=$PATH:/home/lswest/.gem/ruby/2.6.0/bin"
Environment="GEM_HOME=/home/lswest/.gem/"
WorkingDirectory=/home/lswest/Repositories/rails-development/media-tracker
ExecStart=/bin/bash -lc 'bundle exec rails server'
TimeoutSec=30
RestartSec=15s
Restart=always

[Install]
WantedBy=multi-user.target
```

The file I created for a user-specific service looks like the code shown above.

The main difference between this and the system-wide service file is the missing “User” value.

As I said in the section above (for anyone jumping directly to this section):

1. Set the environment variables you’ll need for this service.
2. Set the `WorkingDirectory` option to the project folder.
3. The `/bin/bash -lc` should run the bash shell as a login shell, but, under ArchLinux, this doesn’t seem to fill the environment variables, hence why point 1 exists.
4. The other options in the file are self-explanatory, or, in the case of

`WantedBy`, shouldn’t need to be adjusted.

RUNNING THE SERVICE

Running a user-specific service is as easy as:

```
systemctl --user start media-tracker
```

Note the lack of `sudo`, and the “`--user`” argument. The other commands are all of the same format - `stop`, `enable`, `disable`, and `status`. Just replace the word “`start`” with whatever option you need.

Debug

In case your service fails to run, you can run `journalctl --user -u`



media-tracker to get the output of your service.

CONCLUSION

I hope this article is useful for anyone who, like me, has a custom program they want to run on every login or boot. It seems like a lot of articles on topics like these focus on system-wide services, which is why I also included instructions for user-specific services. If you run into issues, or have improvements to offer me, feel free to send me an email at lswest34+fcm@gmail.com. Similarly, if you have any article requests, direct them to the same email address.

FURTHER READING

<https://wiki.archlinux.org/index.php/systemd> - The ArchWiki article on Systemd

<https://wiki.archlinux.org/index.php/Systemd/User> - The ArchWiki article on user-based Systemd.



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.



There are times when you need a number of variables initialized at startup based on the last time the program ran. In the world of Windows, this is usually saved in a configuration file with a ".ini" extension. There are a number of ways to do this. One would be an XML file or even a database. However, both of these are often overkill and not quick and easy solutions. There is a library that can help us, however, in Python.

INSTALLATION

The library is configparser, and is easily installed with pip:

```
pip3 install configparser
```

Notice that I used pip3 rather than pip. If you are still using Python2, you should use pip, but since Python 2.x is ending life on Jan. 1, 2020, I've decided that I'm going to use only Python 3.x syntax. The current version is 3.7.4 (according to the PyPI site (<https://pypi.org/project/configparser/>) and was last updated on

March 22, 2019. This version is compatible with Python 2.6 onwards, including Python 3.7. There was an earlier version for 2.6-3.5 called ConfigParser which has been around for many years. If you want to use the older version, you can import it directly as a backport:

```
from backports import
configparser
```

Otherwise, you would import it as normal:

```
import configparser
```

USE

The INI file (which is NOT compatible with Windows "official" .ini file format) is a simple text file. You can use the extension .ini, or .cfg, or whatever you want. In this tutorial however, I'll use .ini as the extension. This file consists of key/value entries that are grouped by sections that are blocked with a [section] header. By default, the section names are case sensitive, but keys are not. Leading and

trailing whitespace is remove from both keys and values. The configuration file may also include comments that are, by default, only on an empty line. Inline comments can be used, but might cause problems, so I suggest that you don't use them. Comments, as in Python, start with a "#".

A very gross idea for the layout of the .ini file would be something like this...

```
[Animals]
animal1 = Frog
animal2 = Dog
animal3 = Hog

[Trees]
tree1 = The Larch
tree2 = Elm
tree3 = Ash
```

In this example, we have two sections, 'Animals' and 'Trees'. Each section contains three variables (animal1, animal2, etc) which are our keys and each has a value. You can also define a key without a default value:

```
tree4 =
```

However, when you use the

'tree4' variable, it is a blank string, not None.

The library doesn't try to guess what datatype a variable is. It always stores them internally as string.

THE CODE

Now we'll get into the code that we would use to read, write and create a default INI file. We'll name this program "iniFile.py". We'll start with the imports section:

```
import os
import configparser
```

In this simple demo program, we only need two imports, os and configparser. You'll see why we want the os library in a moment. Now we will define a global variable "iniFileName", which holds the filename of our ini file.

```
global iniFileName

iniFileName =
"MyConfigFile.ini"
```

Now, we'll create a function to



read our ini file:

```
def read_ini():
    global ini, iniFileName
    global tree1, tree2,
    tree3, tree4
    global animal1, animal2,
    animal3
    global theanswer
```

We define a number of global variables just to make things easy. Then, we check to see if the file exists (the `os.path.isfile()` method), and then read the file:

```
if
os.path.isfile(iniFileName):
ini.read(iniFileName)
```

This next bit of code (top right) shows how we can view the various sections and key/value sets.

Now we can assign the values to the proper variables:

```
animals = ini['Animals']
animal1 = animals['animal1']
animal2 = animals['animal2']
animal3 = animals['animal3']
trees = ini['Trees']
tree1 = trees['tree1']
tree2 = trees['tree2']
tree3 = trees['tree3']
tree4 = trees['tree4']
ans = ini['Answers']
```

We can also use the `.get` method of the section object to

assign the value to a variable:

```
theanswer =
ans.get('Life, The Universe
and Everything')
```

Now, we return 'True' to the call to say that the INI file existed. Otherwise, since the INI file doesn't exist at this point, we'll use some default values to create one and return 'False' to force the calling routine to re-read the INI file:

```
return(True)
else:
write_default_ini()
return(False)
```

Now, here is the function that can write to the INI file. In this case, we'll only write one value, but this will show how to do it. Basically, we use the `.set(section,key,value)` to update the ini object, then write it back out properly.

```
def write_ini():
    global ini, iniFileName
    ini.set('Trees', 'tree4',
tree4)
ini.write(open(iniFileName,
'w'))
```

Here (bottom right) is the function to write a default INI file,

```
print(ini.sections())
sections = ini.sections()
for section in sections:
    print("Section: {0}".format(section))
    for key in ini[section]:
        print("Key = {0} - Value =
{1}".format(key, ini[section][key]))
```

just in case it doesn't exist. We use `object.add_section()` to create a section and `object.set()` to add a key/value under that section.

Finally, we write our configuration file to disk, using the global `iniFileName` that we set up earlier at the top of the program:

```
# Writing our
configuration file
with open(iniFileName,
'w') as configfile:
config.write(configfile)
```

This function (next page, top right) simply is used to display all the variables that were pulled from the INI file:

The function `init()` (next page, bottom right) is where the real work is done; we instantiate the config parser object as 'ini', read the INI file and check to see if we get a 'True' (it is there) or 'False' (we needed to create it from defaults), so try to read again, show the variables, then update a value (tree4) then write the changed variable back to the file:

```
def write_default_ini():
    global iniFileName
    config = configparser.RawConfigParser()
    config.add_section('Animals')
    config.set('Animals', 'Animal1', 'Frog')
    config.set('Animals', 'Animal2', 'Dog')
    config.set('Animals', 'Animal3', 'Hog')
    config.add_section('Trees')
    config.set('Trees', 'Tree1', 'The Larch')
    config.set('Trees', 'Tree2', 'Elm')
    config.set('Trees', 'Tree3', 'Ash')
    config.set('Trees', 'Tree4', '')
    config.add_section('Answers')
    config.set('Answers', 'life, the universe and
everything', 42)
```

Finally, we have our "if __name__" entry point to our program which calls the init() function, and when it's done, notify the user that we are all done:

```
if __name__ == '__main__':  
# =====  
# All code is run from the  
# init() function  
# =====  
    init()  
# Notify user that we are  
# done  
    print('Program End')
```

That's all there is to it. Here's what the output looks like the first time the program is run:

```
['Animals', 'Trees',  
'Answers']  
Section: Animals  
Key = animal1 - Value = Frog  
Key = animal2 - Value = Dog  
Key = animal3 - Value = Hog  
Section: Trees  
Key = tree1 - Value = The  
Larch  
Key = tree2 - Value = Elm  
Key = tree3 - Value = Ash  
Key = tree4 - Value =  
Section: Answers  
Key = life, the universe and  
everything - Value = 42  
animal1: Frog, animal2: Dog,  
animal3: Hog  
tree1: The Larch, tree2: Elm,  
tree3: Ash, tree4:  
What's the answer to Life,  
The Universe and Everything?  
42  
theanswer type is <class
```

```
def show_ini_vars():  
    global tree1, tree2, tree3, tree4  
    global animal1, animal2, animal3  
    global theanswer  
  
    print("animal1: {0}, animal2: {1}, animal3: {2}".format(  
                                                animal1,  
                                                animal2,  
                                                animal3))  
  
    print("tree1: {0}, tree2: {1}, tree3: {2}, tree4: {3}".format(  
                                                tree1,  
                                                tree2,  
                                                tree3,  
                                                tree4))  
  
    print("What's the answer to Life, The Universe and Everything? {0}".format(  
        theanswer))  
    print("theanswer type is {0}".format(type(theanswer)))
```

```
'str'>  
Program End
```

Notice that the value for "tree4" is blank. If, however, you look at the INI file, it looks like this:

```
[Animals]  
animal1 = Frog  
animal2 = Dog  
animal3 = Hog  
  
[Trees]  
tree1 = The Larch  
tree2 = Elm  
tree3 = Ash  
tree4 = Birch  
  
[Answers]  
life, the universe and  
everything = 42
```

That's because the last line of the init() function updates the "tree4" variable to Birch in the write_ini() function.

```
def init():  
    # instantiate the ini object  
    global ini  
    ini = configparser.ConfigParser()  
    # call the read_ini function  
    isok = read_ini()  
    if isok:  
        # Call a function that prints out all our variables  
        show_ini_vars()  
    else:  
        isok = read_ini()  
        show_ini_vars()  
    # Change a variable and write it back to the ini file  
    global tree4  
    tree4 = 'Birch'  
    write_ini()
```

I didn't get too deep into the possibilities of this library, but if you want to learn more, you can read about it from the official Python docs on configparser at <https://docs.python.org/3/library/configparser.html>.

I've put the code example for this month on pastebin at <https://pastebin.com/X37remDa>.

Until next time, keep coding and have a great month!



HOW-TO

Written by Hansrudolf Brunner

GRAMPS

During the years, I collected photos and family-documents from my relatives. I am using Gramps to organize the content of several boxes. I am not interested in portraying myself as a descendant of some illustrious person from the past, I think pedigrees are used mainly for horses, dogs and noblemen. Nonetheless, the history of a family may reveal some interesting facts. I spent quite a lot of time

lately with Gramps, and found things about my relatives I didn't know before. Since I have no background in genealogy, the use of Gramps was quite difficult in the beginning. In fact I needed several attempts to get it going. However, if the data is entered in a certain order, you do not run into problems.

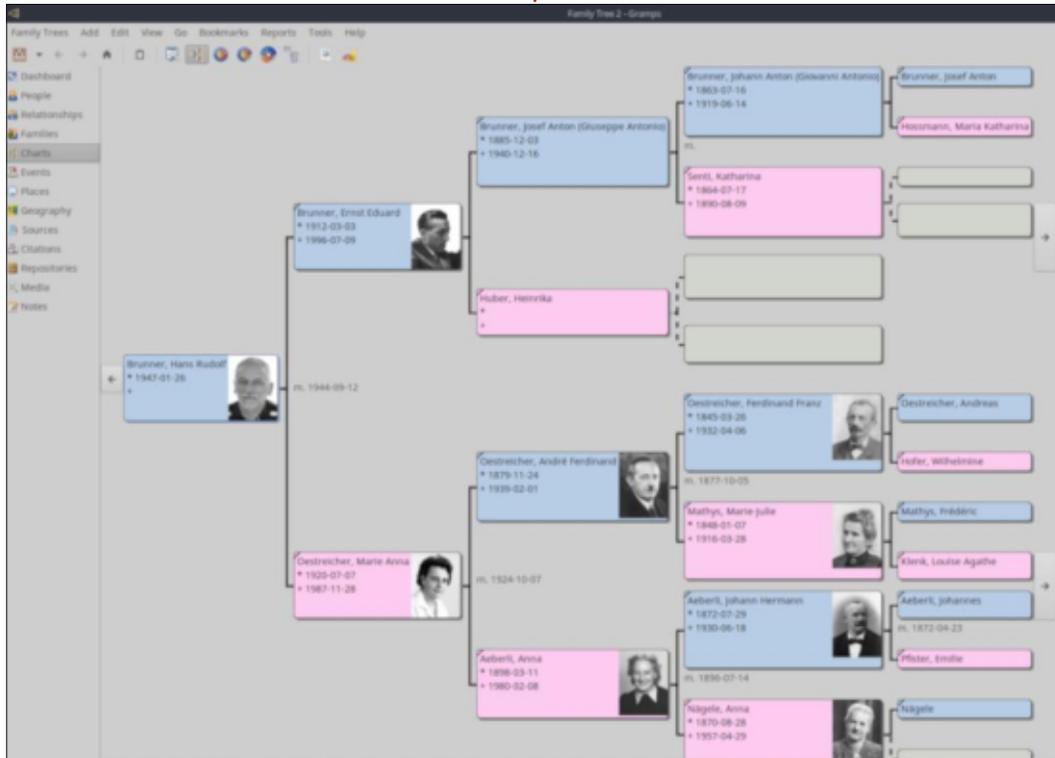
And here is my very short how-to:

- Enter the data for a person: last-name, first-name, dates of birth and death.
- Enter females always with their maiden-names.
- Repeat steps 1+2 for all known persons.
- Define families: select partners from the persons already entered.
 - Family-records are for any kind of relationship, married or not.
 - If someone was married more than once, define several family-

records.

- (Family-records are automatically stored under the name of the male partner)
- Within the families: select children from persons previously entered.

You may add photos, documents, locations, dates and more – but basically you are done with 1-5.



Complete Individual Report

Brunner, Hans Rudolf

Name:	Brunner, Hans Rudolf
Gender:	Male
Father:	Brunner, Ernst Eduard
Mother:	Gstreicher, Marie Anna



Life Events

Birth:	1947-01-26 in Kt. Zürich, Schweiz.
--------	------------------------------------

Families

Wölber, Marianne Barbara

Marriage:	1977-04-22 in Kt. Zürich, Schweiz.
Children:	Brunner, Thomas Philipp Brunner, Beatrice Nora

HOWTO - GRAMPS

I add photos in black&white 1000x750ppi (> Gallery). They will show up in the chart and reports.

There are a lot of reports; for most of them, you have to select first a person as starting point, else you get some other (unwanted) person or branch.

I use only three reports:

> Charts (selected from left panel on screen) =>

- You get a nice report on the screen, unfortunately without print function.

> Reports (select from top) > Text Reports > Complete individual Report =>

- This is a collection of all data for this person with photo (if added in personal record).

> Reports (select from top) > Graphs > Family Lines Graph =>

- You may need to add persons to include (>People of interest) in plot, or

- Restrict the number of persons (>People of Interest).

- Else the graph gets overloaded or ugly.

Maybe you will find other reports helpful, even a website can be created from your data. Give it a

try. At the moment, I stick to the three mentioned reports.

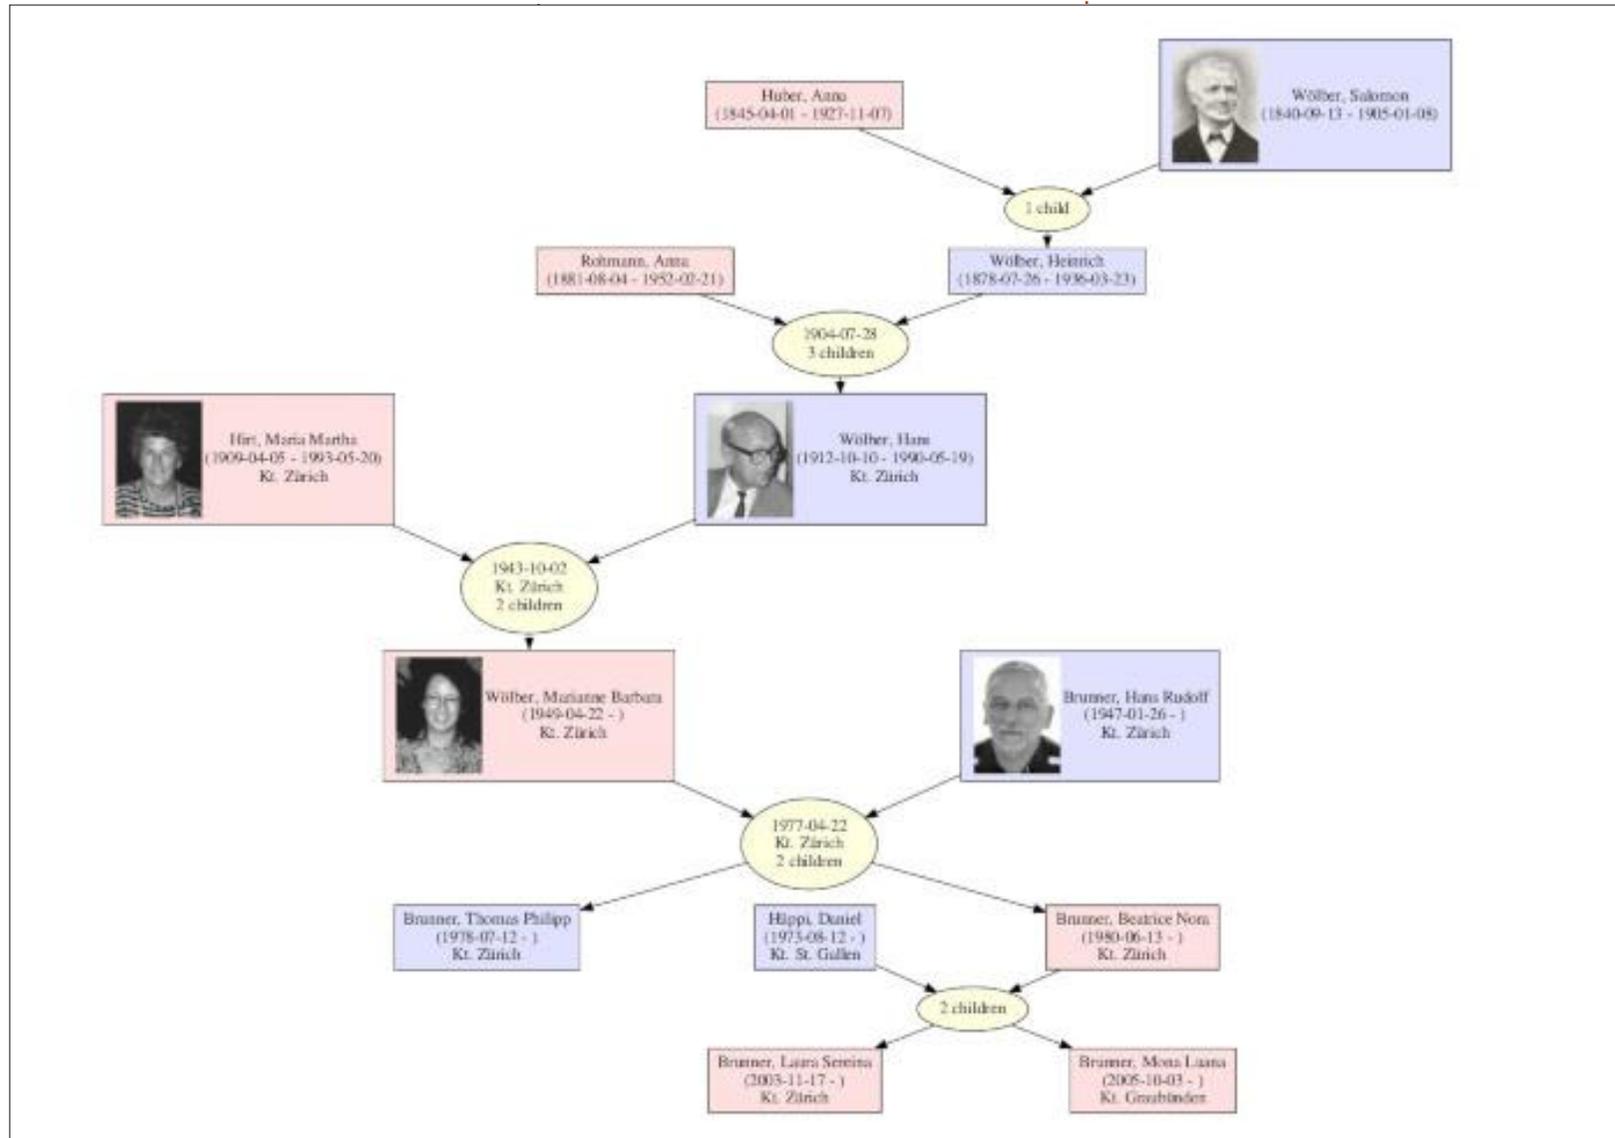
There are some problems with genealogy: there is no end, you will always find another branch to follow; knowledge gets lost when

people die, so ask people as long as they are alive.

And as a last hint: Gramps can be installed for several languages, or you may run gramps in a different language than originally

installed by using (here: swiss-german):

`LANG=de_CH.utf8 gramps`
`"FamilyTreeXXX"`





In issue FCM#144, we had a look at basic netcat. In this issue, we will take a deeper look at netcat. Josh, your wish is our command. Feel free to reach out to us on e-mail or Telegram if there is anything you would like to see more of in FCM. Also, feel free to make a donation to Patreon - <https://www.patreon.com/fullcirclemagazine>. That said, we shall waste no more time and go for the jugular!

PORT SCAN

We will be scanning <http://scanme.nmap.org/>. Please, if you want to scan anything besides this, make sure you own the machine. We do not want our readers in the charge office. So our terminal code should read:

```
nc -vz scanme.nmap.org 22 25 443
```

This scans 3 common ports at once (more, if you would like!). You can refer to our previous article on netcat regarding ports. Instead of me spelling out the options, your

homework will be to look them up in the man page. If you would like to scan a range of ports, you simply type the range:

```
nc -vz scanme.nmap.org 1-100
```

This is not the only way to scan with netcat. You can scan via the service name too. In the terminal, type:

```
nc -vz scanme.nmap.org ssh
```

Like before, you can string multiple service ports together, but you will get results only from open ports.

BANNER GRABBING.

Port scanning is all well and good, for known ports; what about unknown ports? How do we figure out what services are running on those ports? That is where banner grabbing comes in. Banner grabbing nets you information on open ports on a computer. If you are a new administrator, it is a neat way to find out what is happening on the systems you are taking over.

Hackers use it to find exploits on services running on those ports. There is a search engine called Shodan that lets you search computers by type, that searches by banner: <https://www.shodan.io/>

If you look up the -v option in the man page, you see it is verbose. The -z option limits it slightly, so leaving it off will get the banner from that port printed to your terminal. Many times, you can guess the function of the port, as people are lazy and port 22 just gets remapped to port 2222.

Why do you need to know this about your servers? Because this is how servers get hacked. As per Josh's request, I will briefly try to explain how hackers use netcat to actually hack a server. The most common uses for netcat, when it comes to hacking, is setting up reverse and bind shells, piping and redirecting network traffic, port listening, and banner grabbing (which we have covered). Yes, all of this and more is possible with netcat! Once you see and understand how hackers

compromise your systems, you can take preventative measures.

If you do not have a decent internet connection at home, or would prefer to follow along offline, you can download an insecure VM image and practice against that. I use the metasploit one, but you are welcome to use another. Here are a few: <https://pdrsecurity.com/10-sites-find-vulnerable-vm-testing/> or <https://www.vulnhub.com/>

Netcat is best used with other tools like nmap or metasploit, but since we will be focusing on netcat, I will break this information piece into three parts. Next issue we will look at creating backdoors and reverse shells.

AGAIN! I do not know what the laws in your country are, but do not scan outside your network as it may be illegal in your country or state. Yes, just simple port scans can be illegal. If in doubt, use the insecure virtual machine please.





HOW-TO

Written by Mark Crutch

Inkscape - Part 85

A quick recap on where we got to last time: using CSS classes, we were able to toggle the style of a button by clicking on it. But it worked well only for a simple button with no content. If you clicked on the text inside the button, the toggle wouldn't work. And you were still able to select the text, which is less than ideal for a button. What we really want to do is to group together several

objects, and have a click on any of them act as a trigger for the button. And, with a little more CSS, we can deal with the selectable text issue, too. So let's begin in Inkscape, by designing a fancier button...

This button is made up of four objects, each of which has a solid fill color so there are no stroke colors to worry about. When

toggled, we'd like to change the background and text colors, and give the impression of the button being 'pressed in' by making the top/left object dark and the bottom/right one light. We can therefore draw up a small table showing the type of each object and the colors it will adopt in each state.

We're going to put the entire button into a single group. This isn't strictly necessary, as Inkscape's layers are already SVG group objects, but does make it a little easier to work with if we want to add more than one button (or other objects) to a single layer.

Now it's time to set up the CSS classes. The basic idea is that we will give the outer group a class of

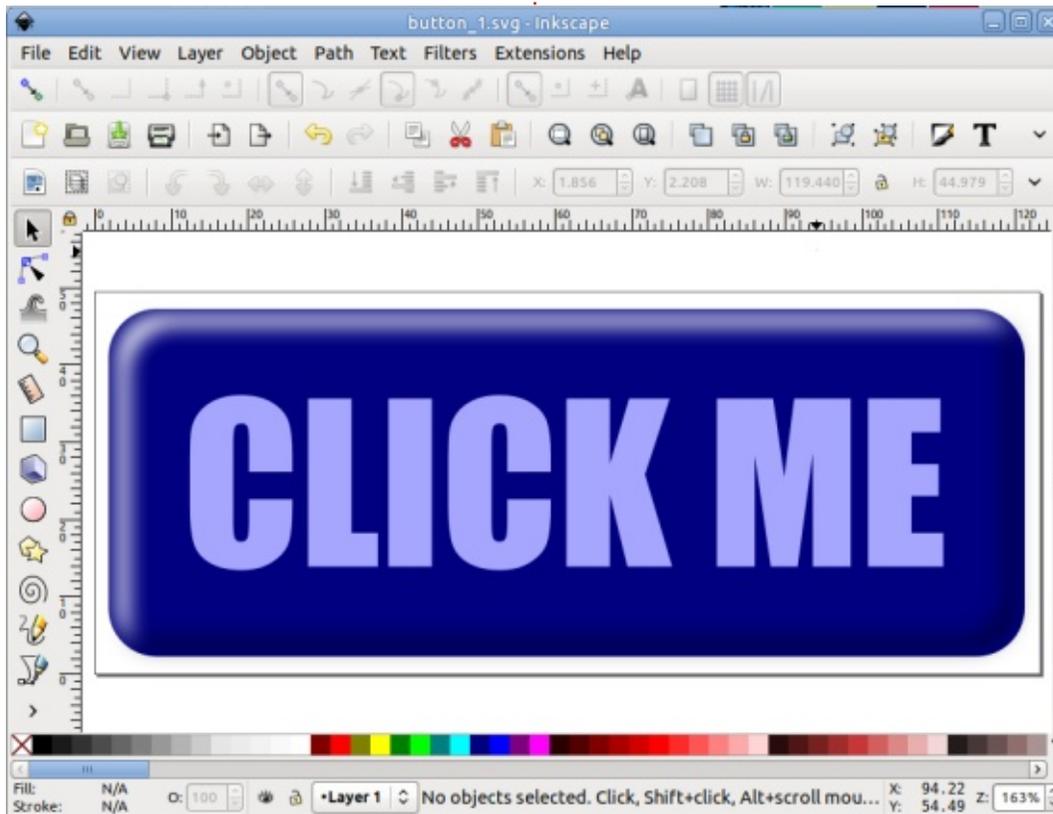
Object	Normal	Clicked
Rounded rectangle	#000080	#800000
Top/left blurred path	#ffffff	#000000
Bottom/right blurred path	#00002e	#ffffff
Text	#a6a6ff	#ffa6a6

'button', and then toggle an additional 'clicked' class on and off using one of the techniques from last time. Let's start by looking at the structure of the button, as created by Inkscape, but with most of the attributes removed for clarity:

```

<svg>
...
  <g inkscape:label="Layer 1">
    <g id="g972">
      <rect id="rect10" />
      <path id="path31" />
      <path id="path874" />
      <text id="text958">
        <tspan
id="tspan956">CLICK
ME</tspan>
      </tspan>
    </g>
  </g>
</svg>

```



It's all pretty much as you might expect: a `<g>` (the layer in Inkscape) containing a `<g>` (our button) which contains the four objects from the table above. There's a slight oddity in terms of the `<text>` object containing a `<tspan>`, but that's just down to the way SVG handles text. If we had multiple lines of text in our button, this might make more sense, as all the lines would then be separate `<tspan>` elements contained within a single `<text>`, but, even with only one line, Inkscape still uses a `<tspan>` even though it's not strictly necessary.

To keep the CSS a little clearer, it's best to change the IDs of some objects, or give them classes to better describe what they do. Otherwise, trying to remember which `<path>` is which at some point in the future becomes a problem. I tend to use classes for these sorts of labels, so they can be reused in other parts of the file as well. After all, you might want a second button to also have a 'top-left' path, so using that string as an ID would become prohibitive (remember, IDs have to be unique in a document, classes don't).

With the addition of a class for

the button group, and one each for the paths, we've basically got this structure:

```
<g class="button">
  <rect />
  <path class="top-left" />
  <path class="bottom-right" />
  <text>
    <tspan>CLICK ME</tspan>
  </text>
</g>
```

Now it's time to add our CSS rules. We'll use the immediate child selector (`>`) to ensure that these rules apply only to elements inside our button, so there's no danger of all the text in the document becoming blue. Compare these rules with the table earlier in this document and you should be able to see what they're doing:

```
<svg>
...
<style>
.button > rect {
  fill: #000080;
}

.button > .top-left {
  fill: #ffffff;
}

.button > .bottom-right {
  fill: #00002e;
}

.button > text {
  fill: #a6a6ff;
}
```

```
</style>
<g inkscape:label="Layer 1">
...
</g>
</svg>
```

Don't forget to also remove any explicit fill and color properties in the 'style' attributes of your elements (including the `<tspan>`), so that the CSS rules you've added aren't overridden. If you load the image into a web browser, it should look the same as the original version in Inkscape. If you're not sure that your styles are working, or that you've removed all the overriding properties on the elements themselves, try temporarily changing the colors in the CSS to other values and confirm that it has an effect when you reload the page.

Take a careful look at each CSS rule to make sure you understand what's happening. Pay particular attention to the difference between an element selector (eg. 'rect') and a class selector (with a dot – eg. '.button'). So, in this case, '.button > rect' (matches any `<rect>` that is an immediate child of an element with the 'button' class), and '.button > .top-left' (matches any element with the

'top-left' class that is an immediate child of an element with the 'button' class). Spotting the difference between a class, ID, and element can be tricky. The syntax for CSS is excessively terse, and less than obvious if you're not used to it, but it's the language we're stuck with so we'll have to make the most of it.

If everything is working okay at this point, it's time to add another set of rules that will apply when the `<g>` has both the 'button' and the 'clicked' classes set. In this case, you just have to concatenate the class selectors – but make sure not to add any spaces between them, as that signifies an ancestor-descendent relationship. Yeah, the syntax of CSS really is that terse.

Here's an example of the additional rule for the `<rect>`, complete with the new color from the table earlier in the article. I'll leave it as an exercise for the reader to create the remaining three CSS rules.

```
.button.clicked > rect {
  fill: #800000;
}
```

You can test your new CSS rules by manually adding an extra

HOWTO - INKSCAPE

'clicked' class to the group (so that it reads class="button clicked"), saving the file, and reloading the page. Remove the extra class and save again before proceeding.

Open the file in Inkscape and select the group that represents the button. In the Object Properties dialog, expand the 'interactivity' section and add the following one-liner to the 'onclick' field (this should be familiar from last month's column):

```
this.classList.toggle('clicked');
```

Save your file, load it into your browser and, if everything has gone smoothly, you should find that clicking on your button toggles it between the two states.



One advantage of wrapping everything in a group, and applying the code to that outer layer, is that clicks on any part of the button are passed through to the enclosing element. This avoids our previous problem whereby clicks on the text didn't toggle the button. But we still have an issue with the text being selectable. We can address this with the 'pointer-events' CSS rule, which lets us tell the browser that all mouse activity over the text – including clicks and selection – should be ignored. Modify your first set of rules so that the last one looks like this:

```
.button > text {  
  fill: #a6a6ff;  
  pointer-events: none;  
}
```

That deals with the selectability problem, but we can go a step further in making our button seem clickable. By adding a 'cursor' property to the group itself, we can make the mouse pointer change when it moves over the button. Add this rule to the CSS:

```
.button {  
  cursor: pointer;  
}
```

Save, reload, and move your

mouse pointer over the button to see the effect.

Previously, I said I would show you how to make clicks on one element have an effect on a completely different one, but I've run out of space in this article, so that will have to wait until next time. Until then, however, you now know how to use CSS classes to style more complex collections of objects, allowing for the creation of much fancier buttons. You also know how to respond to clicks in such a collection simply by putting the code onto a group that wraps all the content.

Why not spend the next month designing ever more impressive buttons? There are plenty of tutorials online for different styles if you're stuck for inspiration. You don't just have to make them toggle when clicked: a common effect is to apply a class in the 'onmouseover' field, and remove it in the 'onmouseout'. You could even go the whole hog and create a button that has both a mouseover effect and a click effect. And next time, I promise, I will show you how to hook your new buttons up to other elements on your page.



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





What is your new year's resolution mate?

I am staying with 1920x1080, why?





After getting my GhostBSD 18.12 fully functional, the new update 19.04 became available. After further reading, Eric Turgeon suggested a full reinstall. Eric modified the True OS base in some form. I did the exact steps as before, and had a successful install. I was quietly jumping for joy.

I rebooted the HP laptop, and attempted to boot to the new upgrade. Yet I did not get past the new GhostBSD splash screen. I simply stalled at this image for well over 20 minutes. I tried the official and community supported editions.

I drifted to their Telegram, and posted my issues. Eric and the entire team were quick to troubleshoot the program. Their early suggestions were to change some of my default boot options, yet at this rate I decided to look for other options. It should be noted that I drifted over to the Project Trident project. It is the graphical version of the former TrueOS desktop. I followed the instructions and placed Project

Trident onto the HP laptop. I booted the computer and installed using the default options. At this time Project Trident was under the 18.04 version. And it is fully installed with the same issues of having to setup my Wifi. However this time my Realtek wireless card was recognized by the installer, and the Lumina DE is surprisingly refined, polished, and resolved.

After speaking to Project

Trident staff and GhostBSD, I made some interesting observations about these 2 similar operating systems, but opposite approaches to creating the BSD end user environment. GhostBSD uses a larger portion of GTK, Project Trident uses Qt. GhostBSD works well with Intel architecture, and Project Trident covers a wider equipment ranges. Unfortunately, my HP laptop is powered by AMD architecture.

Currently, GhostBSD is more popular than Project Trident. Next month I will do a fuller review of Project Trident and a proper review of the Lumina desktop.



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.



What is Project Trident?

Project Trident is a desktop-focused operating system based on [TrueOS](#). It uses the [Lumina desktop](#) as well as a number of self-developed utilities to provide an easy-to-use system that both BSD beginners and advanced system administrators can feel comfortable running ²⁴/₇.



GhostBSD is an Unix derivative operating system based off TrueOS. It is meant to be an easy segway into BSD via graphical installer. This operating system is a pre-configured desktop with common software. The end user can use it for coding, internet surfing, and other daily computer tasks. In the simplest of terms, it is a free and open operating system for any computer user.

When I installed GhostBSD I had difficulty getting my wireless card to work. I reached out on Telegram to ask for help. The lead developer **Eric Turgeon** answered my questions and helped me fix my wireless connection. Eric graciously agreed to a quick interview.

What is your background? How did you get started in this field? Any suggestions to others that want to follow your path?

I did not finish high school, which mean I have no CS degree or software engineering degree. I am a self-taught programmer. I was

hoping to study in computer science at the university, but the irony is that I did quit school after realizing that I was 15 year in school and still had 2 years of French grammar course to complete and it was the only thing keeping me from going to college or university.

On my journey of becoming a hacker/cracker searching for more information and new tools to hack/crack, I did find 'How To Become A Hacker' from Eric Steven Raymond. Eric's essay was coming at the right time because my life was starting to change, I became a Christian, and continuing the path of a hacker/cracker was no longer an option for me. I started to be more interested in Open Source, and I discovered I was able to learn computer science by myself.

I recommend to everyone to read 'How To Become A Hacker' at <http://www.catb.org/esr/faqs/hacker-howto.html> by Eric Steven Raymond. For people to follow my path, they need to be curious and

ready to pass time to read, implement what they learn and start a project that would force them to learn what they need to get things done. Yeah, that sums it for me, because everything I know today is due to the fact that at one point I needed it for GhostBSD, FreeBSD or work.

What lead you to BSD? Why not another operating system platform?

When I did read 'How To Become A Hacker' a couple of times, BSD and BSD UNIX was mentioned, and BSD Unix was sticking in my head because of UNIX. In the essay, it stated 'You can find BSD UNIX help and resources at www.bsd.org.' from there FreeBSD was looking promising. I did download FreeBSD, but with the lack of a GUI I was not able to do much so it did turn me off at first, with more search I did find PC-BSD and installed PC-BSD 1.4, but I was a Gnome2 guy, and at that time PC-BSD was KDE only. I did like Ubuntu a lot, and I thought why is

there no project like Ubuntu in the BSD world, so that was about that time that the idea GhostBSD started. Since only FreeBSD supported native Nvidia drivers, I started there.

Today with FreeBSD, GhostBSD, TrueOS, and FreeNAS, it feels like home. It is hard to imagine to go back to GNU/Linux, and I would say that for me it is appealing to know that BSD is a replacement of Unix and that BSD is derivated from BSD Unix. Also, the BSD license makes much more sense for me compared to GPL.

Whom do you see as a common user of BSD?

I see mainly BSD for servers, but for people that have the right hardware, it can work great for laptop and desktop use. I would say it is probably for people that want a system developed at a safe pace and is not happy with what is happening with other systems. There is also some user that use BSD without knowing it by using

INTERVIEW

iPhone, Mac OSX, PS3, PS4, Netflix, WhatsApp, and so on.

Linux is being consumed by corporate investment, IBM and Redhat for example. Do you think BSD can get a larger foothold of Linux users due to the corporation consumption model in Linux?

FreeBSD is used, by iXsystems for FreeNAS, TrueNAS, TrueCommand, TrueOS. FreeNAS is, in my opinion, the best Open Source Network Attached Storage software out there, it also makes TrueNAS the best commercial solution for companies that need a massive amount of storage and the customers supports is excellent. Apple iPhone and MAC OSX are BSD systems. Sony did use FreeBSD for PS3 OS and PS4 OS. Netflix is using FreeBSD for its Open Connect Appliance. There is a lot of other companies using FreeBSD there is a list at <https://www.freebsd.foundation.org/freebsd>. So FreeBSD as already a foothold and it foothold continue to grow. Has a FreeBSD ports committer, the GhostBSD project leader, founder, and developer and an iXsystems employee, I see FreeBSD gaining more territory every day.

I do not talk about the other BSD's, because I don't know much about them.

What are the best resources for a new BSD user?

For FreeBSD the handbook is excellent, I would say any documentation that a project provides should be good. Absolute FreeBSD, 3rd Edition: The Complete Guide to FreeBSD, FreeBSD Mastery: Jails, FreeBSD Mastery: ZFS are great books also.

Editing configuration files is daunting for new users. Do you see this issue preventing new user adoption?

Yes and no, Yes for users that never use a system at the administration level, but for users that have used Linux to the same level BSD no.

You created GhostBSD. How difficult was it to develop the platform?

Created is a big word let I have put lots of codes together that I did not understand together and call it GhostBSD, at first. After that, I had to take the time to learn how it

works and why it was working. The 2 first release did not contain much code from me, and it was FreeSBIE and code from a guy that did a Gnome live cd with FreeSBIE. I had no experience with programming and did know anything about FreeBSD. So yeah that difficult, but I slowly started to understand FreeBSD and was able to have help at the start to guide me and learned C and Python along the way.

Where do you see GhostBSD in the next 5 years? 10 years?

I would like to have GhostBSD with all the missing bits and tools to make it just work without any effort of the user and to make it OEM ready. My long term vision would be to write a GTK DE that works flawlessly with FreeBSD.

When GhostBSD migrated to TrueOS was it difficult?

No, because it is still FreeBSD, but with OpenRC and with base systems package. From the outside, it looked difficult, but it was reasonably easy. What happened there was a lot of old issues that needed to be fixed, adding all the changes, rolling our

packages and system packages, it was a lot to do for one release. I had many people helping on the GhostBSD Telegram group with testing, with code, and moral support. My manager Joe who rewrite all the code to build GhostBSD, Ken from the Trident Project was very useful also when I had problems. Participating in the TrueOS meeting also helped a lot because I am aware of everything related to TrueOS development. I would say that TrueOS, GhostBSD and Project Trident are working collectively to the greater good of desktop on FreeBSD.

What is the most common use for GhostBSD? What hardware do you recommend for a simple install?

GhostBSD is commonly used for laptop and desktop. It is also great for a workstation with five disks on ZFS RAIDZ3. For me, it is my daily driver for work, and it is the OS I write all the code for all the project I am involved in. I even play games on GhostBSD, Xonotic is my favorite game.

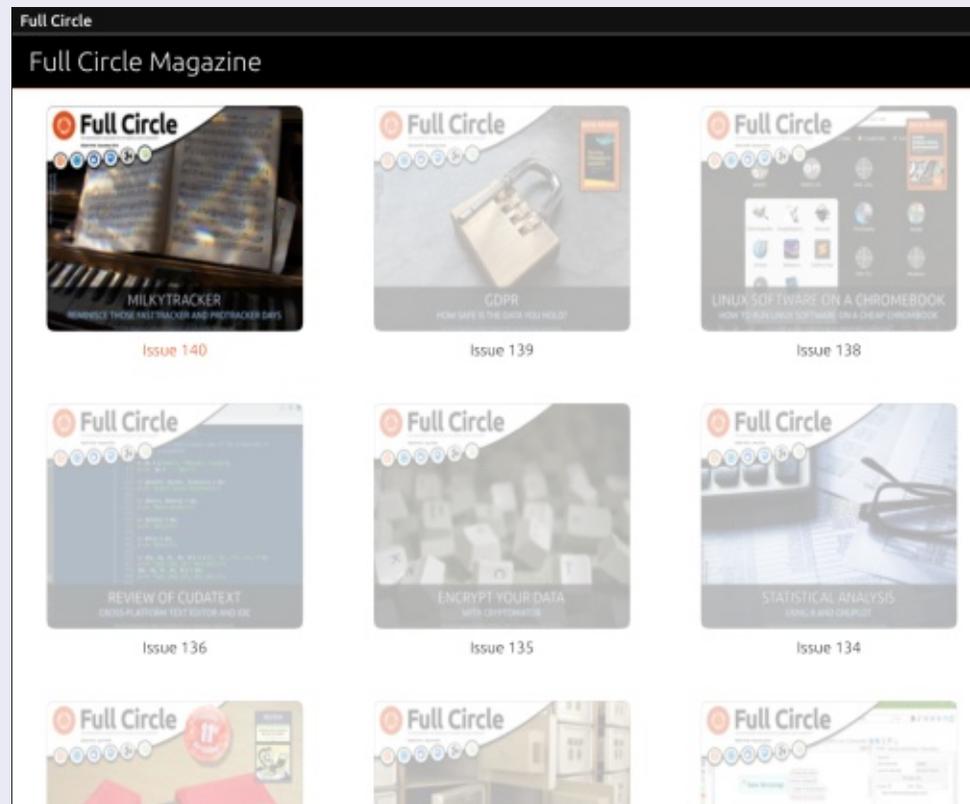
My main desktop is a GIGABYTE X470 AORUS Ultra Gaming with Ryzen 2700X, 64G of ram and Nvidia 1050TI, so it is running fairly

INTERVIEW

good on new desktop hardware, I would say for a desktop watch for a compatible network card with FreeBSD, and the rest should work great. For laptops, most Lenovo ThinkPad laptops are working great. For most people, the problem will be the WiFi, not all Broadcom, Qualcomm is supported, so that is to watch. There is some user that have swap their laptop WiFi cards to be able to use GhostBSD.

Eric, thank you for everything you do with GhostBSD. Your support on Telegram is flawless. Also, I appreciate the chance to speak with you. Congrats on betting out TrueOS for popularity. GhostBSD is a great operating system.

THE OFFICIAL FULL CIRCLE APP FOR UBUNTU TOUCH - UPDATED!



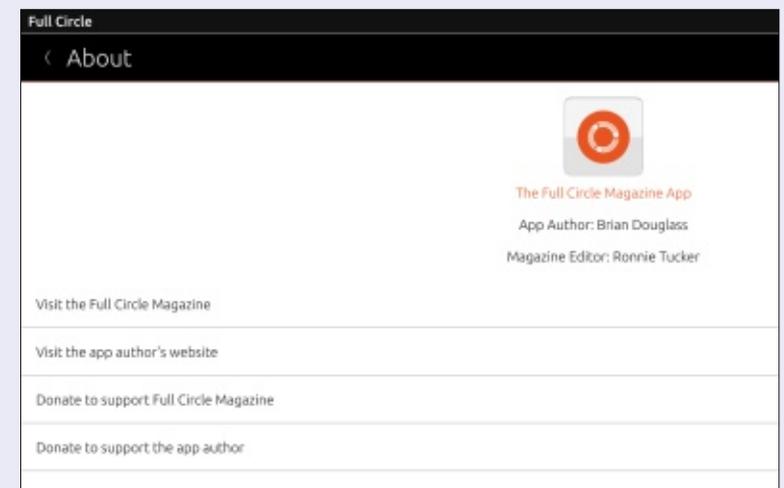
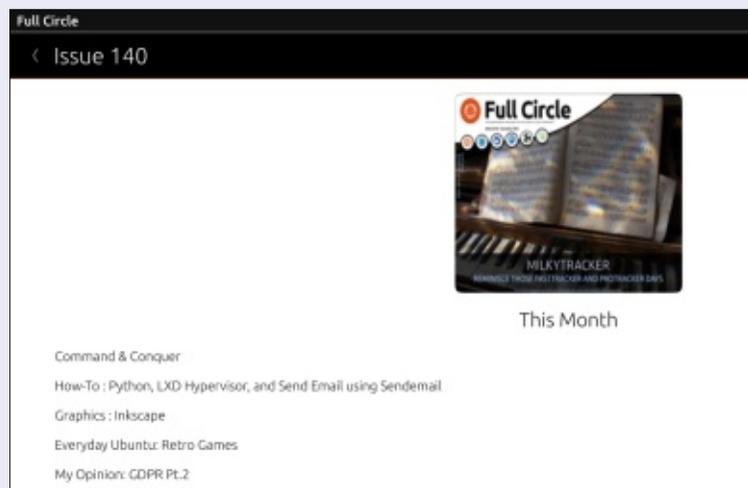
Brian Douglass has updated his FCM app for Ubports Touch devices that will allow you to view current issues, and back issues, and to download and view them on your Ubuntu Touch phone/tablet.

INSTALL

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

<https://uappexplorer.com/app/fullcircle.bhdouglass>

HUGE thanks to Brian for this.





One of the truly wonderful things about Linux is the community that has grown up around it. There is a stereotype of arrogant technophiles who look down on 'noobs' or 'newbies', and there's no sugar-coating it – you will run into some of those online or in real life from time to time, but that's life for you. Most Linux experts I've encountered will at least try to help you if they can, and that is a great thing for anyone trying to learn how to do things in Linux. This month's column will start teaching you how to get your own help on Linux, which in many ways will give you 'the keys to the kingdom'. We'll start with a discussion of 'man pages'.

MAN PAGES

Man pages used to be the foundation of almost all Linux knowledge transfer. Linux itself, and most of the applications in Linux, has man (short for 'manual') pages. I once had an issue with LILO, a common Linux boot loader (this was back around 2003), where

something within LILO had gotten bollixed up enough that I could no longer boot into Linux. Sorrowfully, I resigned myself to having to reinstall the operating system, all my applications, and the loss of all my data.

At that point in time, I was dual-booting my system (as I usually still do), and most of my critical data lived on the Windows side of my computing life, but reinstalling Linux and its applications was still going to be quite a lengthy and tedious process that I was not looking forward to. I posted something online about the issue, and I got back a few responses that basically said haughtily, "Go read the LILO man pages." Well, that wasn't really practical for me, and, for many people, it would be even less so today. The man page would be a complete description of LILO, not the solution to a specific problem, or even solutions to a number of specific problems. They would not be practical for solving a specific issue.

This handful of 'Linux snob's', if

you will, really put me off. Fortunately, I then got a LOT of responses from much more helpful people who let me know that, with the particular problem I had, all I needed to do was boot from a diskette (THAT'S dating me), and actually just run LILO from a command-line (no parameters, just 'LILO' then <Enter>). Sure enough, that solved my problem easily and handily, and illustrated for me that the community was mostly very helpful. Reading through an entire man page to become familiar with LILO top to bottom wasn't really something I had time for, at the time. The Linux community came through for me with a very simple and quick fix.

Oh, and just to date me even more than the floppy disk reference, the flavor of Linux I was running at the time was ... Caldera. Try looking that up online to see why Caldera OpenLinux is mostly remembered by Linux veterans with contempt or outright hatred (or, at least, publisher SCO is remembered that way). It's a long story, but an interesting one, to be

sure.

Ending our little digression here, the man pages can still be a valuable asset in Linux if you're looking to learn (they're less practical for solving a specific current problem, like I had with LILO). To access the man page for a specific function or application (let's say, 'LESS'), go to a terminal window (look for 'Terminal' in the Dash – click the topmost icon in your Launcher and type in 'term', that should suffice, then click the Terminal). At the command-line, type 'man' and the command you want to learn about, in this case:

```
man less
```

Then hit <Enter>. Linux will bring up the man page for 'less', and you can read all the details on how it works. Man pages do tend to be a 'deep dive', but they're still a great resource if you want to learn, and particularly if you're using something on the command-line (which can be, at times, a much more efficient way to do certain things than the GUI is, although

obviously less intuitive). In man, you can use arrow keys to move up and down if the page doesn't fit on one screen (it usually won't). You can hit 'h' on the keyboard for man page help, or 'q' to quit from the man page reading utility and return to a command prompt. Just for fun, you can use

`man man`

to see the man page on the man page reader. Linux recursiveness is a long-standing tradition in the community.

Man pages are also available for many Linux applications. Try:

`man libreoffice`

from the terminal to see one example.

UBUNTU HELP/UBUNTU DESKTOP GUIDE

Ubuntu has a built-in Help utility that is probably much more practical for learning Linux today than man pages. Go to the top right of your screen and click the Ubuntu 'gear' icon:

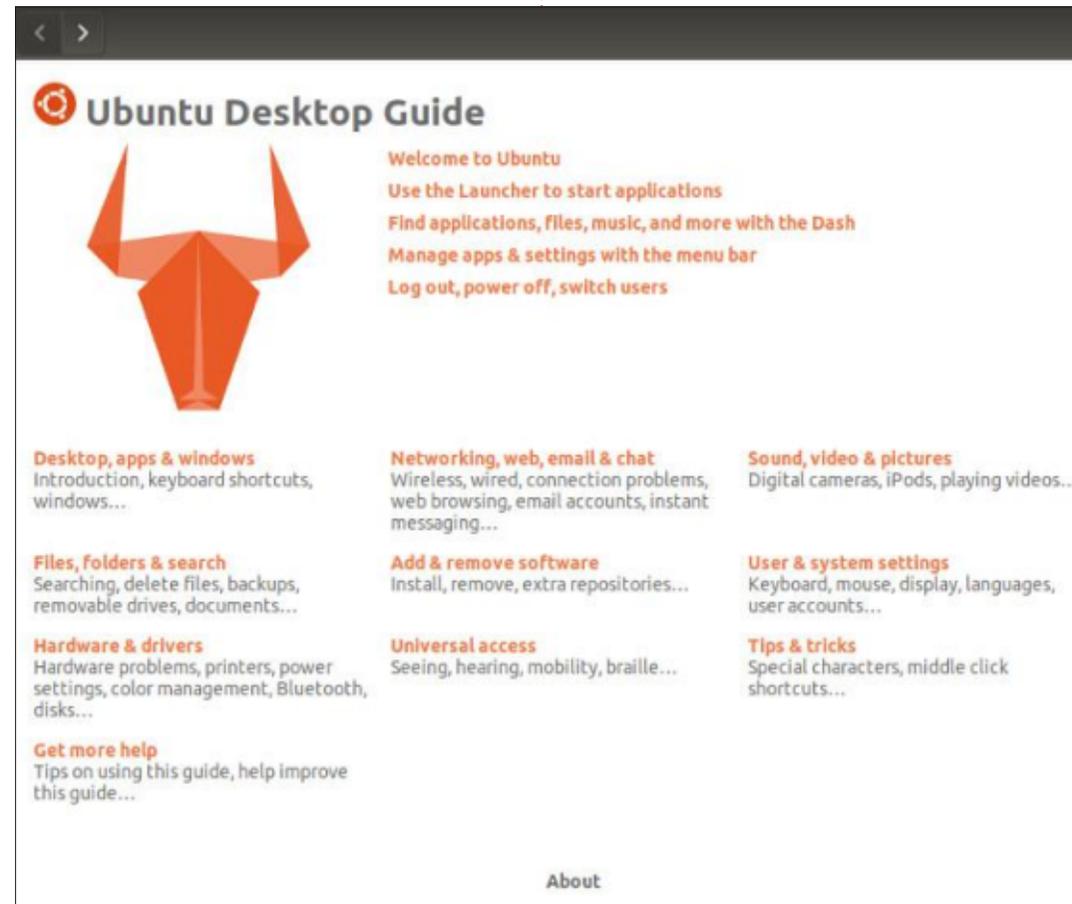
Select 'Ubuntu Help....' and Ubuntu Help will open. It will appear as a box with a question mark inside a circle on the

Launcher if you find yourself needing to switch to/from it. Ubuntu help has a Search function (click the magnifying glass), and a bookmarking utility (click the star to peruse or add bookmarks). You also have arrow keys in the top left to let you navigate back and forth. The Ubuntu Desktop Guide is a great introduction to Ubuntu Linux and is a very good place to start your journey to Ubuntu Mastery.

GOOGLE IS YOUR FRIEND

Given how much traction Linux has online in general, it will come as no surprise that there are almost endless resources available via Google search. You can search for answers to particular questions or problems, you can Google for general information on Ubuntu, Linux, or even Unix (most Unix information will still apply). You can also find a plethora of Linux

```
Terminal File Edit View Search Terminal Help
MAN(1) Manual pager utils MAN(1)
NAME
man - an interface to the on-line reference manuals
SYNOPSIS
man [-C file] [-d] [-D] [--warnings[=warnings]] [-R encoding] [-L
locale] [-m system[,...]] [-M path] [-S list] [-e extension] [-l|-I]
[--regex|--wildcard] [--names-only] [-a] [-u] [--no-subpages] [-P
pager] [-r prompt] [-7] [-E encoding] [--no-hyphenation] [--no-justifi-
cation] [-p string] [-t] [-T[device]] [-H[browser]] [-X[dpi]] [-Z]
[[section] page ...] ...
man -k [apropos options] regexp ...
man -K [-w|-W] [-S list] [-l|-I] [--regex] [section] term ...
man -f [whatis options] page ...
man -l [-C file] [-d] [-D] [--warnings[=warnings]] [-R encoding] [-L
locale] [-P pager] [-r prompt] [-7] [-E encoding] [-p string] [-t]
[-T[device]] [-H[browser]] [-X[dpi]] [-Z] file ...
man -w|-W [-C file] [-d] [-D] page ...
man -c [-C file] [-d] [-D] page ...
man [-?V]
DESCRIPTION
Manual page man(1) line 1 (press h for help or q to quit)
```



tutorial information on YouTube, either via a Google search incorporating 'YouTube' in the search parameters, going to YouTube and searching directly from there, or, as often as not, by simply clicking the Videos tab in your Google search results, where you will usually find plenty of YouTube 'hits'. We'll also talk a bit next month about online forums, which will undoubtedly make up a significant portion of hits that you will encounter in most Google searches on Linux topics. Here are a couple to get you started:

<https://www.youtube.com/watch?v=zA3vmx0GaO8>

<https://www.youtube.com/watch?v=IVquJh3DXUA>

THE UNIX HATER'S HANDBOOK

It's a bit of an 'oldie' at this point, but a very entertaining source of Unix knowledge (almost universally transferable to Linux) is the rather lengthy 'Unix Hater's Handbook'. You can find it available for download here:

<https://web.archive.org/web/20120120031001/http://m.simson.net/ugh.pdf>

The Unix Hater's Handbook takes an amusingly sarcastic look at Linux's quirks and peculiarities, but beyond the humor, there's a LOT of good information in it, presented in a surprisingly entertaining way. It was edited in 1994, from online posts in the Unix-Haters forum, by Simson Garfinkel, Steven Strassmann, and Daniel Weise. The information in it is mostly relevant to command-line functionality in Unix, but it can serve as a surprisingly good foundation of general Linux knowledge. Plus, it's good for more than a few chuckles.

Next month: Getting More Help in Linux (or More Getting Help in Linux, I suppose).



Richard 'Flash' Adams spent about 20 years in corporate IT. He lives in rural northwest Georgia, USA, with his adopted 'son', a cockatiel named Baby.

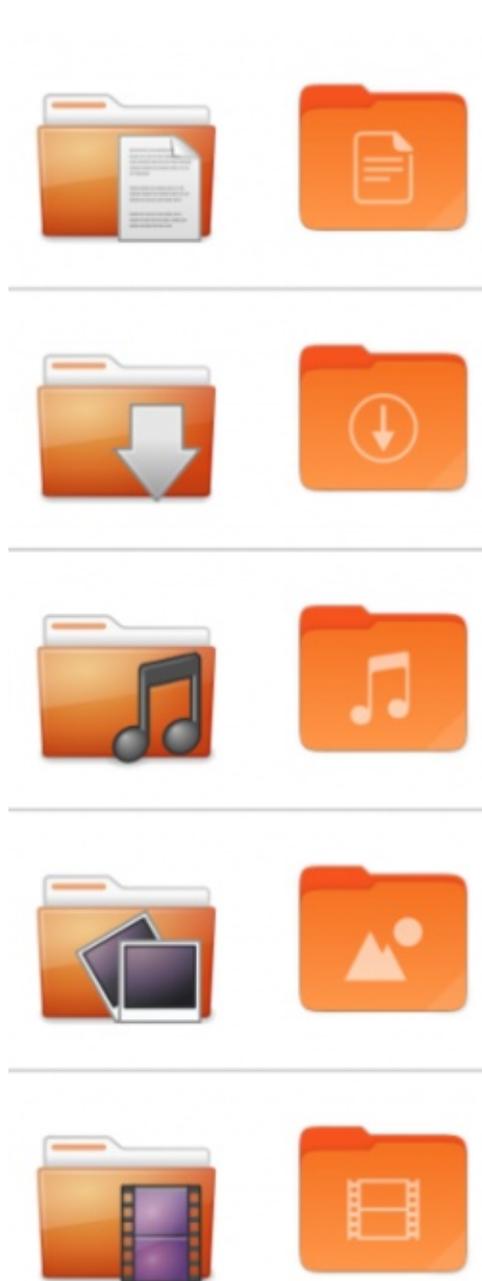


Like OTA-8, OTA-9 is mainly a stability improvement release. We are continuing our progress towards incorporation of upstream technologies into Ubuntu Touch, such as Canonical's Mir 1.x and Jolla's qtcontacts-sqlite. This work will multiply our effectiveness, allowing us to focus less on fundamentals and more on innovation. That doesn't mean you won't be seeing some needed changes, though!

You can find a full changelog for this release in the changelog section below.

A REFRESHED LOOK

Michele Castellazzi merged the work that was done by the Ubuntu Artwork Team between the last release of Ubuntu for Devices and the project being dropped by Canonical. This offers us a trove of new and updated symbols and folder icons, giving a better feeling of cohesion throughout the system.



NEXUS 5 CAMERA FIXES

Rúben Carneiro has fixed one of the most complained-about issues affecting the Nexus 5: the viewfinder freezing after taking a picture. The ability to record video was also fixed by the same commit, but switching between still image and video modes is a bit unreliable and may require switching away from the camera app.

QUIET IMPROVEMENTS: THE QQC2 SURU STYLE

The QQC2 Suru Style is an important project that doesn't get very much publicity. It is a Qt Quick Controls 2 style that follows the Ubuntu Touch design guidelines. This allows developers to use the freely available controls in their QML applications and easily port between different Qt platforms with automatic style changes to suit each.

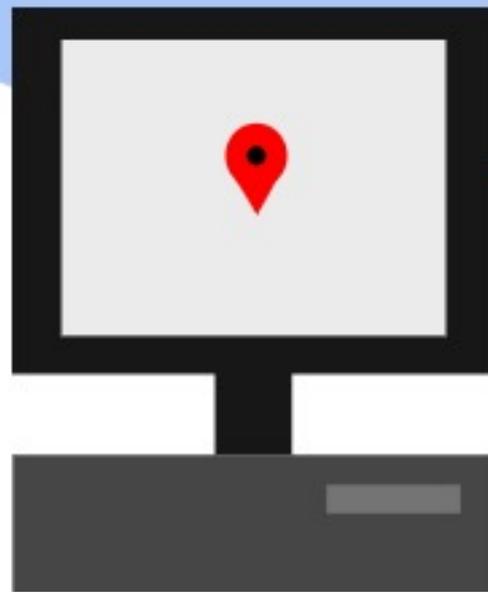
This release includes fixes for the style, including using the system scaling settings, better

detection of the system dark theme, and a new "Busy" indicator. A non-exhaustive comparison of some of the available QQC2 styles can be found in this image gallery.



Meanwhile in the Antarctic...

This is the only place google earth means nothing...



THE DAILY WADDLE
by: ErikTheUnready





LibreOffice Draw has always been the black sheep of the family. Yet this tool is unmissable in the modern workplace. It is mostly used to draw flowcharts. I use them to map out processes for situations, even flowcharts for people who perform certain tasks every day. The idea is to make a few 'templates', and adapt them to your situation. You can map a person's daily routine, and the escalation procedure for each step, and put it behind their office door. That day that person X does not come to work, and the boss needs Y, consult the flowchart behind person X's door and you are on your way. Be sure to add telephone numbers of people to contact if you cannot get hold of person X. It may take you a while to set up, but once you have a bunch of templates, all you have to do is fill in the blanks. Here I made a rough example of such a flowchart in under five minutes. (It is not real, but similar to what you would have in real-life.) You would probably spend 10 minutes on each template in real life – where you align everything, and think

about colours and gradients and arrows.

You can make these and attach them to your IT policies and procedures the users get, so there can be no confusion about how things work. Most of your time will actually be spent sitting with your users and making notes about what and how they do things. LibreOffice Draw is fast and easy to master. Everything you need to draw is on your left, and everything you need to write is on your right. You can even digitally sign each of your templates. There

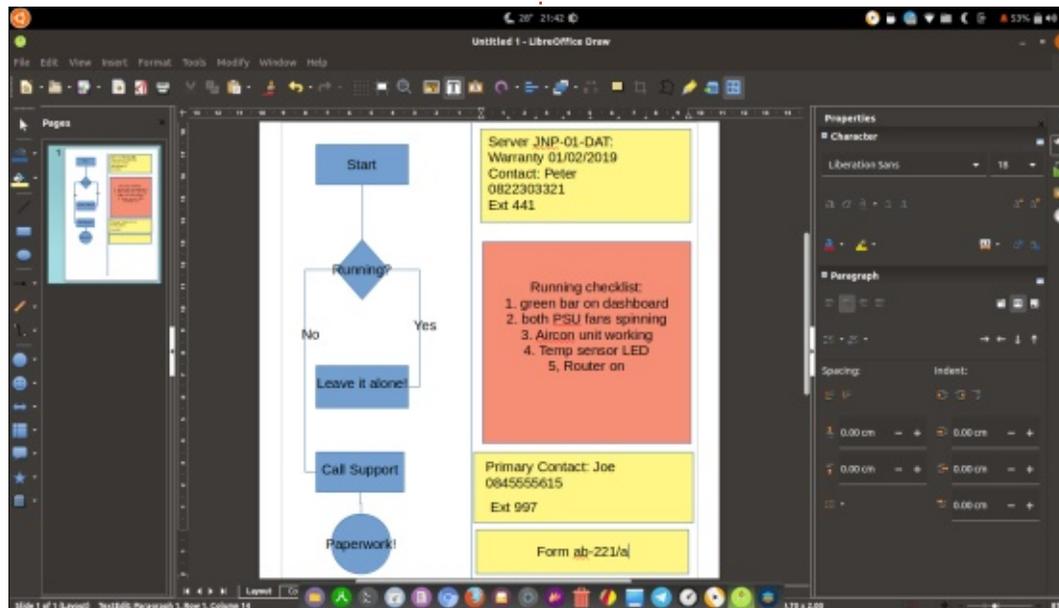
is a tip I can give you regarding use, firstly, press Alt + F12 and edit the settings to use OpenCL as this makes drawing smoother.

Once you are done creating these flowcharts for your users, you can easily print them to PDF and file them immediately. You then clean the test off your 'template', and make the next one. Software like dia, etcetera, specialise in making flowcharts, but this is much faster – with standard A4 layouts that easily slot in with your other work documents. Never let a user tell

you his / her job is too 'complicated'. In my experience, every job can be mapped this way.

This puts onto paper who is responsible for what and who it escalates to. This is also a great tool for hand-overs. Users can explain their jobs to the person taking over, but should there be some problem later, this is a great fallback. This also reduces the time for such hand-overs by half.

I bet you never looked at LibreOffice Draw in this light?



Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.



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://bit.ly/fcmwriting>

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

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

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

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

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

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

TRANSLATIONS

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

REVIEWS

GAMES/APPLICATIONS

When reviewing games/applications please state clearly:

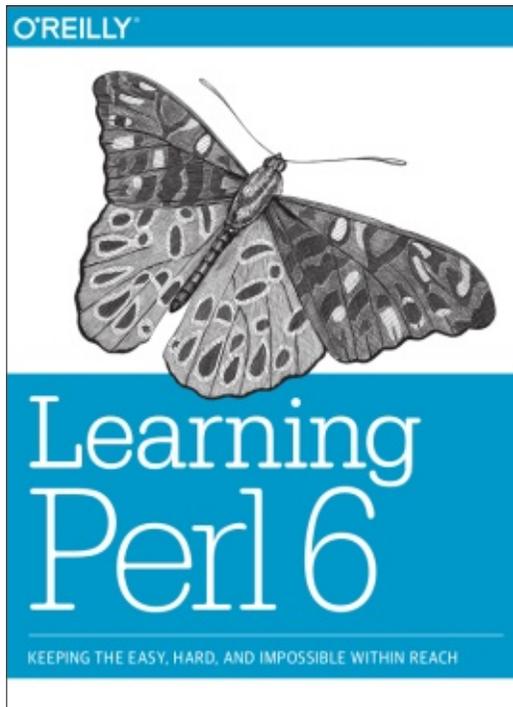
- title of the game
- who makes the game
- is it free, or a paid download?
- where to get it from (give download/homepage URL)
- is it Linux native, or did you use Wine?
- your marks out of five
- a summary with positive and negative points

HARDWARE

When reviewing hardware please state clearly:

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

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



Learning Perl 6

Author: brian d foy (sic)

Price: \$54.82! (Amazon)

From: O'Reily

If you're ready to get started with Perl 6, 'Learning Perl 6' is the book you want, whether you're a programmer, system administrator, or web hacker. Perl 6 is a new language - a modern reinvention of Perl suitable for almost any task, from short fixes to complete web

applications. This hands-on tutorial gets you started.

I must admit I have seen the screen name on perlmonks.org, but I have no idea who the author is. So much for something about the author, as this is obviously a pseudonym.

https://www.perlmonks.org/index.pl?node_id=385334

Perl 6 is the next evolution in the Perl ecosystem and I am holding my breath excitedly until it is included in your standard Ubuntu distribution. You can, of course, try it out already. This is why I decided to get this book. There is something about Perl that makes it sexy. I don't know if it is because perl docs are so much better than awk manpages, or that it is easy to pick up and use.

This book is packed chock full of information from the very first page of the introduction. There are exercises for you on every other page. This is a sink or swim book; do not be fooled by the mild mannered cover. The climb is steep

with numbers and strings, et-cetera, until chapter 5, where blocks are explained as: "Block(s) are the things that group multiple statements into a single thingy." This made me smile. Things ease off a bit, and we are led to believe this will go easier. The terms used are what made me read the chapter again, as I am not a programmer, things like phasers have a completely different meaning to me. The thing I respect most about this book is that is not brimming with acronyms. (A pet hate of mine.) Things are explained clearly and concisely. The odd factoids in-between, with little pictures / icons, break you out of your frown. Yes, frown you will... There are lots of examples, which is a plus. There is a whole chapter on errors - when things go wrong - that does a remarkable job explaining those cryptic error messages Perl loves. Something I have never used in Perl is modules, or to be precise, other people's modules. Chapter 10 will show you how to go about acquiring these modules for your own use. The downloads for the book are

actually kept on Github, so the links in the book are not used. Find it all here:

<https://github.com/briandfoy/LearningPerl6> Downloads

As this is the very first edition, small errors are to be expected. Reading "Perl6, keeping the easy, hard and impossible in reach" is like balancing on a tightrope. There is so much information here, but it always stops short of information overload. At the end of chapter 21, there are almost another 100 pages covering the answers to the various exercises. (Yes that's how many exercises there are!).

I will not spoil the book for anyone - woofles dies in the end - but I CAN say you will not regret reading this. You can read some of it online if you make an account at O'Reilly media. I did not do any of the exercises as I am planning on reading it again, this time doing the exercises - I wanted to get a feel for what I was letting myself in for.

BOOK REVIEW

My recommendation, for anyone reading this book, is to start at the beginning and not skip a single thing, no matter how well you think you know it.

The only negative part is the price of the book, \$54.82 from Amazon and over \$100 locally, which every man in the street can not afford.

The content of the book:

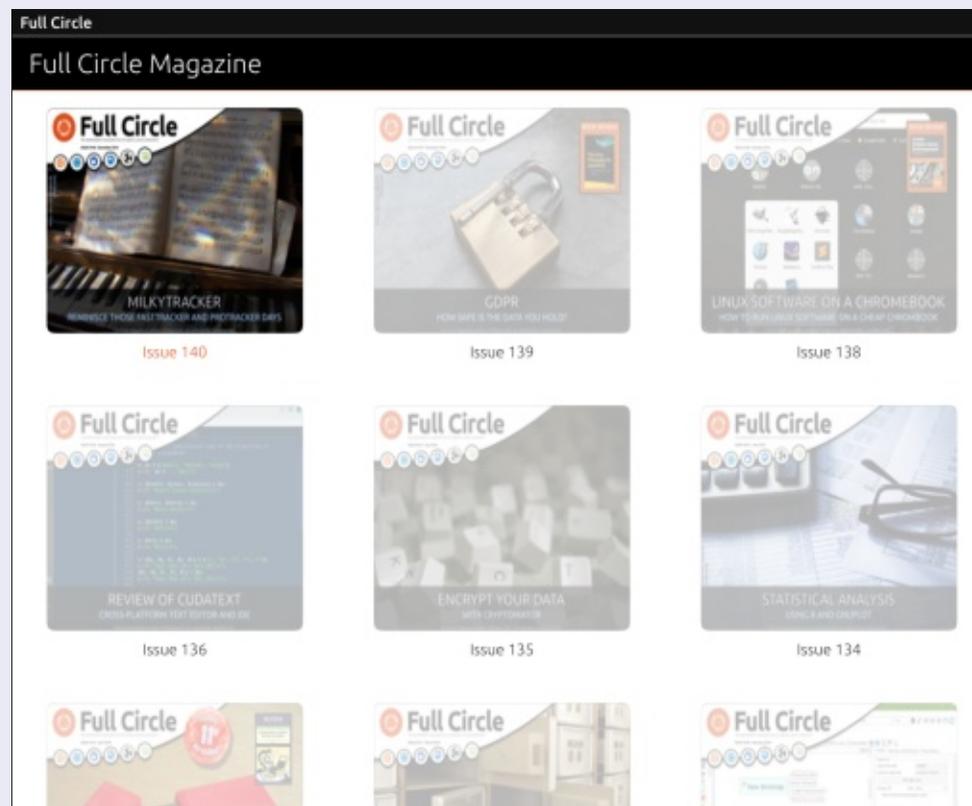


The hefty price gets :



Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.

THE OFFICIAL FULL CIRCLE APP FOR UBUNTU TOUCH - UPDATED!



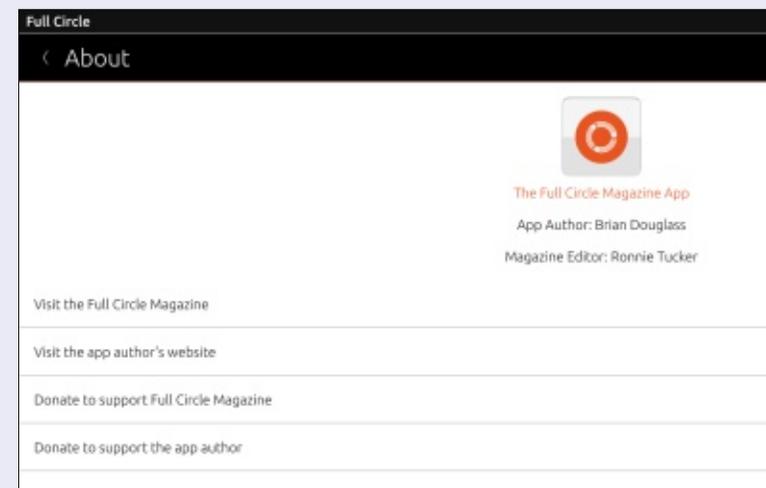
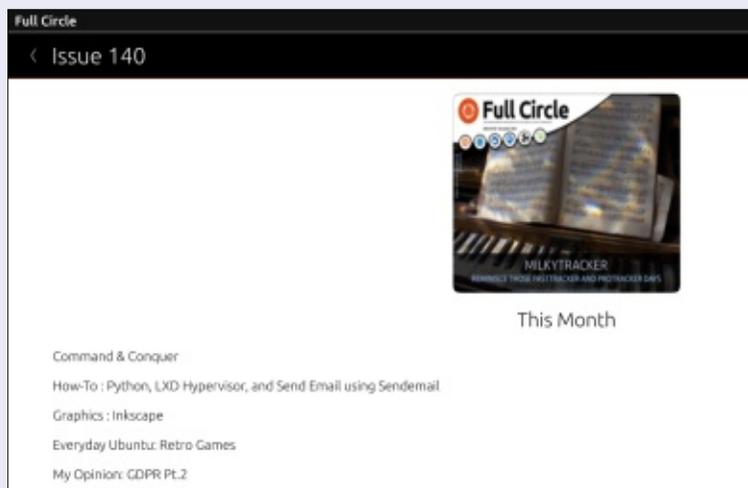
Brian Douglass has updated his FCM app for Ubports Touch devices that will allow you to view current issues, and back issues, and to download and view them on your Ubuntu Touch phone/tablet.

INSTALL

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

<https://uappexplorer.com/app/fullcircle.bhdouglass>

HUGE thanks to Brian for this.





REVIEW

Written by Erik

Quod Libet

Version: 3.9.1 (software centre)
4.2.1 (flatpak & PPA)

Website:
<https://github.com/quodlibet/quodlibet>

From the web:
Quod Libet is a music management program. It provides several different ways to view your audio library, as well as support for Internet radio and audio feeds. It has extremely flexible metadata tag editing and searching capabilities.

This music player is in active development and can be downloaded via your software centre. There is something in the back of my mind that tells me the name is Latin for whatever, but I do not want you to say “whatever” – I want you to give this player a go. It integrated very nicely with most of my themes. It picked up all my music via the library function - it may take a while the very first time. (Unless you have more than 1024 files per directory.) The interface is clean but not customizable. As desktop Linux

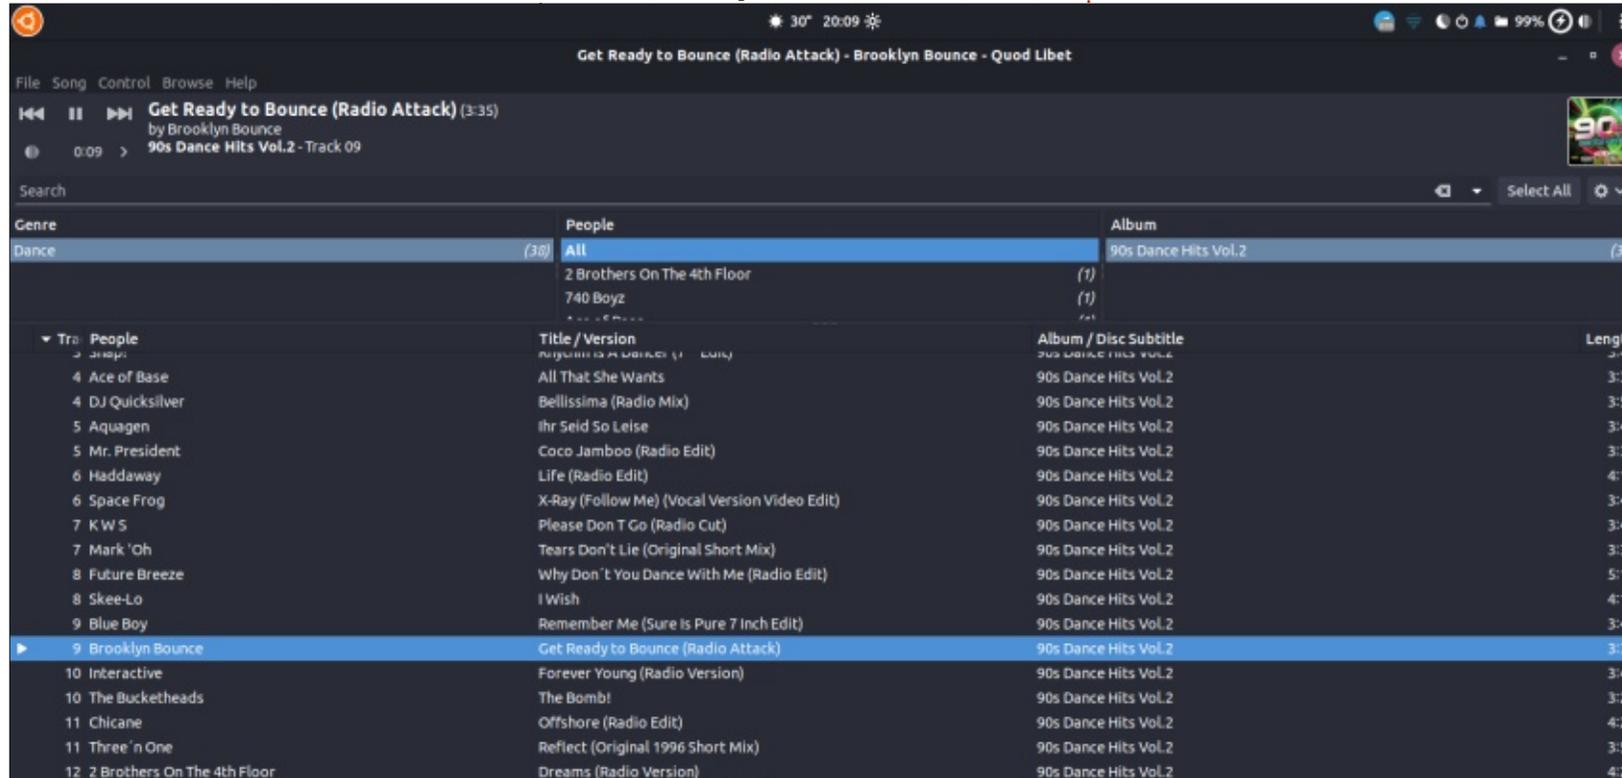
users, we like to customize if nothing else. There is also no spectrum analyser. Some things are a bit arb, for instance, you can enable LastFM scrobbling via the plug-ins, but unless you also enable the LastFM sync, it does not scrobble.

Right, now with the bad out of the way, let’s talk about the good. The settings allow you to set a custom cache size, so there were

no skips when the drive was under load. The player folds up nicely, but does not have a mini player.

Quod Libet supports media keys out the box, which makes using it a pleasure. You can even bookmark your playlists. The filters on right-click of any song are also very handy. Simple key combinations like alt + enter drop you right into the tag editor. Speaking of the tag editor, it may seem minimal, but,

for quick fixes, it does the job and it is quick. Another nice thing about Quod Libet is that it supports a full unicode set, so no funnies with song names. This nifty desktop application supports podcasts and radio streams without hassle. Though it supports control from the tray icon, this did not work on Ubuntu Budgie (no icon). Quod Libet also supports old formats, like tracker files playback. The time tracker is hidden, even in



REVIEW - QUOD LIBET

full screen mode, and you have to click on the time to find it.

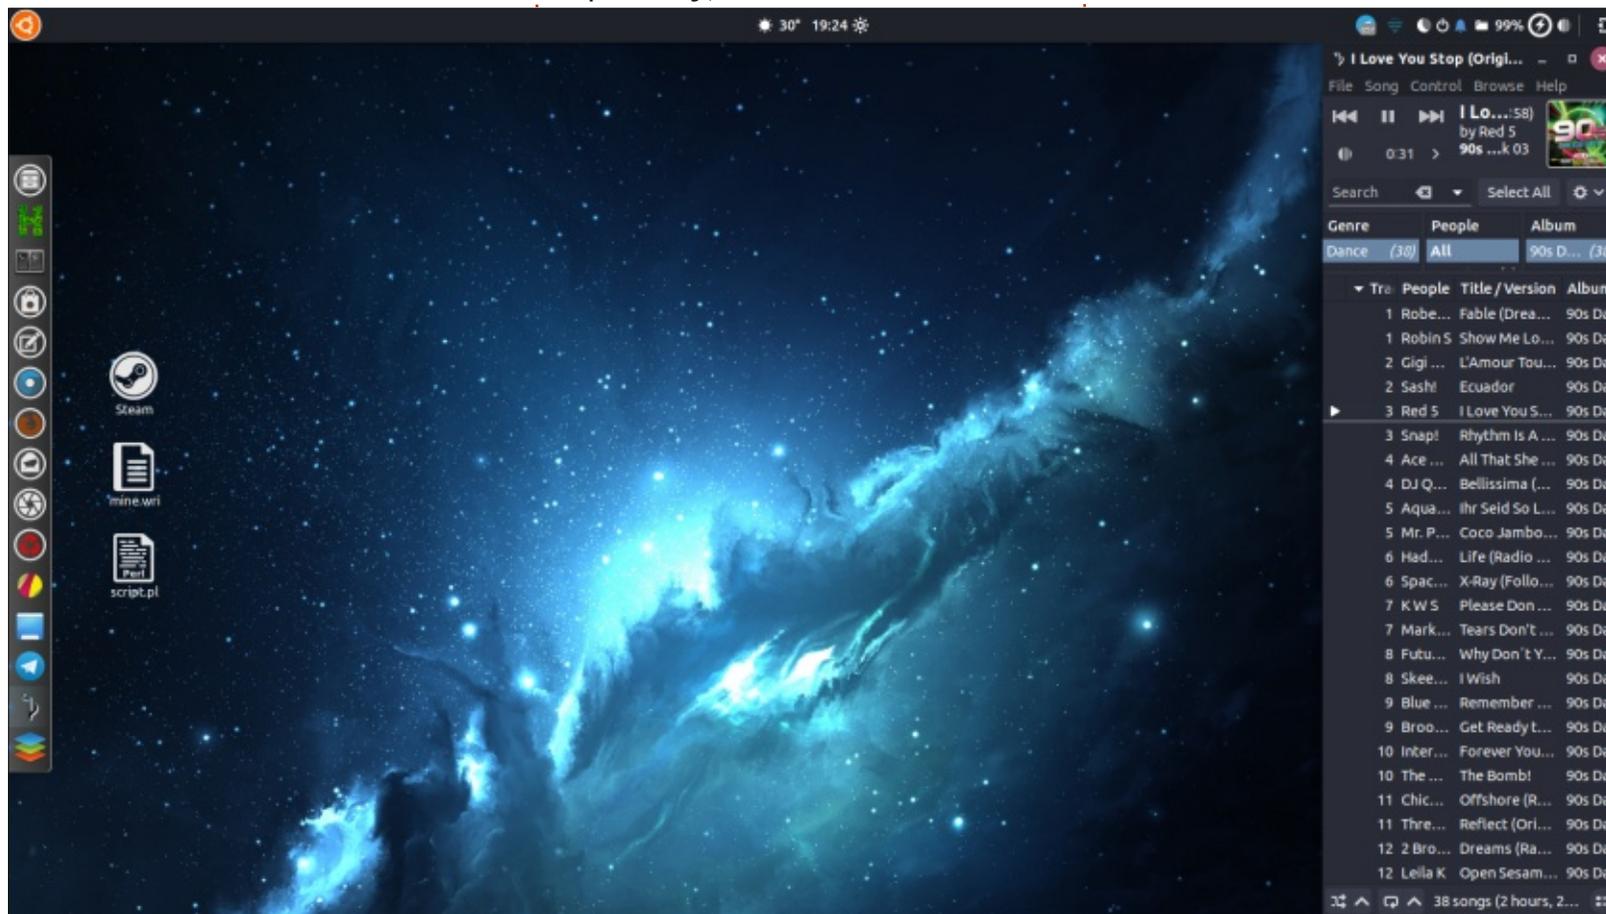
The keyboard shortcuts (Ctrl + ? - but this did not work on Ubuntu Budgie) are very nicely displayed for your perusal. One has to look at just the man page to see that Quod Libet is manageable from the command line. For those who do not want to install it, look here: <https://www.systutorials.com/docs/linux/man/1-quodlibet/>

The player is completely cross platform, so it can be your player in all your operating systems. It can fetch album art for you, or use the embedded album art. Clicking on the album art will produce a bigger picture, depending on whether you right- or left-clicked.

I do encourage you to play with the plug-ins, as there are many and some quite novel ones, like bpm-tap. Lastly, I want to mention that

it has karaoke for all you budding x-factor singers.

Quod Libet feels like a lucky packet; every time you play with it, you discover something new!



Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.



LETTERS

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

Join us on:



goo.gl/FRTML



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



twitter.com/#!/fullcirclemag



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



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

FULL CIRCLE NEEDS You!



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

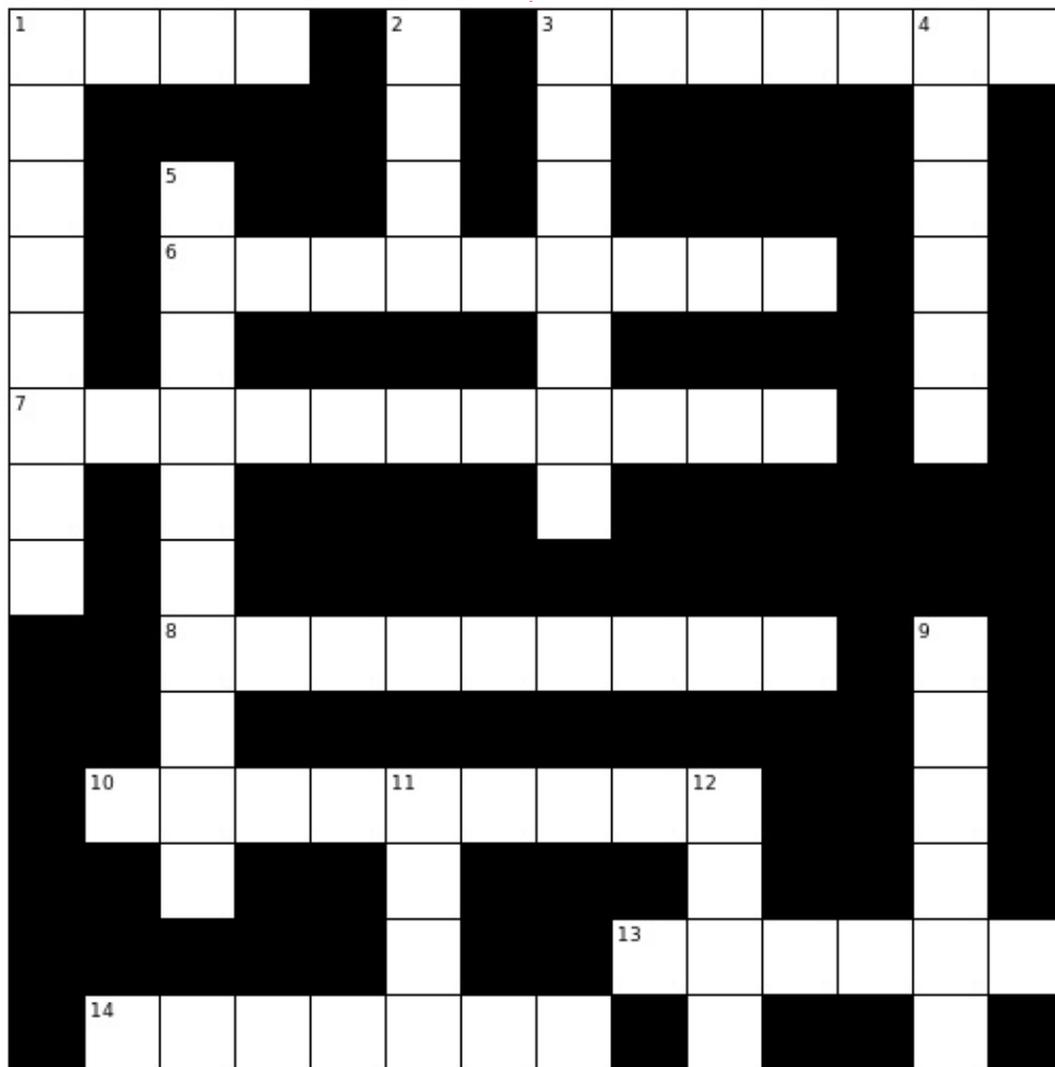
See the article **Writing for Full Circle** in this issue to read our basic guidelines.

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



DistroWatch.com

Put the fun back into computing. *Use Linux, BSD.*



ACROSS

- 1A: A German distro, heavy into greens.
- 3A: From the USA, latest version is code named "Zulu".
- 6A: Originally forked from Mandriva.
- 7A: From New Zealand, based on Crux.
- 8A: Actually BSD, but often found in fossils!
- 10A: I say Patrick Volkerding, you say....
- 13A: Drom France, based on Puppy Linux.
- 14A: Digital Audio player based on Gentoo for headless servers.

DOWN

- 2D: Another German distro - the name is an island in Hawaii.
- 3D: Altispeed just lost all their data.
- 4D: Mark Shuttleworth thinks its cool.
- 5D: The mascot is a tiger-striped Tux!
- 9D: A Pentest distro that now has a "home" edition.
- 11D: Look ma, I am a haxxor...
- 12D: Based on mint, with XFCE, made just for children.

Compiled by Erik

Answers are elsewhere this issue.



Q&A

Compiled by EriktheUnready

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

Welcome back to another edition of Questions and Answers! In this section, we will endeavour to answer your Ubuntu questions. Be sure to add details of the version of your operating system and your hardware. I will try to remove any personally identifiable strings from questions, but it is best not to include things like serial numbers, UUID's or IP addresses.

Back when Pterodactyl abduction was not covered by insurance and my neighbour Barney Rubble got new digs, the company that I worked for was one of the few who did support for Wang equipment (as in Mandarin for "king", not the other meaning). To support high-speed printing, they supplied what was known as 'band printers' (dot-matrix was very slow back then). This consisted of a steel band with letters printed on it, rotating at crazy velocity, and hammers like that of a typewriter striking the letters. A common fault would be for the band to jump off the

spindles or moving out of phase. We would say it slipped. As a joke, we (the FSE's) would enter "knit one, slip one" (and other variations of it) on our reports. One of the companies who had these was a spinning mill and they also made knitted things. Now, imagine the confusion when a new IT manager took over and started looking back at past errors to come up with a plan for the future... Don't be afraid to write notes or make shell scripts for when you reinstall your OS, reinstallation is a fact – like Windows spying on you.

Q: Hello. Something is definitely wrong with my installation, I suspect it is the hard drive of my HP Compaq laptop. It is ancient. How can I check for bad sectors on my drive?

A: I would suggest booting from a USB thumb drive and doing this, but you can do it on a live system too: `sudo badblocks -v /dev/sdaX > badblocks.txt` where X is your drive. You can run: `sudo fdisk -l` to find out what that

number should be.

Q: Hi, I need some help here. I have an old proline monitor that is 1280 x 1024, yet it keeps detecting and switching to 1024 x 768. I prefer the square monitors to the stretched widescreen ones, and would like to keep these. Somebody told me to run `dpkg --reconfigure`, but I broke it before like that. The monitor runs at 60Hz or 75Hz. Please help me with this as it is annoying as hell.

A: Firstly, those old Proline's are not quite HD. The command you are looking for is `xrandr`, XOrg's own monitor configuration utility. First query your monitor with: `xrandr -q`, this will list all the modes. Possibly yours will be: `xrandr --output VGA1 --mode 1280x1024 --rate 60.00` if it is connected to your VGA port. The nice thing is you can put `xrandr` statements (as above) inside your `.xinitrc` and it will work. Reference: <https://xorg-team.pages.debian.net/xorg/howto/use-xrandr.html>

Q: My Ubuntu seems messed up. If Ubuntu wasn't so damn friendly, I'd go back to Debian, but sometimes the quirks drive me mad. Sometimes it connects to wireless, sometimes not. Then there is that 'special' case where the taskbar icon runs up and down through the bars, almost like it is flashing or stuck on initializing. I don't know if it is connected or connecting or what state it is in. I usually just reboot my laptop. I didn't have this on Ubuntu 16.04. Is there a simple way to check?

A: Yes. You can use `nmcli`. It is quite a versatile command. `nmcli --help` will bring up the basics. What you will not see in that list is the 'd' option (read the manpage for all the switches). Use the command: `nmcli d` (no hyphen) This will bring up your devices and type and the state. (connected or not)



Q&A

Q: My Ubuntu Budgie 18.4 has a line in my task bar. How do I get rid of it?

A: This is very little info to go on. A. You could have a LCD malfunction. B. It could be a theme, like Pocillo-light-slim. C. It can even be your wallpaper! Without more information, I just cannot tell.

Reply: *Hahaha yes, you were right, do I feel sheepish!*

(Not sure what we were right about – he never got back to us.)

Q: Before I ask my question below, can I just say that I enjoy you magazine and podcast. I have been reading and listening ever since leaving Windows in 2012 and moving to Ubuntu. It's a great effort. Now to my question: I have an old CanoScan LiDE 60 USB powered flatbed scanner that I have faithfully carted around Australia for years. It just worked with Ubuntu 14.04, 16.04 & 18.04, and it has served us well doing entire photo albums and all our office work. We currently live in a camper trailer so we keep any

essential paperwork digitally - we have no room for filing cabinets and the flatbed without a printer means it is easy to stow. Unfortunately, this old scanner is now playing up and ready for replacement. I tried the CanoScan 300 & 400, but could not get them to work with Ubuntu 18.04. I've checked the hardware compatibility site but it seems out-of-date for flatbeds. Can you help with a flatbed scanner that will work out of the box with Ubuntu 18.04?

A: Thank you for the kind words. The CanoScan 300 is not Linux compatible. That said, a user on the Late night Linux Telegram channel got it to work using this link: <https://askubuntu.com/questions/1107632/install-canon-lide-300-400-on-ubuntu-18-04>

After some investigation (pardon the long wait), it seems that in the new scanners, only the Epson flatbed scanners are SANE compliant out of the box. Tested with SANE drivers and SimpleScan, it works.

Q: Hey guys, I was trying to install Deflemask, but can't get it to work. My computer uses Ubuntu 18.04.2 LTS. Gnome 3.28.2. Their website forum is forbidden 403 for some reason. <http://deflemask.com>. My problem is described in the manual tho. *"This is usually due to a missing library, remember to get OpenGL (sudo apt-get install freeglut3), if you are on a 64-bit distro, you have to Download 32-bit libraries (sudo apt-get install ia32-libs)".* It says it is missing files. I tried to install what they said, but I must be too dumb. I just don't get it. ia32libs says no installation candidate. You gotta help me, I never felt this stupid in windows.

A: I have never even heard of Deflemask, let alone use it. I tried the installation and got the dll shuffle too. Made me feel like Using Fedora 2 back in the day and breaking it. Do not feel 'dumb', rejoice! It means you are learning more and realising there is more to know. Instead of ia32libs, use (t)rusty old dpkg: `sudo dpkg --add-architecture i386` and update your repo's. (That is two dashes, the word processor sometimes turns it into one long dash.)

Then install:

```
sudo apt install libSDL-  
image1.2:i386
```

```
sudo apt install libgtk2.0-  
0:i386
```

```
sudo apt install libglu1-  
mesa:i386
```

(note the ":i386" at the end). It should run fine after that.

P.S. I really like this software, maybe we can do an article on it sometime.

Q: Is this a good laptop for Ubuntu and CAD? <[link to deal removed](#)> (Cherry Trail SoC 2GB laptop, 32GB SSD.)

A: No. CAD requires a graphics card, unless you are willing to suffer in silence; failing that, a very powerful CPU. (Cherry trail is an Atom CPU SoC, like a Raspberry Pi). CAD is memory hungry, 2GB will barely cover it. Try for an HD screen, 1366 is usable, but will strain your eyes on a 14-inch display. As for Ubuntu, not sure, it may need Raspbian...

Q : Hey guys! Can you please help me, I have searched the internet empty, but can't find a solution. Run Xubuntu 18.04 with the latest Firefox, and I have tried all the about:config tricks, but I can't seem to fix the screen tearing when scrolling. I have installed nvidia proprietary drivers, added compton, etc. The screen tears about 2/3 the way down. I want to smash my screen, that's how frustrated I am.

A : Dear frustrated, open Firefox preferences and under general go to performance. Now untick "smooth scrolling" - if this does not fix your problem, get back to me. (yes it sounds counter intuitive, but it works).

Q : My machine is an i5/8GB with Ubuntu 18.04. I want to theme Geany. I have become used to the dark interface of Sublime text. My problem is this: Geany themes only cover the code pane.

A : This one also threw me a curve ball. However, there is a youtube video on this: https://www.youtube.com/watch?v=ln9tvv_qe9l

Q : My question is about Focus Writer. I have installed it from synaptic as I did not want a snap package. I can't edit a theme and I am feeling very stupid. Could you redpill me?

A : The short answer is that you can not edit default themes. You can, however, duplicate a theme and edit that. It will move to "custom", or you can even create your own that will be editable after saving in the 'custom' tab.

Q : hey, my problem is my ubuntu is hanging up. sometimes with graphic corruption. It has never happened before. This is my specs. <output of neofetch removed.> It is not overheating. What could be wrong?

A : If you take a gander at that output, you will see: Disk (/): 3G / 410G (99%) - fourth from the bottom. Boot into 'advanced' or 'recovery' mode and use the make space option. If it fails to make space, boot from a live medium and delete some stuff. It may be wise to install Stacer to monitor

your disk usage and clean up some space.

Q : I am so addicted to 2048-qt game in Ubuntu. How can I hide it so my supervisor can't see what I am doing?

A : The short answer is you can not. However, there is a command-line version (with graphics) that will show up as terminal open and not a game, but you never heard it from me.

Q : Help! I can't log into my Skype any more. It just says that I have been logged out for security reasons. The more I log in, the more it logs me out immediately. Have I been hacked?

A : No, it's Microshaft at its finest. You need to uninstall, reboot, reinstall the latest version, and go through the 'forgot password' routine and reset your password. Only then will you be able to log in. Failing that, you may have to make a new profile.

Q : How can I create a shortcut to point to the default

application for reading mail, without specifying what it is? I know about mailto: but this does not work well. I am creating a vanilla desktop with applications on it. I want to roll it out across multiple machines, but the machines may / may not be the same, but it must be standardized. If that makes any sense?

A : That sounds like an interesting project, will you keep us informed about your progress? You can try: `exo-open --launch MailReader %u` in the command parameters. It is unashamedly stolen from: "`exo-open --launch WebBrowser %u`" (That is two dashes, word processors do funny things.)



Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.





Look at you! Back for more! I am glad you decided to look into getting Linux certified. I have a motto. Linux is not difficult – and I will show you why. So, without further ado, let’s dive into Capacity Planning (Topic 200). The very first part is system availability.

Now you may ask, “what is system availability?” Well, is the system available? Can you log into it? Can you actually use it?

Well, the command targeted first is ‘uptime’. If you do not know what it is, quickly type it in your terminal. You will see some cryptic output :

```
08:11:44 up 11:41, 1 user,
load average: 0.08, 0.13,
0.16
```

The LPI wants you to understand what you are looking at. (Feel free to look in the man page quickly).

Linux is not difficult – let us break it down:

The first part, up to the comma:

	total	used	free	shared	buff/cache	available
Mem:	8093332	755216	5594628	235620	1743488	6827692
Swap:	2097148	0	2097148			

“08:11:44 up 11:41” - indicates the current time and that the system has been up for 11 hours and 41 minutes. (To verify this, enter the “date” command).

The second part, up to the next comma: “1 user” - the number of users logged on to the system. (To verify this enter the “w” or “who” command).

The third part, “load average: 0.08, 0.13, 0.16” - does what it says on the box: The catch is the times, they are one, five and fifteen minutes. (Why fifteen minutes?).

How simple was that?

Let us look at memory. (Measuring and troubleshooting resource usage).

If you ever used DOS, you will know the “free” command. In Linux, it is basically the same, just

more verbose. Please run it in your terminal now. Your output should look something like that shown above.

If you have not done so yet, look at free in the man page.

Linux is not difficult – let us break it down:

The two rows, “mem” and “swap” refer to just that “mem” being RAM and “swap” being swap space. The output may seem strange, but if you had a look at the man page, you would have seen the -m option, which will change the output to megabytes and the -h option to make it “human readable”. If you have not, try them now.

The columns tell us exactly what they are:

On some older systems you

may see buff/cache split into their own columns. Since Linux is very efficient at managing memory, it will cache things to make your system faster, or buffer for future reads.

You may see in the “swap” row that the “used” column is 0. This is because swap does not get used until your memory starts getting full. (Yours may have a value). Swap is like the windows page file.

Now on your desktop you have a task manager or system manager that will display your memory for you, but when you SSH into a server, this is not available. So, you might want to get an update at regular intervals. How would you do that? With the -s option! For example, we want an update every three seconds. We would use: “free -s 3” and, as with most Linux commands, you can stack options or switches, so you could use “free



-h -s 3" and Linux would be happy with that. Try it now. Open a web page and browse in the background and see how free changes in your terminal.

How simple was that?

Now we know there are many ways to skin a cat, in Linux. Free is not the only command available to you. The book wants us to know another, "vmstat". As always, open it in your man page and have a look. Please execute the vmstat command in your terminal. Mmm... What do all those two-letter acronyms mean?

Your output may look like that shown top right.

Linux is not difficult – let us break it down:

The very first line sorts your columns into groups, and it is immediately evident that vmstat is a lot more talkative than free. The first part is self-explanatory, but, under swap, you will see "si" and "so". Easy-peasy – swap in and swap out. Under io you will see "bi" and "bo". Easy-peasy – blocks in and blocks out. (Remember io is to block devices, and you won't forget what they stand for). Under

```
procs -----memory----- --swap-- ----io---- --system-- -----cpu-----
r  b   swpd free  buff  cache   si  so   bi  bo   in  cs  us  sy  id  wa  st
0  0       0 5549944 150920 1626744   0  0    7  5   230  39  1  1  98  0  0
```

system, you will see "in" and "cs". Think about how your system operates, these will be interrupts per second and context switches per second. It is not always the easiest to remember, but, once you understand it, your system operation, interrupting and context switching will spring to mind.

Like free, you can get updates at regular intervals. However, you do not use -s, you can simply type "vmstat 3". Did you look at the man page yet? The first option is "-a" - display active and inactive memory. It will replace "buff" and "cache" and give you even more insight into your memory usage. Should you want a vertical output with even more stats, you can use the -s option. Type vmstat -s into your terminal and peruse the output.

"Old skool" admins won't be trusting some new-fangled dashboard (get off my lawn!!), you can actually attach a timestamp to your readings and use awk to cut it up the way you want it. "How does

one attach a timestamp?", you may ask. Well, that is the -t option. Just like free, Linux will be quite happy with "vmstat 3 -t". Try it in your terminal now. How simple was that?

If you have learned something new, great! If you already knew all this, good, you are prepped for your LPI exam.

Speaking of exams, let us find a question on the internet. I Googled - "201-400-demo.pdf". Here is an example question, don't learn these parrot fashion, as they are probably not real exam questions, only fashioned in the style of exam questions!!

In the following output, the load averages represent the system load averages for what

time frames?

```
12:10:05 up 18 days, 19:00, 2
users, load average: 0.47,
24.71, 35.31
```

- A. 1, 5 and 15 minutes
- B. 1, 15 and 30 minutes
- C. 1, 15, and 30 seconds
- D. 15, 30 and 60 minutes

Hah! We know that one. No explanation needed! (Why fifteen minutes?).

In this example output (below), which descriptions match the purpose of the free, buff and cache columns? (Choose THREE correct answers.)

- A. Used swap space
- B. RAM available for filesystem buffers
- C. Available free RAM
- D. RAM used for buffers
- E. RAM used for filesystem cache

```
# vmstat 1 100
procs ---memory--- --swap-- ----io---- --system-- ----cpu----
r  b   swpd free  buff  cache   si  so   bi  bo   in  cs  us  sy  id  wa
0  0  0 282120 134108 5797012 0 0 0 2 0 0 0 0 100 0
0  0  0 282120 134108 5797012 0 0 0 0 1007 359 0 0 100 0
0  0  0 282120 134108 5797012 0 0 0 0 1117 577 0 0 100 0
0  0  0 282120 134108 5797012 0 0 0 0 1007 366 0 0 100 0
```

If you answered A or B, please go over what we covered.

buff = RAM buffers

cache = RAM cache

free = free RAM

* as free, buff and cache is under the memory column.

In the following output, what is the 5 minute load average for the system?

```
# uptime  
12:10:05 up 18 days, 19:00, 2  
users, load average: 0.47,  
24.71, 35.31
```

- A. 0.47
- B. 24.71
- C. 35.31
- D. There is no 5 minute interval. It is some value between 0.47 and 24.71.
- E. There is no 5 minute interval. It is some value between 24.71 and 35.31.

I'll leave this one to you!

RESOURCES:

<https://www.lpi.org/how-to-get-certified/free-training-materials>

Feel free to contact us with any queries and, as always, you can find me on the FCM telegram group.



Erik has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.

Crossword answers:





It's been over a year since our last Steam Hardware Survey and with all of the changes that have taken place in the world of Steam for Linux, I figured now would be a good time to review our place in the overall gaming world. The last time we had a review of the Steam Hardware Survey here at Full Circle Magazine was back in 2017, so we're definitely due for a refresher. To review our last Steam Hardware Survey issue please refer to FCM#120 from April 2017. I was curious as to whether there would be a noticeable change in the number of Linux users since last time we conducted a Steam Hardware Survey two years ago, especially considering the SteamPlay Proton update from August 2018. For those who don't know about SteamPlay, it's basically Steam providing an emulator which is similar to Wine or Play-On-Linux. An emulator is basically a tool that takes a Windows software and tricks it into believing that it's running on a Windows PC even though it may indeed be a Linux computer. By far, the most popular Windows to

Linux emulator is Wine (Wine is a sort of acronym which stands for: Wine Is Not an Emulator). In the case of SteamPlay, the emulator is basically run by what Valve (the company behind Steam) is calling Proton. To find out more about Proton and SteamPlay refer to FCM#137 from 2018.

Recently, as I was about to play a game for the first time in a few months, I was pleasantly surprised to be welcomed by the Steam message that asks me whether I'm interested in running the Hardware Survey. Of course, I accepted and took screenshots of the whole process for your viewing pleasure. I strongly advise all of you Linux gamers to run the Survey as it's a great way to measure how vibrant Linux gaming is when it's compared to Windows or Mac OS X.

After you accept to run the Survey, there's not much that you need to do except for going through a couple of more windows and clicking on Next until you get to the last screen that says you're

done. So, in other words, the survey is not something that you are actively involved in but it's a rather automatic process which gathers information from your computer and sends the results to Valve. A (sort-of) bug that I encountered immediately after my Survey ended was that when I clicked on the View Results web link, I was taken to the Steam Store page but not the page which actually has the Survey results. So, I had to take the extra step of searching for that page – which can be found here: <https://store.steampowered.com/hsurvey/Steam-Hardware-Software-Survey-Welcome-to-Steam>

Looking at the results, and comparing them to our 2017 results, we can see that the percentage of Linux users has increased slightly from 0.77% in 2017 to 0.81% in May 2019 (and yes, you read that correctly, it's almost 1% but not yet). Although the percentage indicates a slight increase in Linux users, it has been reported by other people that the

actual number of users has gone up a lot because the overall number of gamers (including Linux gamers) has also gone up substantially.

In about a year, I'm hoping to once again be randomly selected to partake in the Steam Hardware Survey. By then, a couple of things should have taken place. First, at this time next year, I should either already be using Ubuntu 20.04 or at worst I'll be planning on updating by June 2020 at the latest. Second, and much more important, by this time next year the number of Linux gamers on Steam could potentially be over the 1% mark. The reason for this is that, in the upcoming months, it is rumored that the next wave of Chromebooks released will be able to run almost all Linux apps (which I would assume would also include Steam). With Chromebooks being more popular than any other Linux laptop, it would not be the least bit surprising if finally the number of Linux gamers playing video-games on Steam finally surpasses the 1% mark. For now, let us

UBUNTU GAMES - STEAM SURVEY 2019

rejoice knowing that we've at least added another 0.05% of total users.

Below are the most relevant results gathered from my hardware after the survey was run:

Computer Information:

Manufacturer: Unknown
Model: Unknown
Form Factor: Desktop
No Touch Input Detected

Processor Information:

CPU Vendor: AuthenticAMD
CPU Brand: AMD FX(tm)-6100 Six-Core Processor
CPU Family: 0x15
CPU Model: 0x1
CPU Stepping: 0x2
CPU Type: 0x0
Speed: 3300 Mhz
6 logical processors
6 physical processors
HyperThreading:
Unsupported

Operating System Version:

Ubuntu 18.04.2 LTS (64 bit)
Kernel Name: Linux
Kernel Version: 4.15.0-48-generic
X Server Vendor: The X.Org Foundation
X Server Release: 11906000
X Window Manager: GNOME Shell
Steam Runtime Version: jenkins-steam-runtime-beta-

release_0.20190320.2

Video Card:

Driver: NVIDIA Corporation GeForce GTX 960/PCIe/SSE2
Driver Version: 4.6.0
NVIDIA 396.54
OpenGL Version: 4.6
Desktop Color Depth: 24 bits per pixel
Monitor Refresh Rate: 60 Hz

VendorID: 0x10de
DeviceID: 0x1401
Revision Not Detected
Number of Monitors: 1
Number of Logical Video Cards: 1
Primary Display Resolution: 1920 x 1080
Desktop Resolution: 1920

x 1080

Primary Display Size: 23.54" x 13.23" (26.97" diag)
59.8cm x 33.6cm (68.5cm diag)
Primary Bus: PCI Express 16x
Primary VRAM: 2048 MB
Supported MSAA Modes: 2x 4x 8x 16x

Sound card:

Audio device: Realtek ALC892

Memory:

RAM: 15,986 MB

Miscellaneous:

UI Language: English
LANG: en_US.UTF-8

Total Hard Disk Space Available: 351,159 MB
Largest Free Hard Disk Block: 64,551 MB
VR Headset: None detected

Recent Failure Reports:



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

Windows, Mac and Linux Use the dropdown to filter by platform or view combined stats.

APRIL 2019 (CLICK LINE ITEM TO SEE MORE DETAIL)

ITEM	MOST POPULAR	PERCENTAGE	CHANGE
OS Version			
Windows		96.12%	+0.21%
Windows 10 64 bit		67.18%	+0.26%
Windows 7 64 bit		24.27%	+0.48%
Windows 8.1 64 bit		2.98%	-0.28%
Windows 7		1.15%	-0.16%
Windows 8 64 bit		0.24%	-0.04%
Windows 10		0.20%	-0.03%
OSX		3.06%	-0.21%
MacOS 10.14.3 64 bit		0.74%	-0.53%
MacOS 10.14.4 64 bit		0.71%	+0.71%
MacOS 10.13.6 64 bit		0.59%	-0.06%
MacOS 10.12.6 64 bit		0.24%	-0.02%
MacOS 10.14.2 64 bit		0.17%	-0.15%



PATRONS

MONTHLY PATRONS

Alex Crabtree
 Alex Popescu
 Bill Berninghausen
 Brian Bogdan
 CBinMV
 Darren
 Dennis Mack
 Devin McPherson
 Doug Bruce
 Elizabeth K. Joseph
 Eric Meddleton
 George Smith
 Henry D Mills
 Hugo Sutherland
 Jack
 Joao Cantinho Lopes
 John Andrews
 John Malon
 John Prigge
 JT
 Kevin O'Brien
 Lee Allen
 Leo Paesen
 Linda P
 Mark Shuttleworth
 Norman Phillips
 Oscar Rivera
 Paul Anderson
 Paul Readovin
 Rob Fitzgerald

Roy Milner
 Scott Mack
 Sony Varghese
 Tom Bell
 Tony
 Vincent Jobard
 Volker Bradley
 William von Hagen
 Taylor Conroy

SINGLE DONATIONS

2019:
 Floyd Smith
 Jack Hamm
 aram v nathan
 Joachim Haupt
 Hari Zafiriadis
 Glenn Heaton
 Adam Gwizdz
 George Parker
 Linda Prinsen
 Frank Dinger
 Graig Pearen
 Stefano Giancarli
 Raymond Meyer
 wil van schaik

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://bit.ly/fcmwriting>

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#146

Deadline:
Sunday 09th June 2019.
Release:
Friday 28th June 2019.



Full Circle Team



Editor - Ronnie Tucker
ronnie@fullcirclemagazine.org

Webmaster - Lucas Westermann
admin@fullcirclemagazine.org

Editing & Proofreading

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

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

Getting Full Circle Magazine:

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/>



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.

FCM PATREON: <https://www.patreon.com/fullcirclemagazine>