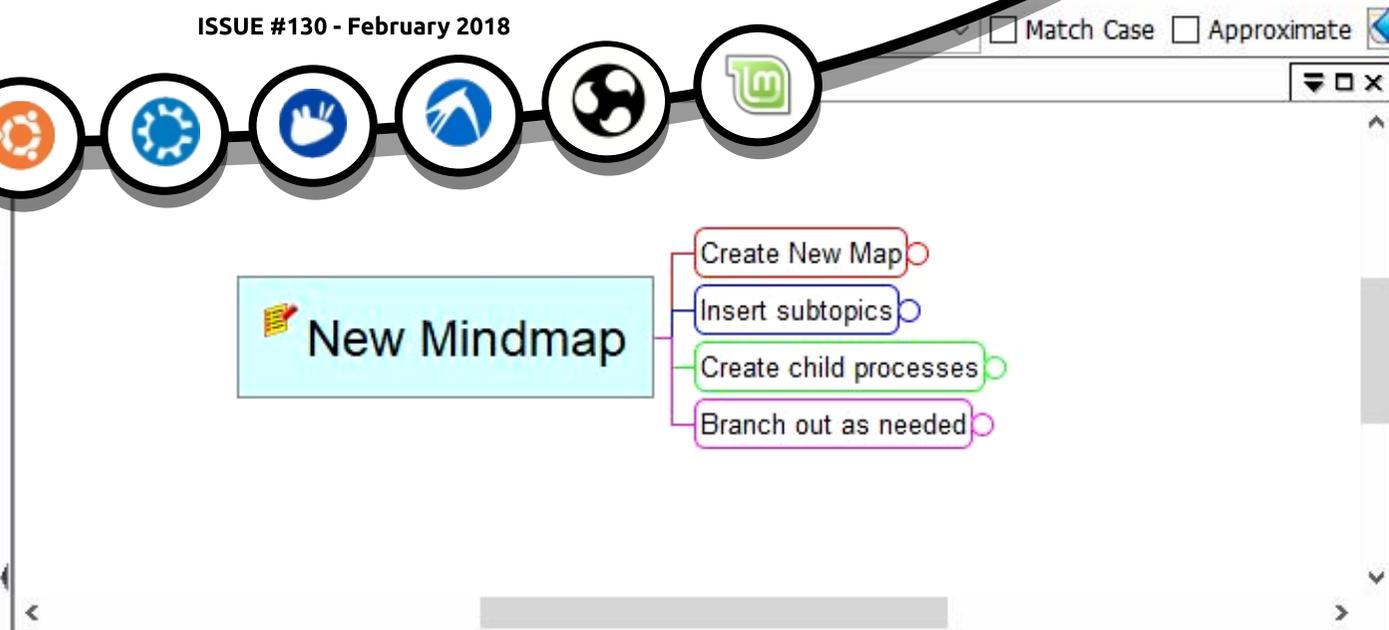




Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

ISSUE #130 - February 2018



Format Calendar and attributes Presentations

Node style

Apply level styles

Automatic edge color

Change Style

Color of whole core

Change Text

Change Background

Core text

Change Format

Change Node numbering

Shape

Node shape

Horizontal margin

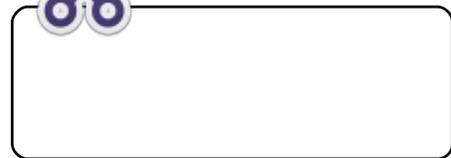


FREEPLANE

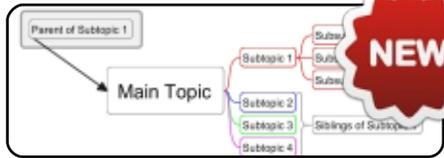
NEW SERIES ON BUILDING MINDMAPS



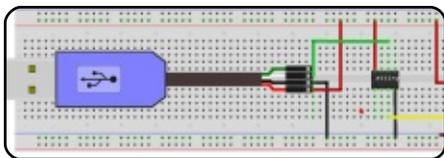
HowTo



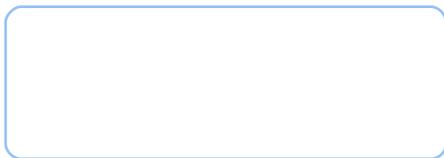
Sphinx p.20



Freeplane p.22



Great Cow Basic p.25



p.XX



Inkscape p.31



Graphics



Full Circle

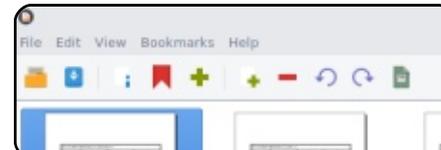
THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY



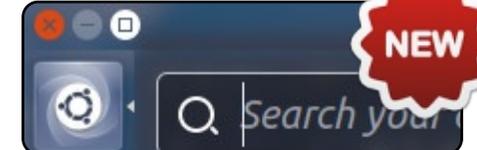
Linux News p.04

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

Command & Conquer p.17



Researching With Linux p.35



Everyday Ubuntu p.36



Linux Labs p.39



My Story p.45



My Opinion p.XX



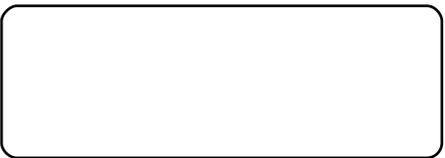
Review p.48



Letters p.XX



KODI Room p.XX



Q&A p.52



Ubuntu Games p.54



My Desktop p.XX



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 THE LATEST ISSUE OF FULL CIRCLE.

This month we have a new series. Elmer Perry (LibreOffice) returns to the fold to show you how to use Freeplane. Freeplane is a mindmapping tool. If mindmapping isn't your thing, don't worry we still have Great Cow Basic, and Inkscape. Regarding Python, I have heard back from Greg who's recovering from an illness. So feel free to send him emails of encouragement.

Elsewhere in this issue we have a review from Lucas on **Able2Extract 12**. I reviewed it way back when it was at version 10 and was very impressed with its conversions. Version 11 seemed to have no Linux edition (which is a bit odd) and while version 12 did need a tweak from Lucas to get it running, it does seem to be going from strength to strength.

With Linux Labs, Charles is looking at desktop memory usage. Just where does all that memory go? While Chris Binnie makes a return looking at nice little BASH time savers.

I don't know about where you are, but here in the UK we're bracing for the impact of the GDPR (General Data Protection Regulations - an upgrade to the UK's aging Data Protection laws) which comes into full force in May this year. It does mean that I've had to look into full-drive encryption for some Windows work machines. So, in true FCM fashion, I'm doing a write up on how I did full-drive encryption on a Windows 7 machine. Fear not, it's the exact same procedure, and software (Veracrypt) for Linux. I'm hoping to be all topical and squeeze it in next month.

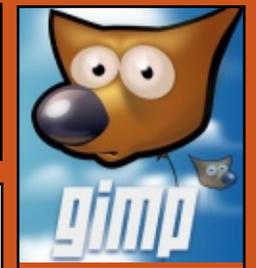
All the best, and keep in touch!

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>



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

Weekly News:



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



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



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



FREE LINUX TOOL MONITORS SYSTEMS FOR MELTDOWN ATTACKS

SentinelOne this week released Blacksmith, a free Linux tool that can detect Meltdown vulnerability exploitation attempts, so system administrators can stop attacks before they take root.

The company has been working on a similar tool to detect Spectre vulnerability attacks. Though free, Blacksmith is not open source. SentinelOne decided to expedite its development in-house to save time, said Raj Rajamani, vice president of project management.

The company has made the tool available to everyone for free in the hope of securing Linux systems while reliable patches are developed, he told LinuxInsider.

The Meltdown vulnerability affects Intel chips and Linux-based systems. A similar design flaw, Spectre, affects AMD and ARM

chips. No comprehensive solutions currently are available for either flaw.

Meltdown is a design flaw in all Intel chips produced in the last decade. It creates a vulnerability that puts Linux, Windows and macOS-powered computers at risk. The flaw is in the kernel that controls the chip performance that allows commonly used programs to access the contents and layout of a computer's protected kernel memory areas.

Two key factors influenced SentinelOne to prioritize the Linux version of the tool. Linux is very susceptible to such attacks, with no comprehensive solution available. Also, Linux is the preferred OS of the world's top supercomputers. That makes Linux a high-value target for attackers.

Those reasons made it clear that it was critical to help secure Linux environments as quickly and effectively as possible, said Migo Kedem, SentinelOne's director of product management.

Source:
<https://www.linuxinsider.com/story/Free-Linux-Tool-Monitors-Systems-for-Meltdown-Attacks-85094.html>

LINUX 4.15: GOOD NEWS AND BAD NEWS ABOUT MELTDOWN AND SPECTRE

Linus Torvalds, Linux's primary creator, had good and bad news about the chip security problems Meltdown and Spectre. The good news is the lead up to the Linux 4.15 was "quiet and small, and no last-minute panics, just small fixes for various issues". The bad news? "It's not like we're 'done' with Spectre/Meltdown."

On the Linux Kernel Mailing List (LKML), Torvalds explained, "The bulk of the 4.15 work is all the regular plodding 'boring' stuff. And I mean that in the best possible way. It may not be glamorous and get the headlines, but it's the bread and butter of kernel

development, and is in many ways the really important stuff."

Torvalds continued, "While Spectre/Meltdown has obviously been the big news this release cycle, it's worth noting that we obviously had all the *normal* updates going on too, and the work everywhere else didn't just magically stop, even if some developers have been distracted by CPU issues. In the *big* picture, 4.15 looks perfectly normal, with two thirds of the full 4.15 patch being about drivers ... not by CPU bug mitigation."

But, trying to mitigate the Meltdown and Spectre problems still ate up a lot of time and the problems are still far from done. First and foremost, like all operating system developers, Linux is waiting on Intel's hardware designers to complete their firmware and microcode patches.

Source:
<http://www.zdnet.com/article/linux-4-15-good-news-and-bad-news-about-meltdown-and-spectre/>

UBUNTU REVERTING TO XORG IN BIONIC BEAVER

Ten years' worth of effort to replace the Xorg graphics framework has been given a "must try harder" mark by Ubuntu, which says its next release will not use Wayland by default.

Ubuntu's desktop engineering manager Will Cooke made the announcement last Friday, saying the decision applies to the Bionic Beaver release due in April.

He listed three shortcomings in Wayland: screen sharing works (for example in Skype, Hangouts and WebRTC) better in Xorg, remote desktop control ditto, and "recoverability from Shell crashes is less dramatic".

Screen sharing is the big killer: in Wayland, both the screen sharing protocol and the GNOME implementation (PipeWire) are still under development. PipeWire was formally unveiled as a project in September 2017.

Once development is completed, Cooke wrote, there's

still going to be a lag while third-party developers integrate the screen sharing.

As for crash recovery: under Xorg, the shell can be recovered independently of the display server and running applications. Wayland hasn't reached that point yet, so if it falls over, so do any applications the user has loaded.

Source:
https://www.theregister.co.uk/2018/01/30/ubuntu_reverting_to_xorg_in_bionic_beaaver/

RED HAT BUYS COREOS FOR \$250M TO EXPAND ITS KUBERNETES AND CONTAINERS LEADERSHIP

Red Hat, Inc., the world's leading provider of open source solutions, announced that it would acquire CoreOS, Inc., a company known for providing the Container Linux operating system (formerly CoreOS Linux), Tectonic for Kubernetes, and Quay Enterprise container registry, for the price of \$250 million USD.

CoreOS joining Red Hat means automated operations are coming to all. In other words, both companies will work together to expand Kubernetes, the open-source system for automating scaling, deployment, and management of containerized applications in business environments, as well as to innovate in containers and distributed systems.

Red Hat's acquisition of CoreOS will also accelerate the adoption and development of the best hybrid cloud platform available to date for modern application workloads, whose demand continues to grow every day. Red Hat is already a leader in enabling enterprises around the globe to embrace container-based apps with its Red Hat OpenShift enterprise-ready and comprehensive Kubernetes platform.

Red Hat expects to close the transaction in January 2018, which won't have a material impact to Red Hat's guidance for the company's Q4 or the fiscal year ending February 28, 2018. On the other hand, CoreOS promises to continue to honor all existing

customer engagements, and send a direct email to all of them about these new changes.

Source:
<http://news.softpedia.com/news/red-hat-buys-coreos-for-250m-to-expand-its-kubernetes-and-containers-leadership-519641.shtml>

LIBREOFFICE 6.0 RELEASED WITH NEW FEATURES

Whenever we talk about Microsoft Office alternatives, LibreOffice manages to come up quite a lot in the conversation. The Document Foundation has been working continuously to improve the overall experience of using this free and open source office suite. The latest LibreOffice 6.0 release is no exception.

As expected, LibreOffice cross-platform release is available for Windows, macOS, and Linux. You can also use its cloud version as well from any computer or web browser. There are many significant changes made to the core engine as well as the Writer,

Calc, Impress/Draw modules.

Notebookbar, which is still in experimental phase, continues to get the attention it deserves. This equivalent to Office ribbon has been tweaked for easier navigation. For Writer, Impress, and Calc, there's Grouped Bar Full variant.

On the security front, OpenPGP keys can now be used to sign ODF documents. There's experimental support for OpenPGP-based encryption as well. To use these features, one needs a third-party PGP tool like Gpg4win.

In cloud version, work has been done to deliver an experience closer to the desktop version. For example, a Save As feature has been added and amount of rows managed by Calc has been increased. Also, a Find and Replace dialog and spell checking has been added to Calc, Writer, and Impress.

Source:

<https://fossbytes.com/libreoffice-6-0-released-features-download/>

“LIGHTWEIGHT” LINUX LITE 3.8 & “MATURE” OPENSUSE LEAP 15 BETA RELEASED

Marking the final release in 3.x series, the Linux Lite developers have released Linux Lite 3.8 operating system. One of the best lightweight Linux distros around, this distro comes with a number of changes since the 3.6 release.

The major changes for Linux Lite 3.8 include better support for LibreOffice, regional DVD support, Font Viewer/Installer, and Google Search-powered homepage in Firefox.

TLP for Laptops has also been added to Lite Tweaks to provide you the benefits of better power management without dealing with all the technicalities. While it's highly customizable, the default configuration is already optimized for battery life.

Following the rolling development model that was used to make Leap 42.3, Leap 15 will be developed until the final build. At

the moment, the development of Leap 15 has reached the beta phase and snapshots are available. The first beta version is Build 109.3, which will be followed by subsequent snapshots.

The latest beta comes with an all-new look, thanks to KDE's next Plasma 5.12 LTS. Leap 15 is currently powered by Linux 4.14 LTS kernel. It's worth noting that rpm 4.14 is a big update that remains to be integrated in the beta builds.

Source:

<https://fossbytes.com/linux-lite-3-8-download-opensuse-leap-15-beta/>

SKYPE RELEASED AS SNAP ON UBUNTU, LINUX MINT

With snap support, Skype can be easily installed on Linux Mint, Manjaro, Debian, Arch Linux, OpenSuSE, Solus, and Ubuntu, and furthermore, users can be provided with automatic updates whenever new versions are ready.

This helps Microsoft make sure that the latest improvements are

always available for all Linux users, and judging from how fast the snaps ecosystem expanded since the launch in 2016, the software giant made the right call by embracing them with Skype.

Microsoft has obviously used this occasion to praise Skype for Linux and snaps, saying in a press release published together with Canonical that the whole purpose of this decision was to make new features available to everyone as smoothly as possible.

Skype suffered a major overhaul in 2017 and improved substantially on the majority of platforms, as Microsoft migrated to new infrastructure whose purpose appears to be transforming the service from a VoIP solution to a more advanced messaging application available cross-platform.

Given the increasing focus on Linux, Microsoft couldn't leave the open-source behind, and today's announcement doesn't come as a big surprise. Times have definitely changed, since former CEO Steve Ballmer once said Linux was a cancer, Microsoft now reiterates its love for Linux with every occasion.

Source:

<http://news.softpedia.com/news/microsoft-loves-linux-skype-released-as-snap-on-ubuntu-linux-mint-519665.shtml>

CHROME OS IS ALMOST READY TO REPLACE ANDROID ON TABLETS

Google's latest Chrome update, version 64, is now making its way to Chromebooks with a number of much-needed, tablet-focused features. Those include a new split-screen feature for multitasking while in tablet mode, and a screenshot feature borrowed from Android, as noted by 9to5Google. While these are relatively standard upgrades, they do paint a picture of the future of Chrome OS as the rightful replacement for Android tablet software. As it stands now, Chrome OS is very close to taking up the mantle there, and features like this push it ever closer to becoming the hybrid OS for all types of Google-powered screens.

This has been in the works for quite a while as Google's Chrome

and Android teams have coordinated closely to ensure the influx of low-cost, hybrid computing devices like 2-in-1 Chromebooks get the best of both worlds. There is, of course, Android app compatibility on Chrome OS, an initiative that first arrived somewhat half-baked last year and has taken months to fully jell as Google worked out the kinks. For instance, just last month Google added the ability for Android apps on Chromebooks run in the background. In July of last year, Google also began embarking on a touch-focused redesign of Chrome OS to make the software more functional in tablet mode.

Google has been merging certain aspects of Android and Chrome OS for years

We're likely not getting the full-blown merging of the two divisions and their respective platforms anytime soon, or perhaps ever, as Google has played with the idea for years without ever seeming to decide that one platform should supersede the other. In essence, however, Android remains Google's dominant mobile OS, while Chrome OS has been taking on more responsibility as Chromebooks

have steadily become more capable and tablet-like.

Source:

<https://www.theverge.com/2018/2/2/16965520/google-chrome-os-64-update-android-tablets-operating-system-replace>

WHY KDE'S PLASMA MOBILE IS THE IDEAL PLATFORM FOR LINUX FANS AND DEVELOPERS

For the last decade, the mobile market has been under complete lockdown. Unless you were Android or iOS, you didn't stand a chance at making much of a run at success. Canonical failed miserably with the Ubuntu Phone. Blackberry had to resort to their own take on Android. Firefox OS couldn't even get off the ground.

And yet, thanks to the Purism Librem 5, there's another attempt at creating an open source mobile platform on the horizon. Many of us prognosticators and pundits have been ansty to see what's to come for this platform, and finally someone has made some headway,

and that's KDE. The platform is Plasma Mobile. From the looks of it, KDE is on to something.

I'm going to preface this by saying I worked with Plasma Mobile on a VirtualBox VM—which is obviously not the target environment. Couple that with the fact this is a very early release and you have one unstable operating system.

However unstable Plasma Mobile may be at the moment, it offers a glimpse into what is in store for the platform, and it's pretty impressive. Where Canonical attempted and failed miserably to complete reinvent the wheel, Plasma Mobile uses what KDE does best and reworks it into a metaphor that functions very well for the mobile space. Of course, by "functions very well" I can only assume that it's not very useful at the moment.

Source:

<https://www.techrepublic.com/article/why-kdes-plasma-mobile-is-the-ideal-platform-for-linux-fans-and-developers/>

MELTDOWN/SPECTRE STATUS FOR RED HAT AND ORACLE

The Red Hat family of operating systems addressed Meltdown and Spectre in its v3.10 kernel quickly, but relied too much upon Intel's flawed microcode and was forced to revert from a complete solution. Oracle implemented alternate approaches more suited to its v4.1 UEK, but both kernels continue to lack full Spectre coverage while they wait for Intel. Conspicuously absent from either Linux branch is Google's retpoline, which offers far greater and more efficient coverage for all CPUs. Auditing this status is a challenge. This article presents the latest tools for vulnerability assessments.

A frenzy of patch activity has surrounded this year's Meltdown and Spectre CPU vulnerability disclosures. Normally quiet microcode packages for Intel chips have seen four updates in the month of January, one of which was finally to roll back flawed code that triggers random reboots. For enterprise-grade hardware, Intel's quality control has left much to be

desired.

It is likely premature to deploy new monitoring and compliance tools, and a final solution for this set of vulnerabilities will wait until correct microcode is obtained. Still, it may be important for many organizations to evaluate the patch status of servers running Linux kernels packaged by Oracle and/or Red Hat.

Meltdown patches exist now and should be deployed immediately on vulnerable servers. Remediating all Spectre vulnerabilities requires not only the latest kernels, but also a patched GCC to compile the kernel that is capable of implementing "retpolines", or compatible microcode from your CPU vendor.

Source:

<http://www.linuxjournal.com/content/meltdownspectre-status-red-hat-and-oracle>

KALI LINUX 2018.1 RELEASED FOR ETHICAL HACKERS

In 2016, Offensive Security—the developer of Kali Linux ethical hacking distro—decided to switch to a rolling release model. However, from time to time, they keep releasing the Kali snapshots with all the latest patches, fixes, and updates. Following the same tradition, the developers have pushed the first snapshot for 2018.

Aptly named Kali Linux 2018.1, this release contains all the fixes and updates released since last November's Kali 2017.3. The team faced challenges like Spectre and Meltdown exploits, whose patches can be found in Linux 4.15 kernel.

This release is powered by the new Linux 4.14.12 kernel, which brings support for newer hardware and improved performance. This will enable the ethical hackers and penetration testers to use Kali in a more efficient manner to safeguard the security.

2018.1 comes with the support for AMD Secure Memory Encryption. It's a new feature of AMD processors that allows automatic DRAM encryption/decryption. This will, theoretically, prevent the

machines from cold-boot attacks.

Kali 2018.1 also supports the increased memory limits. With kernel 4.14, the new CPUs will be able to support 128PB of virtual memory and 4PB of physical memory.

In case you're running Hyper-V to run Kali VM images, there's some more good news. In the new release, Hyper-V integration services are included, which support Network Monitoring, Replication, and Dynamic Memory.

Source:

<https://fossbytes.com/kali-linux-2018-1-released-download-features/>

MOZILLA FIREFOX 58.0.2 RELEASED WITH WINDOWS AND MACOS FIXES

Firefox 58.0.2 comes with fixes for a signature validation issue during update on macOS, as this has prevented some users from installing the latest version. Apple users are now supposed to be provided with a bug-free experience in this new version,

especially if they're using the built-in update system.

Then, Mozilla says that it has addressed blocklisted graphics drivers related to the off main thread painting crashes. This is a problem that was reported in both the main Firefox 58 release and the later update to version 58.0.1, and the parent company claims everything should work correctly now.

Additionally, Firefox 58.0.2 introduces a fix for a tab crash occurring when printing documents. There were only isolated reports in this regard, and printing for example worked correctly on my work system, but with this new version, the task should run even smoother for everyone.

And last but not least, this new update comes to fix clicking links and scrolling emails on Microsoft Hotmail and Outlook webmail. In the previous versions, users were complaining of various link issues when accessing the inbox of the Microsoft emails, so updating to Firefox 58.0.2 should refine the experience.

Source:

<http://news.softpedia.com/news/mozilla-firefox-58-0-2-released-with-windows-and-macos-fixes-519714.shtml>

OPERA 51 RELEASED: IT'S 38% FASTER THAN FIREFOX QUANTUM 58

Mozilla triggered the browser war with the release of Firefox 57 aka Quantum. But comparing the two proved it wasn't easy to defeat king Chrome.

With the release of Opera 51, the Opera browser also finds itself in the battle. Its developers claim that the new version, based Chromium 64, is around 38% faster (Speedometer 2.0 benchmark) than Firefox 58 when tested on an HP Spectre running Windows 10. The reason behind the quickness is Opera 51's Windows version is compiled with Clang.

The new version adds to the browsers feature set which includes a tool to block cryptojacking on people's computers. A click to scroll feature makes things easier and probably removes the need for a mouse on

laptops. A user can click a tab to automatically scroll to the top from the bottom and click it again to return to the bottom of the page.

Opera 51 makes the reset process easier by including a new 'Reset browser settings' button in Settings > Browser. The option deletes all the data excluding bookmarks, internet history, and saved passwords. Earlier, the process involved creating a separate preferences file and removing the entire browser from the system.

The new stable version also marks the return of the bookmark import/export feature. When watching videos, users can pop-out the player in a separate window. Now, Opera has added a 'back to tab' option in the video pop-out to go back to the tab playing the video.

Source:

<https://fossbytes.com/opera-51-released-its-38-faster-than-firefox-quantum/>

KDE SLIMBOOK II PLASMA-BASED LINUX ULTRABOOK LAPTOP IS CHEAPER, MORE POWERFUL

Meet KDE Slimbook II, the second-generation of the KDE Slimbook laptop that emphasizes the widely-used KDE Plasma open-source desktop environment for GNU/Linux distributions. Build for running the KDE Neon Linux distro, the 1st-generation KDE Slimbook laptop was announced a year ago and offered some attractive features, including a 13.3-inch screen, faster SSDs, and latest Plasma desktop.

A year later, KDE Slimbook II is here to redefine what a modern Linux laptop should be and promises to be up to 15 percent faster than its predecessor by featuring either a 2.5 GHz Intel Core i5 CPU with 3MB cache and Turbo Boost up to 3.1 GHz or a 2.7 GHz Intel Core i7 CPU with 4MB cache and Turbo Boost up to 3.5 GHz. It also supports 4, 8, or 16 GB or DDR4 RAM that's up to 33 percent faster than DDR3.

KDE Slimbook II features the

same crisp Full-HD 13.3-inch matte anti-glare screen as the first generation Slimbook, more powerful Wi-Fi antennas, a bigger touchpad with multi-touch and support for various gestures, 3xUSB 3.0 ports, including a reversible USB-C port, a slicker backlit keyboard, as well as dual disk drive bay that lets users insert a second drive to extend the internal storage of the laptop, supporting M.2 SSDs of 120GB, 250GB, 500GB, and 1TB.

Best of all, KDE Slimbook II is cheaper than its predecessor and the price starts from 699,00 € for the Slimbook Katana II with an Intel i5 processor and 799,00 € for the Slimbook Katana II with an Intel i7 CPU.

Source:

<http://news.softpedia.com/news/kde-slimbook-ii-plasma-based-linux-ultrabook-laptop-is-cheaper-more-powerful-519729.shtml>

CLLOUDLINUX'S KERNELCARE PROMISES TO FIX MELTDOWN & SPECTRE FLAWS WITHOUT REBOOTS

Meltdown and Spectre affect the kernel and other components of a Linux-based operating system, including QEMU, Xen, Nvidia graphics drivers, as well as web browsers like Firefox, Chrome, and Opera. To patch your Linux computer against these bugs that affect billions of devices, requires you to reboot your systems, but not with KernelCare, a commercial live patching service from CloudLinux.

According to CloudLinux, KernelCare is now capable of live patching the Meltdown vulnerability and the first variant of the Spectre exploit on CloudLinux 7 series of operating systems, as well as Red Hat Enterprise Linux 7, CentOS 7 and CentOS 7 Plus, and Proxmox Virtual Environment 3.10.

To use KernelCare on your CloudLinux, CentOS, Red Hat Enterprise Linux, and Proxmox VE systems, you'll have to pay from \$2.25 USD per server monthly, but there's also a free trial that supports updates for all Linux kernels on an unlimited number of servers, and CloudLinux even promises to bring support for more

GNU/Linux distributions soon.

These include CloudLinux 6, CentOS 6, Red Hat Enterprise Linux 6, and Virtuozzo 6 series, as well as Ubuntu, Debian, and other distros. However, CloudLinux noted the fact that Xen PV is not and won't be supported on KernelCare, which is widely used among cloud providers and enterprises.

Source:

<http://news.softpedia.com/news/cloudlinux-s-kernelcare-promises-to-fix-meltdown-spectre-flaws-without-reboots-519728.shtml>

NINTENDO SWITCH HACKERS MANAGE TO LOAD LINUX ON THE HYBRID CONSOLE, AND IT 'CAN'T BE PATCHED'

Last month, a group of hackers teased the public with its claims that the Switch can be hacked with a future-proof exploit.

Some advances in the hacking scene were put on display previously and only flaunted homebrew apps that ran on the hybrid game system. However, the

latest one shared by fail0verflow, a hacking group, flaunted a picture of the Nintendo console with a Linux boot screen on display, the photo of the team's achievement was shared via Twitter for the public to see.

Moreover, it was noted that the exploit was supposedly future-proof and cannot be patched out via official software updates.

Upon its reveal, several sources immediately called it out as a possible fake. Current photo-editing software could easily reproduce a doctored image of almost anything people can find on the web. However, most industry experts acknowledge fail0verflow's track record when it comes to hacks released from the Nintendo Wii up to the Sony PlayStation 4, which lends to the group's credibility concerning the Switch exploit.

Although it might seem like the final build completion is just a few months away, past exploits for other game systems did not release until more than a year after. Referring to the image from fail0veflow, it might still be a while, judging by the external

wires and cables.

Source:

<http://www.techtimes.com/articles/220751/20180211/nintendo-switch-hackers-manage-to-load-linux-on-the-hybrid-console-and-it-cant-be-patched.htm>

KDE PLASMA LINUX DESKTOP IS NO LONGER VULNERABLE TO USB ATTACKS

One important security vulnerability fixed in the KDE Plasma 5.12 LTS desktop environment is a USB exploit that could allow a local attacker with physical access to the unpatched computer to execute arbitrary commands if the malicious USB flash drive was mounted via the removable device notifier function and contained certain characters in its volume label.

All KDE Plasma users running a previous version of the desktop environment should update their installations to the latest KDE Plasma 5.12 LTS release as soon as possible. The new version is

already available in the software repositories of popular GNU/Linux distributions like Kubuntu/Ubuntu, Arch Linux, OpenSUSE, and others, so there's nothing holding you back to update it right now.

If you can't update your KDE Plasma desktop to version 5.12, there's a workaround for the USB bug, as you'll have to mount all removable USB devices with the Dolphin file manager instead of using the device notifier. Previous LTS users using Plasma 5.8 can update to KDE Plasma 5.8.9 LTS, which also addresses this security flaw. Other Plasma users can apply the patches in the advisory.

Source:

<http://news.softpedia.com/news/kde-plasma-linux-desktop-is-no-longer-vulnerable-to-usb-attacks-update-now-519767.shtml>

CANONICAL PLANS TO RELEASE UBUNTU 16.04.4 LTS (XENIAL XERUS) ON MARCH 1, 2018

The Ubuntu 16.04.4 LTS maintenance update was

originally scheduled for release today, February 15, 2018, but Canonical decided to delay it last month due to the Meltdown and Spectre security vulnerabilities that have been publicly disclosed last month and found to affect billions of devices.

Furthermore, a new release date was not announced until today when Canonical said that it would release Ubuntu 16.04.4 LTS very early next month, on March 1st, 2018. Until then, the company expects to have all the updates ready in the archive for existing users who would want to update their installations.

Ubuntu 16.04.4 LTS would be the fourth of five scheduled maintenance updates for the long-term supported Ubuntu 16.04 LTS (Xenial Xerus) operating system series, and it will ship with up-to-date kernel and graphics stacks based on those from Ubuntu 17.10 (Artful Aardvark).

Canonical promised to offer five years of support for Ubuntu 16.04 LTS, until April 2021. The operating system was released two years ago, on April 21, 2016, and it received three such maintenance

updates until now. One more maintenance update is planned for the Xenial Xerus series, Ubuntu 16.04.5, which could arrive later in the year.

Until then, Ubuntu 16.04 LTS users using the Ubuntu 16.04.3 milestone released last year would be able to update their installations to the Ubuntu 16.04.4 release on March 1, 2018, without the need to download the new ISO images, which are designed for new installations only.

On April 26, 2018, Canonical plans to release a new long-term supported Ubuntu series, the Bionic Beaver, a.k.a. Ubuntu 18.04 LTS, which will be supported with similar maintenance updates containing the most recent security and software updates until April 2023.

Source:

<http://news.softpedia.com/news/canonical-plans-to-release-ubuntu-16-04-4-lts-xenial-xerus-on-march-1-2018-519817.shtml>

CANONICAL WANTS TO COLLECT SOME DATA FROM UBUNTU USERS TO IMPROVE NEW RELEASES

The information Canonical's Ubuntu Desktop engineers need to improve certain aspects of the Linux-based operating system about includes users' setups, installed software, Ubuntu flavor and version, network connectivity, CPU family, RAM, disk size, screen resolution, GPU vendor and model, as well as OEM manufacturer.

In addition, the company says that it needs to know your location, yet it promises to not store IP addresses of users. Other information that would be collected includes total installation time, automatic login info, selected disk layout, LivePatch enablement, and if you choose to install updates or third-party software during installation.

Canonical says that it plans to implement the new data collection option in the installer through a checkbox named something like "Send diagnostics information to help improve Ubuntu" and enabled

by default. Of course, users will be able to uncheck this box during the installation if they don't want Canonical to collect their data.

However, it's important to know that this would help the Ubuntu Desktop development team to focus their efforts on the things that matter the most to you in future versions of Ubuntu. Also, Canonical said that all the collected data from the installation would be securely sent to a service run by Canonical's IS team via HTTPS.

The data is saved locally, on your computer, and would be sent to Canonical on first boot if an active network connection is detected. Users will be able to access the respective file that contains the collected data and inspect it thoroughly. Canonical said that the results of this data collection would be made public.

Source:

<http://news.softpedia.com/news/canonical-wants-to-collect-some-data-from-ubuntu-users-to-improve-new-releases-519815.shtml>

HATE TO RUIN YOUR DAY, BUT... BOFFINS COOK UP FRESH MELTDOWN, SPECTRE CPU DESIGN FLAW EXPLOITS

When details of the Meltdown and Spectre CPU security vulnerabilities emerged last month, the researchers involved hinted that further exploits may be developed beyond the early proof-of-concept examples.

It didn't take long. In a research paper – "MeltdownPrime and SpectrePrime: Automatically-Synthesized Attacks Exploiting Invalidation-Based Coherence Protocols" – out this month, bit boffins from Princeton University and chip designer Nvidia describe variants of Meltdown and Spectre exploit code that can be used to conduct side-channel timing attacks.

In short, the team have discovered new ways for malware to extract sensitive information, such as passwords and other secrets, from a vulnerable computer's memory by exploiting the Meltdown and Spectre design blunders in modern processors.

The software mitigations being developed and rolled out to thwart Meltdown and Spectre attacks, which may bring with them performance hits, will likely stop these new exploits.

Crucially, however, changes to the underlying hardware probably will not: that is to say, whatever Intel and its rivals are working on right now to rid their CPU blueprints of these vulnerabilities may not be enough. These fresh exploits attack flaws deeply embedded within modern chip architecture that will be difficult to engineer out.

Before you panic: don't. No exploit code has been released.

Source:

https://www.theregister.co.uk/2018/02/14/meltdown_spectre_exploit_variants/

FREE SATURN "RANSOMWARE-AS-A-SERVICE" ALLOWS ANYONE TO BECOME A NOTORIOUS HACKER

Last week, MalwareHunterTeam spotted a new ransomware named Saturn. As per a detailed report from Bleeping Computer, Saturn ransomware is being actively distributed at the moment, but the methods used for its distribution are unknown. Also, this ransomware is not decryptable at the moment.

After infecting the machine, Saturn ransomware executes commands to disable Windows repair and clear Windows backup catalog. It encrypts the files and adds Saturn to their name. The ransomware also leaves a ransom note in each folder, which contains a link to the payment site. The ransom amount is set to \$300 at the moment, which doubles after seven days.

That was a brief introduction of Saturn ransomware. But that's not the end of the story. The creators of Saturn are offering the

ransomware for free via a Ransomware-as-a-Service (RaaS) affiliate program.

The members of the affiliate program need to generate an infection file and distribute it to other users via email or other kinds of campaigns. The infected users end up paying the ransom in form of Bitcoin to the malware creators. The affiliate member of the program gets 70% of the payment and creators get 30%.

This zero buying cost and payment division model could attract lots of notorious players who could try to get their hands on Saturn.

Source:

<https://fossbytes.com/saturn-ransomware-as-a-service-free/>

UBUNTU ADDS NEW "MINIMAL INSTALLATION" OPTION FOR FEWER PREINSTALLED PACKAGES

The development of the next Ubuntu LTS release, i.e., Ubuntu 18.04 Bionic Beaver, is

going on in full swing. The desktop development team has decided to add a new option in the installation process that allows you to perform a lean installation of Ubuntu.

As reported by OMGUbuntu, this option appears during the installation process. If you choose the minimal install option, lots of software packages are stripped away, and you get an Ubuntu installation with a desktop environment, core system tools, web browser, and not much more.

This option will appear in the Ubuntu installer (Ubiquity) in the section where you are asked if you wish to install restricted multimedia playback codecs.

To be precise, this option removes about 80 packages, including the likes of Transmission, LibreOffice, Cheese, Rhythmbox, Thunderbird, etc. So, if you're a person who deletes tons of software soon after a clean installation of Ubuntu, this option will be useful.

Source:

<https://fossbytes.com/ubuntu-new-minimal-installation/>

ORACLE OPEN-SOURCES DTRACE UNDER THE GPL

Oracle appears to have open-sourced DTrace, the system instrumentation tool that Sun Microsystems created in the early 2000s and which has been beloved of many-a-sysadmin ever since.

As noted by developer Mark J. Wielaard, this commit by an Oracle developer shows that something is afoot.

Big Red recently listed DTrace as one of the open source projects it hosts and has also created a new mailing list for discussion of what it's called a "Linux port of the Solaris Dynamic Tracing Framework."

The company has also revealed "The license for DTrace has changed over time. Versions before 0.6.1 include a CDDL kernel module and a proprietary userspace component; versions 0.6.1 and above have a GPL kernel component and a UPL userspace component. The UPL is generally held to be compatible with the GPL, so it therefore looks possible

to include DTrace in Linux. As DTrace aficionado Brendan Gregg told us in 2016, Linux has only recently added tracing tools to compare with Sun's progeny. The prospect of bringing all of DTrace to Linux is therefore rather tantalising.

Source:

https://www.theregister.co.uk/2018/02/19/oracle_open_sources_dtrace_changes_licence_to_gpl/

LIBREOFFICE 6.0 OPEN-SOURCE OFFICE SUITE PASSES 1 MILLION DOWNLOADS MARK

The Document Foundation announced recently that its LibreOffice 6.0 open-source and cross-platform office suite reached almost 1 million downloads since its release last month on January 31, 2018.

That's terrific news for the Open Source and Free Software community and a major milestone for the acclaimed LibreOffice office suite, which tries to be a free alternative to proprietary

solutions like Microsoft Office.

The 1 million downloads mark was reached just two weeks after the release of LibreOffice 6.0, which is the biggest update ever of the open-source office suite adding numerous new features and enhancements over previous versions. These include a revamped design with new table styles, new gradients, new Elementary icons, and improved Notebooks, superior interoperability with Microsoft Office documents, support for new import and export file formats, and much more.

Last week, the Document Foundation published the first point release of the LibreOffice 6.0 office suite to improve the security and robustness of the application on all supported platforms. LibreOffice 6.0.1 is now the latest version, but The Document Foundation will be supporting it with more updates during the past few months.

Source:

<http://news.softpedia.com/news/libreoffice-6-0-open-source-office-suite-passes-1-million-downloads-mark-519831.shtml>

MEET THE ANTSLE: THE PERFECT OUT-OF-THE BOX VIRTUAL MACHINE SOLUTION

I've worked with virtual machines on many platforms from VirtualBox to VMware to KVM to qemu. For the most part, they're all pretty much the same—a nice GUI tool built around powerful command tools, designed to create and manage virtual machines on a host operating system. Of the available options, VirtualBox has been my go-to tool for VMs for some time. Currently I have over 50 virtual machines at the ready. Fortunately, I don't run more than two or three at a time, as that would take a serious cut of my workstation resource pie. The other downfall of my solution is that those virtual machines are taking up precious resources—that I don't always have to spare.

But what if there was an alternative that could make managing your virtual machines incredibly easy? A solution that is not only self-contained, but cost-effective?

That's the Antsle. A private

cloud server, designed for developers, that can serve businesses of all sizes. With this piece of hardware, you can roll out servers, containers, you name it—all from a user-friendly web-based GUI. Starting at \$799 and going to \$4,499 (for base models), you can purchase and have an Antsle built to meet your specific needs.

But what I've found most impressive about the Antsle is how easy it is to spin up an "antlet" (think virtual machine). In seconds, you can deploy a VM based on any one of the built-in templates (Ubuntu, Fedora, Debian, CentOS, FreeBSD, Kali Linux, or Windows Server). If the list of templates doesn't suit your needs, you can upload a virtual appliance or use an ISO to create a new template, which can then be used to deploy "antlets."

I've been working with an Antsle for a week now. It took me roughly an hour to get up to speed on rolling out antlets, and a day or so to realize that this solution would usurp VirtualBox as my go-to VM solution. Considering I've been using VirtualBox for a decade, that's should serve as a

testimony for the ease of use and power to be found in Antsle.

Source:

<https://www.techrepublic.com/article/meet-the-antsle-the-perfect-out-of-the-box-virtual-machine-solution/>

INTEL RETURNS WITH ANOTHER SPECTRE PATCH FOR 6TH, 7TH, 8TH GEN CORE PROCESSORS

Intel hasn't stepped back on their efforts to show how concerned they're for the people affected by the critical Meltdown and Spectre vulnerabilities disclosed ahead of schedule. The first lot of security patches released by Intel ended up rebooting people's machines.

Later, the world's second largest chip maker had to advise users to refrain from applying the updates. For a bug-free future, Intel CEO sort of promised last month that the Intel chips releasing later this year will be free of spectre and meltdown bugs.

Now, Intel has released a new

Spectre security patch for their Skylake, Kaby Lake, and Coffe Lake processor architectures that power the 6th, 7th, and 8th generation Intel Core processor lines. The update also applies to Intel Core X series, as well as, their Xeon Scalable and Xeon D server chips.

The new microcode update includes a mitigation technique called 'Retpoline' for Spectre variant 2 (Branch Target Injection). Developed by Google, it is one of the multiple mitigation techniques Intel has delivered as a part of the security patch that went through "extensive testing by customers and industry partners to ensure the updated versions are ready for production."

Intel will release the fresh microcode to their OEM partners, and it will land on affected machines via system updates. Maybe this time, the update doesn't cause further problems like it did in the past.

Source:

<https://fossbytes.com/spectre-patch-intel-core-chips-skylake-kaby-lake-coffe-lake/>

CANONICAL'S UNITY 8 DESKTOP REVIVED BY UBPORTS WITH SUPPORT FOR UBUNTU 18.04 LTS

As you are aware, last year Canonical decided to stop the development of its futuristic Unity 8 desktop for Ubuntu and the Ubuntu Touch mobile OS. Days after their sad announcement a few community members appeared interested in taking over the development of Unity 8, the most promising one being Yunit.

However, the Yunit project didn't manage to improve Unity 8 for desktops in the last few months as much as the community would have wanted, and, after a long battle, they decided to pass the baton to UBports team, which is announcing the initial build for devs and an official website for Unity 8.

This first Unity 8 build maintained by UBports apparently supports the next-generation Wayland display server through the latest Mir 0.30 display server, which is still developed by Canonical. We believe that UBports

will try to improve Wayland support so you can run your favorite Linux apps soon.

Until then, all developers and anyone who wants to help UBports continue Canonical's Unity 8 and Ubuntu convergence dream is welcomed to join the dedicated UBports Unity8 Developers group on Telegram, and check out the installation instructions below.

If you want to install UBports' Unity 8 desktop environment on your personal computer, you should have either Ubuntu 16.04 LTS (Xenial Xerus) or Ubuntu 18.04 LTS (Bionic Beaver), which is currently in development. Also, you should be aware of the fact that most X11 apps won't work.

Source:

<http://news.softpedia.com/news/canonical-s-unity-8-desktop-revived-by-ubports-with-support-for-ubuntu-18-04-lts-519890.shtml>

BOTCHED NPM UPDATE CRASHES LINUX SYSTEMS, FORCES USERS TO REINSTALL

A bug in npm (Node Package Manager), the most widely used JavaScript package manager, will change ownership of crucial Linux system folders, such as /etc, /usr, /boot.

Changing ownership of these files either crashes the system, various local apps, or prevents the system from booting, according to reports from users who installed npm v5.7.0. —the buggy npm update.

Users who installed this update —mostly developers and software engineers— will likely have to reinstall their system from scratch or restore from a previous system image.

The bug was first reported a week ago but was left without an answer from npm developers. Users filed a new bug report after last night's release, and the npm team has released npm v5.7.1, a version that removes the buggy code.

FreeBSD users have also reported being impacted by the bug. Mac and Windows users didn't experience any issues. The

problem did not affect every Linux user.

Running the npm update commands as root doesn't result in npm trying to reassign root ownership to all files, so the issue appears to affect only npm update operations prefixed by a sudo command.

Npm is the de-facto package manager for all small, medium, and large-scale JavaScript project. Npm is packed with Node.js, and is also the largest package manager on the Internet, hosting libraries and plugins for Node.js, Ember, jQuery, Bootstrap, React, Angular, and many other JavaScript frameworks. You won't find a JavaScript developer that doesn't use nowadays.

Source:
<https://www.bleepingcomputer.com/news/linux/botched-npm-update-crashes-linux-systems-forces-users-to-reinstall/>

CANONICAL OUTS NEW UBUNTU KERNEL UPDATE WITH COMPILER-BASED RETPOLINE MITIGATION

New Linux kernel security updates have been released for Ubuntu 17.10 (Artful Aardvark), Ubuntu 16.04 LTS (Xenial Xerus), Ubuntu 14.04 LTS (Trusty Tahr), and Ubuntu 12.04 ESM (Extended Security Maintenance), adding the compiler-based retpoline kernel mitigation for the Spectre Variant 2 vulnerability on amd64 and i386 architectures.

Canonical fixed the Spectre Variant 2 security vulnerability last month on January 22, but only for 64-bit Ubuntu installations. This update apparently mitigates the issue for 32-bit installations too. Spectre is a nasty hardware bug in microprocessors that use branch prediction and speculative execution and it could allow unauthorized memory reads via side-channel attacks.

In addition to this, the new kernel updates address a race condition (CVE-2017-17712) in Linux kernel's IPv4 raw socket

implementation and a use-after-free vulnerability (CVE-2017-8824) in the DCCP protocol implementation, allowing a local attacker to execute arbitrary code or cause a denial of service. Both security flaws were discovered by Mohamed Ghannam.

Also, the new kernel patches a use-after-free vulnerability (CVE-2017-15115) discovered by ChunYu Wang in Linux kernel's SCTP protocol implementation, which could allow a local attacker to crash the system by causing a denial of service or execute arbitrary code. These security issues affect Ubuntu 17.10, Ubuntu 16.04 LTS, and Ubuntu 14.04 LTS.

Source:
<http://news.softpedia.com/news/canonical-outs-new-ubuntu-kernel-update-with-compiler-based-retpoline-mitigation-519909.shtml>



Over a decade has passed since I bought and built my first desktop PC (until then, I'd used my family's). Included in this build was a monitor, keyboard, and mouse. As time went on, I upgraded the desktop PC, and also the peripherals. Specifically, I eventually bought a second monitor (for accurate color representation, and extra screen real estate). Fast forward to current day - I have finally replaced both monitors (one was 11 or 12 years old, and the other was 7 or 8). Both monitors had issues (random power issues, only 1 input worked on either, etc). Not to mention the first monitor was 1680x1050, and had atrocious color accuracy. It was therefore time to upgrade - but to what? This article is dedicated to my move back to a single, ultra-wide monitor, and some tweaks I've implemented in i3.

THE MONITOR

For those who want the exact model, it's an LG 34UM68-P.

Specs: 34", with a resolution of 2560 x 1080. 99% sRGB coverage.

A FEW OBSERVATIONS RIGHT OFF THE BAT

It wasn't that much smaller than my dual monitor setup (both monitors were under 23"), but due to the fact that this was one panel, the viewing angle was more comfortable. I had previously mounted both monitors on raised arms connected to my desk, but that always resulted in a sharp angle between both monitors, where I tried to avoid too much of a gap due to spacing and bezels.

The new monitor also offered an adjustable stand, meaning I could do away with the monitor arms, and instead just raise/lower and tilt the ultrawide to a comfortable position.

I also instantly noticed that having a browser window open the full width (as tiling window managers do when only one window is open) was

overwhelmingly huge. Thankfully, opening a second window, and adjusting the split a tiny bit, makes for a very comfortable size.

Brightness - I never touched the brightness settings on either of my old monitors (set to 100), and this new monitor was like staring into a small sun when looking at anything with a white background. Dialing the brightness down to 50% helped immensely.

MY REQUIREMENTS

My dual monitor setup was originally like this (monitors are indicated with L for the left-hand 1080p panel, and R for the right-hand 1050p panel).

PS4 -> HDMI Switch -> Monitor (L)

Windows PC -> HDMI Switch -> Monitor (L)

Nintendo Switch -> HDMI Switch (Monitor L)

Intel NUC -> HDMI Switch -> Monitor (L)

Intel NUC -> Mini DisplayPort -> Monitor (R).

The HDMI Switch would automatically swap inputs depending on what was active - meaning my left-hand monitor was used to display everything, while the right-hand monitor only ever showed my NUC. This was great because it let me keep an eye on emails, or watch videos on the right while playing games or testing websites in IE under Windows.

Naturally, I wanted to be able to do something similar with one monitor. Which is why I required the monitor to support Picture-By-Picture. Essentially, the monitor can activate two inputs simultaneously and display them as 1280x1080 side by side. Naturally, when using a PS4 or Nintendo Switch, it will have to scale the resolution of 1920x1080 to be able to see the whole screen. Which results in black bars above and below the image. However, due to the size (and closeness) of the 34" panel, the actual size of

the image isn't much smaller than my 1080p screen originally was. Scaling the resolution down to 1280x1080 under Linux also ensures that watching a video (for example) yields the exact same size bars as the PS4 input.

While I do frequently use both inputs, I will also enjoy focusing on one thing at a time, which this monitor does great as well. On 1920x1080 inputs, there will again be black bars on the left/right of the screen. Personally, this doesn't bother me, but your mileage may vary.

SWITCHING RESOLUTIONS IS A PAIN USING XRANDR

Running i3wm means you will have to manage your resolutions manually. For this, I have always used xrandr. Unfortunately, with this new setup, I found myself changing resolutions frequently. This also resulted in a bug I noticed with polybar (my system panel for i3) where it would not adjust itself to the resolution - only restarting it would fix the issue. So I pulled up a vim window and wrote myself a very small shell script.

```
#!/bin/bash

if [ "$1" == "On" ]; then
    xrandr --output DP1 --mode 1280x1080
    killall polybar
    polybar main_bar &
    disown
elif [ "$1" == "Off" ]; then
    xrandr --output DP1 --mode 2560x1080
    killall polybar
    polybar main_bar &
    disown
fi
```

The reason I needed a shell script and not an alias was that I could not string the polybar main_bar & into the disown command when using an alias. That meant closing the terminal also closed polybar. And since polybar dumps some messages into the terminal, it resulted in some odd looks. Disown essentially disconnects the background commands from a terminal. The file is called splitScreen, and running it as splitScreen On enables the settings for when I'm dividing the screen, and splitScreen Off changes it back.

One thing I noticed - xrandr won't register the 1280x1080 resolution option unless you've already split the screen using the monitor's on screen display. Not a

big deal, but worth noting.

WHY NOT 4K?

I did consider a 4K monitor, but I simply couldn't justify the increase in price. Especially due to the fact that I own a PS4 Pro, I would have wanted to get a 4K, HDR monitor, which are extremely expensive. Personally, I find monitors in Germany to be much more expensive than other countries (such as Canada). I'd much prefer to eventually get a 4K HDR television (cheaper) for my console alone. As for 4K for work - I don't have a lot of situations where I've really wished I had more screen real estate, and wasn't going to spend nearly 2x as much on a monitor that didn't fulfill a palpable need at my end. Tied into this was also a severe lack of

showrooms in my area that actually carried IPS panels, or modern monitors. Most options are geared towards either gamers (poor sRGB coverage, and absolutely no Adobe RGB coverage), or budget-friendly monitors (high refresh time, 1080p resolution, bad color accuracy). If I had been able to actually see and try out a 4K monitor, it may have changed my mind. But you can't miss what you don't know, correct?

ANYTHING ELSE?

The main factor was always cost. I wanted to replace my monitors, but I didn't want to spend more than 400€ if I could avoid it. The monitor I eventually bought was offered on Amazon at around 430€. However, I was fortunate enough to notice an Amazon Warehouse deal down to 380€ - which fell perfectly within my budget. The downside? A single large scratch on the back of the monitor, and a box that looked like it had gone through a war. In my mind? Perfectly acceptable trade-off.

UNFORESEEN COSTS?

I had to buy a mini DisplayPort to DisplayPort cable, mainly because I wanted to get rid of the cheap adapter I had previously been using.

More unforeseen was that I ended up replacing my desk entirely. It was previously a free hand-me-down desk that was quite curved. The problem with the curve was that the relatively wide foot on the new monitor wasn't playing nicely with some of my cable restrictions (the ethernet cable I run through the room to get internet, for example). Since I was generally not too thrilled with the usable space on the old desk, I went to Ikea and bought a new one. The upside here is that I can now make much more efficient use of the space I have, and have nearly doubled my open working area between my main desk (with monitor, PS4, PC, etc) and my second desk (which I use for charging my laptop and as a workbench when doing PC repairs or working on guitars).

Could I have made do without the new desk? Certainly. Am I happy about it? Absolutely - I wasn't even aware at times how cramped I felt sitting in the curve

of that desk. Not to mention the ability to cable-manage much better on the new desk due to the shape and form of the frame.

SURPRISE BENEFITS?

I often play music from digitized sheet music or tablature - I found that by fullscreening my PDF viewer, I can fit 3 pages of music on screen at once. Which means a lot less pausing to advance through pages using my keyboard.

I also find myself vertically splitting windows a lot more than before, meaning I can comfortably fit a video playing in 1/4 of the screen, have a browser open to 50% of the screen, and use the last 1/4 to have a code editor open for tweaking HTML or CSS while I test.

CONCLUSION

I absolutely love my new setup - I have had to get reacquainted with some keyboard shortcuts in i3 to more easily manage my windows (as I previously did only 50/50 splits, or moved windows between monitors). But after doing that (and learning I can

fullscreen a video to the container size by fullscreening the video using the player's control, and then hitting mod+F), I have been able to better utilize the space than ever before.

The better sRGB coverage is also a life-saver when it comes to testing websites - the oldest monitor I had often swallowed small details like shadows and lines. I also find the width of this monitor (I estimate it's maybe 10-15cm narrower than my dual monitor setup, including bezels) to be much more comfortable to look at. I no longer have to actively turn my head or hold it at an odd angle to clearly see what I'm looking at.

RECOMMENDATIONS?

If you're also looking for an ultrawide monitor (or are considering it), I would highly recommend it if you don't mind the following:

- Managing window size
- Black bars around videos
- Dealing with resolutions and their limitations depending on the input device.

If you're looking for a gaming

monitor (as opposed to one for productivity), I'd probably steer you towards something with a high refresh rate/response time. That being said, this monitor does support FreeSync (requires an AMD graphics card). I'm just not sure how comfortable gaming on a 2560x1080 screen would be, assuming you keep the monitor close enough to reach the on-screen controls. You may prefer a 1440p monitor.

Do you have any questions or comments that I haven't mentioned in this article? Let me know by sending me an email at lswest34+fc@gmail.com. Similarly, if you have any requests for things I should cover, please let me know and I'll do my best to cover them in the coming months.



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.



Sphinx is a very nice document processor; it can create html, latex, pdf's, and more. I started using it a few years ago to write a short book about Wood Turning. The reason for choosing Sphinx is its ability to create the different styles by creating only a config file and writing the document in reStructuredText. Here is the documentation:

<http://docutils.sourceforge.net/rst.html>

If you have not used markup, you are missing out. Try it out. All you need is an editor such as vim/gvim or ??? – your choice – do not use a word processor. So here is the basic layout for the markup and how to install Sphinx (by the way, Sphinx was written to create Python documentation).

This was copied from the Sphinx documentation page: <http://www.sphinx-doc.org/en/stable/>

- Output formats: HTML (including Windows HTML Help), LaTeX (for printable PDF versions), ePub, Texinfo, manual pages, plain text

- Extensive cross-references: semantic markup and automatic links for functions, classes, citations, glossary terms and similar pieces of information
- Hierarchical structure: easy definition of a document tree, with automatic links to siblings, parents and children
- Automatic indices: general index as well as a language-specific module indices
- Code handling: automatic highlighting using the Pygments

highlighter

- Extensions: automatic testing of code snippets, inclusion of docstrings from Python modules (API docs), and
- Contributed extensions: more than 50 extensions contributed by users in a second repository; most of them installable from PyPi

OK, let's install Python and Sphinx; most distributions come with Python installed or at least in the package manager. You may still

need to install PyPi to use pip:

```
sudo apt install python-pip
```

Install Sphinx, either from a distribution package or from PyPi

```
$ pip install Sphinx
```

or

```
sudo apt install python-sphinx
```

this works best for Ubuntu

```
.. Woodturning documentation master file, created by
   sphinx-quickstart on Thu Dec  5 22:01:12 2013.
   You can adapt this file completely to your liking, but it should at least
   contain the root `toctree` directive.
```

```
Welcome to Woodturning's documentation!
=====
```

Contents:

```
.. toctree::
   :maxdepth: 2
```

```
chapter1
chapter2
chapter3
chapter4
```

```
Indices and tables
=====
```

```
* :ref:`genindex`
* :ref:`modindex`
```

HOWTO - SPHINX

16.04.01

Then run and answer all questions, be sure to say yes to autodoc extensions. If you want to create a directory first and change to that dir, quickstart will do it for you:

```
$ sphinx-quickstart
```

After this is finished, you will have a `index.rst` file that looks like this, and you will have a `conf.py` file created from the questions you answered in the quickstart.

```
:maxdepth: 2
```

```
intro  
tutorial
```

Just add your chapters etc into this file as shown in the example at the bottom of the previous page.

You may want to look through the `conf.py` file and see if it is what you want, if not make the changes. Then type:

```
make html
```

This will build the documentation you just created in html format under the `_build/html`

directory. If you make for other types of files, you will see the directories created in the `_build` dir. Example:

```
doctrees  html  latex  tex  
xml
```

I hope this is of use to some people who like to try new ways of using Ubuntu/Linux.

THE OFFICIAL FULL CIRCLE APP FOR UBUNTU TOUCH

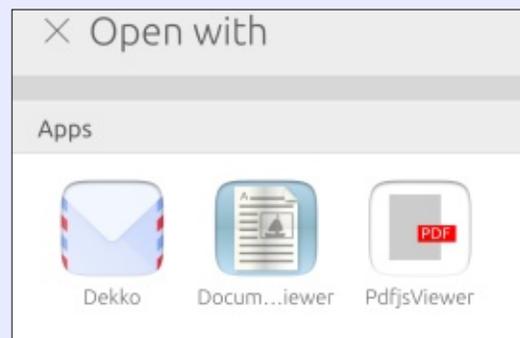


Brian Douglass has created a fantastic app for Touch devices that will allow you to view current issues, and back issues, and to download and view them on your Ubuntu Touch phone/tablet.

INSTALL

Either search for 'full circle' in the Ubports App 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>





HOW-TO

Written by Elmer Perry



Freeplane - Pt1

Freeplane is a program for creating mind maps. The program is versatile enough to fit anyone's mapping style. You can make the maps as simple or as complex as you need for your application. Freeplane is available in the Ubuntu Software Center or online at their web page, https://www.freeplane.org/wiki/index.php/Main_Page.

WHAT IS A MIND MAP?

A mind map is a visual diagram for organizing ideas and information. Use a mind map to brainstorm a new idea, create a process, or organize your thoughts. A mind map can help you keep notes during a lecture and see patterns you might otherwise miss. I've used mind maps through the years to write stories and articles, and create the processes for programs. I also used mind maps to create teaching and training curricula. Yes, I know other tools exist for doing these things, but a mind map gives me something other tools don't. With a mind map, I can see the entire

program structure and functions. Other tools don't give me this broad overview. Being able to see the entire thought process allows me to see connections I would otherwise miss.

The basic structure of a mind map starts with a central topic. This is the idea, class, process, or topic you are mapping. From the central topic, the map branches into several subtopics related to the topic. The subtopics can branch out to other child topics, and those child topics out to other child topics. The branches can go as deep as needed for the topic.

INSTALLATION

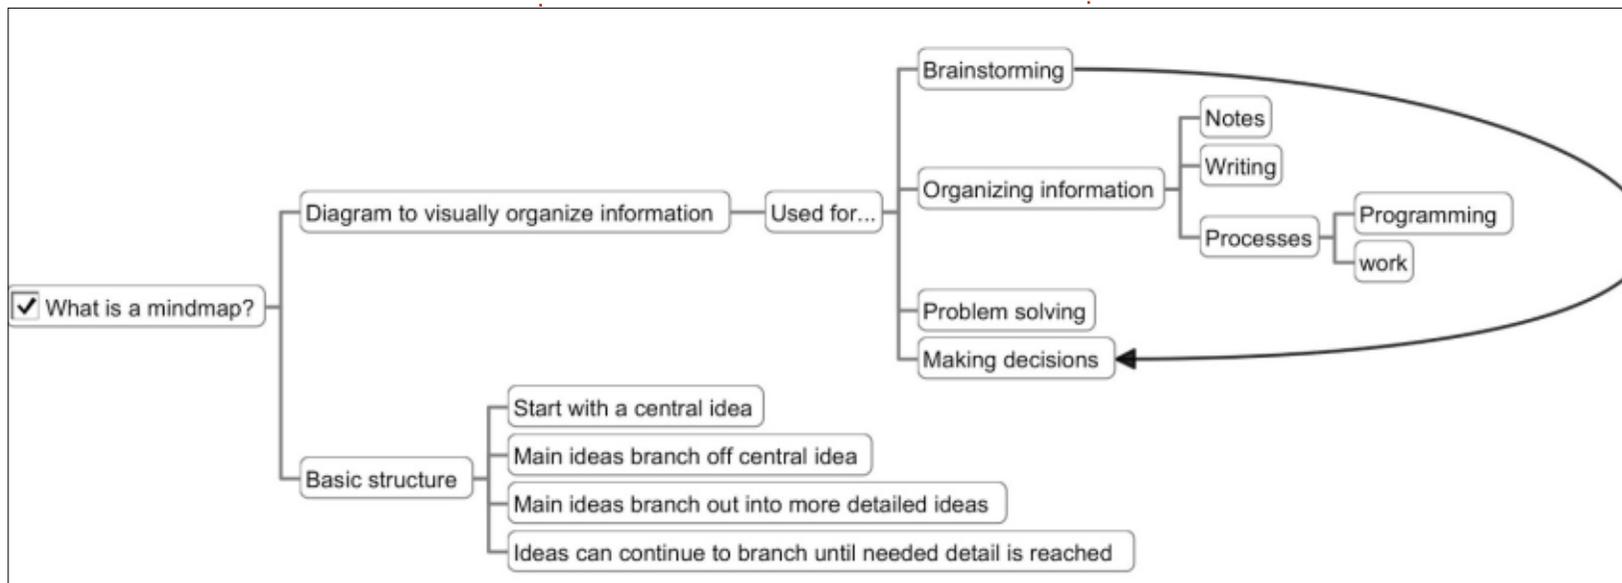
You can download Freeplane from the Ubuntu Software Center. I'm starting with the latest download version, 1.6.10. While you can follow along with the version available in the Software Center, I will discuss new features found in 1.6. You can download version 1.6 from the Freeplane home page, https://www.Freeplane.org/wiki/index.php/Main_Page. Freeplane is a Java program. If you are running the latest Java runtime, you should have no problems running

Freeplane.

PROGRAM WORKSPACE

The program has a standard layout with the menus and toolbars at the top, and a workspace at the bottom. On the left-hand side, you can pull out an Icons toolbar, and on the right-hand side, you can pull out the Tool Panel. By default, you can pop up the Note panel at the bottom. In the Preferences, you can change where the Note panel docks.

The menu bar sits below the



HOWTO - FREEPLANE

title bar. The menus give you access to the program's commands.

The program has several toolbars. The main toolbar shows by default. The main toolbar gives you access to save options and common formatting actions. The Filter toolbar allows you to filter what nodes of the map show,

making it a handy search tool. The F-Bar shows the actions assigned to the F-Keys. You can customize the actions assigned to the F-Keys. The Icon toolbar gives you a set of icons for marking nodes in your maps. And the status bar sits at the bottom of the window. The status bar displays information about the current map or selected node. The Filter, F-Bar, Icon, and status bars

hide by default, but you can activate them through the menu View > Controls.

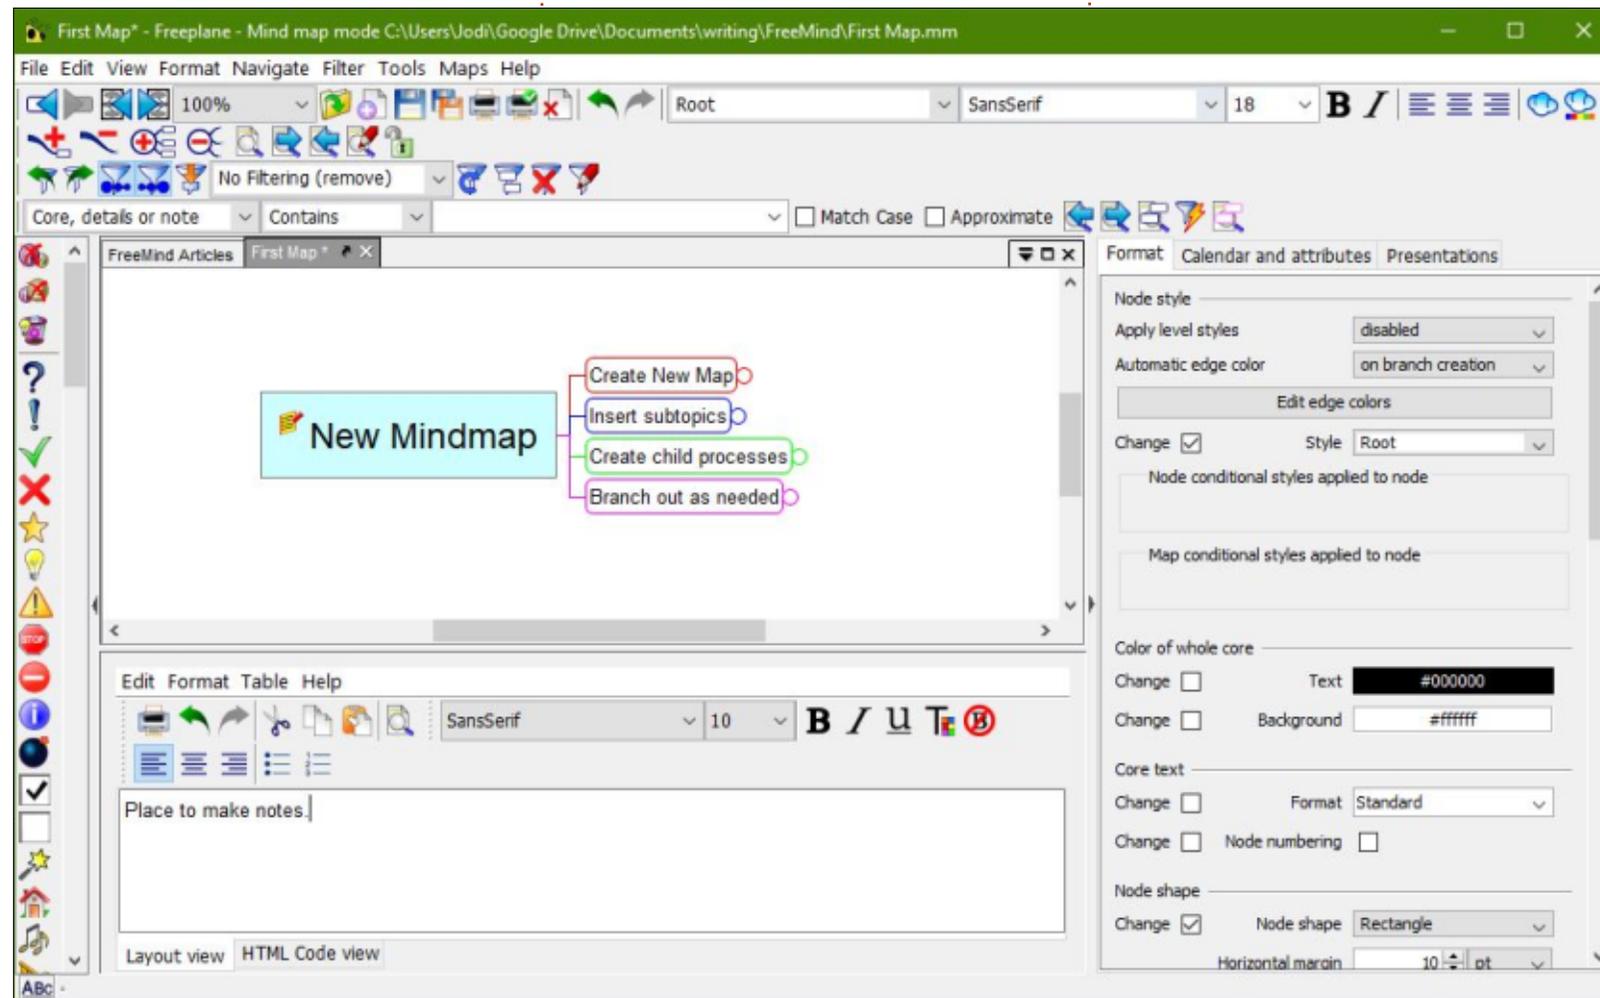
The workspace in the middle is where you will build your map. You can have many maps open at one time, and each one displays a tab at the top of the workspace. Scrollbars hide by default, but you can show them using the View >

Controls menus. You can click and drag on a blank space in the map area to move the map around. The wheel on your mouse will scroll the map up and down. Hold the SHIFT key and scroll the mouse wheel to move the map left and right.

YOUR FIRST MIND MAP

Enough talk! Let's make our first map. If you haven't already done so, install and open Freeplane. Create a new map by selecting File > New Map from the menus. In the template dialog, select the standard template. In the 1.6 version, the name is standard-1.6.mm. This creates a new map with a root topic in the center of the screen. Type in a name for your topic like "My First Map." Press the Enter key to save the name.

To create subtopics for your main topic, press the Insert key on the keyboard. Nodes that are one level deeper than the current node are called child nodes of the current node. Type the text for the subtopic, and press the Enter key. Create more subtopics by pressing the Enter key. Nodes on the same level are called siblings in



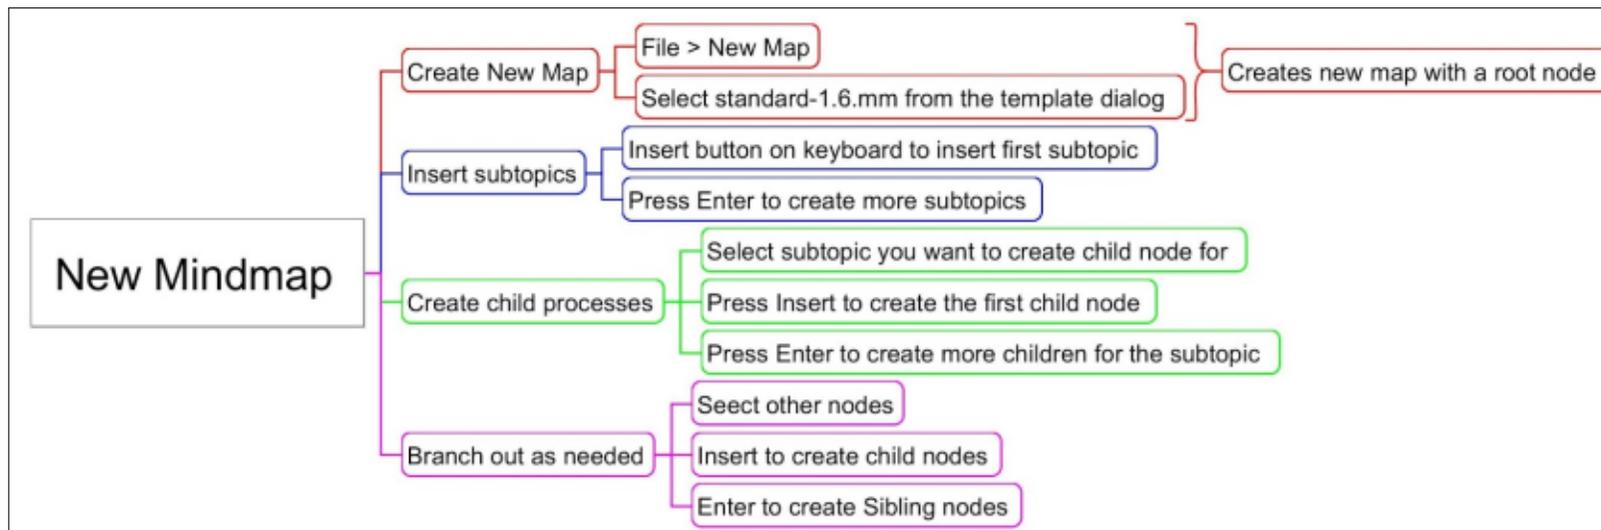
HOWTO - FREEPLANE

Freeplane. Create several subtopics by pressing Enter and typing in some text.

Now, let's create child nodes of a subtopic. Select one of the subtopics and press the Insert key. Type a name for the new node. Press Enter to create siblings of the new node or Insert to create a child of the new node.

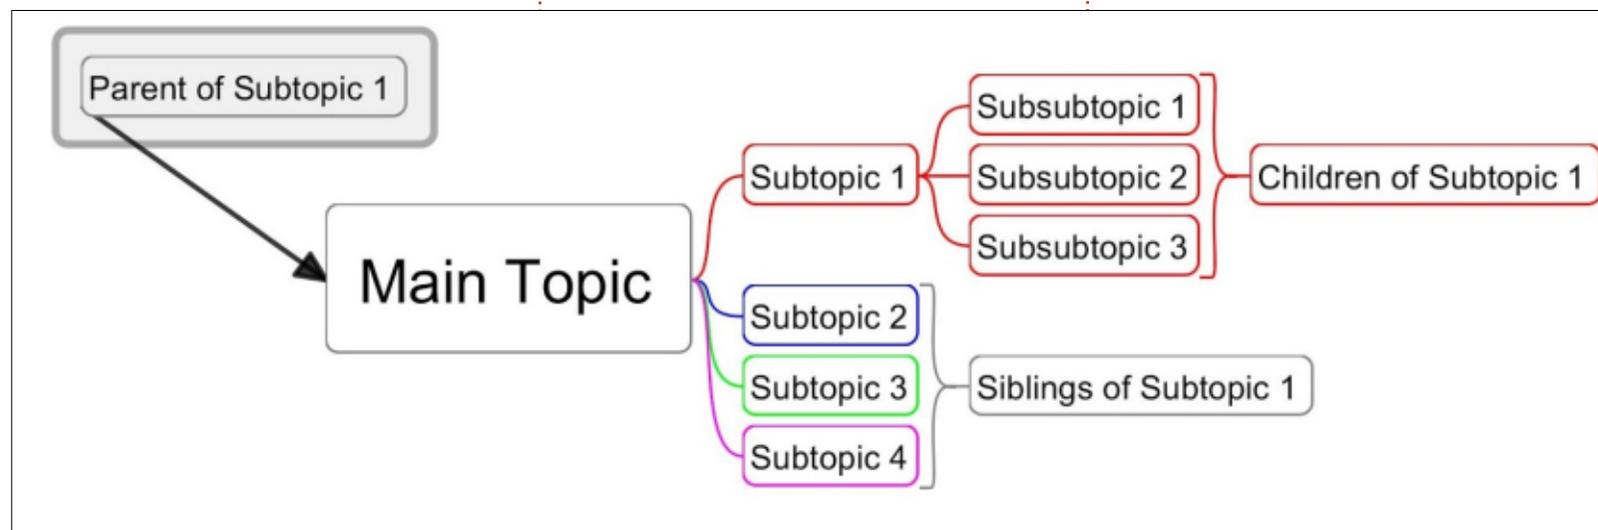
Branch out as far and using as many nodes as you need. The depth of your map is only limited by the needs of your mind and topic.

Next time we will look at some of the core elements of Freeplane, and do more with our map.



As I go through this series, I will explore the capabilities of Freeplane and the world of mind mapping. I will look at the possibilities and benefits of mind mapping. If you've never used mind maps before, I encourage you to join us for this series. It may change the way you process

information.



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



In the last issue, I showed you how to dim an LED with a potentiometer. In this issue, we will read the values of the potentiometer using the analog-to-digital conversion (ADC), and send these values to a serial console. This time, we need additional hardware – a USB-to-serial adapter (it is advised not to use the serial port on your computer as damage may occur to the microcontroller because of high voltage from the computer serial port).

Many USB-to-serial adapters will work. Common chipsets are CH340, CP210x, FTDI, PL2303, but any other chipset will work as well. Importantly, they must provide voltages at 3V or 5V, and a Linux driver. It would also be useful if the USB-to-serial adapter is 'breadboard friendly', which means it has pins to connect to it, whether directly or through cables, to the breadboard. You could use an Arduino as a usb-serial adapter, see <https://oscarliang.com/use-arduino-as-usb-serial-adapter-converter/> for more information.

SERIAL COMMUNICATION AND HARDWARE

A serial communication means that a byte of data is sent bit-by-bit through a dedicated receiving and transmitting line. See the end of this article for more basics of serial communications.

For the program flow, it is further important if the communication is blocking (the program must halt until the data is fully received or transmitted) or non-blocking (the data is buffered and the program need not halt until data is transmitted or received, then the data would be processed). Some microcontrollers have a Universal Synchronous and Asynchronous Receiver-Transmitter (USART) implemented. The ATmega328p (the one on the Arduino UNO R3) for example has a USART on PIN2 for receiving and PIN3 for sending serial data.

Great Cow BASIC has a solution to get the hardware USART to work with a few lines of code.

```
#chip mega328p, 16
#DEFINE USART_BAUD_RATE 9600
#DEFINE USART_BLOCKING 'optional, see non-blocking
mode in help file
DIR PORTD.1 OUT 'Transmitting line of USART
DIR PORTD.0 IN 'Receiving line of USART
DIM beer AS INTEGER
WAIT 500 ms
beer = 99
'Main program
DO
  HSerPrint str(beer) + " bottles of beer on the wall, "
  HSerPrint str(beer) + " bottles of beer."
  HSerPrintCRLF
  beer = beer - 1
  HSerPrint "Take one down ... "
  HSerPrint str(beer) + " bottles of beer on the wall."
  HSerPrintCRLF
  WAIT 500 MS
  IF beer = 1 THEN
    HSerPrint "1 bottle of beer on the wall, 1 bottle of
beer."
    HSerPrintCRLF
    HSerPrint "Take one down ..."
    HSerPrintCRLF
    WAIT 500 MS
    HSerPrint "There are no more left :("
    HSerPrintCRLF
    WAIT 500 MS
    HSerPrint "Went to the store..."
    HSerPrintCRLF
    beer = 99
    WAIT 500 MS
  END IF
LOOP
```

Additionally, Great Cow BASIC supports some functions to send and receive data in different data types. See the Great Cow BASIC help for more information, but, to

use the USART in the ATmega328p, and send some data, you would start with the following lines. (In real life you would, of course, send more meaningful data):



SETUP THE HARDWARE USART FOR SERIAL OPERATIONS

An example of using the hardware USART in the ATmega328p is shown on the previous page.

SETUP FOR SERIAL OPERATIONS USING SOFTWARE

Some microcontrollers do not have a hardware USART; which is true for the ATtiny13a I am using for these examples. Another reason could be the need of an additional serial connection without changing the microcontroller. In order to receive and send data, the serial communication could be achieved in software as well. For me, this would be quite a hard task to program, but, luckily, Great Cow BASIC has two methods for serial communication. For comparison, I will provide an example for each, and observe the differences in the two methods. Great Cow BASIC also has two methods for Serial operations using software, and

these are both shown below. For simplicity, I provide the examples for sending values over the serial line. If you want to receive data, go right ahead and add the few missing lines of code.

OPTION 1 - 'LEGACY SOFTWARE SERIAL METHOD':

This legacy method was implemented in 2007 - and is a popular method for software serial operations. This method (top right) can handle up to three different serial lines, with baud rates between 300 and 19,200.

OPTION 2 - 'OPTIMISED METHOD':

This method (bottom right) uses an additional library for Great Cow BASIC from Frank Steinberg, and is an optimised software serial driver written in portable assembler. Portable assembler means this optimised method supports AVR and PIC microcontrollers. This method provides higher baud rates (eg: an AVR clocked at 1 MHz can handle up to 28.800 baud), and the

```
#CHIP tiny13a, 1.2      'see my advice in the conclusion
                        section below
#define RS232Out PORTB.1

'Config Software-UART
#define SendAHigh Set RS232Out ON
#define SendALow Set RS232Out OFF
#define SerPrintCR
#define SerPrintLF

DIR RS232Out OUT
DIR PortB.4 IN
InitSer 1, r9600, 1+WaitForStart, 8, 1, none, normal
DIM POTI AS BYTE

DO
    POTI = READAD(ADC2)
    SerSend 1, POTI
    WAIT 500 MS
LOOP
```

```
#CHIP tiny13a, 1.2
#OPTION Explicit
#include <SoftSerial.h>      'include the library. This is
                             required.
#define ADSpeed MediumSpeed

'Config Serial UART:
#define SER1_BAUD 9600      'baudrate must be defined
#define SER1_TXPORT PORTB 'I/O port (without .bit) must
                             be defined
#define SER1_TXPIN 1      'portbit must be defined

DIR PortB.4 IN
DIM POTI AS BYTE

'Main program
DO
    POTI = READAD(ADC2)
    Ser1Print POTI      'send the value
    Ser1Send 13      'new line in terminal
    Ser1Send 10      'line feed in terminal
    WAIT 100 MS
LOOP
```

HOWTO - GREAT COW BASIC

program generated is small enough to fit on the tiny13a microcontroller.

If you compare the code examples, keep in mind that the ATmega328p on the Arduino comes with a 16 MHz Resonator, so the system clock is at 16 MHz. The ATtiny13a comes with a 9.6 MHz integrated systems clock; but, by default, the clock is divided (read: limited) by 8 to a mere 1.2 MHz. To change that, it would be necessary to change the fuses, which I did not want to do in this set of articles – I wanted to keep things simple.

SOFTWARE PREREQUISITES TO WORK WITH THE SERIAL ADAPTER

If you did not have to use a serial adapter, nor serial terminal software, on your computer, here are some brief explanations to get the serial communication to work. The chances are good that your computer has the necessary drivers right out-of-the-box:

Plug in your serial adapter to your USB-Port, and type:

```
dmesg | grep tty
```

You should see a line like this:

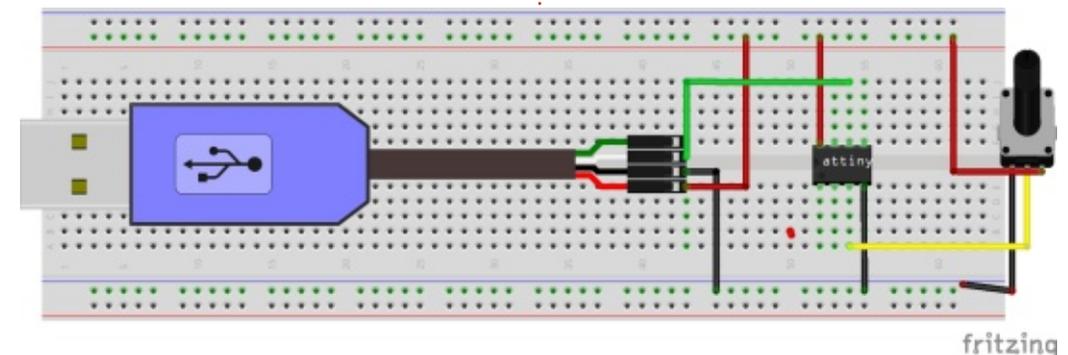
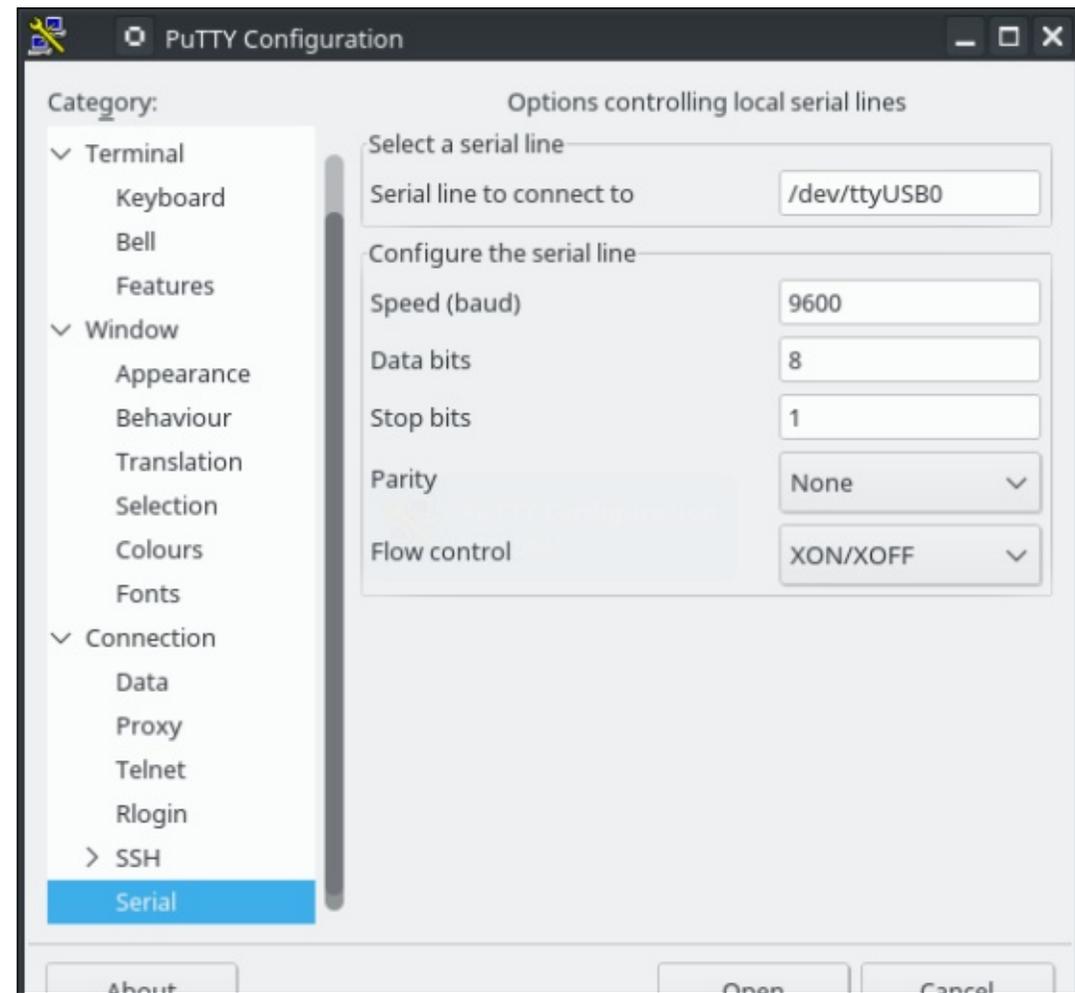
```
[...] usb 6-2: cp210x converter now attached to ttyUSB0
```

Now, you know the chipset of your adapter is supported (in my case it is a cp210x), and it was assigned as device ttyUSB0. Then give your Linux user the permission to use this device. Type:

```
sudo chmod 666 /dev/ttyUSB0
```

Give your superuser password if prompted.

Then, if you do not have a serial terminal installed, I suggest putty, as it is easy to use and easily installed. A `sudo apt install putty` is all you need to do. Then start putty and change the serial line to `/dev/ttyUSB0`, the baud rate according to your chosen one in the microcontroller program (I chose 9600 baud for the examples). Make sure the other parameters of the serial line are correct, scroll down the category tab and look for the entry Serial. Make sure the data, stop, and parity bit settings are correct. The entry flow control is not important for the time being - leave it as is.



HOWTO - GREAT COW BASIC

Then, click 'Open' - and you should see the awaited data scrolling along the terminal.

BREADBOARD CIRCUITRY

Compile the program and flash it to the ATtiny13a with your hardware programmer of choice. The USB-to-serial adapter should have +5V or +3V, and GND, to power up the microcontroller. Then, hook up the receiving line from the adapter to the sending PIN (PB1) of the ATtiny13a. Hook up the Poti to +5V or +3V, and GND as well, and the middle pin to the PIN ADC2 (PB4).

CONCLUSION

After some experimentation, it shows that the legacy software serial method produces a program that is large - too large for the tiny13a; The version of my code shown above compiles to 1.016 Bytes. Why is this a problem? First, you have no more program space left for anything else. Second, and more importantly, the readings on the terminal were not human readable as we are limited to the raw values being displayed. I would suggest, if using the legacy serial

method, to utilise a microcontroller which has more program flash memory, eg, the ATtiny85.

But, all is not lost! Serial communication on the ATtiny13a is still feasible. The optimized serial software compiled to around 368 bytes and worked like a charm. With the optimised software, we have a working solution to get serial I/O to work; even with the small program memory device like the ATtiny13a!!

Serial communications is an important method for exchanging data. Great Cow BASIC makes the configuration and use very simple. Serial communications is easy - it is useful for debugging and to communicate key messages to the users. It can also be used for inter-microcontroller communications, albeit there are better methods like I2C and SPI.

GITHUB REPO

Evan Venn (Anobium...my mentor with respect to these articles) suggested to open a GitHub repository for the source code, and, I must admit, this was a

good idea. If you want to download the sources instead of copy-pasting it, you can now check it out with git or an SVN client. See <https://github.com/Anobium/Great-Cow-BASIC-Demonstration-Sources/tree/master/Publication%20Solutions/Full%20Circle> for more information.

REFERENCES

serial communication
<https://www.teachmicro.com/microcontroller-serial-communication/>

ACKNOWLEDGEMENT

I wish to thank Evan Venn (Anobium) from the Great Cow BASIC Team for his insights and valuable hints.



Boris holds a bachelor degree in business administration and works for an insurance company. While not working, he is a family person and enjoys playing with his kids or tinkering with his personal projects. Contact info and additional material at his site: <https://www.evil-publishing.de/fcm>

CORRECTIONS

In the last issue a formula got lost, for completeness:

$$ADC \text{ Value} = \frac{Volt}{255}$$

SERIAL COMMS - A FEW DEFINITIONS AND NOTES:

by Michael Kennedy and Gord Campbell

These notes focus on some of the main features of the common RS-232 implementations, but a few other variations are also mentioned. They're intended to supplement the story in FCM #130 (etc) on Great Cow BASIC, by Boris.

"RS-232" (Recommended Standard 232) is a popular communications system. Originally (1960), it was used for connections with old-style telephone/terminal/typewriter systems.

The simplest RS-232 implementations handle two-way communications between just two

devices (say, Dev-A and Dev-B). The cabling has just three wires: ground, data pulses going from Dev-A to Dev-B, and data pulses from Dev-B to Dev-A. The cable is normally up to about 15-20 metres, but could be up to a few hundred metres, max. The hardware at both ends could be microcomputers, or dedicated, relatively smart, RS-232 chips, or just suitable line drivers - where the entire implementation is handled in software.

The speed of the data-pulses must normally be set, in advance, on both devices, and must be based on how quickly both devices can transmit and receive the pulses, on the cable length and quality, the overall electrical environment, etc. Typically, it might be set at 300 baud, 1,200 baud, 2,400, 4,800, 9,600, 14,400, 19,200, or maybe even up to 115,200 baud and beyond. A "baud" refers to the width of each "bit/pulse" on the cable, so 4,800 baud is 4,800 bits per second.

To send a single "byte" from Dev-A to Dev-B:

- Dev-B must be ready to accept the byte (obviously!).
- Dev-A sends out one "Warning"

pulse, called a Start-Bit, to wake up Dev-B!

- Dev-A then sends the byte's bits. Usually eight, but sometimes maybe as few as five are adequate.
- Dev-A might send a single Parity bit (Odd/Even/None).
- Dev-A sends one, or one-and-half, or two, "Stop"-bits.

So, typically, a single byte might involve the transfer of about 10-12 bits/pulses. And in that case, a speed of 4,800 baud might result in a maximum transfer of about 400-480 bytes per second - assuming there's no wasted time between each "byte" - which would never be reached in reality!

UART: The above approach is termed "Asynchronous", because there can be an indeterminate amount of time between bytes. The devices are termed UARTs (Universal Asynchronous Receiver and Transmitter).

SIMPLEX/DUPLEX: Where data pulses are allowed on only one of the two data-wires at any one time, this is defined as Half-Duplex. If Dev-A may send data to Dev-B, and Dev-B may send data to Dev-A, at the same time, this is Full-Duplex. There's also a "Simplex" approach,

where only one device does all the sending, and the other all the receiving - with just a single "data"-wire, for example, a wireservice sending stories to a newsroom.

Further, less common, variations on RS-232...

USARTs: Two wires could be used to send data-bits from Dev-A to Dev-B, and two more wires for the data-bits from Dev-B to Dev-A. The second wire (of each pair) contains clock pulses - that match the data-bits on the first wire. Then, the extra Start/Stop bits may not be needed for each byte, and it may not be necessary to "lock" the two devices into matching speeds - because the clock pulses determine the widths of the data-pulses. This approach should result in slightly higher communications speeds and higher quality. The devices at both ends are termed USARTs (Universal Synchronous and Asynchronous Receiver and Transmitter); they can handle both protocols.

Additional Control wires: In some implementations, extra wires are used ([D]CD Carrier Detect, RTS Request to Send, CTS Clear to Send, DTR Data Terminal Ready,

DSR Data Set Ready, RI Ring Indicator, RTR Ready To Receive, etc), which are used by each device to tell the other device that they are Busy, Ready, Unavailable, etc. Often called "Handshaking" or "Flow-Control" signals.

RS-422, Longer cables: Slightly different devices and cabling arrangements can be adopted where longer cable runs, and/or higher speeds, are needed. The cable contains two wires for each of the data-pulses runs, and two more wires for each of the clock-pulses - like standard twisted-pair LAN cables.

RS-485, Multiple devices: Different ICs can be deployed where, for example, more than two (maybe 10, 20, 30...) devices can share the same cable(s), and, for example, all of them can talk with a single "server/master" device.

Error Handling!: If you're writing "serial" software, and/or taming serial hardware, for private educational or entertainment purposes, then you probably won't be over-worried if a few errors arise now-and-then. For example, if you're sending data, serially, to a display, or controlling a few LEDs,

etc, you might chose to not bother if the device is present or not, or powered on, or fully “reset” after a power-on, or ready to send/receive serial data, or that its serial configuration matches ours, or that it’s faulty, or has a hundred other issues....

However, if your RS-232 app is used in an MRI/CAT-scanner, and it’s controlling doses of radiation, positions of motors, etc, and you’re the patient, then you’ll definitely have to handle all of the above conditions! Then, you’ll probably have to use some of the additional hand-shaking signals/wires that I mentioned above, and you’ll probably have to introduce a whole “protocol” and timings to control all your comms. Eg, when a sender/transmitter wishes to send some data to a receiver, then, minimally:

- Both devices would use the available control-wires to ensure that both are ready-and-willing to communicate.
- Dev-A says: “Hello, I wish to send you an alert/reading... Are you ready, and can you take it now?”.
- Dev-B: Yes, send it on.
- Dev-A: Here it is..... xxxxx....
- Dev-B: Got that, 123 bytes in total, checksum=xxx, thank you.

- Dev-A: Excellent – thank you, over-and-out.

At all of the above stages:

- The sender of each packet might have timeouts running, to check if the transmission choked...
- The receiver would have timeouts running, to ensure that the packet arrived on time, and fully...
- All those hand-shaking lines would be continually checked, at both ends, to ensure that nobody cut a cable, or switched the power off...
- If minor errors arose, retries would probably be attempted, automatically, for a while...
- The handling of all major/catastrophic errors would have to be very carefully designed, and implemented, and tested.

So, to ring that bell, or turn on that emergency light, or rotate that stepper motor, a single “Print.....” statement might evolve into a very complicated 500-line “LED_On()” function! ;-)

Finally: For lots of good, general, details, see: <https://en.wikipedia.org/wiki/RS-232>. And do check all the other 18,000,000 websites ;-)

If you’re implementing RS-232 comms on a specific device (UART, uP, CPU, etc), then look up the details in the manufacturer’s datasheets of those devices. Normally, the devices will work exactly as per the specs. However, if you’re stressing the devices, and/or writing low-level drivers, etc, you might sometimes observe “inexplicable” results. When stressed, some of these devices have nasty bugs, so you should run some web searches, and, hopefully, others will have posted explanations and suggestions.

Good Luck!

Michael Kennedy and Gord Campbell



HOW-TO

Written by Mark Crutch

Inkscape - Part 70

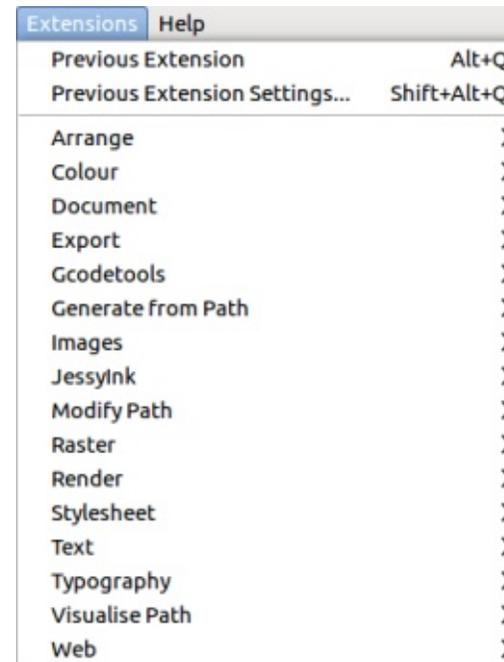
With 69 previous instalments of this series behind us, it's clear that Inkscape has a lot of features and functionality, despite being slightly shackled by the limitations of the SVG file format. But there are some tasks which don't warrant inclusion in the main Inkscape application. To support these, the developers added a simple extension mechanism to Inkscape, which allows it to hand the document off to an external program for further manipulation, and receive the modified document back in response.

The program that performs the modification could be anything from a small shell script to a fully blown, compiled C++ application. In practice, most take the form of simple Python scripts. Because the program just needs to manipulate a text document, just about any language could be used, but the output must still be a well-formed SVG file, so using a language with good XML libraries certainly makes the job of writing an extension easier.

Because extensions are

external programs, they can be shipped independently of Inkscape. If a manufacturer wants to create an extension to convert documents into the right format for their plotter or vinyl cutter, they can do so. With the right programming skills, users can even create their own extensions. But before considering third-party extensions, it's worth looking at the ones that ship with Inkscape by default.

Extensions can be found, unsurprisingly, under the Extensions menu. Here you'll initially find a couple of shortcuts – the first to re-run the last used extension with the same settings you used previously, and the second to open the settings dialog for the last used extension, if it has one (otherwise it also just re-runs the extension). The rest of the menu is taken up with sub-menus, each of which holds menu items to launch the extensions themselves – and, in some cases, further sub-menus before you get to the actual launchers.



If you take a couple of minutes to scroll through the available extensions, you'll realise that there are a lot of them. Over 150 of them on my default installation of 0.92! That should keep this column full for the next few years, so let's press on with the first one...

...or maybe not. I'm not so cruel as to go through each and every extension in detail, but will pick out a few examples in order to show the common UI features that you'll find. Since extensions are

just normal programs, they can accept parameters, in the same way that running a tool on the command-line often requires additional arguments. The exact arguments that are needed are defined in a configuration file for the extension (this also includes other details, such as which sub-menu to put the launcher in). This file defines not only the names of any additional arguments, but also the type of value the argument expects. This allows Inkscape to generate a simple dialog, ensuring that the right type of UI widget is used for each parameter. You can tell in advance if an extension will prompt for additional parameters by looking at its name in the menu: as is the convention in computer programs, entries ending in an ellipsis ("...") will display a dialog, while those without will have an immediate effect.

As an example of an extension with no UI, let's consider the Color > Brighter extension. As you might expect this makes the selected objects brighter, which it does by altering their fill and stroke colors

HOWTO - INKSCAPE

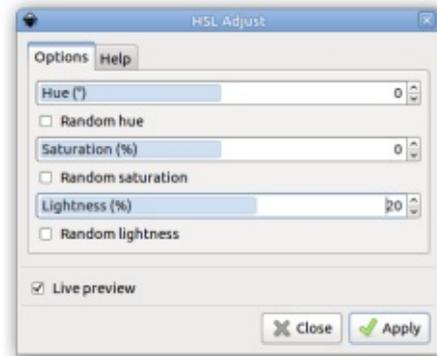
– albeit by only small amounts at a time. Here’s a before-and-after image, having applied this extension many, many times to the object on the right:



Because this extension has no UI, its effect takes place immediately, so you might think that using the “Previous Extension” menu entry, or, better still, its keyboard shortcut (ALT-Q) would be a fast way to incrementally increase the brightness of an object. Unfortunately using an extension de-selects everything in your drawing, so repeatedly calling the same extension also involves an intermediate step of re-selecting the object you wish to operate on. In many cases, including the Brighter extension, failing to select an object results in the effect being applied to every element in the drawing. If you forget this, and just hammer ALT-Q a few times you’ll quickly find that everything gets brighter, not just the element

you had selected when you first ran the extension.

A better approach to achieve a similar effect is to use the Color > HSL Adjust... extension. As the ellipsis indicates, this extension displays a user interface, so you can adjust the amount of brightening you wish to apply before the extension actually runs. Furthermore, recent versions of Inkscape add a “Live Preview” checkbox to the extension dialog, allowing you to see the effect of your changes before they’re finally applied.

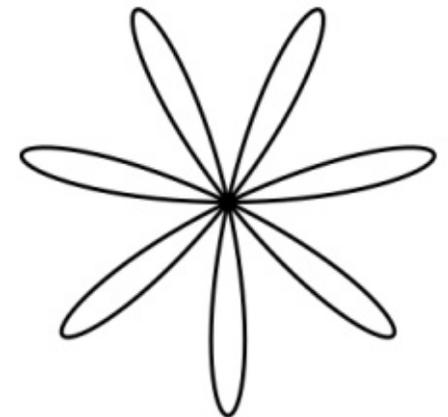


On the surface it might seem like the “Live Preview” option is a bit of a no-brainer. Why would you ever not want that checked? But consider that each extension is a separate program that needs to be launched, receive a copy of the entire Inkscape document, process it, return the entire document back to Inkscape, and then close. And this process will happen for every little change you make to the settings in the dialog. With a large document or a complex extension it can take several seconds, or even minutes, to preview the changes. Un-ticking the checkbox lets you change several settings at once without this overhead. If you already know the values you need to enter, or just want to preview after making a number of changes, being able to turn off this checkbox is vital.

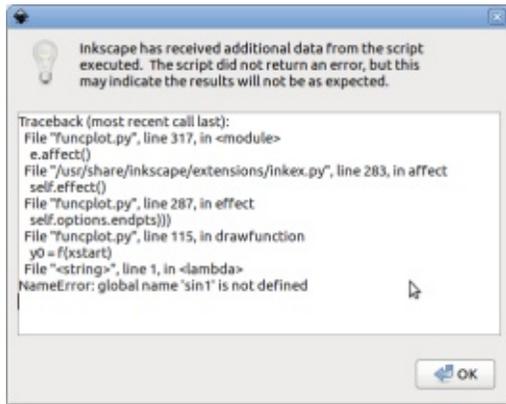
This dialog also shows a few other items worth noting. Inkscape has displayed the boolean parameters as checkboxes, and the numeric parameters as the GTK3 style hybrid spinbox and slider that is used elsewhere in the program. Furthermore, the sliders have different ranges – the Hue input runs from -360 to +360, whereas the Saturation and Lightness run

from -100 to +100. By using the right field types, and limiting the possible values, well written extensions can ensure that users are protected from entering nonsense that the extension will then have to deal with.

But Inkscape’s selection of UI widgets is limited to a few basic types, and doesn’t even allow the developer to specify a template or regular expression to validate free-text fields. For most extensions, this isn’t a problem, but some do have specific requirements for the data that you enter into their fields. For example, Render > Function Plotter... lets you draw mathematical curves by entering a function into a text field. Typing “sin(x*7)” and setting the checkbox to use polar coordinates, for example, renders this seven-lobed flower.



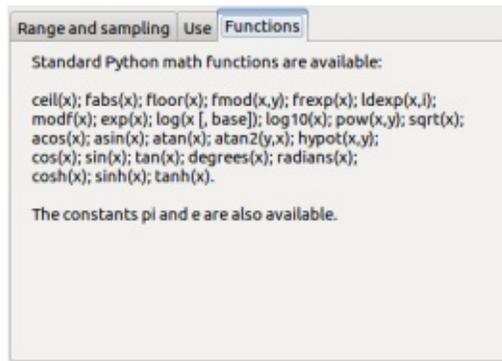
Type an invalid function into the input field, however, and you're likely to see something like this instead:



The problem is that Inkscape has no way to validate the input, and the extension has decided to send the whole Python error back to Inkscape, rather than a more friendly "The formula you entered is not valid". When you are working with a free-text field that requires specifically formatted content, it's worth disabling the live preview until you've finished editing your data.

One other thing to look out for is a Help tab, or similar, which will often contain further information about what sort of values are valid in the fields. The Function Plotter, for example, has a "Functions" tab

which details the Python mathematical functions that can be used in the text field:



Whilst experimenting with different settings in an extension's dialog, you'll probably discover that you can't zoom or pan the canvas, or change the selected objects. This is particularly frustrating when using something like the Function Plotter, as some combinations of parameters can lead to shapes being drawn that are too large or small for the current zoom level. The cause of this restriction is actually the preview mode – simply un-tick the Live Preview checkbox and you'll be able to make changes on the canvas, before ticking it again to restore the preview.

Once you've finished playing with the Function Plotter, it's worth exploring some of the other

extensions in the Render submenu. Given the ubiquity of barcode scanning software on phones now, the Render > Barcode > QR Code... extension could be a useful tool when designing a poster, flyer or other promotional material.



This menu also includes extensions to draw calendars, grids (including logarithmic and polar grids, which Inkscape can't do natively), printer's registration and color marks, and charts (though you're probably better off using a spreadsheet or dedicated graphing program for anything other than the simplest of charts). On the more frivolous side of things, there are also extensions to create spirograph-style images, simple fractal trees, and fake 3D objects.

The last extension to look at this month is one that you'll find either absolutely invaluable, or you'll probably never need to use at all. It's also one of the extensions with a less than descriptive name: Render > Hershey Text. What this extension does is to render some text using a so-called Hershey font. At first glance, this might not look terribly different from normal Inkscape text in a similar font:

This is normal text
This is Hershey Text

Zoom in, however, and things start to look a little different.

This
This

Switch them both to a minimal stroke, and no fill, and the

difference really becomes apparent:



The image shows two versions of the word 'This'. The top version is a clean, simple outline with smooth curves. The bottom version is a Hershey-style outline made of straight lines, with a small square dot on the 'i'.

As you can see, the normal text has a clean and simple outline with curves where needed. Conversely the Hershey text is made up of straight lines that sometimes overlap awkwardly – and just what is that little square doing in the dot on the i? The answer is that Hershey text is intended for use with pen plotters, laser engravers, and similar devices.

Suppose you're preparing a file for use with a laser engraver, but you want some solid text. Such devices really care only about outlines, not fills, so using a normal font would just result in an outline version of your text. If you wanted to fill it, you would need to draw a hatch pattern inside the text, resulting in the laser going

over adjacent parts of your design repeatedly. At best, this might cost you more, as the job will take a lot longer. At worst you might find that these parts of the design become excessively scorched by keeping the beam in the same area for so long.

Hershey text, on the other hand, is not designed to be used with a fill. Rather the letters should be left as outlines, with the thickness of the beam or pen being used to provide any "fill" as the outline is drawn. With that knowledge, the little square makes a lot more sense.

If you print your designs using only an inkjet, laser printer, or even a professional printing press, you'll be fine using normal text. But if you decide to use one of the increasing number of laser cutting and engraving bureaus that accept Inkscape files, or if you purchase one of the hobbyist pen plotters that work with the program, this little extension could save you a lot of time, money or ink.



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

**FULL CIRCLE
WEEKLY NEWS**



Join our new hosts Wayne and Joe as they present you with a short podcast (<10min) with just the news. No chit-chat. No time wasting. Just the latest FOSS/Linux/ Ubuntu news.

RSS:

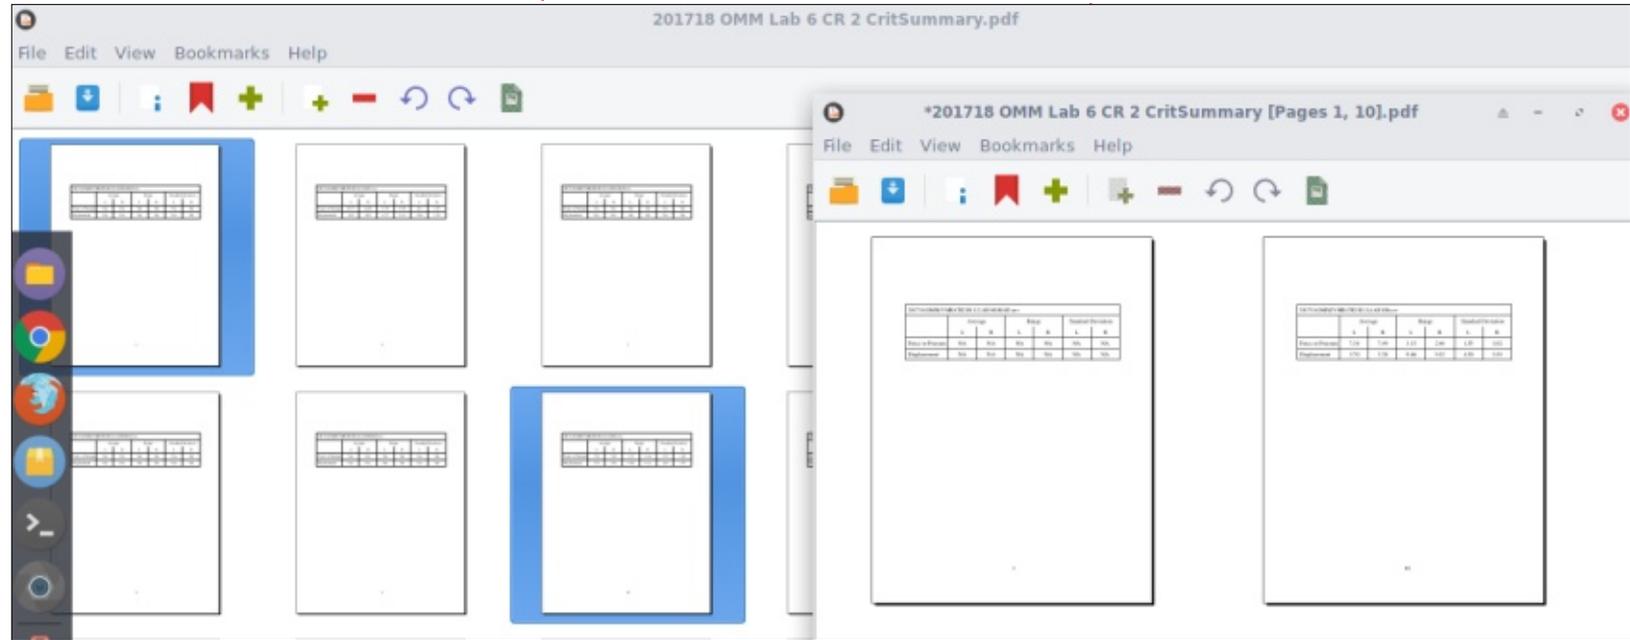
<http://fullcirclemagazine.org/feed/podcast>





This past month I needed a simple but effective pdf extractor. This past month I finished processing a learning lab data set. My Perl batch files can effectively gather my finished data files, and transition it into a simple statistical plot. The file is outputted as large PDF file as seen below.

Yet I need to take the pages of this file and send it to the learning lab students. This lab is a simple review of therapeutic manipulation employed by student doctors. So I needed another way to send the students their data. The data needs to be separated from the bulk file. I wanted an easy and reliable program to remove the pages and create a new PDF. I wanted a simple point, click, and save program.



I found that program, and it is called PDFmod. The program is available for Ubuntu and Fedora. I installed it via dnfdragora.

It has an extremely simple interface. You open a file and the

pages populate in the field within the box. The pages are simple thumbnails.

However when you click on the thumbnail, the PDF page

populates into a larger window. I knew what pages I needed, so simply point and selected. The pages become blue, and by selecting a small green plus a new PDF is created. You would then go and save the file as needed. I have not reviewed the other options or

functions to PDFmod. Yet I positive this program can be used by anyone. I am sure PDFmod will become another tool I use for research.

201718.OMM.VMB.CR2.R1.L2.AE101.csv						
	Average		Range		Standard Deviation	
	L	R	L	R	L	R
Force or Pressure	10.55	14.48	8.79	8.90	4.11	4.72
Displacement	4.24	4.89	6.37	6.43	2.60	3.32



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.



So, the first step on the path to Linux joy is to install your distribution. If you're reading this magazine, it's very likely you've already chosen Ubuntu, Kubuntu, or some Ubuntu derivative. There are many, many articles on installing Linux, and it's ground pretty well-trodden by now, so we'll skip ahead just a little. Once you've got Linux itself installed,

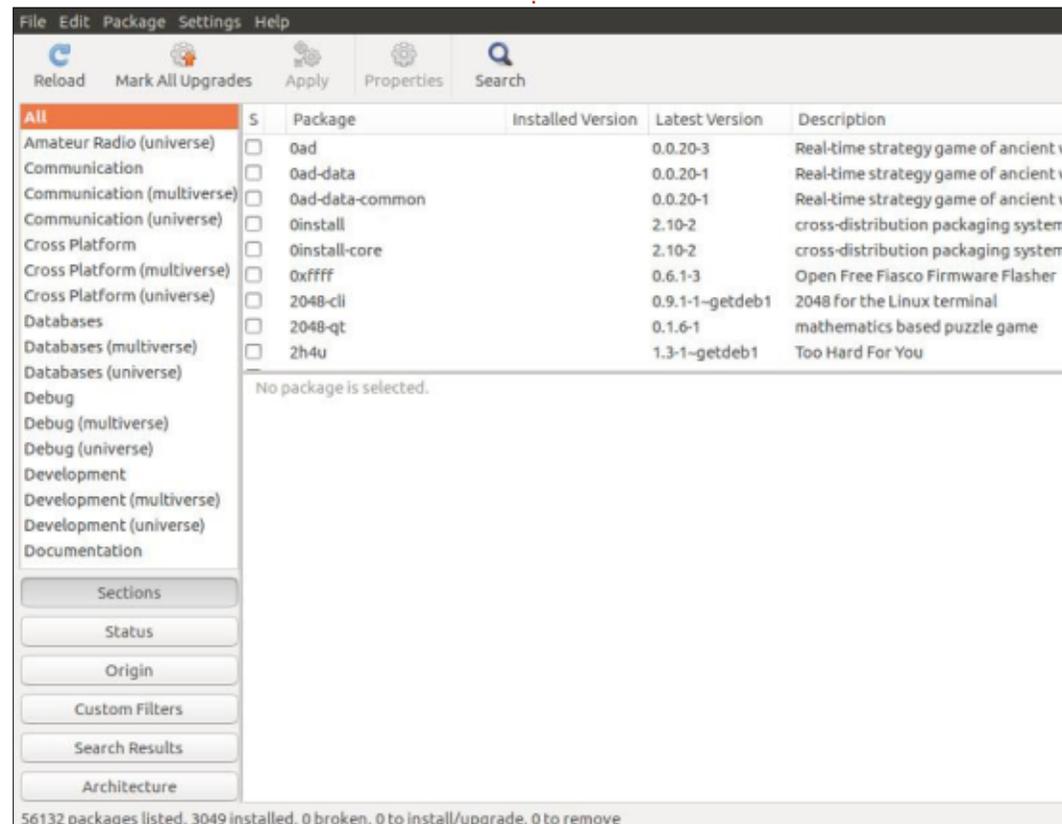
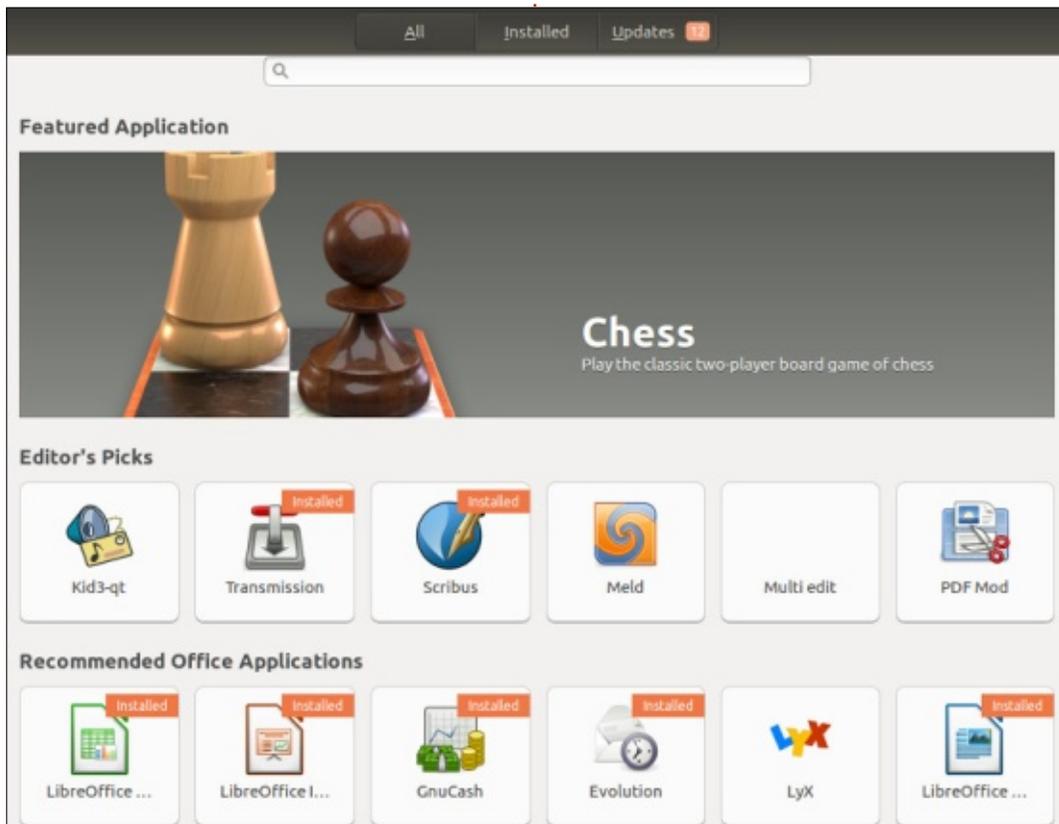
what next?

Well, Linux all by itself is obviously not all that useful - you need software to accomplish much of anything. But installing software itself is not quite as straightforward as it is on Windows, Mac OS, or Android. Your first resource will probably be the Ubuntu Software Center (below):

While the Ubuntu Software Center is well-designed and capable, I personally find it's missing a lot of packages I want to install, so one of the first things I install on every Linux system I set up is Synaptic Package Manager. Synaptic and Ubuntu Software Center together cover most of what I need to install. I like the interface of Synaptic a little better,

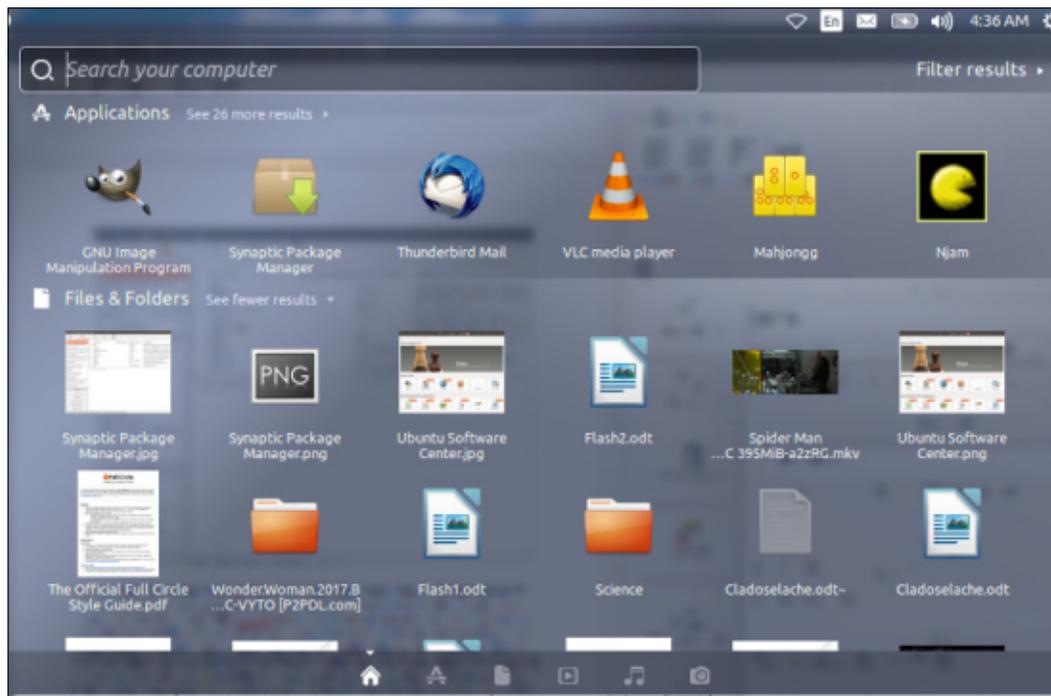
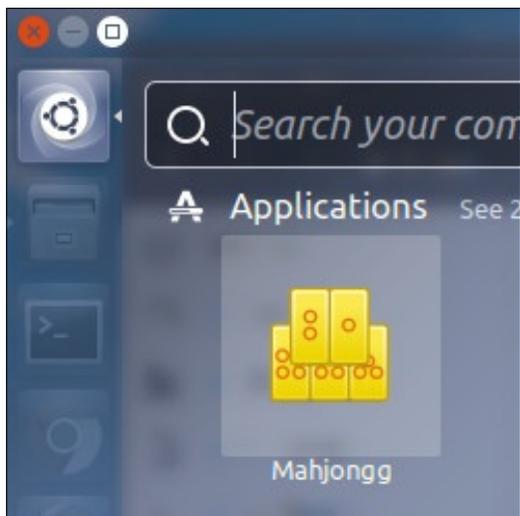
and find it more intuitive, but your mileage may vary. Once installed, Synaptic Package Manager looks like the image shown below right.

There's another great resource built into Ubuntu that I really like for software installations; it's the apt-get command. Modern Linux distributions are much more functional through the GUI's



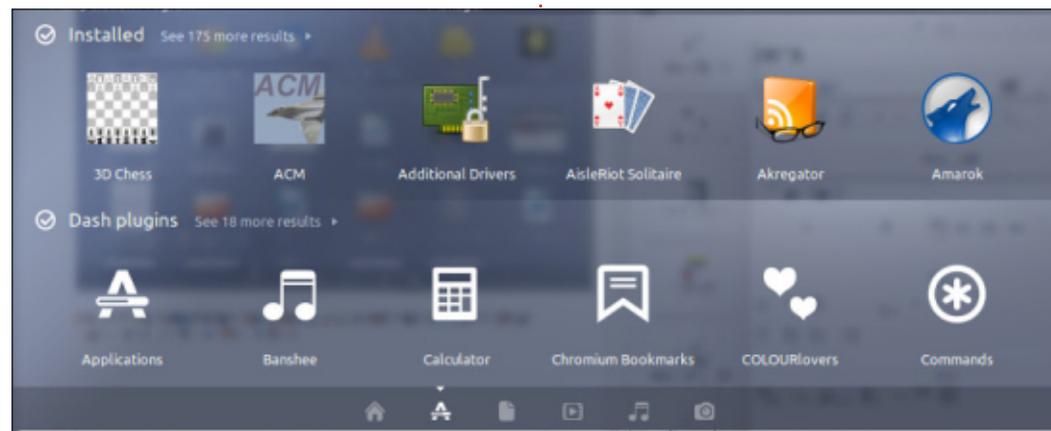
EVERYDAY UBUNTU

(Graphical User Interfaces) used for desktop environments than Linux used to be, but there are still some things that can be done more easily, or with more power and control, through the CLI, or Command Line Interface. This is accessed through the software 'Terminal' (although there are other ways to get to the command line, this is the commonest). You can go to the Dash in Unity to find the terminal (we'll also talk about the Dash more in a later column, but, for now, just note that it's a main, key component to the Unity interface). The Dash will be at the top of the default control strip (Launcher) on the left-hand side of the screen. It has the Ubuntu 'wheel' design, seen here:



Click on it. You'll see the 'dash' shown above.

The 'A' at the bottom is for Applications. Click it, then click 'Installed' to see all the applications on your machine.



The Terminal should be listed (the list will sort alphabetically by default):

Click on the Terminal to launch it. Before you do that, though, it's an excellent idea to add the

Terminal to the Launcher, if it's not already there by default. Click and hold the Terminal icon, then drag-and-drop it leftwards onto the area of the Launcher where you want it to be. Of course, you can also do this with other applications that you use frequently; it will probably be one of the first customizations you want to do to your desktop. If you have an idea which applications you're likely to use most often, go ahead and drag them onto the Launcher.

Once you are in the Terminal, you'll see a prompt that will say:

```
username :~$
```

Where username is the user name you set up on the machine during initial installation. At the prompt, type:

```
sudo apt-get install njam
```

then hit ENTER. Terminal will prompt you for the administrative password you set up during initial installation - type it and hit ENTER. Confirm that you do want to download and install "Njam". Now apt-get will get and install the latest version of "Njam", a fun PacMan-style game. Most of the

software names are fairly intuitive and easy to guess for apt-get (try `sudo apt-get install gimp` or `sudo apt-get install libreoffice`, for example). These will install The Gimp photo editing software, or the LibreOffice office suite, respectively.

The `sudo` part of the command is an abbreviation for 'super user do', and tells Linux to allow your user account to do things normally restricted to administrator, or super user, accounts. Sudo is a command you will probably use often in Ubuntu, so make note of it.

I like the apt-get command because it is easy and fast (or, put another way, quick and dirty). It does a good job of checking for dependencies, which are programs or software that a given application needs, other than itself, to run properly.

Finally, it's a really good idea to set up repositories to aid in future installations. Repositories are online resources for Linux that store software. Ubuntu will default to using its own repositories, which are very good, but you may want to add some more. Some common

repositories in use include:

Playdeb for games:
<http://www.playdeb.net/updates/Ubuntu/17.04#how to install>

The page has complete instructions for installing the repository, but just downloading and installing the DEB package is the easiest way. Save the DEB to a folder on your machine, then run it once the download is completed.

Getdeb for other applications:
<http://www.getdeb.net/updates/Ubuntu/17.04#how to install>

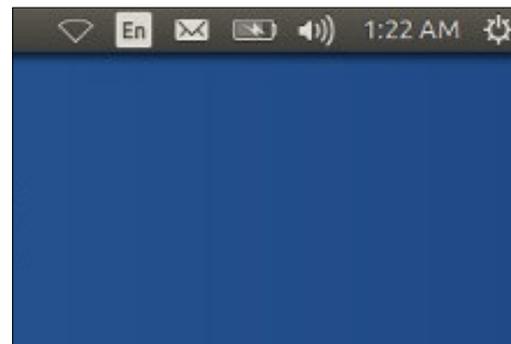
Also with a DEB package that can simply be downloaded and installed.

There are other repositories that allow you to install software that may not be freeware, or may not be open source, or that simply aren't in the default Ubuntu repositories. You'll have to decide for yourself if it's important for you to use only completely open source or freeware software.

Other repositories are added in Unity, through Software and Updates, in the System Settings. To get there, click the 'gear' icon at

the top right:

In the menu that opens, click System Settings. This will bring up System Settings (much more on it in a later column):



Click Software and Updates, then the Other Software tab, and you'll see this screen:

Click Add and you can add the APT line (the repository's location) for any new software repository. You'll have to look to the hosting website for the APT line. Many will require a signing, or GPG, key, as well - look on the hosting site for specific instructions for adding a key, if necessary.

Next time: Suggested software applications.

I invite feedback on easier/better ways to do things. Any such submissions in response

to articles or content will be considered the property of Full Circle Magazine for publication purposes, without remuneration, unless the writer/commenter specifies otherwise. That said, commentary and feedback are heartily encouraged and appreciated, at acer11kubuntu@gmail.com.



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.



Last month, I retweeted a link to an It's FOSS article by Abhishek Prakash on How Much Swap Should You Use in Linux? <https://itsfoss.com/swap-size/> The short answer put forth in the article was "it depends on a number of factors."

This article got me wondering about how much RAM and swap would I use on an average day at work? A few details about my work machine: Intel Core i7-3770 @ 3.9GHz, 28GB RAM, Intel 530 180GB SSD + Seagate ST1000DM003 1TB hard drive.

For the test, I ran Linux Mint Cinnamon 18.3 – with the latest updates. I started by rebooting the machine to eliminate anything that might be running in memory. At the LightDM login prompt, I switched to a virtual terminal and ran `cat /proc/meminfo`, the result was:

MemFree 28,225,148 kB
SwapFree 23,999,484 kB

Before going further, it's worth mentioning that at no time during

the testing did SwapFree ever change. This seems to confirm Abhishek's theory that, with a lot of RAM, you might not ever touch swap space.

I started by logging in and loading the Cinnamon desktop. I had my phone connected to my PC in charging mode (since boot), but nothing else was running. I ran the same `cat /proc/meminfo`. The result was:

MemFree 27,752,840 kB

Just over 472,308 kB difference. Firefox is undoubtedly

one of the most used applications on Ubuntu, so I started by launching it. I'd heard in the past that Firefox tends to "leak" RAM. With Firefox launched and sitting at the default page, RAM usage was as follows:

MemFree 27,183,420 kB

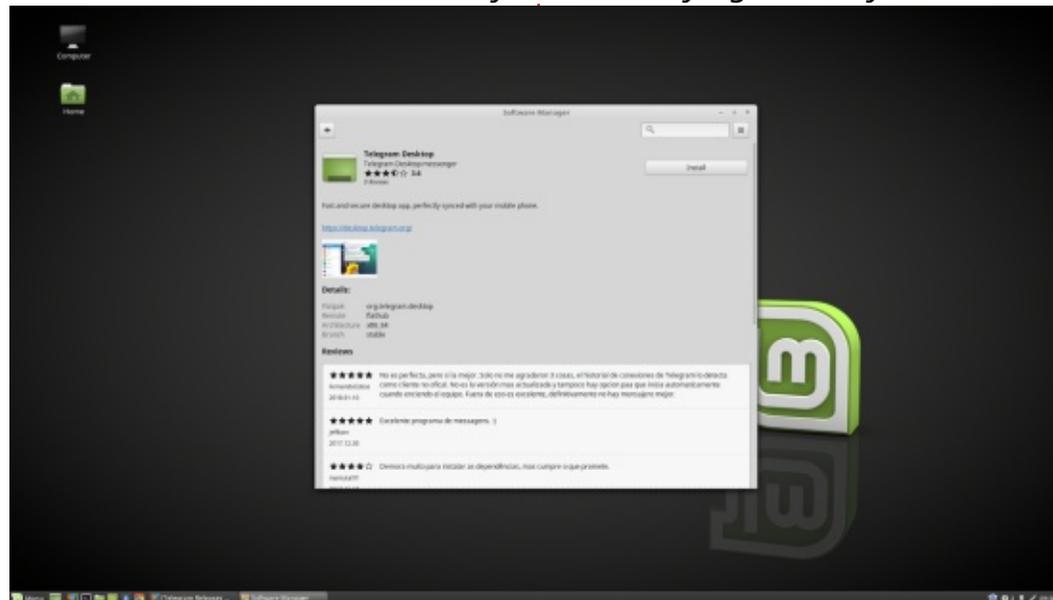
This is a 1,041,728 kB difference. On a system with 28GB of RAM, this difference is barely noticeable, but on a system with only 2GB of RAM, the RAM available to the rest of the applications and operating system is already significantly reduced.

Keeping in mind that a little less than half of that number is the Cinnamon desktop, it's an argument for using a lighter-weight desktop like LXDE, iceWM, or Enlightenment - at least on older hardware. It also illustrates the fact that modern applications simply need a bit more RAM.

Of course no one opens Firefox and just lets it sit at the home page. I opened 5 tabs with the following web pages: my own blog (some photos, no videos), the Full Circle Magazine website, Slashdot (remember when that was a thing), Distrowatch, and OMG Ubuntu UK. Memory usage was as follows:

MemFree 26,620,820

Keep in mind that, at this point, there are no other desktop applications open, only Firefox with 5 tabs. Memory usage ran up to 1,604,328 kB or 562,600 kB more than a single tab opened. It occurred to me, while opening all the tabs, that some websites have a lot more content than others: Yahoo, for example, tends to have



LINUX LABS

a lot of ads and Javascript. I closed all the tabs and opened only Yahoo. It took almost a minute of waiting on a fast connection before the Yahoo page completely loaded (all scripts). The amount of RAM required just to load Yahoo alone is a bit astonishing:

MemFree 26,736,172 kB

In other words, Yahoo by itself takes 447,248 kB, roughly the equivalent of 3-4 tabs. I picked Yahoo on purpose because I see a lot of people using Yahoo as a starting page (or for email). I'm not sure why people pick such a heavy site as a home page in this day and age, but it still seems to be relevant for some people.

Earlier in the article, I mentioned that I had heard Firefox "leaks" memory. It seems reasonable to assume that all the RAM doesn't get flushed when you simply close an application. On closing Firefox the amount of RAM available was:

MemFree 27,452,500

When I first logged in to Cinnamon, the memory free was 27,752,840 kB. Loading Firefox

with several pages, then closing it, seems to have eaten an extra 300,340 kB, confirming that yes, applications seem to eat a bit more RAM even when they're closed. I confirmed no Firefox processes were running by running `ps aux | grep firefox`. The result was my own `grep` session, no instances of Firefox loaded.

So what about that other magical browser from Google:

Google Chrome? Google Chrome (not Chromium) took up a bit more RAM than Firefox. The `cat /proc/meminfo` result was:

MemFree 27,024,536 kB

I figured some might argue the point that other applications had already been loaded into RAM and this RAM was being used still by the residue of those applications so I rebooted the system once

more, opened a virtual terminal, and ran `cat /proc/meminfo` once more. This time (before logging in to Cinnamon) the memory usage was slightly less:

MemFree 28,232,364 kB

I loaded up Google Chrome to the default page and ran `cat /proc/meminfo` again and got:

MemFree 27,324,268 kB



Indeed Google Chrome seems to be a bit lighter on memory usage (than Firefox), coming in at 908,096 kB with the default site opened.

Before I rebooted and tested Chrome, I decided to test memory usage with multiple applications opened. I opened Firefox to the Full Circle Magazine website. I then opened the Mint Software Manager, the Nemo file manager, the Background (wallpaper) application, GIMP (with a screenshot I'd previously captured), and set VLC to stream an Internet radio station. With all these applications opened, I switched to the virtual terminal and ran `cat /proc/meminfo`. The result was:

MemFree 25,353,268 kB

On a machine with only 2GB of RAM, we would be well into the land of swap (2,871,800 kB), but, on this work machine, the RAM usage barely registers a blip. I closed all applications except VLC (which was still streaming audio) and memory usage dropped:

MemFree 26,392,388 kB

At this point, I wondered if the issue was VLC or whether it was all remnants of Firefox and the other applications? And if it was these other applications, how could I "flush" the RAM without restarting (either Linux or the Cinnamon desktop)?

I found an article on Techmint on How to Clear RAM Memory Cache, Buffer and Swap Space on Linux at

<https://www.tecmint.com/clear-ram-memory-cache-buffer-and-swap-space-on-linux/>.

The Techmint article outlines the way to clear PageCache, dentries, and inodes, and all three together, but it doesn't really outline what dentries and inodes are, though it does warn not to clear dentries and inodes on a production server. To clear only PageCache, the article suggests

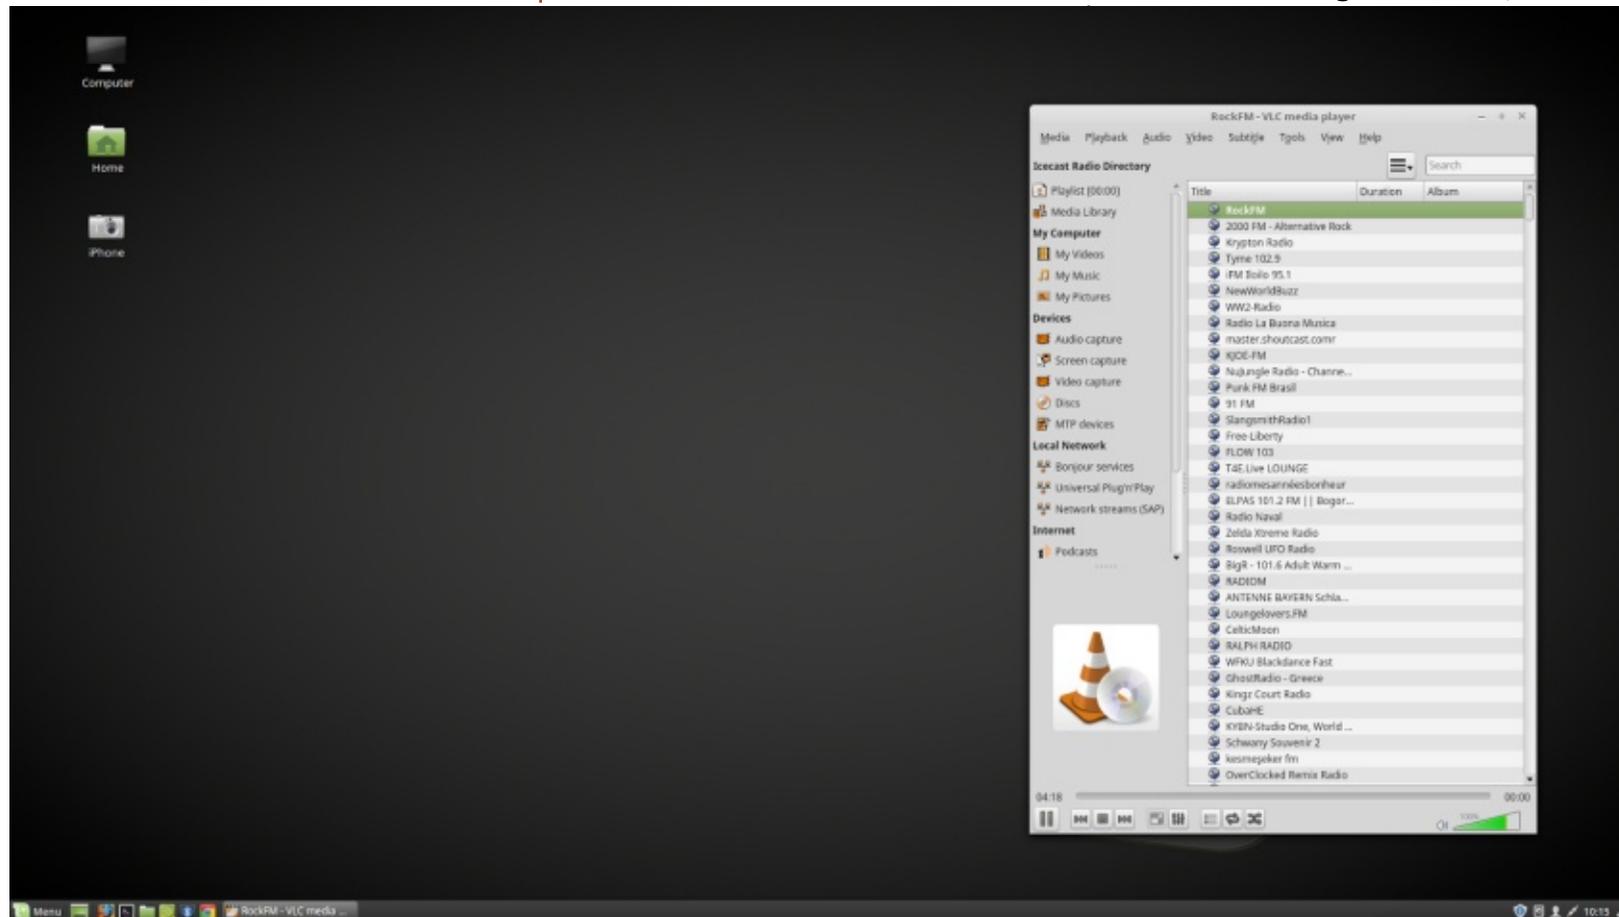
running (as root):

```
sync; echo 1 > /proc/sys/vm/drop_cache
```

It's worth mentioning that I tried running this command using `sudo` and it didn't work, you must actually log in as the root user, which you can do by typing:

```
sudo su -
```

After running 'sudo su -', and



entering in your password (as long as you belong to the group that can use sudo), you'll see the root prompt. From the root prompt, the 'sync; echo' command appears to work. Because I wasn't running any kind of production server or anything critical, I ran the command to flush all three items: PageCache, dentries and inodes:

```
sync; echo 3 >
/proc/sys/vm/drop_cache
```

Again a warning not to run it on a production server. For a better outline of what dentries are, check our Bruce Fields' Virtual File System article here:

<http://www.fieldses.org/~bfields/kernel/vfs.txt>

At this point, I still had VLC streaming the same radio station. The freed up memory result was surprising:

```
MemFree 27,475,792 kB
```

It's a bit more memory used than when we first logged into the Cinnamon desktop, but it does appear to free up a significant amount of memory (1,083,404 kB). This seems to indicate that VLC was using just 277,048 kB more

RAM than when we first logged in. VLC continued to stream in the background while I checked the RAM usage.

What does this all mean? It seems to mean that applications do "leak" RAM, and, even if you've closed an application completely, there may still be some residual RAM usage. When using a web browser, it's much like using any other application that loads a lot of data; a single large file, or web page, may use up a lot of RAM compared to many smaller pages/files. While we all love the mantra of reuse before recycle, old computers with 2GB or less RAM might struggle when used for even light surfing if you surf large sites with lots of content, or open multiple tabs - 4GB of RAM really starts looking like a minimum even for basic web surfing. Lastly, if you are stuck using a system with a small amount of RAM for surfing, you can clear up some of that leaked RAM by clearing the PageCache, dentries and inodes cache.

LINKS OF INTEREST:

Bruce Fields' Virtual File System article:

<http://www.fieldses.org/~bfields/kernel/vfs.txt>

Techmint - How to Clear RAM Memory Cache, Buffer and Swap Space on Linux:

<https://www.tecmint.com/clear-ram-memory-cache-buffer-and-swap-space-on-linux/>

It's FOSS - How Much Swap Should You Use in Linux?

<https://itsfoss.com/swap-size/>



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



When you're putting Bash scripts together, and portability with older systems and other types of shells isn't that important to you, then you might be pleased to discover that it's easy to make light work of your scripts by using a feature called Parameter Expansion. Over time, I've compiled a little list of the most common uses I've had for this feature, and I thought some people might find them helpful. I have to admit these examples are a little erratically listed, so, if you find that they're of interest, I'd suggest collecting your own notes.

In case you're new to Shell Scripting, many of the examples below are a modern shorthand, if you like. Just avoid using them on ten-year-old Bash and KornShell

versions, and the chances are that they will work as expected on your friendly, neighbourhood Linux system. Nonetheless, Unix purists should probably stop reading at this point!

I seem to remember that this functionality was first introduced in Bash version 4.2, and is available in versions from that point release onward, but I might be wrong. You will probably receive an odd-looking error if your shell version doesn't support this functionality, but it's unlikely that anything will break horribly. Figure One shows how to find out which version of Bash you're running on a machine, executed as any unprivileged user – as are the other commands.

Let's not beat about the bush any longer. These examples are hopefully self-explanatory. There's plenty more reading online such as on the GNU site (https://www.gnu.org/software/bash/manual/html_node/Shell-Parameter-Expansion.html), and the Bash Hackers Wiki (<http://wiki.bash-hackers.org/syntax/pe>). The following examples all start simply with the command you would type, and underneath is what your shell will output (in most cases, version dependent) having hit the Enter key. Integrate these snippets into your scripts in any way that you see fit.

FOOD FOR THOUGHT

This short list is really just a

taster of this powerful functionality to whet your appetite.

- Find chars, two chars in length, after position three:

```
$ me="1234567890"; echo
${me:3:2}
```

45

- Ignore the first three chars:

```
$ you="1234567890"; echo
${you:3}
```

4567890

- Replace chars after matching them:

```
$ them="1234567890"; echo
${them/456/xxx}
```

123xxx7890

- Remove chars from string:

```
$ blue="1234567890"; echo
${blue/456}
```

1237890

```
chris@Judy:~$ bash --version
GNU bash, version 4.3.48(1)-release (x86_64-pc-linux-gnu)
Copyright (C) 2013 Free Software Foundation, Inc.
License GPLv3+: GNU GPL version 3 or later <http://gnu.org/licenses/gpl.html>

This is free software; you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.
chris@Judy:~$ █
```

- Display chars up to a delimiter:

```
$ red="1234567890"; echo  
${red%5*}
```

1234

- Show chars after a delimiter:

```
$ yellow="1234567890"; echo  
${yellow#*5}
```

67890

- Discover a variable's length
(note the unusual use of hash
here):

```
$ orange=123;  
length=${#orange}; echo  
$length
```

3

- Remove pattern (the front
part of a variable):

```
$ green="/etc/resolv.conf";  
echo ${green#/etc/}
```

resolv.conf

- Remove pattern (the end of a
variable):

```
$  
pink="chris_secret_file.tar.g  
z"; echo ${pink%.tar.gz}
```

chris_secret_file

- Find and Replace:

```
$ chris="Containers are great  
!"; echo  
${chris/great/fantastic}
```

Containers are fantastic !

- Discover a substring:

```
$  
chrisbinnie="www.devsecops.cc  
"; echo ${chrisbinnie:4:9}
```

devsecops

THE END IS NIGH

Hopefully you agree that, with the addition of some of the examples mentioned on the URLs above, these commands make it possible to significantly speed up your scripting. I hope you enjoy putting them to good use either as one-liners on your command-line, or in shell scripts (as I mostly use them).



Chris Binnie is a DevSecOps consultant who currently works with Docker, Kubernetes and OpenShift security. Further reading can be found here:

<http://www.containersecurity.net>



MY STORY

Written by Rob Lindsay

Linux - Still Number 3 (Maybe 4)

Previously, I looked at my decision to buy an HP Chromebook 11 G5 and exploration of Crouton and Ubuntu MATE. The Chromebook purchase had been a result of having an aging Macbook Pro and wondering what to replace it with. I decided to give a Chromebook a try and bought an HP Chromebook 11 G5.

Early experimentation with Chromeos and Crouton was replaced with focus on Chromeos after I acquired a Raspberry Pi 3B. Here I look at the Chromebook/Raspberry Pi symbiosis as a possible Macbook Pro replacement.

Ok, you may ask, why replace one computer with two?

Well, using the Chromebook by itself led almost immediately to Developer Mode, and Crouton enabling me to run Ubuntu Mate and various other flavours of Ubuntu on the machine.

This was fun, but removed one

of the most attractive aspects of the Chromebook... The validation of the operating system on bootup.

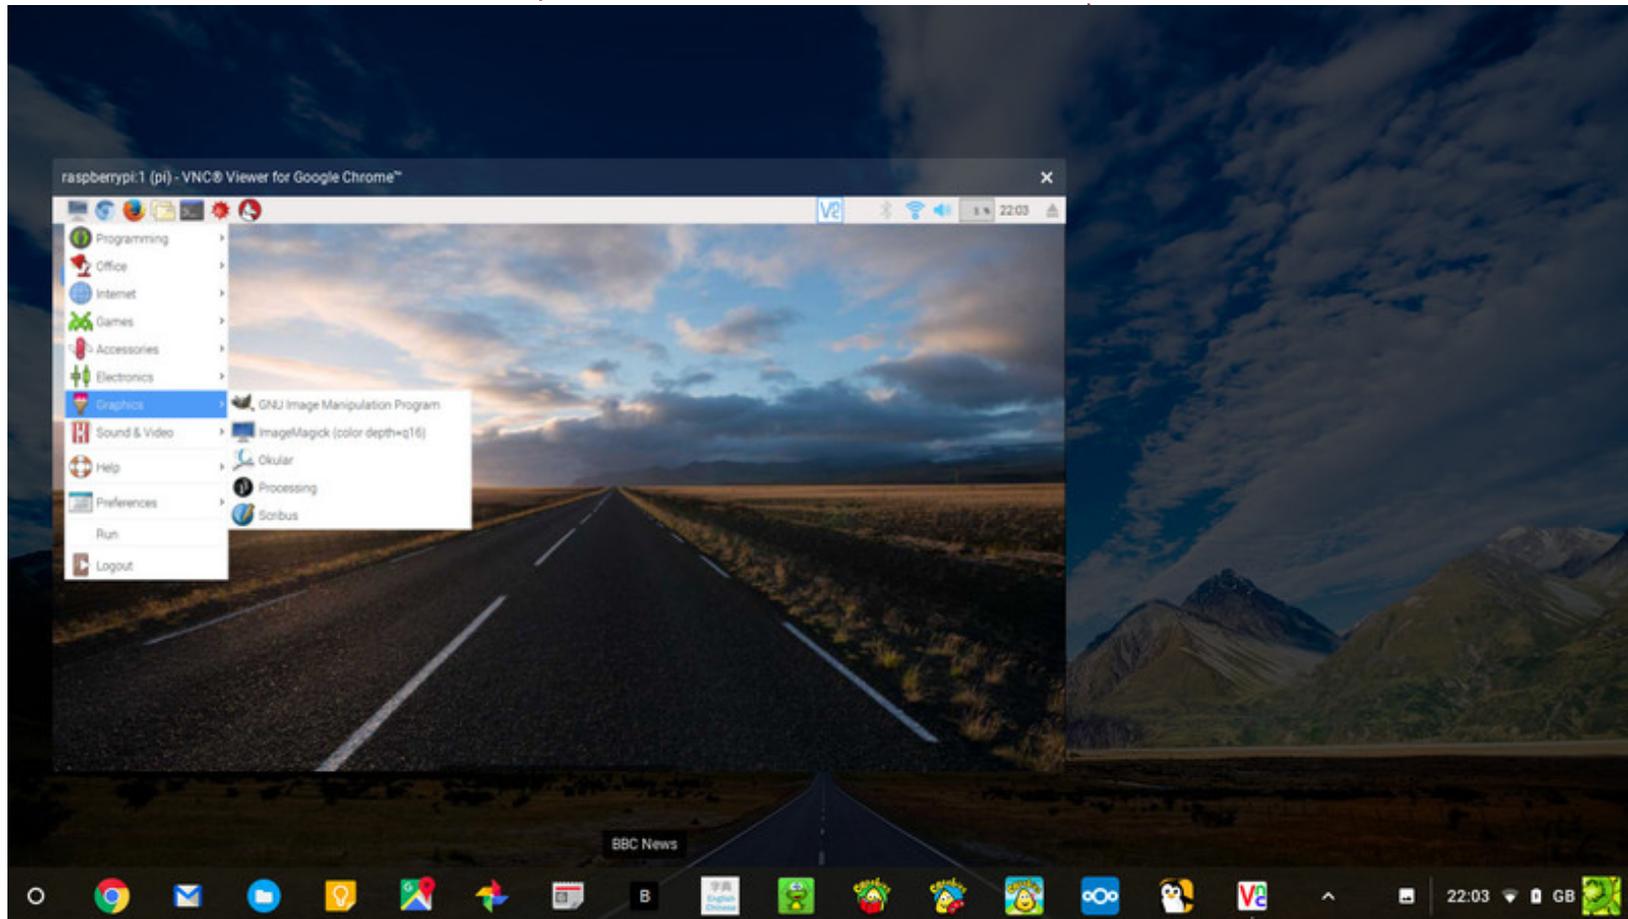
I ended up believing that one should run a Chromebook in normal mode, and look for other ways of accessing traditional

software when needed. I didn't want to pay for a rollApp or similar accounts. Without really thinking of any connection between the Chromebook and anything else, I decided to buy a Raspberry Pi 3B and see what it could do.

Once I had got the machine, I

spent some time investigating its capabilities.

My final decision was to use it as an extension to the Chromebook, acting as a server, and, amongst other things, providing access to traditional software like LibreOffice, Scribus



MY STORY

and Gimp. I also thought that it would be nice to have it act as a DropBox-like file server to myself and other members of the family, both on the local home network and via the internet.

This was achieved by loading Nextcloud on the RASPI connected to a 2TB powered external drive.

To access the RASPI from the Chromebook locally, I use a combination of VNC Client and Secure Shell SFTP. The VNC Client gives me the RASPI desktop on the Chromebook, and SFTP allows transfer of files between the two.

If away from home, or overseas, and wanting to do something on the RASPI, I use VNC Connect which is free to RASPI users, and Cloud Commander, to transfer files between the two. If I want to access my Nextcloud installation when away from home, I fire up Dataplicity and access Nextcloud through my personal wormhole.

To use Cloud Commander, I start two instances of Dataplicity. In the first I type "cloudcmd" and in the second [Cloud Commander uses port 8000], I initiate port 8000 in ngrok and copy the https url into

my browser.

My experience is that VNC Connect cannot be set up headlessly on the RASPI. You have to plug it via an HDMI cable to a TV, and use a USB mouse and keyboard to initiate the link to the RealVNC server. Once this has been done, and port 5900 has been port forwarded in your router, you are able to access the RASPI GUI from anywhere.

There are a number of security considerations to be mindful of when connecting to the RASPI over the internet, the most important being to change the default RASPI password when initially configuring Raspbian.

POSTSCRIPT

Having mentioned rollApp earlier in this piece, I have just discovered OffiDocs. Using this when away from home might be better than using VNC Connect.

What is interesting is that, since my purchase of this Chromebook nearly a year ago, the options for access to traditional software like LibreOffice and Gimp via the

browser seem to have multiplied, and I think this trend will continue.

The OffiDocs pathway gives browser access to a fully configured Linux virtual machine.



Rob is a retired teacher and has an ongoing interest in OSes, particularly Linux, that started with RedHat and Mandrake 15 or so years ago.

FULL CIRCLE WEEKLY NEWS



Full Circle Weekly News

Join our new hosts Wayne and Joe as they present you with a short podcast (<10min) with just the news. No chit-chat. No time wasting. Just the latest FOSS/Linux/ Ubuntu news.

RSS:

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





GUIDELINES

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

RULES

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

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

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

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

document.

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

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

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

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

TRANSLATIONS

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

REVIEWS

GAMES/APPLICATIONS

When reviewing games/applications please state clearly:

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

HARDWARE

When reviewing hardware please state clearly:

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

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





REVIEW

Written by Lucas Westermann

Able2Extract 12

I was recently given a product key for Able2Extract 12, a PDF converter & editor. Previously, I've always done these sorts of tasks using various command-line tools. As extracting text from PDFs, or editing them in any way, is not a task I do too frequently, I cannot promise that I've tested everything the software has to offer. That being said, here are my experiences and thoughts.

COMPATIBILITY

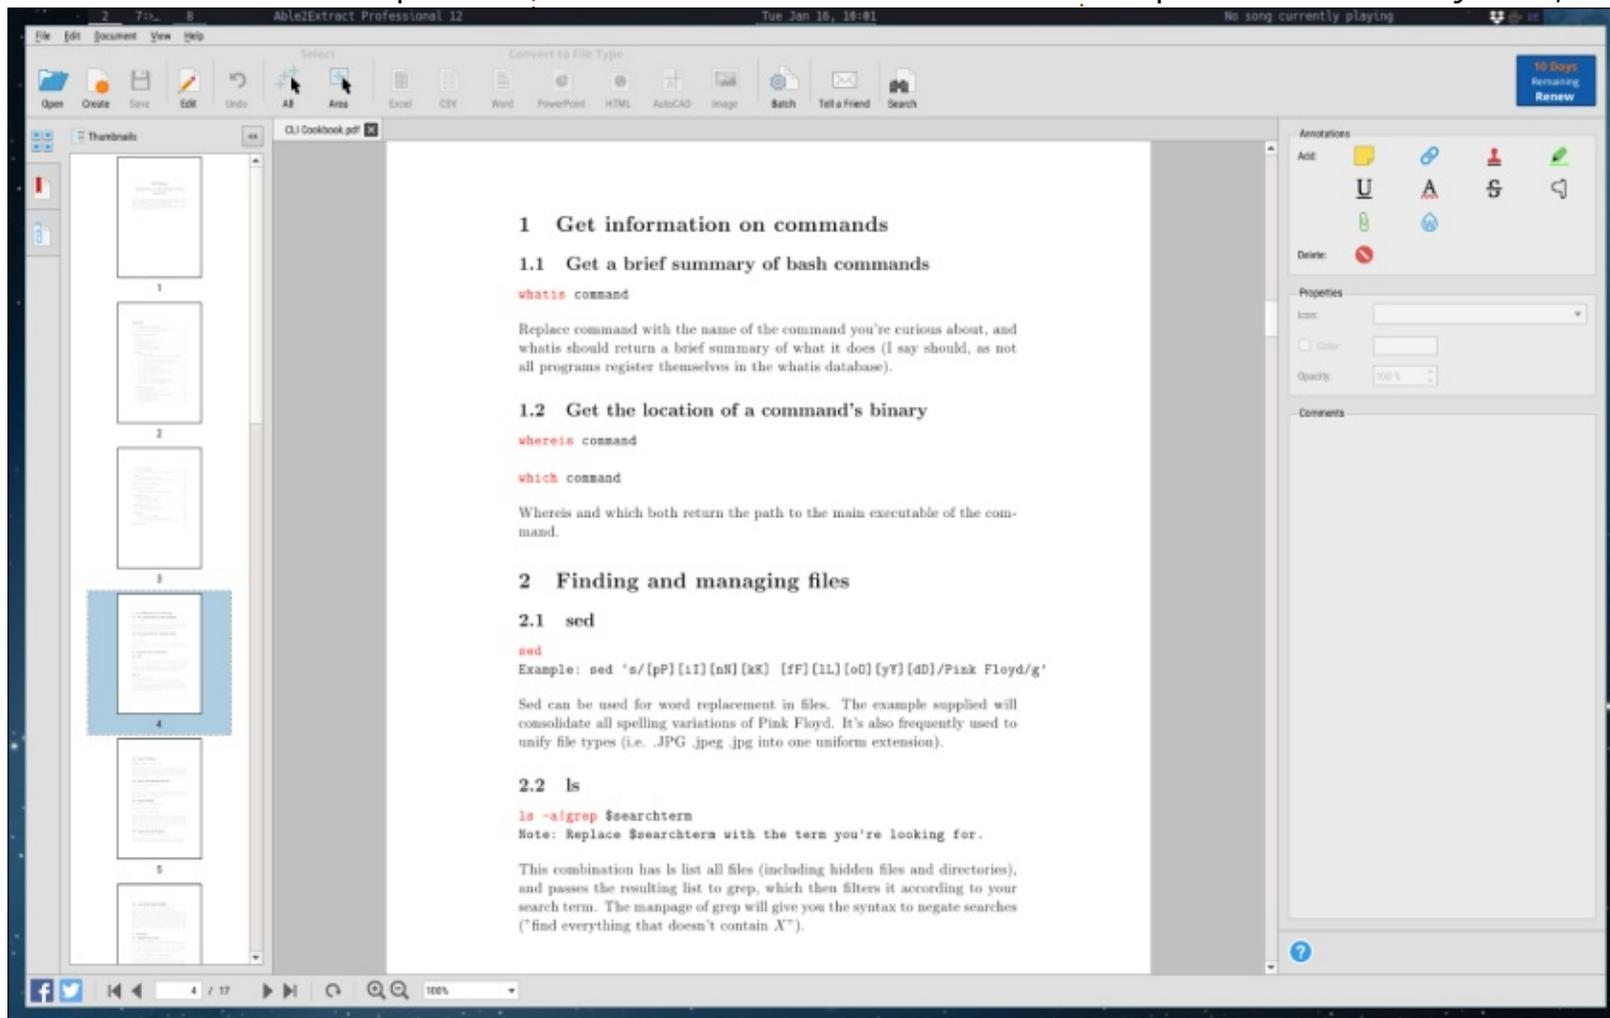
While the software offers packages for only Ubuntu and Fedora, I was able to create a PKGBUILD that correctly installed and runs the Ubuntu .deb file under ArchLinux.

I did, however, run into an issue in Ubuntu 16.04, Ubuntu 17.10, and ArchLinux. Specifically, the application would crash with an error about the QT Fonts location. After contacting the company, we were able to resolve the issue. Apparently, the application requires the variable

\$QT_QPA_FONTDIR to be set to the root path. Instead of defining this system-wide in /etc/environment or in my user's .bashrc, I instead created a bash script that sets the variable and runs Able2Extract. The script is:

```
#!/bin/bash
export QT_QPA_FONTDIR=/
/opt/investintech/a2ep/bin/Able2ExtractPro
```

I went this route because the Able2Extract package does not seem to add the bin to your \$PATH variable, meaning it can be run only from the folder, or the .desktop file. After dropping the script into a folder on my PATH, I



REVIEW

am able to run it as normal. This has the added benefit of not interfering with any other applications, should they want the same variable.

APPLICATION INTERFACE

The layout of the application itself is very familiar (after having used software such as Adobe Acrobat), and it offers some helpful (non-intrusive) tips when starting it for the first time.

FEATURES

The application allows you to create, edit and convert PDF files. Part of the conversion process utilizes OCR technology in order to convert PDFs to editable files such as documents (.odt), or presentation slides. It also offers the ability to create spreadsheets, CSV, HTML, images, and AutoCAD files.

I tested the Word, Excel, and HTML modes on a few recipe scans I have. Some of these files were taken via smartphone camera, and others were scanned on an actual flatbed scanner. The OCR system worked well for most of the files I

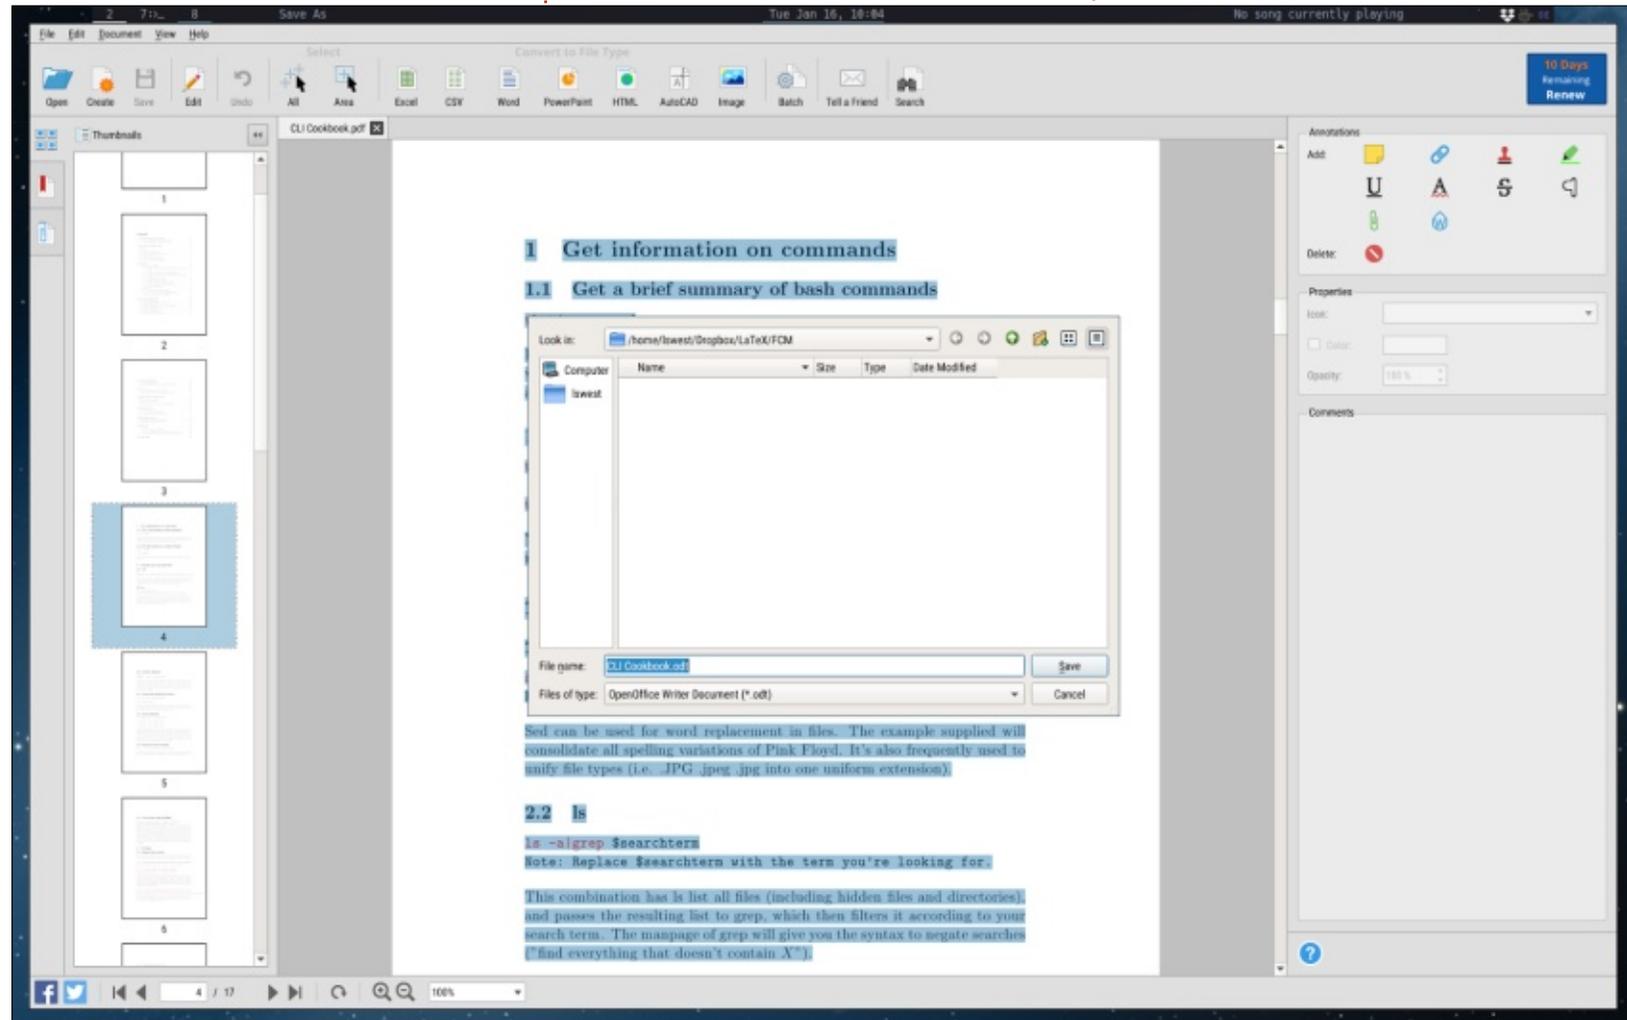
tried, although one particularly badly photographed image had a few gaps where light reflections obscured the text. That being said, I could have filled in the blanks using logic, or adjusting the contrast of the image to make it more legible for myself. I was most impressed by the HTML results, as it actually added plenty of styling

to the text to make it look clean and legible. If you were planning on turning PDFs into unstyled HTML files to add to a website, you should have a plan in place to strip out the inline styles. I did not see any options for the HTML converter.

The conversion options offered

do allow you to handle things such as missing or unrecognized glyphs, or to set the file format for Word and Powerpoint conversions (on my system, it defaulted to OpenOffice). You can also do some document styling such as margins.

The creation tool selects an image file and turns it into a PDF - I



REVIEW

did not see an option to select text documents or word documents (though you could create PDF files using a PDF printer or something like LaTeX). The editing tools include things like adding stamps, highlights, text, comments, etc. They also include things like redacting sections of files, deleting PDF pages, extracting specific pages, and adjusting text styles. The text style adjustment appears to work only on some PDFs - in my tests these options were grayed out. They probably work only on PDFs that were created from a text document, as opposed to image-based scans.

RESULTS

As noted in the previous section, almost every attempt I

made yielded a complete copy of the PDF. In some cases (low contrast, poorly lit, etc), there were some gaps in the resulting file. These could relatively easily be corrected or filled out (especially if you have access to the original document). The worst result came from a recipe that was in 3 columns - while the OCR system managed to correctly separate the columns (I've experienced some that treat 3 columns as 1 line), the character recognition of the actual text was not that impressive. The font in the PDF file was very small, and quite faint, which could have added to the lack of accuracy. The resulting file would have definitely needed proofreading and correcting (though most OCR files should be checked before deeming it finished).

Overall, the results I've experienced using Able2Extract 12 rivals any other OCR software I've ever used, and is much better than other Linux-based alternatives I've tried so far. Is it always perfect? No, but in every test I ran, it yielded a file that would have reduced the effort required to copy the file by hand by at least 50-60%. In most cases it would have required only a few small corrections.

CONCLUSION

If you do a lot of PDF work (splitting documents, OCR scans, etc), and don't have an application for Linux to do this in, I would highly recommend giving Able2Extract a shot. Even if you have an application you use, you

may not be happy with the OCR results - and then I would recommend you try Able2Extract as well.



It's almost a perfect score - if the package worked out of the box, and if there were extra options for HTML conversions, I'd be happy to give it a 5.



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.





Join us on:



goo.gl/FRTML



facebook.com/fullcircle-magazine



twitter.com/#!/fullcirclemag



linkedin.com/company/full-circle-magazine



ubuntuforums.org/forum-display.php?f=270

FULL CIRCLE WEEKLY NEWS



Join our new hosts Wayne and Joe as they present you with a short podcast (<10min) with just the news. No chit-chat. No time wasting. Just the latest FOSS/Linux/ Ubuntu news.

RSS:

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



FULL CIRCLE NEEDS You!



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

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

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



Q&A

Compiled by Gord Campbell

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

Q When does Ubuntu Mate 16.04 support end?

A The "Mate" parts will get updates for three years, the underlying Ubuntu parts will be updated for five years. This is typical of the *buntu distros.

Q Is it safe to write to an NTFS partition from Linux?

A Some people report various problems in Windows after writing to the C: drive from Linux. However, there is agreement that it's safe to write to other NTFS partitions. Sometimes, you really need to do it, for example to copy a large file to an external drive so you can take it to friends who use Windows.

Q I'm a Data Warehouse developer, mostly Informatica and SQL. I would love to have a local install of ORACLE and/or DB2 and SQL Server for that matter. I would prefer an older version, not cloud based. If any of you could

recommend a source for a local installation of a DB, I would appreciate it.

A (Thanks to *SeijiSensei* in the Ubuntu Forums) I've used PostgreSQL for a couple of decades now and would not choose any other SQL server for Linux. PostgreSQL has excellent command-line utilities to create users and databases – tasks that are unnecessarily complex in MySQL. It has a solid ODBC driver so it can communicate with Windows software. I have a couple of PG databases that I manage with Microsoft Access that way.

TIPS AND TECHNIQUES

MEAT AND POTATOES

Many of the questions and answers in this column come from the Ubuntu Forums. However, they are not representative of what appears in the Forums.

I try to avoid duplication, so the most popular themes from the Forums -- the meat and potatoes -- have appeared long ago. The most popular themes are:

- My Wi-Fi doesn't work,
- Sound doesn't work the way I expect it to,
- I installed an update and now something is broken.

There are a few very patient people in the forums who help with Wi-Fi questions. They help people to identify the actual Wi-Fi chip in their computers, and, in most cases, how to get it working. If you have an issue with Wi-Fi, I urge you to search in the Forums first, because you will probably find the solution is already there. You don't need to ask a question.

The key piece of information is the identification of the Wi-Fi chip. In most cases, looking at the output of `lsusb` and `lspci` will show you what it is. Knowing that you have a Dell Dimension is only marginally helpful, but the useful information may look like this:

Ethernet controller: Qualcomm Atheros AR2413/AR2414
Wireless Network Adapter [AR5005G(S) 802.11bg] (rev 01)

So my Google search would be: Wi-Fi AR5005G ubuntuforums

More than with Wi-Fi, there are a host of possible issues with sound. I plug in my earphones and expect the sound to switch from the built-in speakers. How can I boost the output from my microphone? Why doesn't my HDMI connection include sound? And many more.

Once again, identifying the sound chip is useful. For example, `lspci` output:

Audio device: Advanced Micro Devices, Inc. [AMD/ATI] SBx00 Azalia (Intel HDA)

And once again, there's a good chance you can find the answer to your issue by doing a search.

For problems caused by an update, there are things you can do ahead of time which may help.

First is having a detailed identification of the components of your system. My suggestion is to run these commands:

```
cd Desktop
```

```
sudo lshw -html > config.htm
```

Now copy config.htm from your desktop to a flash drive, so you can view it on another computer.

Second, figure out how to use grub to boot into a kernel which is not the latest one. Here's a web page to read:

<https://askubuntu.com/questions/82140/how-can-i-boot-with-an-older-kernel-version>

Finally, have a Boot Repair DVD ready to go, before you need it.

Instructions:

<https://help.ubuntu.com/community/Boot-Repair>

Between the Ubuntu Forums and Askubuntu, the answer to any problem you encounter is likely sitting there waiting for you. All you need is a good Google search to find it.



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





Linux has never been known as a gaming platform, and, as a result, there really hasn't been anything resembling a decently sized collection of games available. Multiple factors have contributed toward this lack of games, but, lucky for us, there have also been lots of solutions and companies that have been battling it out to bring us top quality games to Linux. A great example of a company fighting to expand Linux gaming is Feral Interactive, a cross-platform video game publisher.

Although Feral Interactive has been around since 1996, it wasn't until 2014 that they ventured into Linux. Originally based out of London, UK, they focused exclusively on games for macOS. Then in 2014, Feral Interactive released 'XCOM: Enemy Unknown' for Linux, and, since then, it has continued to bring games to the Linux platform. In addition, Feral Interactive has also expanded into iOS since 2016. Considering how long they have been around, you would imagine that their Linux

gaming library would be relatively small, but you'd be wrong. Feral Interactive has actually been one of the most active publishers bringing games over to Linux. Perhaps one of the reasons why they have been able to port as many games as they have could be attributed to the fact that they constantly work with other publishers such as Warner Bros, Square Enix & 2K Games, among others.

Their relationship with Warner Bros. has brought us two of the best games I've played on Linux: 'Mad Max' and 'Middle-Earth: Shadow of Mordor'. Both of these

games feature some of the best graphics you'll see in a Linux game. 'Middle-Earth: Shadow of Mordor' (reviewed back in our Full Circle Magazine 103 from November 2015) won many awards after being released, including GameSpot's 'Game of the Year' Award. One of its core features (that was critically acclaimed) is what they began calling the Nemesis system – which basically makes the enemies evolve very much the same as any hero would evolve in any role-playing game. In other words, your decisions will directly influence what happens to your enemies ever-evolving character traits. This makes

Shadow of Mordor a fun game to play because it personalizes the enemy for you. Mad Max is another award-winning game that was the result of the Warner Bros. & Feral Interactive relationship. Mad Max was released one year after Middle-Earth and was also very well received by critics and fans alike. You can find a review of Mad Max in Full Circle Magazine 118.

Besides Warner Bros., Feral Interactive has also had a successful relationship with Square Enix. Some of the gems that have been the product of the Square Enix/Feral Interactive partnership have been titles such as Tomb



Raider, Hitman 2016, and Life Is Strange.

Tomb Raider was released for Linux back in April 2016, and you can read all about it in Full Circle Magazine 111. Tomb Raider was also a big hit among critics and fans. Everything from the graphics to character development resulted in Tomb Raider being one of the best games from the entire Tomb Raider series, as well as one of the best games of the last decade.

Life Is Strange was a hidden gem that proved to be a most awesome game as it featured a new playing style in which the main component is the ability to go back in time and relive specific moments again in hopes of achieving a new outcome. Life Is Strange was reviewed in Full Circle Magazine 117.

Then, in 2016 we were finally treated to games being released on Linux at the same time as they were released for Windows PC & macOS. One of these games was Hitman 2016. Although not a game from the main series, Hitman 2016 is a sort of standalone title from the Hitman universe but it brings with it everything we've come to

love from the Hitman games. The main selling point from these games being their stealth playing style, Hitman 2016 proved to be a great addition to this genre. Hitman 2016 was reviewed in Full Circle Magazine 119.

Feral Interactive has worked with other publishers in addition to Warner Brothers & Square Enix. Its partnership with 2K Games and Sega has also produced a growing number of titles that have only made the Linux game list more rich both in quality as well as quantity. There's pretty much a game for everybody, from strategy games like any from the Total War Series, to racing games like Grid: Autosport, Dirt Rally as well as F1 2015 and F1 2017.

The Feral Interactive web site is a good place to find lots of great Linux games. The site is also the Feral Interactive store which means that any title published by Feral Interactive can be purchased from them direct. One cool page I like to look at from time to time is the Upcoming Titles Radar which is sort of a game onto itself. This page always shows you the name of the latest Feral Interactive release. However, the radar on the

page gives you clues as to what games will be released in the future. At the time of this writing, the latest title to be released by Feral Interactive is Deus Ex: Mankind Divided. Then, as you work your way out from the radar's center are four categories (including the center itself): Out Now (at the radar's center), Very Soon (usually games that are just about to be released), Soon (games coming in the next 6-12 months), and Quite Soon on the outer circle of the radar (games that are in production but not to be released anytime soon).

Something interesting about Feral Interactive is also their weekly games broadcast series, #FeralPlays, which is a weekly live-stream on [twitch.tv/feralinteractive](https://www.twitch.tv/feralinteractive) and features a different game being played live, usually on Tuesdays 6PM GMT, and anyone can watch. This twitch channel is also considered the official Twitch channel for Feral Interactive.

There are other publishers out there who have brought great games to Linux, but I have wanted to acknowledge Feral Interactive for quite a while, and to give them

the credit they deserve. More than anything, since Feral Interactive came on-board, their presence has been felt all across the Linux gaming spectrum. Not only have they managed to port lots of games over to Linux, but also these games are critically acclaimed titles which have been ported and released relatively bug-free. The quality of a game that is ported to a different platform from the one it was written for, can vary greatly from title to title, yet Feral Interactive has managed to maintain a high level of performance on the titles they've brought over to Linux. Great game performance goes a long way in the world of Linux gaming. I've played lots of buggy games, but none of those came from Feral Interactive. Successfully porting a game to Linux is an art all by itself, and Feral Interactive seems to have nearly mastered such art.



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



PATRONS

MONTHLY PATRONS

2016:

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

JT

Elizabeth K. Joseph
 Vincent Jobard
 Chris Giltane
 Joao Cantinho Lopes
 John Andrews

2017:

Matt Hopper
 Jay Pee
 Brian Kelly
 J.J. van Kampen

2018:

John Helmers
 Kevin O'Brien
 Kevin Raulins
 Carl Andersen
 John Helmers

SINGLE DONATIONS

2017:

Linda Prinsen
 Shashank Sharma
 Glenn Heaton
 Frank Dinger
 Randy E. Brinson
 Kevin Dwyer
 Douglas Brown
 Daniel Truchon
 John Helmers

Ronald Eike
 Dennis Shimer
 Iain Mckeand
 Jaideep Tibrewala
 Kevin Dwyer

2018:

Yvo Geens
 Graig Pearen
 Carlo Puglisi
 James A Carnrite
 John Holman
 P G Schmitt
 Robert Cannon
 Thomas A Lawell
 Ronald Le Blanc
 Luis Eduardo Herman
 Glenn Heaton
 Peter Swentzel

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

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

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

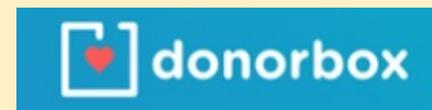
A big thank you to all those who've used Patreon and the PayPal button. It's a HUGE help.



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



<https://paypal.me/ronnietucker>



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



HOW TO CONTRIBUTE

FULL CIRCLE NEEDS YOU!

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

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

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



FCM#131

Deadline:
Sunday 11th Mar. 2018.
Release:
Friday 30th Mar. 2018.



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.



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