REVIEW

UBUNTU 20.10

Photo: Derek (Flickr.com)

DEMYSTIFYING NMAP

NEW

THE NETWORK MAPPER
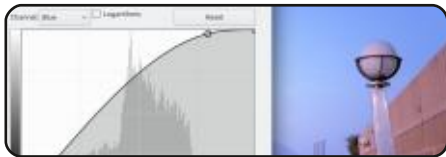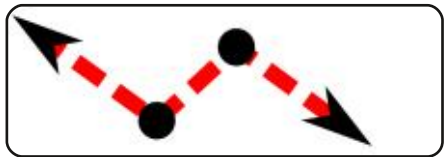
Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

EDITORIAL

WELCOME TO THE LATEST ISSUE OF FULL CIRCLE

This month, we have the usual suspects of Python, Inkscape, Krita and Rawtherapee, Podcast Production but Erik begins a new series (in Command & Conquer) about Nmap. To paraphrase Groucho Marx: "*it may look complicated, it may sound complicated, but don't let it fool you. It really is complicated*". But this month brings the Krita series to an end.

I'm sure you know by now that the next version of Ubuntu, and it's cousins, have been released. Well, Adam Hunt springs to the rescue with a review of the brand spanking new Ubuntu 20.10 with reviews of the other flavors in the coming months.

Elsewhere, we have an excellent piece from Hüseyin about how, in 2015, a municipality in Turkey converted to using Pardus Linux and LibreOffice. It just shows how it can be done!

That's it for this month. Stay safe and well!

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

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

Find Full Circle on:

facebook.com/
fullcirclemagazine

twitter.com/#!/fullcirclemag

http://issuu.com/
fullcirclemagazine

https://mastodon.social/
@fullcirclemagazine

Weekly News:

http://fullcirclemagazine.org/
feed/podcast

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

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

LIBREOFFICE CELEBRATES 10 YEARS OF THE PROJECT:

28.09.2020

The LibreOffice community has celebrated ten years since its inception. Ten years ago, the leading developers of OpenOffice.org formed a new non-profit organization, The Document Foundation, to continue the development of the office suite as a project independent of Oracle, not requiring developers to transfer property rights to the code and making decisions based on the principles of meritocracy.

The project was created a year after the takeover of Sun Microsystems due to dissatisfaction with the tight control of development by Oracle, which prevented interested companies from connecting to collaboration.

https://
blog.documentfoundation.org/
blog/2020/09/27/libreoffice-10th-
anniversary/

OPENSSH 8.4 RELEASED:

28.09.2020

After four months of development , the release of OpenSSH 8.4 is presented, an open client and server implementation for the SSH 2.0 and SFTP protocols."OpenSSH is a 100% complete SSH protocol 2.0 implementation and includes sftp client and server support."

Just a long list of changes:

https://lists.mindrot.org/pipermail/
openssh-unix-dev/2020-September/
038826.html

DSL (DOS SUBSYSTEM FOR LINUX) PROJECT FOR RUNNING LINUX APPLICATIONS FROM MS-DOS ENVIRONMENT:

09/28/2020

Charlie Somerville, who develops the CrabOS operating system in Rust as a hobby, presented a comic, but quite working draft DOS Subsystem for Linux (DSL). It is positioned as an alternative to the WSL (Windows Subsystem for Linux) subsystem developed by Microsoft for those who prefer to work in DOS. Like WSL, the DSL subsystem allows you to directly launch Linux applications, not from Windows, but from the MS-DOS or FreeDOS command shell. The subsystem sources are distributed under the AGPLv3 license.

A DOS environment with a DSL layer can be run in the form of a QEMU virtual machine, or installed on real hardware. Linux programs are launched using the dsl utility, similar to the wsl utility.

https://twitter.com/charliesome/
status/1308061921032110085

RELEASE OF UBUNTU GAMEPACK 20.04:

09/29/2020

Ubuntu GamePack 20.04 is available for download, which includes tools for running more than 85 thousand games and applications, both specially developed for the GNU / Linux platform, and games for Windows launched with PlayOnLinux, CrossOver and Wine, as well as older games for MS-DOS and games for various game consoles (Sega, Nintendo, PSP, Sony PlayStation, ZX Spectrum).

The distribution kit is built on Ubuntu 20.04 and includes all updates, as of September 2020.

https://ualinux.com/ru/news/reliz-
ubuntu-gamepack-20-04-
platformy-dlya-zapuska-igr

MIR DISPLAY SERVER 2.1 RELEASED:

09/29/2020

The release of the display server Mir 2.1 is out, the development is continued by Canonical, despite the abandonment of the development of the Unity shell and the Ubuntu edition for smartphones. Mir remains in

demand in Canonical projects and is now positioned as a solution for embedded devices and the Internet of Things (IoT). Mir can be used as a composite server for Wayland, which allows any application that uses Wayland (for example, built with GTK3 / 4, Qt5, or SDL2) to run in Mir-based environments. Installation packages are prepared for Ubuntu 18.04-20.10 (PPA) and Fedora 30/31/32 . The project code is distributed under the GPLv2 license.

https://discourse.ubuntu.com/t/mir-2-1-0-release/18527

ScummVM 2.2.0 is out:
09/29/2020

The release of the free cross-platform interpreter of classic games, ScummVM 2.2.0 was released, which replaces executable files for games and allows many classic games to run on platforms for which they were not originally intended. The project code is distributed under the GPLv2 license.

In total, it is possible to launch more than 250 quest games and about 1600 interactive text games, including games from LucasArts, Humongous Entertainment, Revolution Software, Cyan and Sierra, such as Maniac Mansion, Monkey Island, Broken Sword, Myst, Blade Runner, King's Quest 1-7 , Space Quest 1-6, Discworld, Simon the Sorcerer, Beneath A Steel Sky, Lure of the Temptress, and The Legend of Kyrandia.

https://www.scummvm.org/news/20200927/

Mesa 20.2.0, a free implementation of OpenGL and Vulkan released:
09/29/2020

Represented by release of free API OpenGL implementation and Vulkan - Mesa 20.2.0 . In Mesa 20.2 implemented full support for OpenGL 4.6 for GPU Intel (i965 drivers, iris) and AMD (radeonsi), OpenGL support 4.5 for GPU AMD (r600), NVIDIA (nvc0) and llvmpipe, OpenGL 4.3 for virgl (virtual the GPU Virgil3D for QEMU / KVM), as well as Vulkan 1.2 support for Intel and AMD cards.

https://lists.freedesktop.org/archives/mesa-announce/2020-September/000600.html

OBS Studio 26.0 Video Streaming System Released:
30.09.2020

OBS Studio 26.0 for streaming, streaming, compositing and video recording is out. The code is written in C / C++ and is distributed under the GPLv2 license. Binaries are generated for Linux, Windows and macOS.

The development goal of OBS Studio is to create a free version of the Open Broadcaster Software application, not tied to the Windows platform, supporting OpenGL and extensible via plugins.

https://github.com/obsproject/obs-studio/releases/tag/26.0.0

Xen Hypervisor Supports Raspberry Pi 4 Board:
30.09.2020

The developers of the Xen project announced the implementation of the possibility of using the Xen Hypervisor on Raspberry Pi 4 boards. Adapting Xen to work on previous versions of Raspberry Pi boards was hampered by the use of a non-standard interrupt controller that does not have support for virtualization. The Raspberry Pi 4 used the usual GIC-400 interrupt controller supported in Xen, and the developers expected that there would be no problems starting Xen. But in reality, everything turned out not to be as expected, and in order to make Xen work on the Raspberry Pi 4, noticeable changes had to be made to the code to work with memory. The fixes were included in the Linux 5.9 kernel, and now Xen is recognized as ready for virtualization on Raspberry Pi 4 boards.

https://xenproject.org/2020/09/29/xen-on-raspberry-pi-4-adventures/

Pale Moon 28.14 Browser Released:
30.09.2020

The Pale Moon 28.14 web browser has been released, a fork of the Firefox codebase to provide better work efficiency, preserve the classic interface,

minimize memory consumption and provide additional customization options. Pale Moon binaries are built for Windows and Linux (x86 and x86_64). The project code is distributed under the MPLv2 (Mozilla Public License) license.

The project adheres to the classic organization of the interface, without switching to the Australis interface integrated into Firefox 29, and providing ample customization options. Remote components include DRM, Social API, WebRTC, PDF viewer, Crash Reporter, code for collecting statistics, parental controls and people with disabilities. Compared to Firefox, the browser retains support for XUL technology and retains the ability to use both full-fledged and lightweight themes.

https://forum.palemoon.org/viewtopic.php?t=25306&p=200477#p200477

OPENCL 3.0 FINAL SPECIFICATIONS PUBLISHED:
09/30/2020

Khronos, which is responsible for the development of specifications for the OpenGL family, Vulkan and OpenCL, announced the publication of the final specifications for OpenCL 3.0 , defining APIs and extensions of the C language for organizing cross-platform parallel computing using multi-core CPUs, GPUs, FPGAs, DSPs and other specialized chips. From those used in supercomputers and cloud servers to chips that can be found in mobile devices and embedded technology. The OpenCL standard is completely open and does not require royalties.

At the same time, an open source OpenCL SDK with tools, examples, documentation, header files, bindings for C ++ and C libraries for developing applications compatible with OpenCL 3.0 was published.

https://www.khronos.org/blog/opencl-3.0-specification-finalized-and-initial-khronos-open-source-opencl-sdk-released

APACHE NETBEANS IDE 12.1 RELEASE:
01.10.2020

Apache Software Foundation Organization introduced Apache NetBeans 12.1, which provides support for Java SE programming languages, Java EE, PHP, C / C ++, JavaScript and Groovy. This is the seventh release made by the Apache Foundation since Oracle donated NetBeans code.

https://blogs.apache.org/netbeans/entry/announce-apache-netbeans-12-1

UBUNTU 20.10 BETA RELEASED:
02.10.2020

The beta release of Ubuntu 20.10 "Groovy Gorilla" is available for testing, which marked a complete freeze of the package base and moved on to final testing and bug fixes. The release is scheduled for October 22nd. Pre-built test images are created for Ubuntu, Ubuntu Server, Lubuntu, Kubuntu, Ubuntu Mate, Ubuntu Budgie, Ubuntu Studio, Xubuntu and UbuntuKylin.

https://lists.ubuntu.com/archives/ubuntu-announce/2020-October/000262.html

FEDORA LINUX EDITION FOR SMARTPHONES INTRODUCED:
02.10.2020

After ten years of inactivity, the Fedora Mobility team is back on track to develop an official edition of the Fedora mobile distribution. Fedora Mobility is designed to be installed on the PinePhone smartphone developed by the Pine64 community. Future editions of Fedora and other smartphones, such as the Librem 5 and OnePlus 5 / 5T, are expected to emerge with support in the stock Linux kernel.

Currently, a set of mobile packages has been added to the Fedora 33 (rawhide) repository , which includes a touch-controlled custom Phosh shell. The Phosh shell is developed by Purism for the Librem 5 smartphone, uses the Phoc composite server running on top of Wayland, and is based on GNOME technologies (GTK, GSettings, DBus).

https://www.mail-archive.com/devel-announce@lists.fedoraproject.org/msg02379.html

NEXTCLOUD HUB 20 COLLABORATION PLATFORM RELEASED:
03.10.2020

Presented the release of the Nextcloud Hub 20 platform, which provides a self-contained solution for organizing collaboration between employees of enterprises and teams developing various projects. Simultaneously published underlying Nextcloud Hub Nextcloud cloud platform 20, allowing the cloud to expand support synchronization and data exchange, providing the ability to view and modify data from any device at any point in the network (using a web-interface or WebDAV). The Nextcloud server can be deployed on any hosting that supports PHP scripts and provides access to SQLite, MariaDB / MySQL, or PostgreSQL. The Nextcloud sources are distributed under the AGPL license.

https://nextcloud.com/blog/nextcloud-hub-20-debuts-dashboard-unifies-search-and-notifications-integrates-with-other-technologies/

NMAP 7.90 NETWORK SECURITY SCANNER RELEASE:
10/04/2020

More than a year after the last release, the release of Nmap 7.90, designed to audit a network and identify active network services. The composition included 3 new NSE-scripts to provide automation of various actions with Nmap. They added over 1200 new signatures for detecting network applications and operating systems.

https://seclists.org/nmap-announce/2020/1

ELBRUS LINUX 6.0 RELEASED:
10/05/2020

Russian MCST Company presented Elbrus Linux 6.0 distribution, built using the Debian GNU / Linux and the LFS projects. Elbrus Linux is not a rebuild, but an independent distribution kit developed by Elbrus architecture developers. Supported systems with Elbrus processors (Elbrus-16S, Elbrus-12S, Elbrus-2S3, Elbrus-8SV, Elbrus-8S, Elbrus-1S +, Elbrus-1SK and Elbrus-4S), SPARC V9 (R2000, R2000 +, R1000) and x86_64. Assemblies for Elbrus processors are supplied on a commercial basis, and the edition for x86_64 systems is declared as freely distributed and free

http://mcst.ru/elbrus_linux_6_release_statement

THE FREE SOFTWARE FOUNDATION TURNS 35:
10/05/2020

The Free Software Foundation is celebrating its 35th anniversary. The celebration will take place in the form of an online event, which is scheduled for October 9. Among the ways to celebrate the anniversary, it is also suggested to experiment with installing one of the completely free GNU / Linux distributions, try to master GNU Emacs, switch to free versions of proprietary software, participate in promoting freejs, or switch to using the F-Droid directory of Android applications .

In 1985, a year after the GNU Project was founded, Richard Stallman founded the Free Software Foundation. The organization was formed to defend against dubious companies caught in the misappropriation of code and trying to sell some of the first GNU Project tools developed by Stallman and his associates. Three years later, Stallman drafted the first version of the GPL license, which defined the legal framework for the free software distribution model. On September 17 last year, Stallman stepped down as president of the Free Software Foundation, and Jeffrey Knuth was elected two months ago.

https://www.fsf.org/blogs/community/fsf-at-35-join-us-in-celebrating-the-incredible-community

UBUNTU RESCUEPACK, LIVE DISTRIBUTION:
05.10.2020

Ubuntu RescuePack is available for download, designed to detect malware and cure infected computers. Includes antivirus packages of ESET NOD32 4,

BitDefender, COMODO, eScan, F-PROT and ClamAV (ClamTk). The distro is also equipped with tools for recovering deleted files. The size of the bootable Live image is 2.6 GB.

The disk runs, without launching the main operating system installed on the computer (MS Windows, macOS, Linux, Android, etc.), to conduct a full anti-virus scan to detect and remove viruses, Trojans, rootkits, worms, spyware and ransomware in the system. The use of an external disk prevents malware from interfering with the neutralization and restoration of the infected system. Data validation is supported in FAT, FAT32, exFAT, NTFS, HFS, HFS +, btrfs, e2fs, ext2, ext3, ext4, jfs, nilfs, reiserfs, reiser4, xfs and zfs file systems.

https://ualinux.com/ru/ubuntu-rescuepack

RELEASE OF THE EPHEMERAL 7 BROWSER, DEVELOPED BY ELEMENTARY OS:
06.10.2020

The release of the Ephemeral 7 web browser, developed by the elementary OS development team specifically for this Linux distribution, has been published . For development, the Vala language, GTK3 + and the WebKitGTK engine were used (the project is not a branch of Epiphany). The code is distributed under the GPLv3 license. Ready-made binaries are made only for elementary OS (the recommended price is $ 9, but you can choose an arbitrary amount, including 0). The browser can be compiled from the source code for other distributions as well.

https://cassidyjames.com/blog/ephemeral-7-developer-tools/

NASNAS 2D GAME DEVELOPMENT FRAMEWORK INTRODUCED:
06.10.2020

The NasNas project is developing a modular framework for developing 2D games in C ++, using the SFML library for rendering and focusing on games in the pixel art style. The code is written in C++ 17 and distributed under the Zlib

license. Supported work in Linux, Windows and Android. There is a binding for the Python language. An example is the History Leaks game created for the GameBoy JAM competition .

https://github.com/Madour/NasNas

RED HAT ENTERPRISE LINUX 7.9 AND ORACLE LINUX 7.9 RELEASED:
08.10.2020

Red Hat has released Red Hat Enterprise Linux 7.9 (the new version was announced a week ago only on the access.redhat.com portal, the announcement did not appear on the mailing list or in the press release section). RHEL 7.9 installation images are available for download by registered Red Hat Customer Portal users only and are prepared for x86_64, IBM POWER7 +, POWER8 (big endian and little endian), and IBM System z architectures. The package sources can be downloaded from the CentOS project Git repository .

The RHEL 7.x branch is being maintained in parallel with the

RHEL 8.x branch and will be supported until June 2024. The first phase of support for the RHEL 7.x branch, which includes functional improvements, is complete. The release of RHEL 7.9 follows the transition to the maintenance phase, in which the priorities have shifted towards bug fixes and security, with minor improvements related to supporting critical hardware systems.

https://access.redhat.com/announcements/5440131

CALLA, AN AUDIO / VIDEO CONFERENCING PLATFORM IN THE FORM OF AN RPG GAME AVAILABLE:
08.10.2020

The Calla project is developing an audio and video conferencing system that allows multiple participants to speak simultaneously. Usually, during online conferences, only one participant is allowed to speak, and simultaneous discussions are problematic. In Calla, to organize natural communication, in which

several people can speak at the same time, they propose to use navigation in the form of an RPG game. The project is written in JavaScript, uses the free Jitsi Meet platform and is distributed under the MIT license.

https://www.calla.chat/

THE RELEASE OF THE MPZ MUSIC PLAYER 1.0

08.10.2020

The first stable release of mpz music player optimized for large local music collections. Mpz's approach is inspired by the "album list" function in Foobar2000. The main feature is a three-pane interface where you can create playlists from catalogs and switch between playlists. During playback,

the audio codecs installed in the OS are used (connected via QtMultimedia). The code is written in C++ using the Qt library and is distributed under the GPLv3 license. Binaries are prepared for Windows and Linux distributions of openSUSE, Debian, Fedora, Ubuntu, CentOS and Mageia. Features include the ability to use Internet radio with playlists in m3u and pls formats, CUE support, the ability to remotely control the player using the MPRIS protocol, log playback, settings in yaml format.

https://github.com/olegantonyan/mpz/releases/tag/1.0.0

RELEASE OF GIMP 2.10.22:

08.10.2020

GIMP 2.10.22 , which continues to hone functionality and

increase the stability of the 2.10 branch, is out . A flatpak package is available for installation (the snap package has not been updated yet). GIMP 2.10.22 is a bug fix release, which for once contains mostly maintenance changes in core code. See:

https://www.gimp.org/news/2020/10/07/gimp-2-10-22-released/

SURICATA INTRUSION DETECTION SYSTEM 6.0 RELEASED:

08.10.2020

After a year of development, the Open Information Security Foundation (OISF) has released Suricata 6.0, a network intrusion detection and prevention system

FULL CIRCLE WEEKLY NEWS

Join our new host Leo Chavez as he presents 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

that provides a means of inspecting various types of traffic. Suricata configurations can use the signature base developed by the Snort project, as well as the Emerging Threats and Emerging Threats Pro rulesets . The source code of the project is distributed under the GPLv2 license.

https://suricata-ids.org/2020/10/08/suricata-6-0-0-released/

FREETYPE 2.10.3 FONT ENGINE RELEASE:
11.10.2020

Introducing FreeType 2.10.3 , a modular font engine that provides a single API for unifying the processing and output of font data in various vector and raster formats.

https://www.mail-archive.com/freetype-announce@nongnu.org/msg00124.html

REDO RESCUE 2.0.6 RELEASE, BACKUP AND RECOVERY DISTRIBUTION:
11.10.2020

The Live-distribution, Redo Rescue 2.0.6 , designed to create backups and restore the system in case of failure or data corruption, is out. State slices created by the distribution can be fully or selectively cloned to a new disk (creating a new partition table) or used to restore system integrity after malware activity, hardware failures, or accidental data deletion. The distribution uses the Debian codebase and partclone toolkit from the Clonezilla project . Redo Rescue's own work is distributed under the GPLv3 license. The size of the iso image is 410MB.

Backups can be saved both to locally attached media (USD Flash, CD / DVD, disks), and to external partitions, which are accessed via NFS, SSH, FTP or Samba / CIFS (an automatic search for shared sections). It supports remote management of backup and restore using VNC or web interface.

https://github.com/redorescue/redorescue/releases/tag/2.0.6

ONLYOFFICE 6.0 RELEASED:
19.10.2020

A new release of the OnlyOffice 6.0 suite is available , designed for working with text documents, spreadsheets and presentations. The editors are desktop applications that are written in JavaScript using web technologies, but combine client and server components in one set, designed for self contained use on the user's local system, without accessing an external service. The project code is distributed under the free license AGPLv3 .

OnlyOffice claims full compatibility with MS Office and OpenDocument formats. Supported formats: DOC, DOCX, ODT, RTF, TXT, PDF, HTML, EPUB, XPS, DjVu, XLS, XLSX, ODS, CSV, PPT, PPTX, ODP. Functionality of the editors can be expanded through plugins, for example, plugins are available for creating templates and adding videos from YouTube. Binaries for Windows, macOS and Linux (deb- and rpm-packages, packages in Snap, Flatpak and AppImage formats will also be generated in the near future).

At the same time, a new release of OnlyOffice DocumentServer 6.0 was announced with the implementation of a server for working with OnlyOffice online editors from a web browser for collaboration. You can also use the Nextcloud Hub platform for collaboration at your facilities, which provides full integration with OnlyOffice.

https://www.onlyoffice.com/blog/2020/10/onlyoffice-desktop-editors-v6-0-with-pivot-tables/

JAMI DECENTRALIZED COMMUNICATION CLIENT RELEASED:
20.10.2020

The first release of the Jami decentralized communications platform , codenamed "Together", was announced. Jami is one of the GNU projects and was previously developed under the name Ring (previously SFLphone), but was renamed in 2018 to avoid trademark infringement. The code is distributed under the GPLv3 license. Binaries are prepared for

GNU / Linux (Debian, Ubuntu, Fedora, SUSE, RHEL, etc.), Windows, macOS, iOS, Android and Android TV. Several options for interfaces based on Qt , GTK and Electron are being developed . Unlike traditional communication clients, Jami is able to transfer messages without contacting external servers through establishing a direct connection between users using end-to-end encryption (end-to-end, keys are present only on the client side) and authentication based on X.509 certificates. In addition to secure messaging, the program allows you to make voice and video calls, create teleconferences, exchange files, organize shared access to files and screen content.

https://jami.net/together-the-new-version-of-jami-and-a-new-step-forward/

GIT 2.29 DISTRIBUTED SOURCE MANAGEMENT RELEASED:

20.10.2020

The Git 2.29.0 is now available . Git is one of the most popular, reliable, and high-performance version control systems, providing flexible non-linear development tools based on branching and merging. To ensure the integrity of history and resistance to changes "retroactively", implicit hashing of the entire previous history in each commit is used, it is also possible to certify with digital signatures of the developers of individual tags and commits. Compared to the previous release, 627 changes were adopted in the new version, prepared with the participation of 89 developers, of which 24 took part in the development for the first time.

https://lkml.org/lkml/2020/10/19/692

SCIENTIFIC LINUX 7.9 RELEASED:

10/21/2020

Scientific Linux 7.9, built on Red Hat Enterprise Linux 7.9 and supplemented with tools oriented for use in scientific institutions is available. The distribution is supplied for the x86_64 architecture, in the form of DVD images (9.9 GB and 7 GB), plus a reduced image for network installation (620 MB). Live images are delayed.

The main differences from RHEL come down to rebranding and cleaning up the bindings to Red Hat services. Scientific-specific applications, as well as additional drivers, are offered for installation from external repositories such as EPEL and elrepo.org . Before upgrading to Scientific Linux 7.9, it is recommended to run 'yum clean all' to clear the cache.

https://www.mail-archive.com/scientific-linux-announce@listserv.fnal.gov/msg00054.html

NETBSD 9.1 RELEASES:

21.10.2020

NetBSD 9.1 operating system is out. There are 820 MB installation images prepared for download , available in assemblies for 57 system architectures and 16 different CPU families.

https://blog.netbsd.org/tnf/entry/netbsd_9_1_released

VIRTUALBOX 6.1.16 RELEASES:

10/21/2020

Oracle has published a patch release for VirtualBox 6.1.16 , which includes 15 fixes .

https://www.mail-archive.com/vbox-announce@virtualbox.org/msg00208.html

TRUENAS 12 NETWORK STORAGE:

21.10.2020

iXsystems introduced TrueNAS 12 , the distribution for the rapid deployment of network-attached storage (NAS, Network-Attached Storage) . TrueNAS 12 is based on the FreeBSD 12 codebase, with integrated ZFS support and web-based management built using the Python Django framework. FTP, NFS, Samba, AFP, rsync and iSCSI are supported to organize access to the storage, software RAID (0,1,5) can be used to improve storage reliability, LDAP / Active Directory support is implemented for client authorization. The size of the iso image is 827MB (x86_64).

The TrueNAS distribution is the result of the merger of the FreeNAS distribution and the commercial TrueNAS platform. For historical reasons, FreeNAS and TrueNAS were developed, tested and released separately, despite the large amount of common code (95%). Users are offered two editions of TrueNAS CORE and TrueNAS Enterprise. TrueNAS CORE is similar to FreeNAS and is free. TrueNAS Enterprise is paid and includes additional enterprise capabilities.

https://www.ixsystems.com/blog/library/truenas-12-0-release/

ASTERISK 18 COMMUNICATION PLATFORM AVAILABLE:
10/21/2020

After a year of development, a new stable branch of the open communication platform Asterisk 18 was released, which is used to deploy software PBXs, voice communication systems, VoIP gateways, organization of IVR systems (voice menu), voice mail,

telephone conferences and call centers. The source code of the project is available under the GPLv2 license. Asterisk 18 (LTS), for which updates will be issued for five years instead of the usual two years. Support for the past LTS branch Asterisk 16 will last until October 2023, and Asterisk branch 13 until October 2021. LTS releases focus on stability and performance optimization, while regular releases prioritize functionality enhancements.

https://www.mail-archive.com/asterisk-announce@lists.digium.com/msg00779.html

THUNDERBIRD MAIL CLIENT 78.4.0 UPDATE:
10/22/2020

The main changes in the new version are related to improved support for add-ons (MailExtensions). Among the non-add-on features, the migration of users of Yahoo and AOL mail services from password authentication to the OAuth2 mechanism is the other notable change.

For add-on developers, two new APIs are proposed browser.tabs.sendMessage for passing messages between add-on components and messageDisplayScripts for connecting content processing scripts that work in the context of the displayed message. Added support for working with multiple highlighted messages in the messageDisplay API . The compose.begin function implements the ability to create a message with attached attachments.

https://www.thunderbird.net/en-US/thunderbird/78.4.0/releasenotes/

BIND 9.16.8 DNS SERVERS AND POWERDNS RECURSOR 4.4.0 RELEASED:
10/22/2020

Bug fix update to DNS-server BIND 9.11.24 and 9.16.8 stable branches, as well as the development of experimental branches 9.17.6. In new versions, in accordance with the

recommendations promoted by the DNS flag day 2020 initiative, the default buffer size for EDNS has been reduced from 4096 to 1232 bytes. The change will eliminate IP fragmentation issues when processing large UDP messages, the processing of which often leads to packet loss and timeouts on the client side.

In addition, a release of PowerDNS Recursor 4.4 caching DNS server is available, which is responsible for recursive name resolution. PowerDNS Recursor is built on the same codebase as the PowerDNS Authoritative Server, but PowerDNS recursive and authoritative DNS servers evolve through different development cycles and are released as separate products. The project code is distributed under the GPLv2 license.

The server provides tools for remote statistics collection, supports instant restart, has a built-in engine for connecting handlers in the Lua language, fully supports DNSSEC, DNS64, RPZ (Response Policy Zones), and allows you to connect blacklists. It is possible to record the results of resolving as BIND zone files. To ensure high

performance, modern mechanisms for multiplexing connections in FreeBSD, Linux and Solaris (kqueue, epoll, / dev / poll) are used, as well as a high-performance DNS packet parser capable of processing tens of thousands of parallel requests.

https://www.mail-archive.com/bind-announce@lists.isc.org/msg00586.html

Ubuntu 20.10 released:
10/22/2020

The release of Ubuntu 20.10 "Groovy Gorilla" is available, which is classified as an interim release, updates will be available for 9 months (support will be carried out until July 2021). Pre-built test images are created for Ubuntu, Ubuntu Server , Lubuntu , Kubuntu , Ubuntu Mate , Ubuntu Budgie , Ubuntu Studio , Xubuntu and UbuntuKylin (China edition).

https://lists.ubuntu.com/archives/ubuntu-announce/2020-October/000263.html

RIAA blocks youtube-dl repository on GitHub:
10/23/2020

GitHub has blocked the repository and all mirrors of the youtube-dl project, which is developing a command line utility for downloading videos from YouTube and other sites. It was blocked under the US Digital Millennium Copyright Act (DMCA) following a complaint from the Recording Industry Association of America (RIAA).

The claims boil down to the presence in the repository of code with examples of downloading licensed content from YouTube. In particular, in the code of scripts for extracting video, there is a section with tests of correctness of work, in which links to materials whose copyright belongs to RIAA participants are indicated as examples to be checked. Tests are also used to verify that the age restriction and encrypted content download (cipher_signature) bypass is working, and these bypass targets are explicitly indicated in the test notes.

https://github.com/github/dmca/blob/master/2020/10/2020-10-23-RIAA.md

GDB 10 debugger released:
10/25/2020

GDB 10.1 is out. (first release 10.x series branch 10.0 was used for development). GDB supports source-level debugging for a wide variety of programming languages (Ada, C, C ++, Objective-C, Pascal, Go, Rust, etc.) on a variety of hardware (i386, amd64, ARM, Power, Sparc, RISC- V, etc.) and software platforms (GNU / Linux, * BSD, Unix, Windows, macOS).

https://sourceware.org/pipermail/gdb/2020-October/048987.html

Pop!_OS 20.10 released:
10/25/2020

System76, a company specializing in the production of laptops, PCs and servers shipped with Linux, published Pop!_OS 20.10. It is developed for delivery to System76 products and comes with a redesigned desktop environment. Pop! _OS is based on Ubuntu 20.10 base . The developments of the project are distributed under the GPLv3 license. ISOs are generated for x86_64 architecture with NVIDIA (2.2 GB) and Intel / AMD (2.6 GB) graphics chips .

Pop!_OS comes with a modified GNOME Shell, the original system76-pop theme, its own set of icons, different fonts (Fira and Roboto Slab), changed settings and an extended set of drivers.

https://blog.system76.com/post/632781631953027072/whats-new-in-popos-2010

Linux kernel 5.9 released:
10/12/2020

After two months of development, Linus Torvalds has released Linux kernel 5.9. Among the most notable changes: limiting the import of symbols from proprietary modules into GPL-modules, speeding up context switching operations using the FSGSBASE processor instruction

set, support for compressing the kernel image using Zstd, redesigning the prioritization of threads in the kernel, support for PRP (Parallel Redundancy Protocol), throughput scheduling in the deadline scheduler, paging ahead of time, CAP_CHECKPOINT_RESTOR capability flag, close_range () system call, dm-crypt performance improvements, code removal for 32-bit Xen PV guests, new slab memory management mechanism, Btrfs "rescue" option, ext4 and F2FS inline encryption support.

The new version received 16074 fixes from 2011 developers, the size of the patch is 62 MB (the changes affected 14,548 files, 782,155 lines of code added, 314,792 lines removed). About 45% of all changes introduced in 5.9 are related to device drivers, approximately 15% of changes are related to updating code specific to hardware architectures, 13% are related to the network stack, 3% are related to filesystems and 3% are related to internal kernel subsystems.

https://lkml.org/lkml/2020/10/11/275

PITIVI VIDEO EDITOR 2020.09 RELEASED:
12.10.2020

After two years of development, the release of the free, nonlinear video editing system Pitivi 2020.09 is available , providing such features as support for an unlimited number of layers, saving a complete history of operations with the ability to rollback, displaying thumbnails on a timeline, and support for typical video and audio processing operations. The editor is written in Python using GTK + (PyGTK), GES (GStreamer Editing Services) and can work with all audio and video formats supported by GStreamer, including MXF (Material eXchange Format). The code is distributed under the LGPL license.

https://blog.pitivi.org/2020/10/11/pitivi-2020-09-hocus-focus/

OPENOFFICE.ORG TURNS 20:
10/13/2020

The free office suite OpenOffice.org turns 20 - on October 13, 2000, Sun Microsystems released the source code for the office suite StarOffice, which was created in the early 90s by Star Division, under a free license. In 1999, Star Division was taken over by Sun Microsystems, which took one of the most important steps in open source history — making StarOffice a free project. In 2010, Oracle got its hands on OpenOffice along with other Sun Microsystems projects, but after a year of trying to develop OpenOffice.org in-house, it handed the project over to the Apache Foundation. The last patch release of Apache OpenOffice 4.1.7 was a year ago, and there have been no major releases in 6 years. The initiative to develop a free office suite was intercepted by the LibreOffice project, which was created in 2010 due to dissatisfaction with Oracle's tight control over OpenOffice.org development, which prevents interested companies from connecting to collaborative work.

LibreOffice developers published an open letter in which they called on the Apache OpenOffice developers to work together, since Apache OpenOffice has long been in deep stagnation, and all development in recent years has been concentrated in LibreOffice. Compared to OpenOffice, LibreOffice has such features as export in OOXML (.docx, .xlsx) and EPUB formats, digitally signed documents, significant Calc performance optimizations, a redesigned NotebookBar interface, Pivot Charts, watermarks, and secure mode (Safe Mode).

Despite stagnation and a virtual lack of support, the OpenOffice brand remains strong with downloads still in the millions and many users are unaware of LibreOffice. The LibreOffice developers have invited the OpenOffice project to inform its users about the existence of an actively maintained and more functional product that continues the development of OpenOffice and includes new features that modern users need.

http://www.openoffice.org/

NEWS

THE KUBUNTU PROJECT INTRODUCED THE SECOND MODEL OF THE KUBUNTU FOCUS LAPTOP:
10/13/2020

The developers of the Kubuntu distribution have announced the availability of the " Kubuntu Focus M2 " laptop branded project and offering a pre-installed desktop environment based on Ubuntu 20.04 and the KDE desktop. The device was released in collaboration with MindShare Management and Tuxedo Computers. The notebook is aimed at advanced users and developers who need a powerful laptop, supplied with a Linux environment optimized for the proposed hardware. The device is priced at $ 1,795. As a base, they used the CLEVO PC50DF1 gaming laptop, also produced under the TUXEDO Book XP15 brand.

https://kubuntu.org/news/kubuntu-focus-model-2-launched/

PORTEUS KIOSK 5.1, DISTRO FOR INTERNET KIOSKS:
13.10.2020

Porteus Kiosk 5.1.0 , based on Gentoo and intended for equipping stand-alone Internet kiosks, demonstration stands and self-service terminals, is out. The bootable distribution image is 110 MB.

The basic image includes only the minimum set of components required to run a web browser (Firefox and Chrome are supported), which is cut in its capabilities to prevent unwanted activity in the system (for example, changing settings is not allowed, downloading / installing applications is blocked, only access granted to selected pages). Additionally, specialized Cloud imaes are offered for comfortable work with web applications (Google Apps, Jolicloud, OwnCloud, Dropbox) and ThinClient for working as a thin client (Citrix, RDP, NX, VNC and SSH) and Server for kiosk network management.

https://porteus-kiosk.org/news.html#201012

KDE PLASMA 5.20 DESKTOP RELEASED:
13.10.2020

The release of the custom KDE Plasma 5.20 shell is available, built using the KDE Frameworks 5 platform and the Qt 5 library using OpenGL / OpenGL ES to accelerate rendering. You can evaluate the work of the new version through the Live build from the openSUSE project and the build from the KDE Neon User Edition project . Packages for various distributions can be found on this page: https://community.kde.org/Get_KDE_Software_on_Your_Linux_Distro

https://kde.org/announcements/plasma-5.20.0/

NPM 7.0 PACKAGE MANAGER AVAILABLE:
13.10.2020

The release of the package manager NPM 7.0 , included in the Node.js distribution and used to distribute JavaScript modules,

has been published. The NPM repository serves over 1.3 million packages, which are used by about 12 million developers. About 75 billion downloads are recorded per month. NPM 7.0 was the first significant release made after the purchase of NPM Inc by GitHub. The new version will be shipped with a future release of the Node.js 15 platform , expected on October 20. To install NPM 7.0 without waiting for a new version of Node.js, you can run the command "npm i -g npm @ 7".

https://blog.npmjs.org/post/631877012766785536/release-v700

CROSSOVER 20.0 RELEASED FOR LINUX, CHROME OS AND MACOS:
10/14/2020

CodeWeavers has released Crossover 20.0, based on Wine code and designed to run programs and games written for the Windows platform. CodeWeavers is one of the key contributors to the Wine project, sponsors its development and returns to the project all the innovations

implemented for its commercial products. The source texts of the open components CrossOver 20.0 can be downloaded on this page: https://www.codeweavers.com/crossover/source

https://www.codeweavers.com/support/forums/announce/?t=24;msg=231110

KRITA 4.4.0 RELEASED:

10/14/2020

The release of the raster graphics editor Krita 4.4.0, intended for artists and illustrators, is available for download . The editor supports multi-layer image processing, provides tools for working with various color models and has a wide range of tools for digital painting, sketching and texture formation. Self -contained images in the AppImage format for Linux, experimental APK packages for ChromeOS and Android, as well as binaries for macOS and Windows are prepared for installation .

https://krita.org/en/item/krita-4-4-0-released/

VIVALDI 3.4 BROWSER RELEASE FOR PC AND ANDROID:

10/15/2020

The simultaneous release of Vivaldi 3.4 browsers for PC and Android took place. Vivaldi is developed on the base of the Chromium engine and continues to develop the ideas of the classic Opera browser, providing a wide range of features, including a tab grouping system, a sidebar, a configurator with a large number of settings, a mode for blocking images and unwanted content, a note-taking system, and a horizontal tab display mode. The browser interface is written in JavaScript using the React library, Node.js framework, browserify and various ready-made NPM modules. Vivaldi builds are prepared for Linux, Windows, Android and macOS. For past releases, the open source project distributes source code of changes to Chromium. The Vivaldi interface implementation is written in JavaScript, available in source code, but under a proprietary license.

https://vivaldi.com/blog/

RELEASE OF RESCUEZILLA 2.0 BACKUP DISTRIBUTION:

15.10.2020

A new release of the Rescuezilla 2.0 distribution is out, designed for backups, system crash recovery and diagnostics of various hardware problems. The distribution is built on the Ubuntu base and continues the development of the "Redo Backup & Rescue" project, the development of which was discontinued in 2012. Rescuezilla supports backup and recovery of accidentally deleted files on Linux, macOS and Windows partitions. It automatically searches for and connects network partitions that can be used to place backups. The graphical interface is based on the LXDE shell. Live assemblies for x86 64-bit systems (771MB) are offered for download . Build for 32-bit systems is temporarily suspended.

https://rescuezilla.com/

THE OPENPRINTING PROJECT BEGAN A FORK OF THE CUPS PRINTING SYSTEM:

16.10.2020

The OpenPrinting project, supported by the Linux Foundaton, has begun a fork of the CUPS printing system . The most active part in the development of the fork is Michael R Sweet, the original author of CUPS.

Since 2007, following the takeover of CUPS by Easy Software Products, CUPS development has been completely controlled by Apple. In December 2019, Michael Sweet, founder of the CUPS project and Easy Software Products, resigned from Apple. The vast majority of changes to the CUPS codebase were made personally by Michael Sweet, but when he announced his departure, Michael mentioned that there are two engineers left at Apple who will provide CUPS maintenance. However, after Michael's dismissal, the CUPS project stopped developing - during 2020, only one commit was added to the CUPS codebase with the elimination of

vulnerabilities.

https://openprinting.github.io/

RELEASE OF UBUNTU CYBERPACK (ALF) 2.0, A FORENSIC ANALYSIS DISTRIBUTION:
10/17/2020

The UALinux project has published the "Ubuntu CyberPack (ALF) 2.0" distribution for collecting and obtaining digital evidence, as well as for conducting security tests. The set includes the tools necessary for obtaining and fixing digital evidence in the investigation of cybercrimes, as well as for the primary forensic analysis of the evidence base in the "field" conditions. The system also contains software for conducting security tests, penetration tests and vulnerability assessments, which makes it an analogue of such distributions as Kali Linux and Parrot. The distribution is based on the Ubuntu 20.04 codebase. The size of the iso image (x86_64) 4.1GB. Both installation and live modes are supported.

https://ualinux.com/ru/news/dostupna-dlya-zagruzki-forensic-sistema-ubuntu-cyberpack-alf-2-0

RELEASE OF ANTIX 19.3:
10/17/2020

AntiX 19.3 , designed for installation on outdated hardware, has been released. The release is based on the Debian 10 (Buster) package base, but comes without the systemd manager and with eudev instead of udev. You can choose between runit and sysvinit for initialization. The default user environment is built using the IceWM window manager, but Fluxbox, JWM and herbstluftwm are also available to choose from. Midnight Commander, SpaceFM and ROX-Filer are available for working with files. The distribution kit works on systems with 256 MB of RAM. Size of iso images are: 1.2 GB (full), 766 MB (base), 470 MB (stripped down), and 166 MB (network install). The images are prepared for x86_64 and i386 architectures.

In the new release, the "antiX wifi switch" utility has been added to select a wireless connection

manager (now connman is used by default, and its own ceni cli interface is an option). Updated package versions, including Linux kernel 4.9.235, Firefox 78.3.0, LibreOffice 7.0.2.2, IceWM 1.8.3. Fresh firmware has been ported from Debian Sid.

https://antixlinux.com/antix-19-3-manolis-glezos-bug-fix-upgrade-isos-available/

RELEASE OF SEQUOIA 0.20.0, A RUST IMPLEMENTATION OF OPENPGP:
17.10.2020

Sequoia 0.20.0 , which offers a library of functions and command line tools that implement the OpenPGP standards (RFC-4880), is out. Sequoia is being developed by three members of the GnuPG project, who decided to create a new free implementation of OpenPGP in the Rust language to improve the security and reliability of the code base, which provides tools for safe working with memory. The project code is distributed under the GPLv2

+ license.

In addition to improving security, Sequoia is also trying to fix the flaws in GnuPG, which cannot be fixed in the main project without breaking compatibility or significantly reworking the code base. For example, GnuPG has a fairly tight relationship between components that makes it difficult to make changes, obfuscates the codebase, and prevents the creation of a complete unit testing system. The gpupg command line toolkit is not synchronized in functionality with the function library and some actions can only be performed using the utility itself.

https://sequoia-pgp.org/blog/2020/10/14/202010-sequoia-0.20.0/

THE BRITISH BROADCASTING CORPORATION HAS INTRODUCED THE SECOND EDITION OF THE MICRO: BIT BOARD:
18.10.2020

The British Broadcasting Corporation (BBC) has presented the second edition of the open micro: bit board, developed for teaching programming and stimulating children's creativity (training courses have been developed for age groups from 7 to 11 and from 11 to 14 years old). The board will go on sale in November. The price has not changed at $ 15.

Design documentation, development-related technical specifications, schematics , bill of materials and developer documentation are published under the Solderpad license (a variant of the Apache License 2.0 for Open Hardware). Design files are available in formats suitable for processing by Eagle, Altium and KiCad.

https://microbit.org/news/2020-10-13/meet-the-new-bbc-microbit/

RELEASE OF THE PROGRAMMING LANGUAGE NIM 1.4.0:
18.10.2020

The system programming language Nim 1.4 is out. Nim is statically typed and built with Pascal, C++, Python, and Lisp in mind. Nim source code is compiled into a C, C++, or JavaScript representation. In the future, the resulting C / C++ code is compiled into an executable file using any available compiler (clang, gcc, icc, Visual C++), which allows achieving performance close to C, excluding the cost of running the garbage collector. Similar to Python, Nim uses indentation as block separators. Metaprogramming tools and capabilities for creating domain-specific languages (DSL) are supported. The project code is supplied under the MIT license.

https://nim-lang.org/blog/2020/10/16/version-140-released.html

OPENBSD 6.8 RELEASED:
18.10.2020

Introducing a release of the free cross-platform UNIX-like operating system OpenBSD 6.8. The release is made on the day of the twenty-fifth anniversary of the project, which was founded by Theo de Raadt in 1995 after a conflict with the NetBSD developers, as a result access to the NetBSD CVS repository was closed for Theo. After that, Theo de Raadt and a group of like-minded people created a new open operating system based on the NetBSD source tree, the main development goals of which were portability (13 hardware platforms supported), standardization, correct operation, active security and integrated cryptographic tools. Full install ISO size base system OpenBSD 6.8 is 530MB.

https://marc.info/?l=openbsd-announce&m=160303500224235&w=2

The VirtualBox Networking Primer
Connecting and Configuring Virtual Machines

The VirtualBox Networking Primer is a no-nonsense guide for the VirtualBox user taking their next steps into virtual networks.

While Oracle VM VirtualBox is a great free tool, the real power of virtualisation comes when you start connecting virtual machines to each other and to the wider world. Software development, sales, education and training are just some of the areas in which network access to virtual machines offers endless opportunities. But the world of computer networks is filled with complex technical jargon.

Complete with principles, practice, examples and glossary, The VirtualBox Networking Primer takes the frustration and confusion out of connecting real-world projects.

Author: **Robin Catling**
Publisher: Proactivity Press
ISBN13 : 9781916119482

Amazon US link:
https://www.amazon.com/dp/1916119484?ref_=pe_3052080_397514860

Amazon UK link:
https://www.amazon.co.uk/VirtualBox-Networking-Primer-Connecting-Configuring/dp/1916119484/ref=sr_1_1?dchild=1&keywords=virtualbox+networking+primer&qid=1600253699&s=books&sr=1-1

Kobo:
https://www.kobo.com/us/en/ebook/the-virtualbox-networking-primer

https://nmap.org/

Blurb: "*Nmap ("Network Mapper") is a free and open source (license) utility for network discovery and security auditing. Many systems and network administrators also find it useful for tasks such as network inventory, managing service upgrade schedules, and monitoring host or service uptime. Nmap uses raw IP packets in novel ways to determine what hosts are available on the network, what services (application name and version) those hosts are offering, what operating systems (and OS versions) they are running, what type of packet filters/firewalls are in use, and dozens of other characteristics. It was designed to rapidly scan large networks, but works fine against single hosts. Nmap runs on all major computer operating systems, and official binary packages are available for Linux, Windows and Mac OS X. In addition to the classic command-line Nmap executable, the Nmap suite includes an advanced GUI and results viewer (Zenmap), a flexible data transfer, redirection, and debugging tool (Ncat), a utility for comparing scan results (Ndiff), and a packet generation and response analysis tool (Nping)."*

Nmap is a tool I find very handy - unmissable, but people often say it is difficult to use. For those people, there is Zenmap. The command-line nmap has a bit of a learning curve, but once you understand it, it is really great.

I sometimes have to get my bearings on a network before working. The more you know, the better your diagnostics will be. However, some clients have shifted to this global IT nonsense, where you cannot do anything on the network without first creating a ticket, waiting three days and supplying blood and stool samples for each thing you want to do. If you do not, the IDS and firewall will shut you down. On top of that the global IT gets it wrong anyway, as they have phone monkeys on the other end and not real IT people.

This is where nmap shines. We can look at how to use nmap, how to sneak past IDS and defeat firewalls.

Let me scare some of you – type nmap in your terminal and just press enter. Look at all the options and read what it says. The rest of you say: "Challenge Accepted!" with me. It goes without saying, that you will be scanning your own home network and not fbi.gov.

Before we talk about nmap's options, I want you to head over to this website and read what a three-way handshake is: https://www.geeksforgeeks.org/tcp-3-way-andshake-process/ - The better that you know the underlying process, the easier you will pick up nmap. Feel free to read all the sections down the side of that page.

Now, nmap is known as a port scanner, but it is more than that. Most of the time, I simply use it to scan for live hosts and pipe the output to a text file. Security researchers and penetration

```
gewgaw @ Aangeneem  ~
 └ $ ▶ nmap -p- 192.168.1.1

Starting Nmap 7.60 ( https://nmap.org ) at 2020-02-07 17:53
Nmap scan report for _gateway (192.168.1.1)
Host is up (0.0016s latency).
Not shown: 65524 closed ports
PORT       STATE     SERVICE
22/tcp     open      ssh
23/tcp     filtered  telnet
53/tcp     open      domain
80/tcp     open      http
443/tcp    open      https
631/tcp    filtered  ipp
3000/tcp   open      ppp
5916/tcp   open      unknown
8081/tcp   filtered  blackice-icecap
20248/tcp  open      unknown
45998/tcp  open      unknown

Nmap done: 1 IP address (1 host up) scanned in 648.50 seconds
gewgaw @ Aangeneem  ~
 └ $ ▶
```

testers will use it in another way. The most common "port scan" I use is port 53 -DNS. Here is a list of common ports:

https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml

or:

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

Why am I telling you to look at port numbers? Well, with nmap, you may just want to scan for a specific port, say SSH on port 22, or the lazy admin obfuscation port 2222. Also 'port' is easy to remember as it is -p. Thus nmap -p 22 192.168.0.0/24 will scan for that port only on my network. By contrast, scan your same network without that "-p 22" and see how much longer it takes. You can also scan a range of ports, simply by specifying the start and end port. Like so: "-p 22-53". By default nmap only scans the 1000 most common ports, so if you would need to scan them all, you would use the "-p-" switch before the host IP. It is a good idea to do this against your IOT devices, to see what is open

(that should not be). This is a long scan, so it is best to do it against one IP at a time. Speaking of time, there is a time option in nmap too. You use the "-T" switch to set the time to wait for a response. However, the range is from 0-5, zero to five, five waits the least time (5ms) for an answer. In other words, 0 waits almost indefinitely, just kidding, but it feels that way. Most of the time I use T4 (recommended by nmap themselves) and usually skip the others. Feel free to try them too. To pipe your output to a file, simply add the redirector and the name of your file to the end of your query. (Example: nmap -p 22 192.168.1.1 > router.txt).

This article is pretty much hands on, you need to open a terminal and scan your network, router, or another live PC on your network.

By now you have tried:

```
nmap

nmap -p 22 <your IP range
here>

nmap -p 22-53 <your IP range
here>

nmap -T5 <your single IP
address here>

nmap -T4 <your single IP
address here>

nmap -p- <your single IP
address here>
```
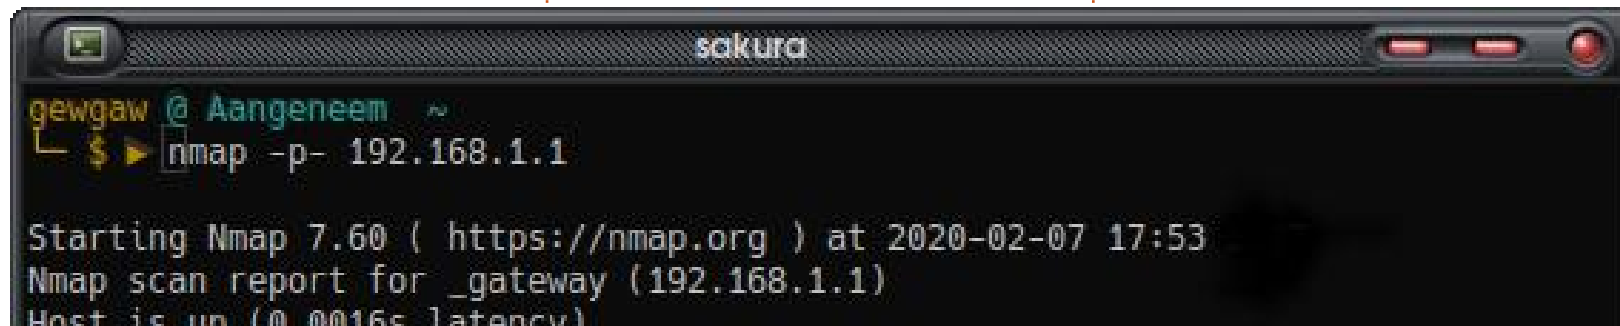
There is another "time" you need to know about and that is the Round Trip Time (RTT). This is the time it takes your packets to make a round trip. To know what this is, simply ping any website. The first value is TTL (Time To Live); if you are unsure what that is, it is the time in seconds, given to the data, to live, so it doesn't bounce around the network forever and ever. The next value you get is a time in milliseconds; that is what you are after. You are able to set this time too, extending the time you wait

for your packet to return. You use the rtt-timeout switch, but you need a prefix. For now I will cover only the "initial" timeout. Therefore, the switch is: --initial-rtt-timeout (two dashes at the beginning), and you need to give the value in milliseconds. Example: nmap --initial-rtt-timeout 200ms canonical.com You can also set the maximum timeout, but we will cover that later.

Let's use the five whys:

• **Why would we use nmap in the first place?**
• You need to reconnoitre the network.

• **Why would we need to reconnoitre the network?**
• We need to make informed problem determinations.

• **Why do we need to make informed problem determinations?**

```
 gewgaw @ Aangeneem  ~
 └ $ ► nmap -p- 192.168.1.1

Starting Nmap 7.60 ( https://nmap.org ) at 2020-02-07 17:53
Nmap scan report for _gateway (192.168.1.1)
Host is up (0.0016s latency)
```

• We can guess as to what the problem is, but it is always better to be sure. Sometimes one error can be very similar to another and some equipment can behave erratically and flood your network.

• Why do we make sure?
• Problems with networks can do many things; if a client cannot access their website, it could be anything from DNS to a network loop. Time is money and the sooner it is taken care of, the sooner business can resume.
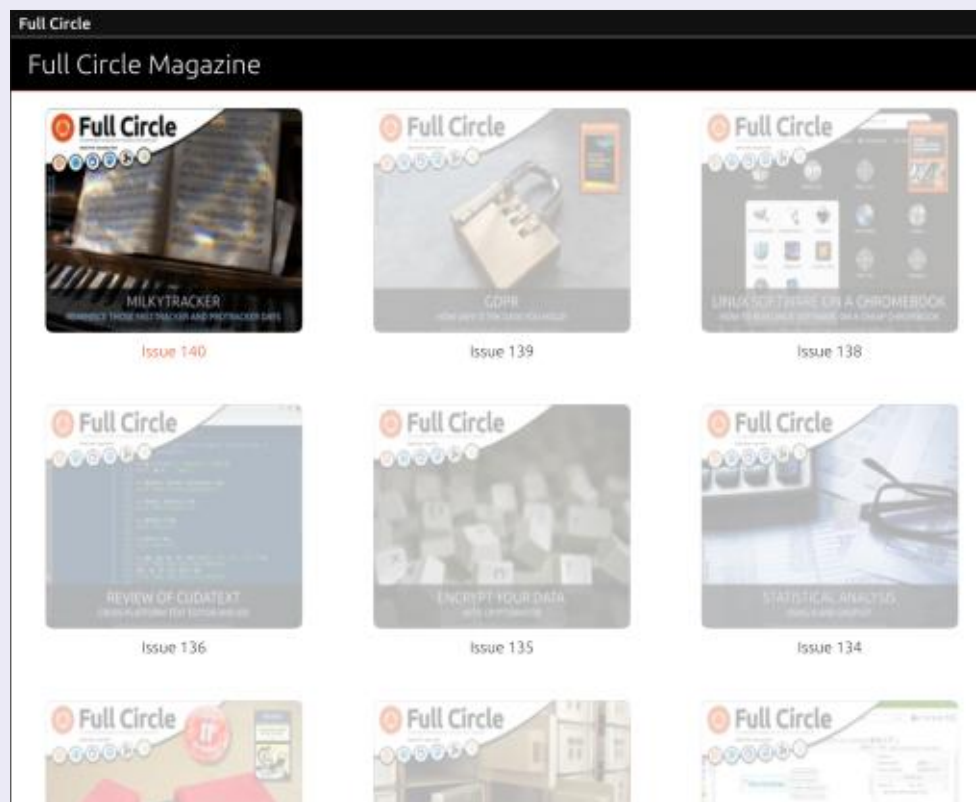
• Why are there so many vectors?
• That's why we have nmap!

Well I hope you enjoyed our toe-wetting session in nmap. Should you have any queries or comments, e-mail:
misc@fullcirclemagazine.org

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

## THE OFFICIAL FULL CIRCLE APP FOR UBUNTU TOUCH - UPDATED!

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

INSTALL

Either search for 'full circle' in the Open Store and click install, or view the URL below on your device and click install to be taken to the store page:
https://uappexplorer.com/app/fullcircle.bhdouglass

HUGE thanks to Brian for this.
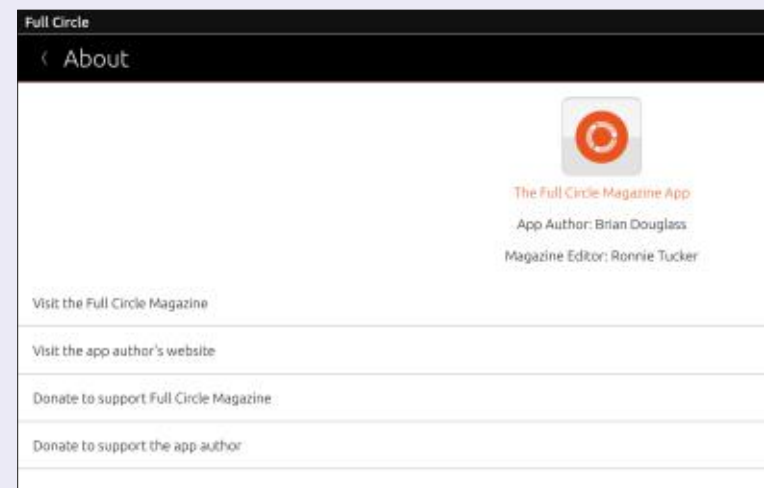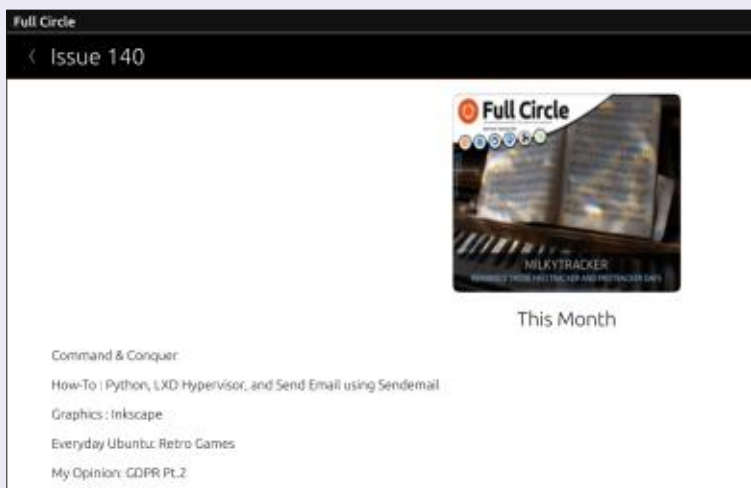
Last month, a PAGE user came to both Don and myself for some help with a fairly new Python library that implements a spreadsheet like table widget written completely in Python. He was fairly new to PAGE and was having trouble getting the table to initialize properly. I thought the library has a tremendous amount of possibilities, so I thought I'd share the information with you. Just so you know, you don't have to use PAGE to implement the sheet widget, but you do have to use Tkinter at least.

To get started, you need to install the library and, as usual, we can use Pip to do it.

```
$ pip install tksheet
```

Once you've done that, you might want to consider downloading the entire source code from the github repository at https://github.com/ragardner/tksheet. I suggest this since there are so many options that this library has, you might not catch some without digging into the source code. The documentation for the library is currently being written, so it's easy to suss out how to do some of the options by having the code available.

We'll use a modified version of the demo program to start with, which does not require PAGE. You can find the full version on the homepage of the repository.

```
from tksheet import Sheet

import tkinter as tk
```

As almost always, we want to import the important libraries, in this case, the tksheet and tkinter libraries.

Now (top right) we'll create a class called demo. This will hold all of the code for the demo program that deals with the sheet widget. As you go through the code, you

```
class demo(tk.Tk):
    def __init__(self):
        tk.Tk.__init__(self)
        self.grid_columnconfigure(0, weight=1)
        self.grid_rowconfigure(0, weight=1)
        self.frame = tk.Frame(self)
        self.frame.grid_columnconfigure(0, weight=1)
        self.frame.grid_rowconfigure(0, weight=1)
```

```
        self.sheet = Sheet(
            self.frame,
            page_up_down_select_row=True,
            column_width=120,
            startup_select=(0, 1, "rows"),
```

will see that many of the functions later on are simply there to allow you to customize the builtin functions that the library has.

As you already know, the __init__ function sets up various parameters and defaults. In the next lines, the sheet object is simply an instance of Sheet (above).

Here (below) is where you load the dummy data into the sheet.

Next (next page. top right), we enable the various bindings that are already coded in the library.

The next two lines insert the frame and sheet objects into a Tkinter grid.

```
self.frame.grid(row=0,
column=0, sticky="nswe")
```

```
        data=[[
            f"Row {r}, Column {c}\nnewline1\nnewline2" for c in range(50)
        ] for r in range(1000)],  #to set sheet data at startup

        height=500,   #height and width arguments are optional
        width=1200    #For full startup arguments see DOCUMENTATION.md
    )
```

```
self.sheet.grid(row=0,
column=0, sticky="nswe")
```

As you can tell from the commented code (bottom right), the next set of lines show ways to control various things like the theme, highlighting columns, rows and cells, etc.

The rest of the code (next page, right side) overrides the default events and bindings, and provides examples of how to handle them in your own code.

Finally, the required instantiation of the demo class object and the call to the Tkinter mainloop.

```
app = demo()
```

```
app.mainloop()
```

When you fire up the program, you should see something that looks like the image on the next page (top right).

While that wasn't too hard at all, I believe that using PAGE makes the process much easier. I've created a VERY simple PAGE GUI to show how quickly and easy it is.

Create a folder to hold your

```python
self.sheet.enable_bindings((
    "single_select",  #"single_select" or "toggle_select"
    "column_select",
    "row_select",
    "arrowkeys",
    "row_height_resize",
    "double_click_row_resize",
    "right_click_popup_menu",
    "rc_select",
    "rc_insert_column",
    "rc_delete_column",
    "rc_insert_row",
    "rc_delete_row",
    "edit_cell"))
```

```python
"""_____ EXAMPLES _____ """
"""_____"""
# _____ CHANGING THEME _____
#self.sheet.change_theme("light green")

# _____ DISPLAY SUBSET OF COLUMNS _____
self.sheet.display_subset_of_columns(indexes=[0, 1, 2, 3, 4, 5],
                                    enable=True)

self.sheet.insert_column(idx=0)
self.sheet.insert_columns(columns=5,
                        idx=10,
                        mod_column_positions=False)

# _____ HIGHLIGHT / DEHIGHLIGHT CELLS _____
self.sheet.highlight_cells(row=5, column=5, fg="red")
self.sheet.highlight_cells(row=5, column=1, fg="red")
self.sheet.highlight_cells(row=5,
                        bg="#ed4337",
                        fg="white",
                        canvas="row_index")
self.sheet.highlight_cells(column=0,
                        bg="#ed4337",
                        fg="white",
                        canvas="header")

# _____ CELL / ROW / COLUMN ALIGNMENTS _____
self.sheet.align_cells(row=1, column=1, align="e")
self.sheet.align_rows(rows=3, align="e")
self.sheet.align_columns(columns=4, align="e")

# _____ ADDITIONAL BINDINGS _____
```

project and start up PAGE within that folder. Expand the default Top level form just a bit and place two buttons near the very top of the form, one on the left and one on the right. The left one should have the text "Load CSV", and the one on the right "Exit". In the command attribute for the left button enter "on_btnLoad" and for the right button "on_btnExit". (We've covered PAGE programming so often, this should be very obvious, but if it isn't, look at one of my previous articles about using PAGE. My article in FCM#155 should be a good easy reference).

Next, place a frame in the form that takes up pretty much of the rest of the form. Then place a custom widget into that frame and expand it to fill the frame. (Mouse-3 (Right-click for most people)) | Widget | Fill Container). Save your project as tksheetGUI.tcl and generate your python code.

Bottom right is what your project hopefully looks like. If not, that's ok. You get the general idea.

Now open tksheet_support.py in your IDE or editor, and let's get to work.


```python
def all_extra_bindings(self, event):
    print(event)

def begin_edit_cell(self, event):
    print(event)   # event[2] is keystroke
    return event[
        2]  # return value is the text to be put into cell edit window

def window_resized(self, event):
    pass
    #print (event)

def deselect(self, event):
    print(event, self.sheet.get_selected_cells())

def rc(self, event):
    print(event)

def ctrl_a(self, response):
    print(response)

def row_select(self, response):
    print(response)

def column_select(self, response):
    print(response)
```

PAGE only creates a single import line, since that's all that you need to get started. Add the following lines to the import section of the code to support our needs. The line that PAGE gives us is not in bold.

```
import sys
import platform
import os

# Third party libraries
from tksheet import Sheet
import pandas as pd
```

We also need to modify the import section that PAGE provides for us. Normally, you would not need to do this, but we are going to add support for the Tkinter messagebox and the filedialog sub-systems.

```
import Tkinter as tk

import tkFileDialog as
filedialog

import tkMessageBox as
messagebox

except ImportError:

import tkinter as tk

from tkinter import
messagebox

from tkinter import
filedialog
```

Note that we actually don't do anything with the messagebox but it's there for your possible future development.

Now (top right) we need to add a couple of lines to the init function, that again, PAGE provides for us.

As I'm sure you are aware, this (middle right) will call the startup() and init_custom() functions before the form is shown to the user. Let's look at the startup function first.

Yes, there is some code that isn't strictly needed, but I like to add it. Basically, we get and provide the python version, program path and program version, and print all of that to the terminal window. Then we set up some global default values for use with the tksheet library.

In the init_custom() function, we will be initializing various settings of the tksheet for our use.

```
def init(top, gui, *args, **kwargs):
    global w, top_level, root
    w = gui
    top_level = top
    root = top
    startup()
    init_custom()
```

```
def startup():
    global version
    pv = platform.python_version()
    print(f"Running under Python {pv}")
    # Set the path of the program
    global progpath
    progpath = os.getcwd()
    print(progpath)
    version = "0.0.1"
    print(f"Version: {version}")
    # ================================================
    # Some default global values
    # ================================================
    global defaultColumns, defaultRows, useBoldFont, defFont
    defaultColumns = 30
    defaultRows = 50
    useBoldFont = False
    defFont = "Arial"
```

It took me a fair amount of time to narrow down the actual settings needed for proper demo use. There are many other settings as you saw in the earlier demo program, but, for us, these are needed. One thing to note is the theme setting (see below). There are four default themes defined by default. The first demo used the light green theme. For the PAGE version, we will use the dark blue theme.

```
def init_custom():
    w.Custom1.page_up_down_select_row = True
    w.Custom1.column_width = 120
    w.Custom1.startup_select = (0, 1, "rows")
    w.Custom1.theme = "dark blue"
    w.Custom1.insert_column(idx=0)
    w.Custom1.insert_columns(columns=5, idx=10, mod_column_positions=False)
```

Now (top right) we need to enable the bindings that we want so that the functionality is what you would normally expect for a spreadsheet demo. These are things like selecting a single cell, a row or column and using the arrow keys as well as allowing for "right-click" (mouse 3) context menu support.

The next few lines (bottom right) allow you to use the popup context menu. Some of them are pretty obvious, but it's not as clear why the last 6 lines would be needed.

The last part of the init_custom function (bottom right) deals with binding the Mouse-3 button to the "rc" routine and sets extra bindings for us to be able to override the default functions within the library.

Now (next page, top right) we'll create the four callback functions that we just defined. I must admit that I created the first three for a different demo that printed the information to a status label, but I felt that, for this very simple demo, printing to the terminal was ok.

I created a simple class to handle the actual tksheet as a

```python
w.Custom1.enable_bindings(
    (
        "single_select",   # "single_select" or "toggle_select"
        # To support column/row select
        "column_select",
        "row_select",
        # To support column/row resize
        "column_width_resize",
        "double_click_column_resize",
        "row_width_resize",
        "column_height_resize",
        "arrowkeys",
        "row_height_resize",
        "Double_click_row_resize",
```

```python
        "right_click_popup_menu",
        "rc_select",
        "rc_insert_column",
        "rc_delete_column",
        "rc_insert_row",
        "rc_delete_row",
        "hide_columns",
        # The following 5 lines must be included for popup support
        "copy",
        "cut",
        "paste",
        "delete",
        "undo",
        "edit_cell",
    )
)
```

custom control. See next page, bottom right.

Almost at the end now. We still have to assign the library as the custom control. PAGE supplies a line that says "Custom = tk.Frame", but we need to replace tk.Frame with our external tkinter control library. You can add the following

```python
w.Custom1.bind("<3>", rc)
w.Custom1.extra_bindings(
    [
        ("cell_select", on_cellSelect),
        ("column_select", on_columnSelect),
        ("row_select", on_rowSelect),
    ]
)
```

line and comment out the 'Custom = tk.Frame' line.

```python
Custom = sheet
```

Finally, we need to provide the callback and support functions for the buttons that we defined when we designed the GUI (next page, top right). PAGE generated the first two functions for us as skeletons. The final function helps the import of the CSV files. You already have seen the proper way to end a PAGE application, but I've included it here for your convenience.

Now (next page, middle) we have the callback for the load csv file button. It's a little bloated, but I thought it would be easier for you to follow the code this way rather than write it more elegantly.

Last but not least, here (bottom right) is the load_csv_file function. Here we are using pandas to read the CSV file and convert it into a list. That way, the tksheet library can utilize it.

Be sure to save your file.

Last but not least (again!), you need to know where a good CSV file is. For this demo, I have used the titanic3.csv file which is provided with the source code for the Pandastable Library. I'm certain it's been released into the public domain. I copied it and put it into

```python
def on_columnSelect(p1):
    print(f"on_columnSelect Triggered {p1}")
    print(p1)
    selcol = p1[1]
    colname = headers[p1[1]]
    disp = f"Column Selected – Col: {selcol}    ColName: {colname}"
    print(disp)


def on_rowSelect(p1):
    print(f"on_rowSelect Triggered {p1}")
    print(p1)
    selrow = p1[1]
    disp = f"Row Selected – {selrow}"
    print(disp)


def on_cellSelect(p1):
    global headers
    print(f"on_cellSelect triggered: {p1}")
    selrow = p1[1] + 1
    selcol = p1[2] + 1
    colname = headers[p1[2]]
    disp = f"Cell Selected – Row: {selrow}    Col: {selcol}    ColName: {colname}"
    print(disp)


def rc(event):
    print(event)
```

```python
class sheet(Sheet):
    def __init__(self, parent, **kw):
        Sheet.__init__(
            self,
            parent,
            page_up_down_select_row=True,
            column_width=120,
            font=("Arial", 12, "bold"),
            theme="dark blue",
            height=500,  # height and width arguments are optional
            width=1200,  # For full startup arguments see DOCUMENTATION.md
        )
```

the development folder for quick access.

Once you run the program, it should look something like this:

I hope that you enjoyed this month's project as much as I did creating it. This library has a lot of potential and will be great for quick display of things in table or sheet format. However, there is no code to write back to a file (that's simple enough for you to work out) and there are no calculation functions at this point. But this is a fantastic library and something that has been needed for a very long time!

As I did once before, I've created a repository on github to hold the code and images from this article. You can find it at https://github.com/gregwa1953/FCM162 .

```python
def on_btnExit():
    print("tksheet_support.on_btnExit")
    sys.stdout.flush()
    destroy_window()
```

```python
def on_btnLoad():
    print("tksheetdemo1_support.on_btnLoadCSV")
    sys.stdout.flush()
    global headers, progpath
    filename = filedialog.askopenfilename(
        initialdir=progpath,
        title="Select file",
        filetypes=(("CSV files", "*.csv"), ("all files", "*.*")),
    )
    data = load_csv_file(filename)
    # print(headers)
    w.Custom1.headers(headers)
    # Now load data into sheet
    w.Custom1.set_sheet_data(data)
    # Get the number of rows and columns
    global totalrows, totalcols
    totalrows = w.Custom1.total_rows()
    totalcols = w.Custom1.total_columns()
    # Print the information to the terminal
    print(f"Total Rows: {totalrows}")
    print(f"Total Cols: {totalcols}")
```

```python
def load_csv_file(filename):
    global headers
    # Load the local csv file into a pandas dataframe
    df = pd.read_csv(filename)
    headers = list(df.columns)
    # Convert it into a list
    dl = df.values.tolist()
    # Return it to the calling function
    return dl
```

As always, until next time; stay safe, healthy, positive and creative!

So this month we will cover how to create your RSS feed, to push your audio directly to the pod-catchers of those people who subscribe to the feed to receive updates as you publish new content.

First, a warning about setting up the feed for your audio. Although the audio is going to be hosted elsewhere (and not actually on the Blog servers), you are limited to a grand total of 1MB of data (yes, less than an old 3.5" floppy's amount of data) for each post being fed via feedburner. If you go over this, the post will not get fed to people subscribed to the feed; it is easy to breach the 1MB limit even with just a long set of show notes without any pictures. This caused me endless frustration when I started to publish the Distrohoppers Digest pages. I was trying to push the feed without success – until, after quite a bit of Googling, I was able to find out what the problem was. As a result, I now publish the show notes as a separate page, with a link after the first paragraph to ....'more' on the

actual Blog post. This allows more text and images if needed, and as this is not on the feed, the page size should be well under the 1MB limit, allowing the feed to be processed.

OK, let's look at setting up your Feed. First, you will need a link to where your audio is stored, either on your self-hosted server, or to the audio stored on archive.org as we do with Distrohoppers and mintCast. If using archive.org, you may want to consider a small monthly donation to them to offset some of their hosting costs; while it is not compulsory to do this, if you are donating 1£/2$ per month, this is 12£/$24 per year, and that is good value for hosting; other service providers start at around $5/month, $60/year. Setting up an account with archive.org is as simple as providing an email address and password, and, if using it for podcast hosting, you can use the email created for the Blogger account. We will look at this further in a future article.

So, to be able to publish the

feed, you need to link your Blog to Feedburner, so navigate to feedburner.com from your google account and you should come to this page. As you can see, there are already a number of feeds in the list for my blogs, but we want to create one for our 'Test Blog' which has the address of 'testtonyh1212.blogspot.com'. To create a feed address for this blog, we need to create a new feed. Under the list is the message:

"Burn a feed right this instant. Type your blog or feed address here:"

in the box below this, put in the blog's web address and tick the box that says 'I am a podcaster'.

Click next, and you will see this box with the options to make an atom or RSS feed, check the RSS feed.

Then click next, and you should have successfully created your feed address which will be showing in the box on the screen.

If you wish, you can just skip straight to managing the feed. At this point, while you are here, you might as well set up the finer points of the feed such as how iTunes will see the feed and key words to assist people finding your content. The picture shown right shows this process in one image including now being able to see the feed on your Feedburner home screen; Test is now the Second Blog listed.

So you now have feedburner set up with an RSS address that you will use later to set up the feed on your Blog. In my case the test Blog feed reads like this: http://feeds.feedburner.com/blogspot/

PSYbSL, with the end PSYbSL being the unique identifier for the feed. It is this full http://feeds... address that we will be using when we create the Feed 'Gadget' for your Blog's home page, so our next step is to navigate to our Blogger account and start setting up the feed.

On the left-hand of your Blogger page will be the menu with all the options to create a new post, etc; navigate to the 'Layout' and open this; you will be presented with something like this (shown above), you may need to scroll down a little to see the right-hand sidebar.
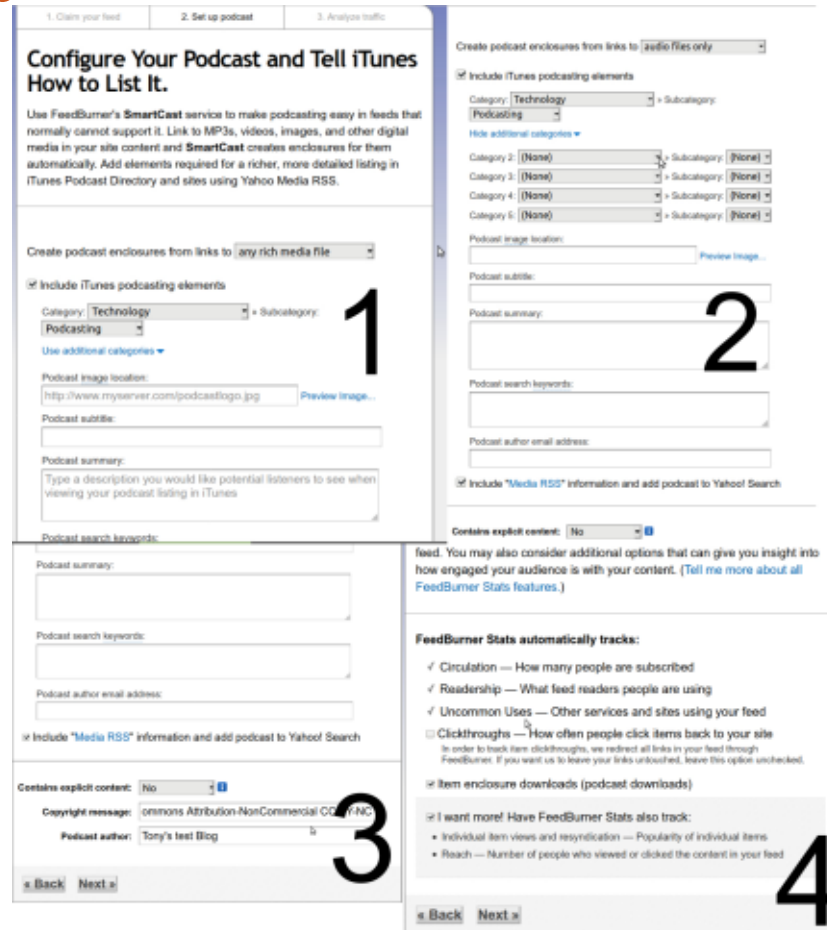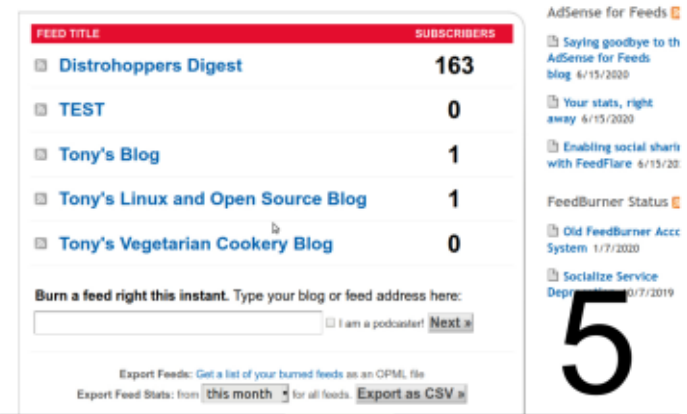
HOWTO - PODCAST PRODUCTION

Click on the 'Add a Gadget' and this screen (shown right) opens.

Again you may need to scroll down to see the 'Feed' section. Click the '+' to open the dialogue box to add your feed details. The following picture is a compilation of following the screens to create the Gadget. After putting in the Feedburner address we created earlier, we then put in our title for the feed and preference for the number of posts that will be seen by pod-catchers, the maximum being five. Save this and the Gadget now shows in the Layout; if you now navigate to the Blog home page, you should see the title you gave the feed next to RSS.

So now anyone with the feed URL that you put into the Gadget will be able to subscribe to that in their pod-catcher, and every time you create a new Blog with audio linked in it, they will have that audio or a link to it pushed to their pod-catcher.

Next time, I will go through uploading your audio to archive.org and then linking this in your blog post each time you record a new episode.
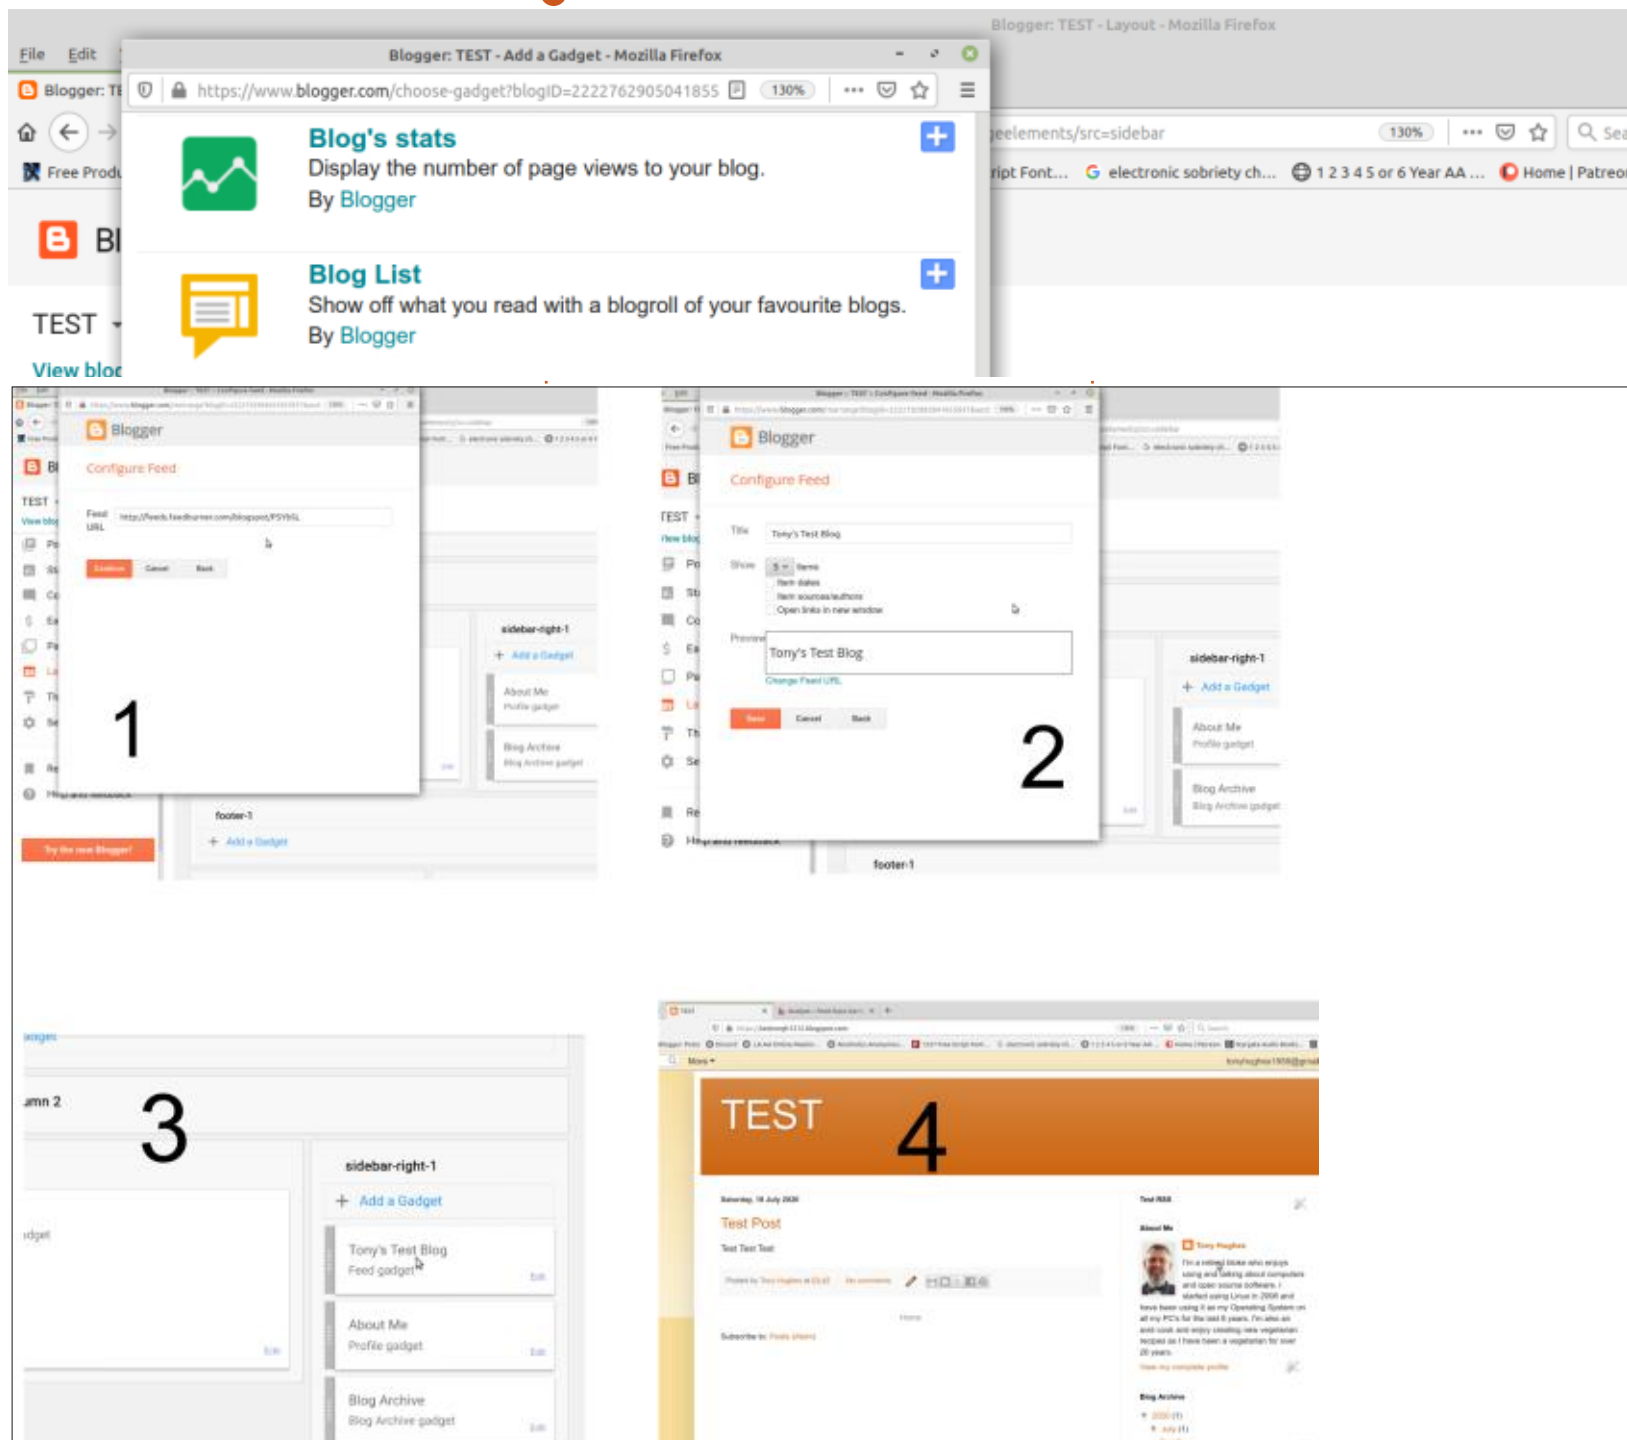
To take a photo, the ideal is for the light to be behind you. We know that this is not always possible: taking a picture of your dog, when he is being adorable, but with the sun behind him, or your friends goofing around at sunset or sunrise. We have all been there. That photo may not be a total waste after all.

In a recent discussion on black-and-white photography, the discussion on applications, like Snapseed's (- not picking on Snapseed, it is just the most well-known in our circle) ability, or inability, to make black-and-white photos look real, came up.

Just taking the color away from a photo does not make a nice black-and-white photo. Doubly so when you have applied post-processing to it. Too often, we see people posting pictures in forums where the edges are darker than the body.

How do these two topics fit together, you may ask? The answer is Rawtherapee.

When you have a photograph where the sun is somehow shining into your lens, be it behind a cloud or peeking over a rooftop, it is not always the easiest to recover. This can be compounded when you took the picture in jpeg-format. In this issue, we want to tackle a photograph like that. Just to go from bad to worse, I will be using a jpeg instead of a RAW image.

This (bottom left) is what we will be working with. As you can see, the sun decided to peek out from behind the cloud as I tried to take a picture of the corpuscular rays. Since there is nothing of interest in the photograph, I would usually just send it straight to the bin. Let's assume there was something in the image that I wanted. If we look at our histogram, you will see that black is actually squashed up to the edge as well as white.

If I turn on clipping, you can see the problem immediately. Hectic loss of detail. (See the right-hand side and the bottom of the photograph.)

Turning this image into black-and-white should lessen the loss of

detail. Alt+C, or clicking the third icon, will bring up the color menu. Find black-and-white, the fourth sub-menu and turn it on. This does not remove the loss of detail we can see with the clipping filters on, but it puts it into the realm of absolutes of light and shadow. Now we have a black-and-white picture that is still boring, but the loss of detail had become deep shadows.

The next tool we are going to use is "local contrast". This falls under detail, so Alt+D or click the second icon.

A word of warning though, improper use of the local contrast can make your black-and-white photograph look extremely "shopped".

The key here is the first two sliders, radius and amount. The larger the radius, the more feathering you will get from the edges. However, radius alone does nothing. This is where the amount slider comes in. If you are in the artistic mood, slide both the radius and amount to maximum. This will bring some drama to the photo, unless you have uniform dark areas that are not black. If you have those, it will look like a kid took a marker to the edges of the picture, before coloring it in with a lighter color. I actually encourage you to do this, then play with lightness level and darkness level to acquire the desired effect. Just keep an eye on your histogram, as we already have clipping that we are trying to hide. You do not want to compound your problems.

Drama!!

But let's assume you want clarity, rather than drama? We need to increase our zoom to 100%.

Tip: don't fret over details in the photo that are going to be cut in the end. By this I mean, if you are planning to crop your picture and you were going to cut off the left hand side, don't worry about how your edit affects the object on the left. It is going to end up on the cutting room floor, so to speak.

At 100%, I can zoom into the roof at the bottom and show you the difference between a radius of 20 (minimum) and a radius of 80. This is what you have to look out for when you use local contrast on a picture.

Compare these two images (below)

Pay attention to the amount of visible tiles on the roof, from the edge. When using this effect on color photos, you may see a light discoloration on the side of the light, and dark discoloration on the side where it is dark. This is one very quick way to tell if a photo has been put through a filter. This is
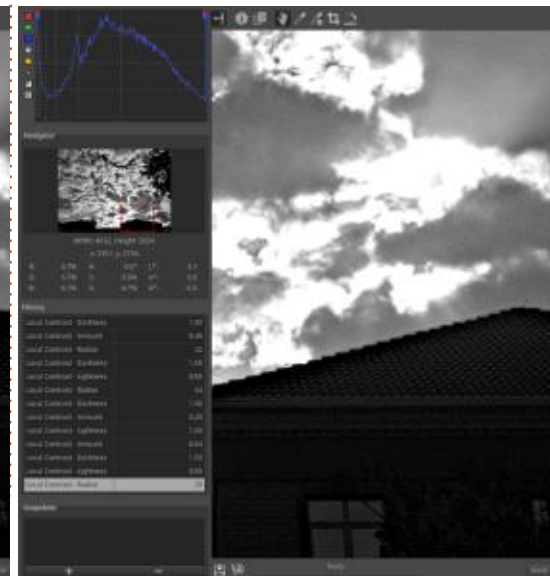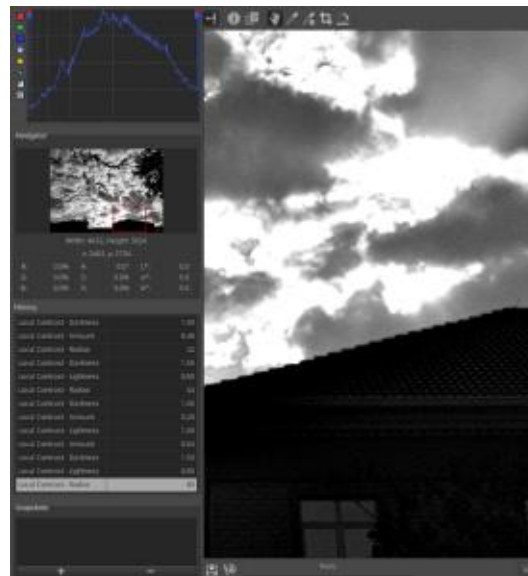
also the difference between editing by hand and using an application like Snapseed. Once your eye is trained to this, it is very hard to "unsee". You will see this in some so-called "professional" photographs on the web, or even magazines. A good rule of thumb is that, when you see the dark edges start to appear, you can either decrease the radius and/or decrease the lightness, to "smooth things out".

I urge you to take a few photos and play with only this effect. It does not have to be black-and-white. That part was just to show you that you can save boring photographs or even a clipped one, by making it black-and-white, where that clipping does not matter as much any more.

Tip: Turn the effect on and off while being zoomed in, as well as being zoomed out.

Local contrast is for sharpening; you do not want to add rifts or halos to your image.

Once you are done, crop your image with the "C" key and you are good to go.

I hope you had fun if you followed along and I hope you learned something.
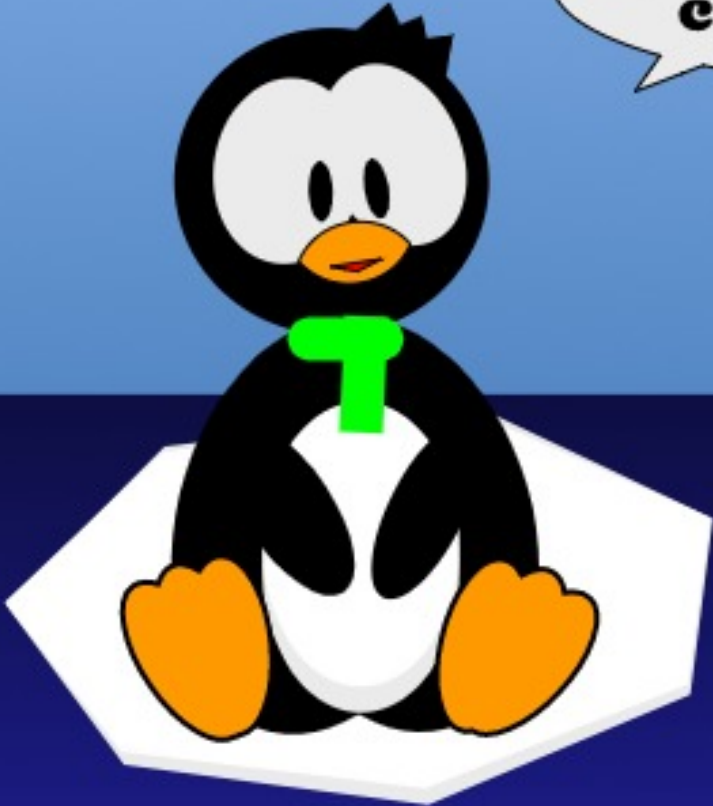
Disclaimer: In no way am I an expert in photo manipulation, I just know a guy who had a dog that gave me some fleas. I thought I may help you scratch an itch also.

As always, if you have any questions or comments, send them to: misc@fullcirclemagazine.org

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

As you may have noticed in the News section of the previous edition of FCM, a new minor release of Inkscape is out. As is often the case with these things, the announcement came only a couple of days after the deadline for my previous article. As it turns out, version 1.0.1 doesn't fix any of the issues I've raised in recent months, or change any of the features that I've described. Phew! No corrections required.

So what is new in the latest release? It's mostly bug fixes and stability improvements, though there is a whole new "Selectors and CSS" dialog that is most likely to be of interest to web developers. I'll cover it in due course, but it's not the topic for this month's instalment. Instead, I'll cover some changes to the 'Path > Stroke to Path' function. To better understand this menu entry, however, I'm going to start with its sibling.

'Path > Object to Path' is a mainstay of any experienced Inkscape user. As the name suggests, it converts your object to a path, and is therefore commonly used when you want to break out of the design shackles imposed by an object's native type. Whilst a rectangle can have only its width, height and corner radius modified, converting it into a path lets you move individual nodes, add more, or delete others. Since paths are so flexible, it can open up a world of design possibilities – albeit at the expense of losing access to the specialized editing tools for the original object type.

There's nothing controversial about Object to Path. The end result maintains its fill and stroke properties, so there's no visual change as a result of the operation. All that's happened is that your native Inkscape object becomes a generic path of the same size, shape and appearance.

Despite the similar name, however, Path > Stroke to Path is an entirely different beast. At its core it converts any stroke that you may have on your object into a new filled path which matches the stroke's original outline. If it sounds confusing, perhaps some examples will help. Let's start with the simplest example possible: a straight line.

The top line is our original path, consisting of two nodes, no fill, and a thick, red stroke. As a rule it's easiest to understand what Inkscape is doing when applied to thick strokes, but everything I describe can also be done with thinner strokes if that suits your requirements.

The bottom line is the result of the Stroke to Path operation. You can see that what we now have is a filled path consisting of four nodes arranged to match the original stroke's outline. The fact that it's a filled path is a really important point to grasp: the original shape had a stroke but no fill, whereas the new shape has a fill but no stroke. The fill color of the new shape is the same as that of the original's stroke color in order to produce a result that is visually identical.

Why would you want to perform such a conversion? Consider trying to make a line that looks a little more hand-drawn. SVG has no support for variable width strokes, but you can fake it by converting your stroke to a path and then tweaking the shape.

These days Inkscape can simulate variable width strokes using Live Path Effects, but that wasn't always the case. The LPE approach also supports varying the

thickness only symmetrically, whereas this manual approach lets you achieve effects such as thickening the stroke on one side of the center line whilst thinning it on the other.

Let's look at a more complicated example. This time we'll jump straight to a multi-segment line with a dashed stroke applied. Once again the original line is at the top, and the bottom shows the effect of converting the stroke to a path.
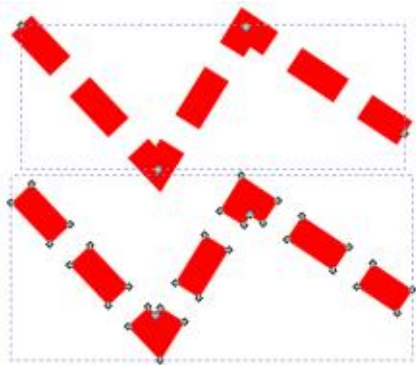
Our converted line is starting to look a bit more interesting. Instead of a simple filled path, we now have a complex path, consisting of a number of filled sub-paths which match the visual appearance of the original line.

At first, this can be a little tricky to get your head around. Because

the resultant shape still looks the same, it's not obvious that it's actually now a filled path that has no stroke. Like any other path, you can actually add a stroke to it; here's the same result but with a thin, dashed, black stroke added to the converted line:
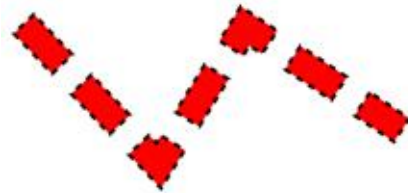
Hopefully it's now pretty clear what Stroke to Path does when presented with the simple case of an object with just a stroke applied. But what happens when your object is a bit more complex? Here's a line with a stroke applied, but also with markers at the start, end, and at each intermediate node.
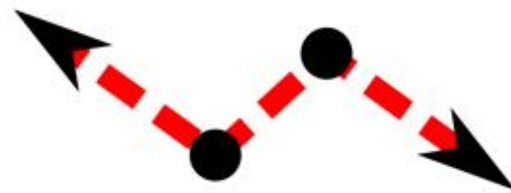
What do you think should happen when Stroke to Path is applied to this object? To my mind there are three possible

alternatives:
• Remove the markers and convert the stroke as before.
• Convert the stroke, but also turn the markers into filled paths.
• Convert the stroke, but copy the markers to the new paths, so that each sub-path ends up with multiple markers.

Really old versions of Inkscape went with option 1. Stroke to Path converted the stroke to a path, and ignored everything else. But since version 0.44 (released in 2006), Inkscape uses the second option (no version uses the third option). This, however, is where not-so-subtle differences in behaviour for 1.0.x start to creep in.

Consider the shape above. Up to version 0.92 using Stroke to Path on this would have created a group containing five objects: the complex path generated by converting the stroke, and a separate path for each marker. In version 1.0.x the result is somewhat different. Now the output is a group containing two elements: the complex path, plus a nested group which contains four more deeply nested groups – one for each marker. Each of those marker groups contains two paths,

one each for the fill and the stroke of the marker. Yes, I said the stroke of the marker. I know you probably weren't even aware of markers having a separate stroke, but apparently they do and they're now converted into a path of their own.

If this sounds a little confusing, perhaps an image of the structure will help. On the left we have the Objects dialog from 0.92, and on the right we have one from 1.0.1 (excuse the different themes – the 1.0.1 snap is still broken with the standard theme, so I'm using the symbolic theme):
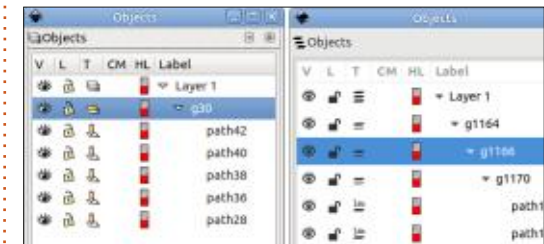
As you can see, the structures are substantially different. If you do want to convert a 1.0 arrangement to the old structure, it's quite simple: you just need to use Path > Union on each pair of marker paths (this is easier if you select them in the Objects dialog rather than on-canvas), then select the group that contains the

markers (i.e. not the top-level group, but the one just inside that), then repeatedly ungroup until you're left with a structure that looks like the one in the 0.92 dialog. Yes, "it's quite simple" was sarcasm.

For most users, this may never present a problem. The sort of use-case that requires you to add markers to your path does not generally intersect much with the use-cases for converting the stroke to a path. So, let's look at a far more common scenario: a path with both a stroke and a fill, but no markers. Here's what happens in 0.92.x, again with the original shapes at the top, and the Stroke to Path conversions at the bottom:
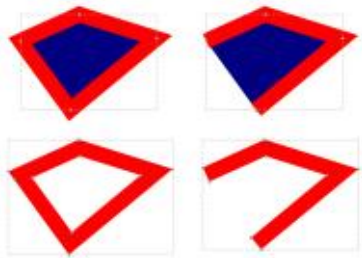
The obvious take-away from this is that 0.92 removes the fill entirely before the stroke is converted to a path. I've shown both closed and open shapes to make it clear that the behaviour is the same in both

cases. Structurally, you end up with a single filled path, just the same as you would if you'd performed the conversion on a shape with no fill or markers, the same as the first example I showed in this article.

Since version 1.0, however, performing a Stroke to Path operation on a shape with a fill results in a group that contains two filled paths: one is the usual path following the shape of the original stroke, whilst the other is a path representing the fill (i.e. it's just a copy of the original path, but without the stroke). The end result is visually identical to the original object.

As I'm sure you've guessed by now, performing a Stroke to Path on a shape with the holy trinity of stroke, fill and markers, results in a group containing:
• A path for the stroke
• A path for the fill
• A group containing an individual group for each marker, with each of those groups containing a path for the marker's stroke and a path for the marker's fill.

There are a few things to unpick from these changes. First of all, if you have an object with just a

stroke then the behaviour is the same as it has always been: you end up with a single filled path, which will have sub-paths if the stroke was dashed.

If your path has a stroke and a fill, however, you'll now end up with a group, rather than the fill being thrown away. This is particularly important to note if following an older tutorial. Many of them either rely on the fill being removed, or instruct you to duplicate the object before applying Stroke to Path so that you don't lose the fill. To get the same result as earlier versions, you can do one of two things:
• Remove the fill before using Stroke to Path.
• After using Stroke to Path, ungroup the result and delete the new path containing the fill (or just keep it, if you need it for subsequent steps).

Despite the confusion this change has caused with some new users, I generally think it's an improvement. Consider the case of drawing a simple cartoon character, with thick black outlines (strokes), and colored clothes and skin (fill). If you want to add some variety to the stroke thickness, the new

behaviour makes it easier to do so without having to duplicate every object, then remove its stroke, just to keep a copy of the filled shape.

Markers are another matter, however. While dealing with an unwanted fill just requires a couple of extra steps, getting the 0.92 structure when markers are in play requires several steps – multiplied by the number of markers – with no easy shortcuts. Although the new functionality may technically be more flexible, in practice there are few use-cases for separating the fill and stroke on a marker. It would have been nice to either have the old method available as a preference, or to include a function or extension that would make it easy to union deeply nested paths and recursively pop them out of their groups.

The take-away from all this is that Object to Path still behaves the same way it always did, Stroke to Path may need an extra step or two if you have a fill, but if you have markers as well then good luck to you!

Next time, I'll start to look at the changes to core drawing tools made with version 1.0.

This series is aimed at learning to make something of the old photos in my possession, and others in the public domain due to their age. You, the reader, are welcome to tag along, and, I hope, glean some small insight and perhaps an idea or two from time to time. No promises are made as to quality of the content, or potential errors and omissions. I am a computer scientist, not a true artist or a professional of image restoration. So please take all this as a best effort, but with no firm guarantees — much as is the case of most open-source software.

In the previous part of this series, we worked on a technique that is often used in museum photo recreations, and consists of bringing out the main subjects of an image in color, while the background remains in black and white. In this final part of this series on using Krita to rework old photos, we will stay with early digital images and work on color density and enhancement.

Exhibit A for this article is a photo I took in Hong Kong, just after the Olympic Games back in 2008. At that time, digital cameras were beginning to get better, and problems with lack of resolution were starting to be resolved (but can one really ever have sufficient resolution?). On the other hand, modern image enhancement techniques were absent, such as High Dynamic Range (HDR) in which several pictures of the same scene are taken at varying exposure points, then combined to build a single photograph with more details both in the shade and in highly illuminated parts of the image. Other color enhancement schemes are now commonplace.

Some of these techniques would have been useful to me back in 2008. In this image, we get some interesting colors from the setting sun (to our left, outside the field of view), but at the same time there are sharp shadows in which all detail has been lost. This is perhaps most noticeable in the buildings at the far left, whose facades are in-shadow, and present a perhaps slightly ominous and not very attractive backdrop to the harbor.

The first thing to do, as usual, is take a look at the histogram.
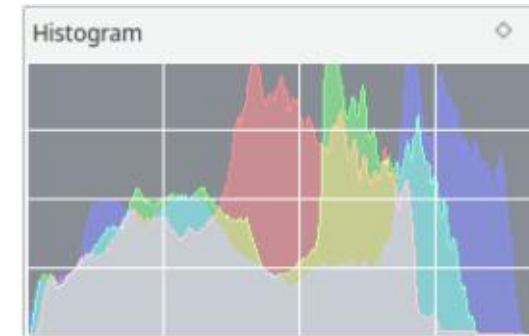
It is interesting to note the contribution of different areas in the image to the overall pixel count. For instance, the clear blue sky and part of the water would seem to make up a large proportion of lighter pixels. However, the large blue pixel count seen to the right of the histogram actually denotes a lack of blue and

excess yellow in lighter colors of the histogram. On the other hand, there is a large amount of reddish pixels in the mid-tones in the histogram, that in fact correspond to an excess of blue in these shades. So, to correctly interpret this histogram, we need to think in terms of the water and sky actually being situated in the mid-tones of our histogram -- they are not actually the lightest pixels of our image. These would be, rather, the yellowish reflections of the chromed railing and the reddish wall to the right of the image.

Finally, most dark pixels would seem to have some lack of blue to them, which is coherent with the darker shades of water and the buildings in the shade -- shadows often contain a slight tinge. However, there is a fairly lower amount of dark pixels, than light.

To correct these aspects and extend, as much as can be, the dynamic range of our image, it would be nice to accentuate the gradient of colors in the lower part of the histogram, giving lighter shadows a tad more detail while retaining some darker shades as such. On the other hand, it may also be worth extending some mid-range pixels up to lighter colors, though only for the red and green channels. The red channel already extends right up to the right-hand limit, and cannot be adjusted further. So, let us go into the menu option Filter, Adjust, and Color adjustment curves. Selecting the Lightness channel, let us move up a tad the lower curve section:

This takes care of the lower (darker) section of the histogram. Now, let's go back once more into Color adjustment curves, and now choose to modify only the Blue channel. When a single primary channel is selected in this tool, a specific histogram of this color is presented within the adjustment curve, and we can indeed see that there is a distinct lack of blue pixels in the lighter (right) part of the curve. Let us adjust this, by bringing the top right-hand extremity of the curve to the left, until it is above the right-most edge of the histogram edge. A similar
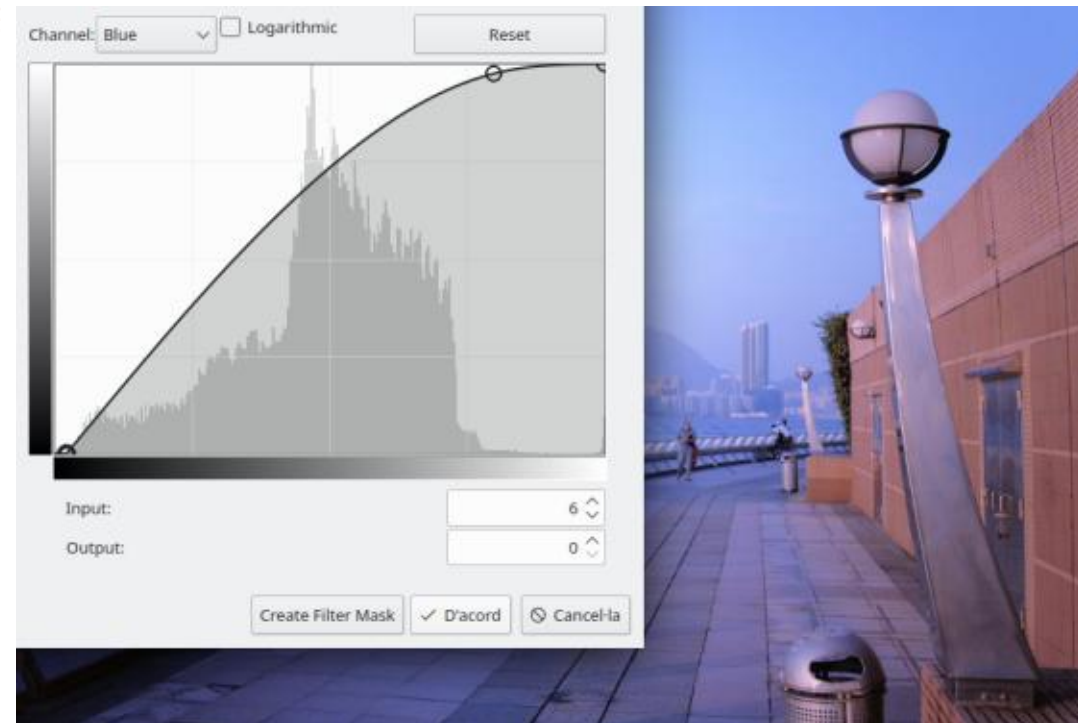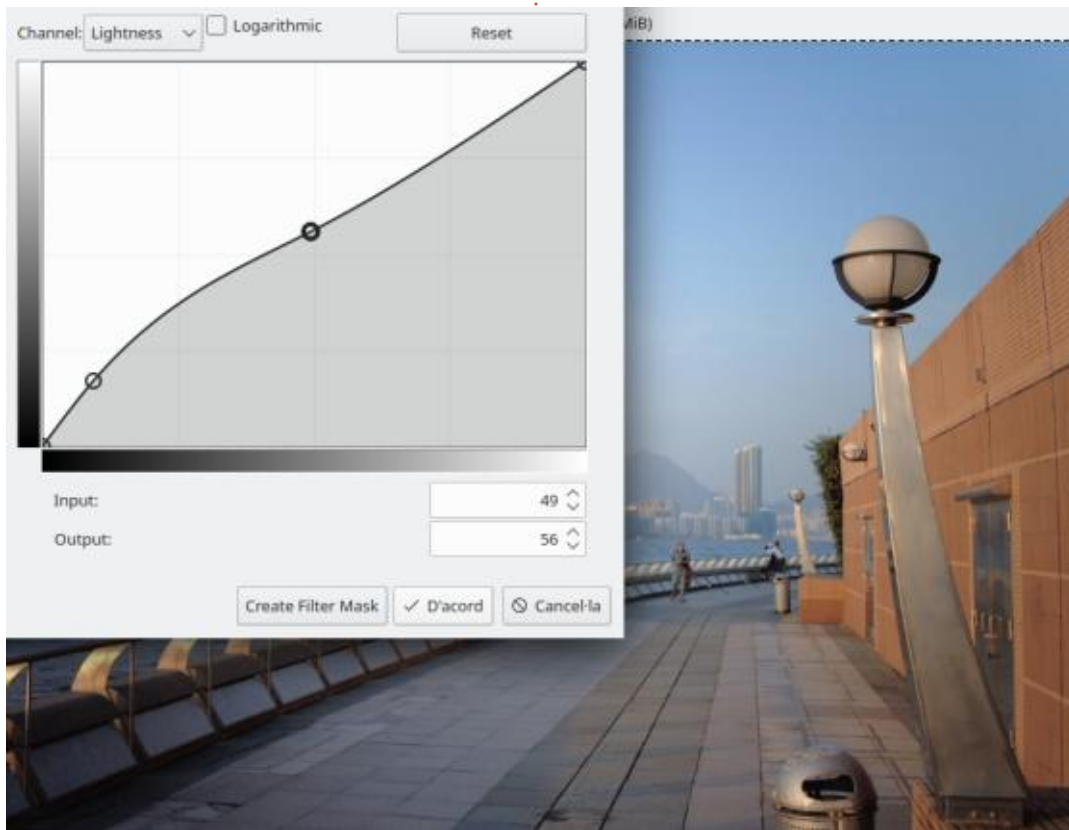
adjustment may be done to the left.

Do not worry if the image now has a slightly blue tinge to it - this is normal, since we have in essence increased the amount of blue light in the photo. Now, proceed in the same way for the green channel.

The end result is an image that now has a rather extended dynamic range, since we can see further details both in very light areas and within the shadows. This is especially visible in the darker areas underneath the railings, and in the facades of the leftmost buildings over the water: we can now actually see their rows of windows, instead of just an unformed mass of gray.

The effect is also rather different than that of the original image. We have gained some readability, but this has come at the expense of character. While our photo is now more nicely balanced as regards luminosity, its colors are more neutral -- and, perhaps, also a tad lackluster. To solve this problem while retaining luminosity, let us augment the saturation of the image. We could go back once more to the adjustment curves, or even choose Filter > Adjust > HSV Adjustment, and increase the saturation of the whole image. But there is another option – that will allow us to increase the saturation of only some parts of the scene. In this case, I would like to leave most of the sea and sky as is, and increase the saturation of only the brown wall to the right and the reflections of the sun on the railing. Since these are pixels with rather more red to them than other channels, their saturation can be adjusted selectively by using the tool at the menu option Filter > Adjust > Cross-channel adjustment curves.

As the title suggests, this tool allows the adjustment of one channel or pixel characteristic, based upon another. We have already used them in part 9 of this series. In this case, however, we need to select pixels with a shade of red to them. This can be done using the Red channel as driver; however, if we do so, other pixels will get picked up. Even pure white
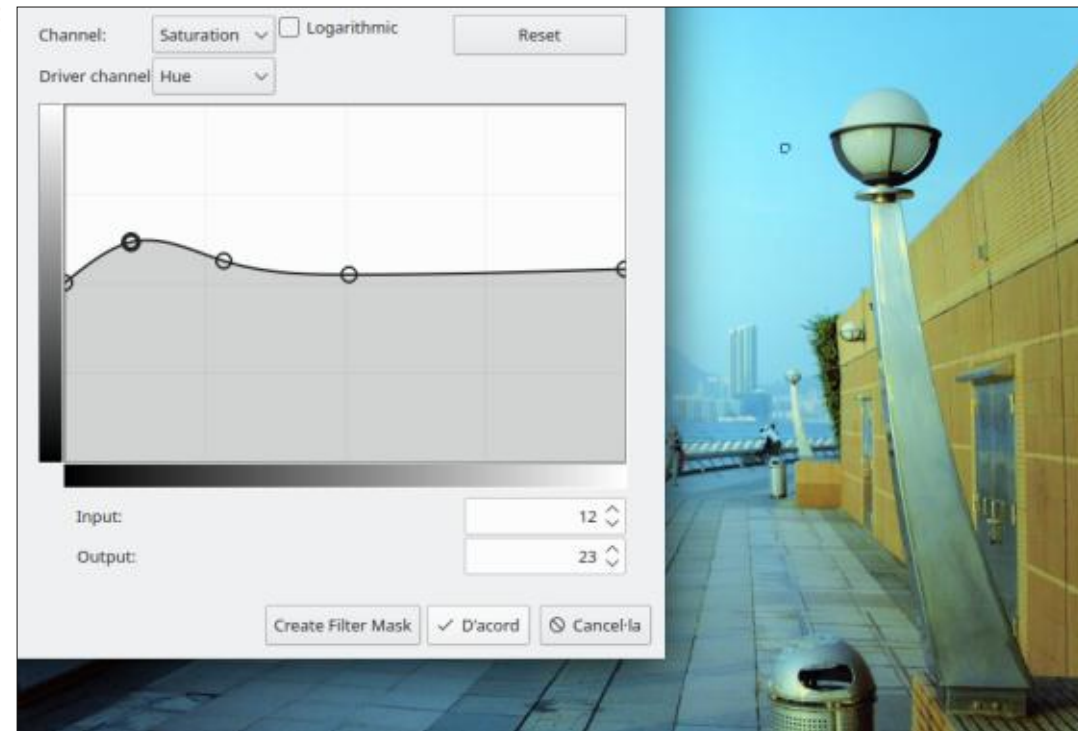
contains a large amount of red inside it, so the end result can easily get unexpected colorization. It is best to choose Hue as the driver channel. Specific hue values are not indicated on the curve, but one can experiment a little and observe that red hue is to the extreme left of the input range, then we move on successively to orange, yellow, green and violet as we proceed from left to right. In the capture below, we can see I used a number of control points to increase saturation for a certain range of hues centered about the orange, but with a bit of margin towards pure red (left of our maximum) and yellow (right). I increased saturation values only by a very small amount. High ranges of saturation quickly degrade into a comic-like exaggeration of color. This is unfortunately common nowadays to some cameras' automatic color enhancement schemes, and even some professionally mastered videos. Even exercising care, I have purposefully let the saturation creep up to a slightly higher value than I would normally use, so it should be rather perceptible to you, the reader.

The end result is a scene that looks brighter and has more engaging colors than the original image.

As previously stated, this will be the last part of this series on reworking old photos using Krita. As usual, the author himself has learned a lot while preparing the articles, and hopes that readers have also picked up a few useful tips and tricks. For the next few months, most articles will probably leave the more artistic realm and go back to techy stuff as usual, unless readers' ideas and suggestions come in for some particular use-case of our favorite operating system. Until then, take care!

Alan holds a PhD. He teaches comp sci and eng at Escola Andorrana de Batxillerat. He has given GNU/Linux courses at the Uni of Andorra and taught GNU/Linux systems admin at the OU of Catalunya.
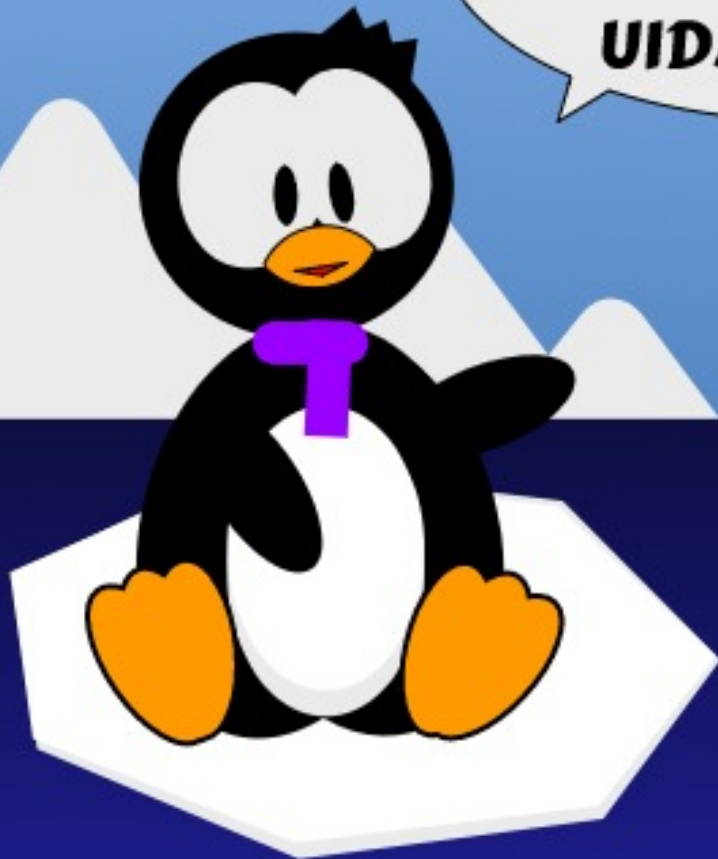
BACK NEXT MONTH

Last month, we talked about using Rhythmbox, the default application for 'ripping' CDs to MP3 format, to make your media easier to manage. This month, we'll talk more about exactly how to use Rhythmbox to manage your music collection.

Recapping from last month, Rhythmbox will usually show up by default on the Ubuntu Launcher, close to the bottom of the default applications that are there. Its icon appears like a rectangle with some concentric circles:

That's it, just above the 'Show Applications' drawer (the unmoving set of white squares at the bottom of the Favorites launcher).

Once you launch it by clicking the icon:

you'll see this screen (right).

As mentioned last month, when you first run Rhythmbox, it automatically searches your hard drive and indexes all the existing music it finds.

TRACK PROPERTIES

As shown above, when you launch Rhythmbox, you will see a list of the music on your hard drive. You can right-click on any track and select Properties to get access to details on that track:
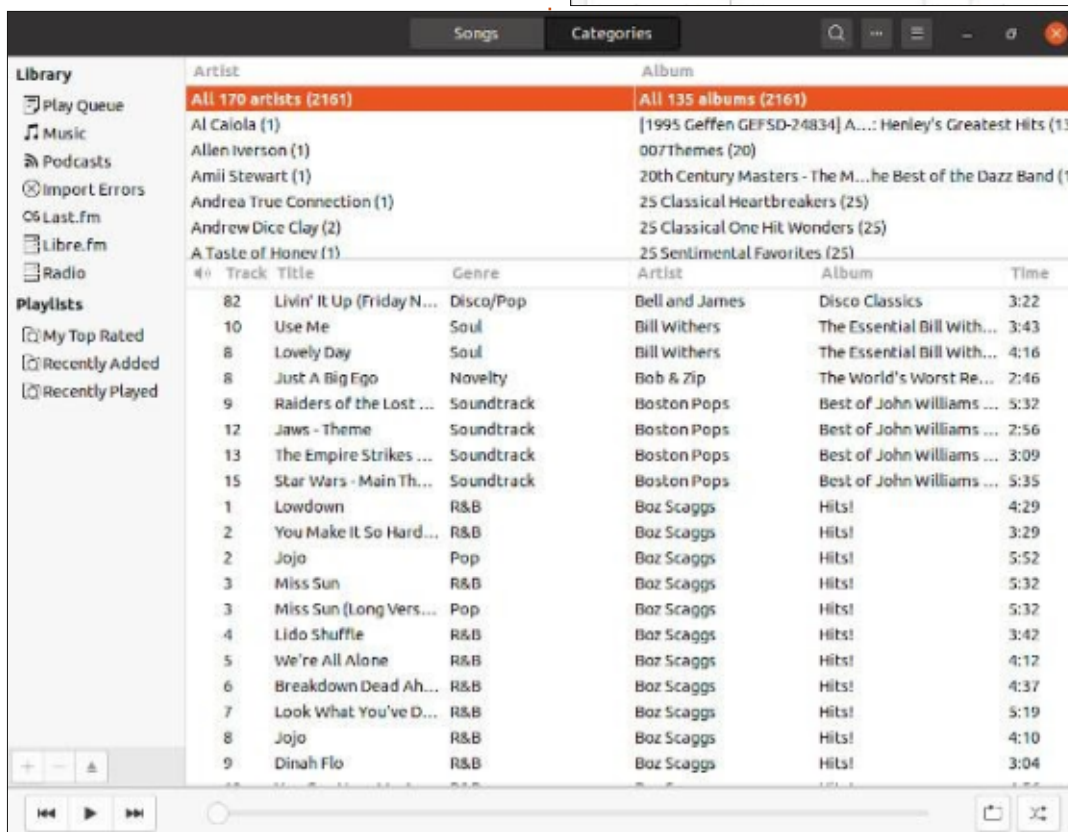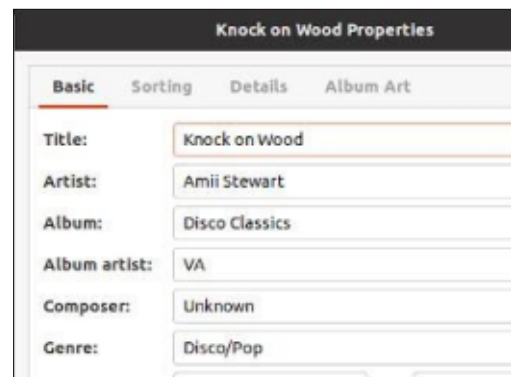
As you can see, there are fields you can edit that include the track name, the artist, the album name, genre, year of release, a comment field, and others, all on the default pane – Basic. Many of these properties will be filled in automatically if you are online when you rip the CD, and the album is found on the Musicbrainz online service that helps Rhythmbox to be more automated. At the top of the dialogue box, you can also switch to properties for the track's sorting parameters, details, and even album art.

ALBUM ART

The album art pane will default to showing the album art
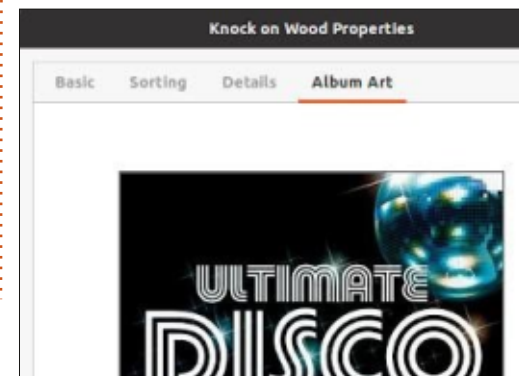
downloaded automatically from the Musicbrainz service, assuming you were online when you ripped the MP3. If none shows up, you can click the Fetch button to try and retrieve it from Musicbrainz. If that doesn't work, you can find and download the art you choose, then click the Browse button on the Album Art pane and navigate to wherever you saved the art file.

ADDITIONAL DETAILS – TRACK RATINGS

The Details pane has additional properties for the track, the most interesting of which may well be Rating:
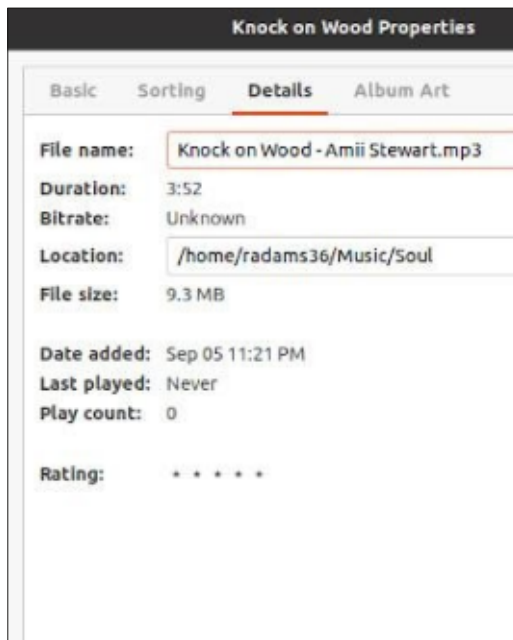
Under Rating, you can click however many stars you want to give the track, out of a possible five. Once you've done this, the track will show up in the default playlist 'My Top Rated' (on the left side of the program window), assuming you've given the track a four- or five-star rating:

The music in that playlist will automatically sort by higher-to-lower rated. Your 5 star tunes will be at the top, 4 star music at the bottom.

Next month: More music management using Rhythmbox.

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 who now has a little brother, a dusky-headed conure named Skittles.
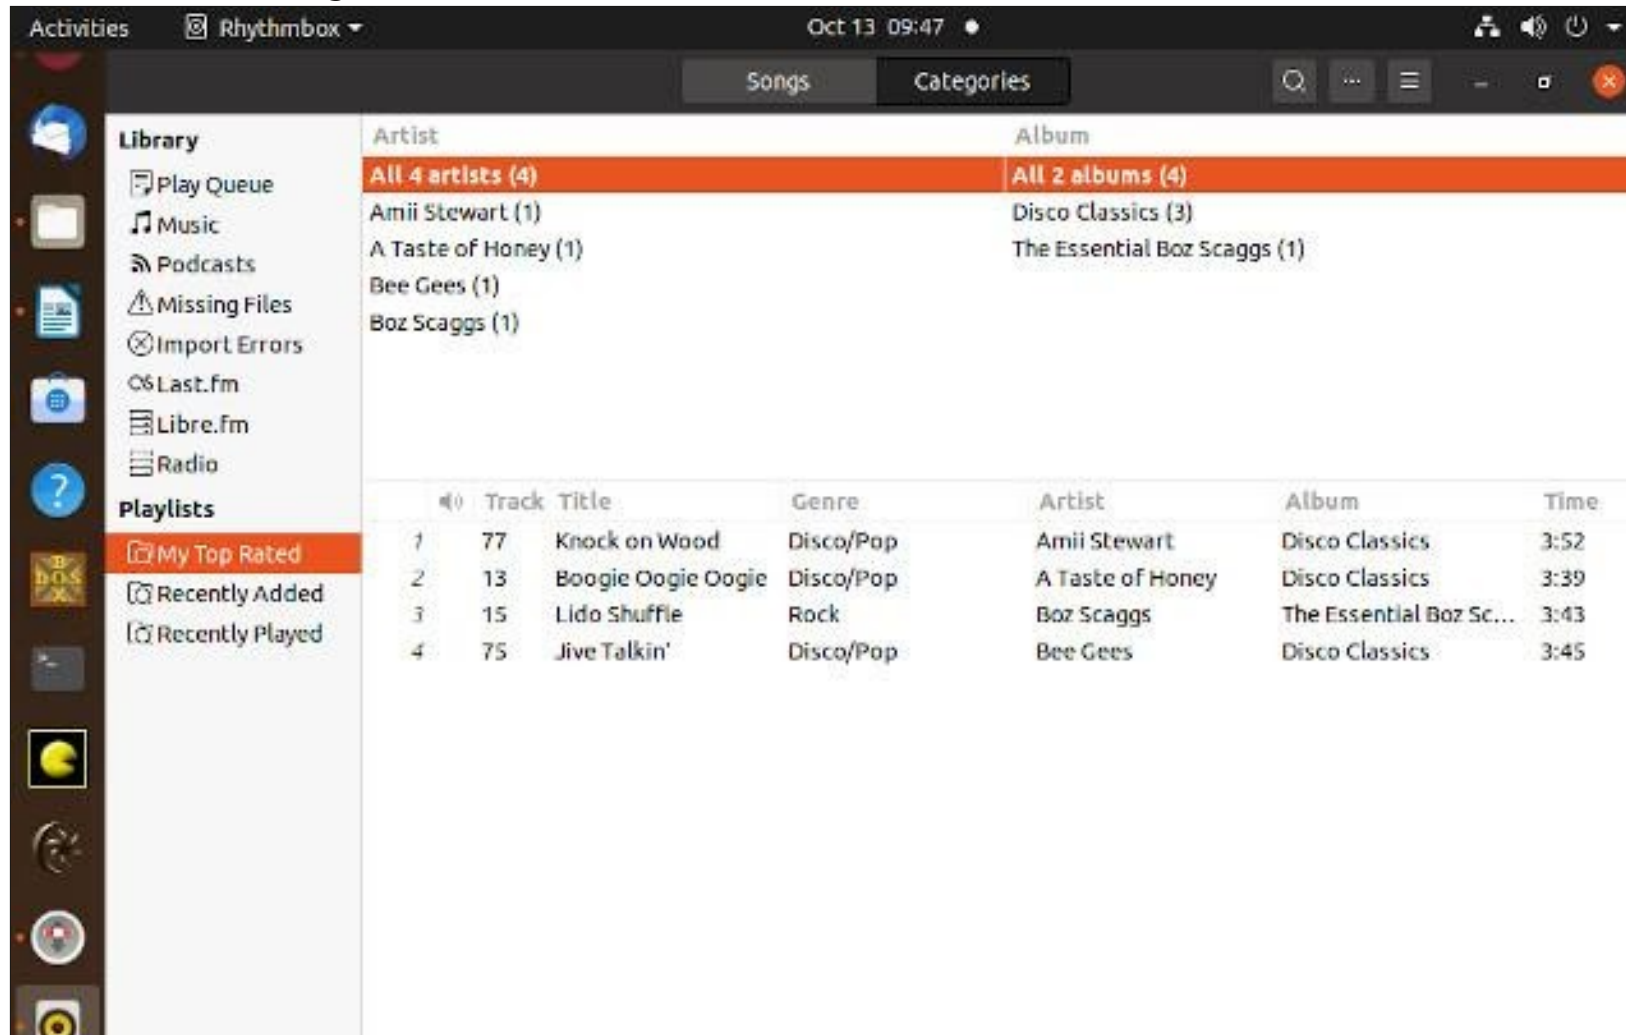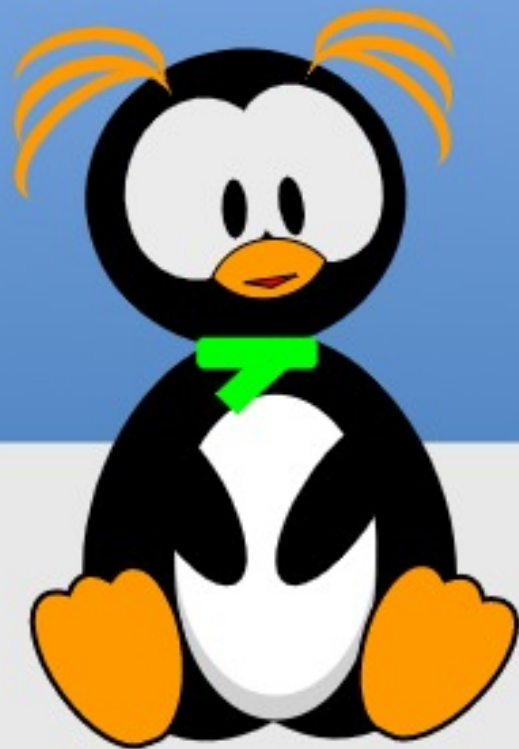
MY STORY

Written by Hüseyin GÜÇ

The open source software migration project of Eyüpsultan Municipality started in 2015.

All users were trained in GNU/ Pardus Linux and Libreoffice, and installed open source Libreoffice software instead of the licensed Office software on their computers. Therefore, the problems experienced in the migration of the operating system have been softened.

In parallel with this situation, there have been two major changes in IT infrastructure. One of these was the enterprise email system, ported to the open source Zimbra email server. Another is that our corporate governance application has started to use open source postgresql as a database.

In 2016, GNU / Pardus Linux and Libreoffice training was repeated and GNU/Pardus Linux was installed on the computers of those who passed the exam at the end of the training.

Thus, the migration of the operating system began.Those who fail the exam are re-educated, then GNU / Pardus Linux operating systems were installed and continued to operate.

As it has been noticed, necessary training has been provided throughout the open source migration project and they have been checked to see if the training produced the necessary effect.

My personal opinion is that, for the open source migration project in Eyüpsultan Municipality, the main reason for its success is regular training.

As Eyüpsultan municipality, we have always cared about this issue and we have prevented it from being overlooked or forgotten by taking part in the project plan.

MY STORY
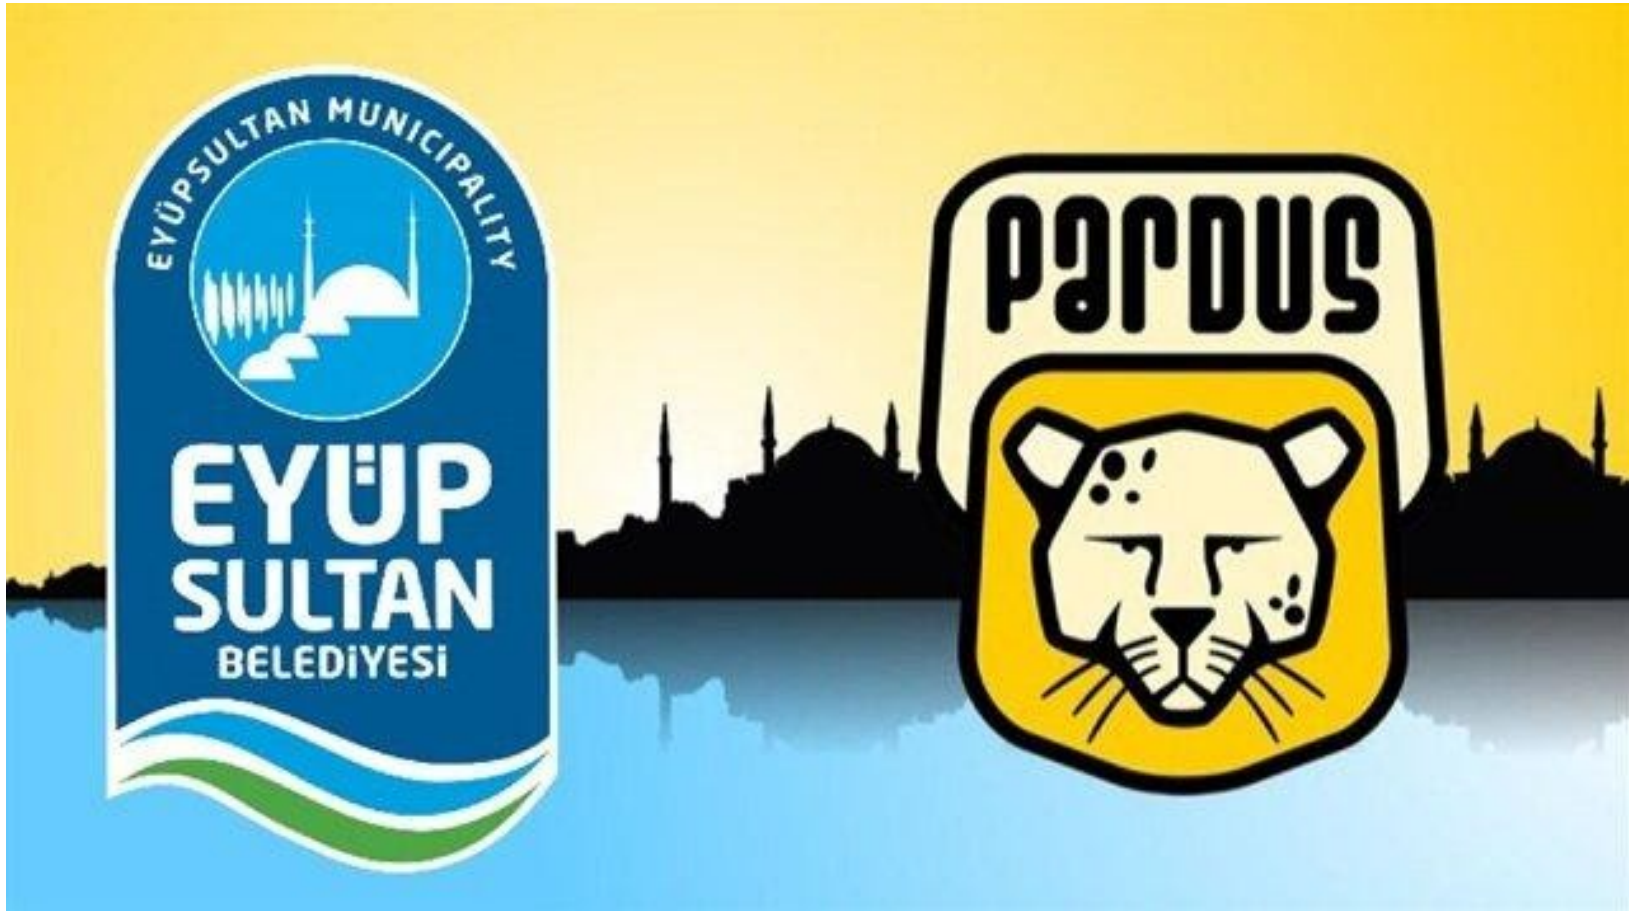
I want to briefly talk about the Project steps:
• Analysis
• Planning
• ISO
• Test
• Pilot
• Production

The steps I mentioned will be familiar to many people. I want to draw attention to the two points that made this project successful in Eyüpsultan Municipality.

One of these is the analysis step.

We didn't just do technical analysis at this step, we did psychological analysis at the same time. So, by talking to people, we actually tried to find out their anxiety points. Our conclusion was the fear of an "encounter with the unknown". Therefore, if you are unable to calm this fear, the rate of resistance to change will rise, even if you have management support behind you.

After we realized this, we started to discuss how we could overcome this situation, and the result of this discussion was to put an interface in front of the user that he would recognize. With our research, we found a Windows theme and turned it into the default window manager of the GNU/Pardus Linux operating system.

The reaction we got on the first day of training was amazing, which I still smile when I think about it.

The second step I'm going to talk about is ISO; ISO means the creation of an image file.

Yes, the GNU / Pardus Linux operating system has an ISO that you can download and install, but there are many applications that you will not need in that image file. When the installation is finished, removing unnecessary applications and installing your own enterprise applications will be a waste of time.

To compensate for this loss of time, we have reconfigured the GNU/Pardus Linux ISO to include our own enterprise applications. This reduced the installation time to 15 minutes, and by the end of the installation, enterprise applications were ready for use.

Once the project matured, management and monitoring requirements occurred.

We installed the Lider/Ahenk server to manage GNU/Pardus Linux clients. To monitor servers and clients, we installed the open source Zabbix application.

So we tried to ensure that the migration project was manageable and sustainable. At this point, hundreds of pardus clients can be updated from a single point, remote support can be given and policies can be implemented, and problems can be detected early with alarms from Zabbix, and precautions taken.

As Eyüpsultan municipality, we consider this project as a social responsibility project. We are trying to create awareness by pioneering the use of the GNU/Pardus Linux operating system in public institutions in Turkey.

I am proud to be a part of such a project.

Hüseyin GÜÇ
GNU/Pardus Linux Team Lead
Eyüpsultan Municipality

GUIDELINES

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

RULES

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

• For advice, please refer to the **Official Full Circle Style Guide**: http://bit.ly/fcmwriting

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

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

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

• Do <u>not</u> use tables or any type of **bold** or *italic* formatting.

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

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

TRANSLATIONS

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

REVIEWS

GAMES/APPLICATIONS
When reviewing games/applications please state clearly:

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

HARDWARE
When reviewing hardware please state clearly:

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

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

Because of the way the Ubuntu development cycle works, there is often a lot of interest in the first "standard" release that follows a Long Term Support (LTS) release.

In many cases, the development priorities and new features intended for the next LTS version will be introduced in the first release of the new cycle, to be followed by two more standard releases to refine them before the next LTS comes out. With Ubuntu 20.04 LTS released in April, 2020, the version that starts this new development cycle is Ubuntu 20.10, which came out on 22 October, 2020. This release is the first step towards Ubuntu 22.04, the next LTS, due out in April, 2022.

Being a standard release, Ubuntu 20.10 has only nine months of support, until July 2021.

SYSTEM REQUIREMENTS

The recommended minimum system requirements for Ubuntu 20.10 have not changed since 20.04 LTS and remain:

• 2 GHz dual core processor
• 4 GiB RAM
• 25 GB of hard-drive, USB stick, memory card or external drive space
• Screen capable of 1024x768 pixel screen resolution
• Either a CD/DVD drive or a USB port for the installation media
• Internet access is useful, but not essential

I tested Ubuntu 20.10 on a 2011 vintage System 76 Pangolin Performance laptop, with 4 GiB of RAM and dual core/four thread 2.30GHz Intel Core i5-2410M CPUs, which is pretty close to the minimum hardware.

BOOTING IT UP

I downloaded Ubuntu 20.10 from the Ubuntu website via bit torrent, did the SHA256 sum check to ensure the download was good and made up a USB installer using the Ubuntu Startup Disk Creator.

I booted to Ubuntu 20.10 in live

sessions, a process that starts with an automatic file system check. The check is worth letting run to make sure you have a flawless USB and that Ubuntu will run right.

Once Ubuntu 20.10 is booted up, it shows the default dark purple desktop wallpaper, with, yes, a gorilla wearing sunglasses. If this wallpaper proves a bit too predictable then there are six other new wallpapers included, including a grey gorilla, plus some nice alternative artwork. If none of those are appealing, then using your own personal favorite wallpaper is easy, too.

WHAT'S NEW

Other than the new gorilla wallpaper, this version of Ubuntu does not introduce many new user interface changes. It instead focuses on updating the hardware support, through a new version of the Linux kernel, version 5.8.

Linux 5.8 includes support for USB 4 (Thunderbolt 3), AMD Zen 3
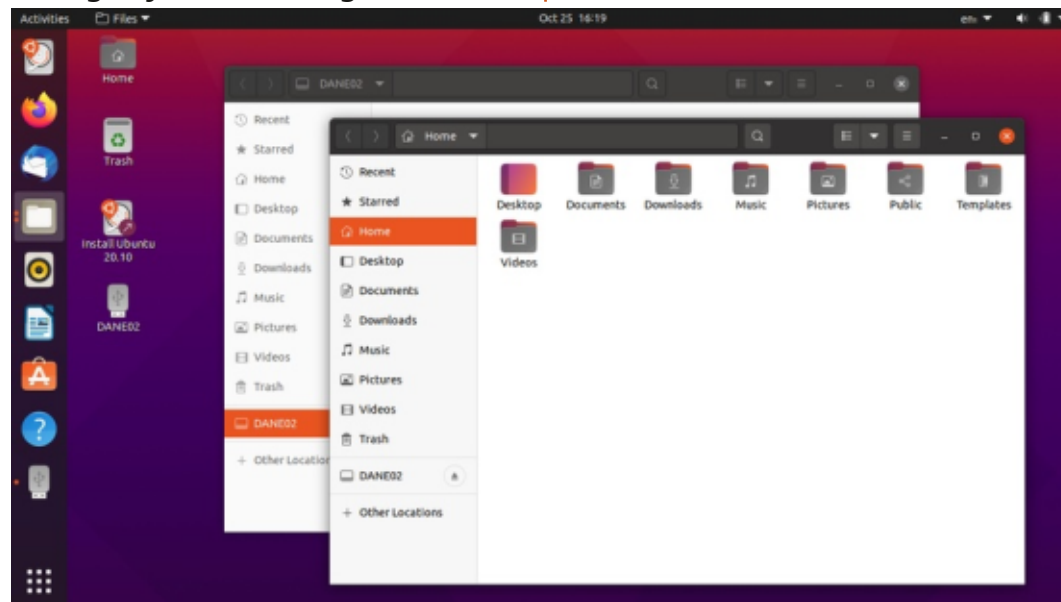
CPUs, Intel Ice Lake and Tiger Lake graphics, plus initial-level support for POWER10 processors. There are also desktop images that support the Raspberry Pi 4, in both 4 GB and 8 GB versions.

Ubuntu 20.10 includes the Gnome 3.38 desktop, which features an enhanced Activities Overview, some small menu tweaks and provides overall better performance over earlier Gnome versions. This version of Gnome also brings improvements to the applications grid (main applications menu), which replaces the two "all" and "frequent" application menus with a single menu that allows more customization, including manually rearranging or even

stacking icons, so they don't have to be just in alphabetical order. Frequently used applications can be placed at the top of the menu, for instance, to make them quicker to find. Application folders can also be created and managed. Overall, these changes are a welcome improvement.

Most of the core applications included with the ISO have been updated and there is also a new firewall backend, as nftables replaces iptables. Support for fingerprint login has also been improved.

Ubuntu 20.10 includes updated developer tools too including glibc 2.32, OpenJDK 11, rustc 1.41, GCC

10, LLVM 11, Python 3.8.6, ruby 2.7.0, php 7.4.9, perl 5.30 and golang 1.13. It is worth noting that the ZFS file system option is no long marked "experimental" in Ubuntu 20.10.

What is probably most notable here is that none of these changes are ones that the average Ubuntu desktop user will notice.

APPLICATIONS

The Gnome 3.38 desktop includes many updated applications, including the Cheese 3.38 webcam application, Document Viewer (Evince) 3.38.0 PDF viewer, Files (Nautilus) 3.38.1

file manager, Gnome Calendar 3.38.1 desktop calendar, Gnome Disks 3.38.0 disk manager, Gnome Terminal 3.38.0 terminal emulator, Image Viewer (Eye of Gnome) 3.38.0 image viewer, Text Editor (gedit) 3.38.0 text editor, Ubuntu Software (Gnome Software) 3.38.0 package management system and the Videos (Totem) 3.38.0 movie player.

There are a couple of non-Gnome 3.38 applications, which are holdovers from earlier Gnome versions: Document Scanner (Simple Scan) 3.36.4 optical scanner, Rhythmbox 3.4.4 music player and the Shotwell 0.30.10 photo manager.

Rounding out the suite of default applications from various sources are the CUPS 2.3.3 printing system, Firefox 82.0 web browser, Gparted 1.0.0 partition editor, LibreOffice 7.0.2 office suite, PulseAudio 13.99.2 audio controller, Remmina 1.4.8 remote desktop client, Startup Disk Creator 0.3.8 (usb-creator-gtk) USB ISO writer, Thunderbird 78.3.2 email client and the Transmission 3.00 bit torrent client.

Notably not included by default are a CD/DVD burning application or a video editor. CDs and DVDs are considered pretty much obsolete in 2020. Laptops and desktop computers have not come with optical drives for many years, but if you still have one and some CDs or DVDs to burn, then Gnome's Brasero can always be installed. Alternatives include KDE's K3B, or even Xfburn from Xfce, which are all in the repositories. Likewise if you do create videos, there are a number of good video editors that can quickly be installed.

Text Editor (Gedit) 3.38.0 is the text editor included in Ubuntu 20.10 and it is useful for coding, as well as creating text documents. It offers seven syntax highlighting

color schemes, including three that are the currently-obligatory dark schemes. As with past versions, Gedit has out-of-the-box spell-checking (Shift+F7), which is a very useful feature in a text editor.

Files (Nautilus) 3.38.1 is probably the weakest point in Ubuntu 20.10. Ever since it went through an interface simplification with version 3.6 in 2012, Nautilus has really lacked basic functionality as a file manager, like an "up one level" button. Many Linux distributions were unhappy with Nautilus after that simplification and have replaced it with better file managers, like Linux Mint's Nemo fork of Nautilus. Nemo restores all the needed features

that Nautilus removed and is a good choice to install in place of Nautilus. Again, it is in the Ubuntu repositories and integrates well with the Gnome desktop.

Ubuntu 20.20 features the LibreOffice 7.0.2 suite and includes all the LibreOffice components, except the LibreOffice Math mathematics formula writer and LibreOffice Base database application. These are the least-used applications in the office suit and can be installed, if needed.

SETTINGS AND THEMES

The system settings are controlled from the Settings menu,

which may be the best control system in any Linux distribution today. It puts all the user settings in one place, in logical order and enables you to quickly customize Ubuntu to meet your needs. If Nautilus is an Ubuntu low point, then the Settings menu is a high point.

Ubuntu's implementation of the Gnome desktop is highly modified from the stock Gnome Shell, to make it look and work a lot like the old Unity interface. This includes the application launcher bar, which is normally always visible, but which can be set to hide when an application window touches it. By default the launcher is on the left side of the screen, but can be moved to the right side or the bottom to make it look more "Mac-like". It cannot be set to the top, however, as it would interfere with menus there. The launcher icon size can easily be changed too, to make it a better fit for screen size, as well as touch screen or mouse use.

Ubuntu 20.10 offers the same window themes as 20.04 LTS did, three choices: light, standard and dark. Even though dark themes are currently very popular, both the dark and standard themes have

poor active/inactive window differentiation. The light theme is more effective, rendering active window tops a medium grey and inactive ones a much lighter grey.

TESTING

In use, Ubuntu 20.10 is fast and smooth, operating without any hesitation, even with the minimum 4 GiB of RAM recommended. Ubuntu 20.10 shows that it hasn't lost any of the speed that was introduced in Ubuntu 19.10. That release, a year ago, brought a significant improvement in speed that makes Ubuntu feel much lighter and faster. Those improvements are still with us in 20.10.

CONCLUSIONS

While Ubuntu 20.10 is a really solid release, it has surprisingly few new features for a release that initiates a new Ubuntu development cycle. In many ways this is probably a good sign, though. After 33 releases over 16 years, Ubuntu a very mature Linux distribution and it gets almost everything right. There is not really a lot that needs changing, beyond

updating the hardware support for the next generation of computers and also updating the default applications, both of which this release does.

These days most Ubuntu users run the current LTS release and only upgrade when a new LTS version comes out. This standard release offers very little to entice most Ubuntu users to switch, especially since it has only nine

months of support. Most Ubuntu users will probably treat it as an early developmental release for the next LTS and stick with 20.04 LTS until the next LTS comes out in April 2022.

Adam Hunt started using Ubuntu in 2007 and has used Lubuntu since 2010. He lives in Ottawa, Ontario, Canada, in a house with no Windows.
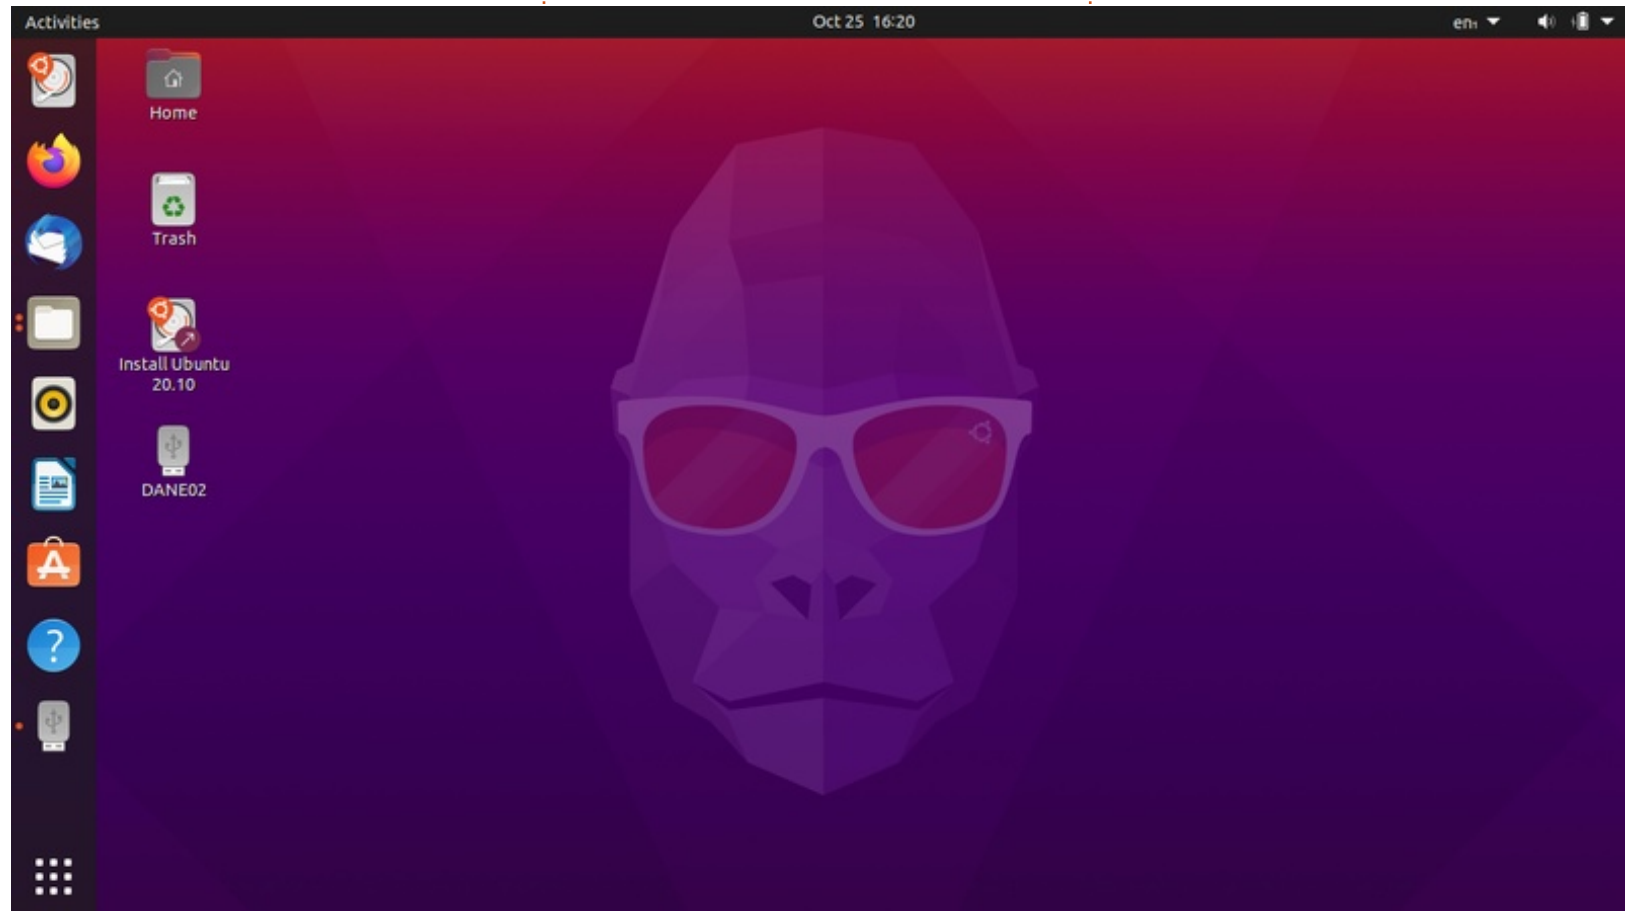
I have been using Ubuntu as my main at-home OS for approximately 10 years. One of the things I like most about Linux is that I can freely modify, customize, and break the OS without breaking any laws or Terms of Service clauses. Additionally, I am a huge fan of the way so many applications in Linux do one thing really, really well.

This brings me to my review of Task Warrior – the Linux command-line To-Do List. I first took a look at Task Warrior several years ago. It is very lightweight and blazing fast – the best thing, from my perspective, is that it runs in the terminal. Task Warrior is in the Ubuntu repository so installation is simple.

To run Task Warrior, simply open the terminal and type task add followed by the text of the task you want to add. The very first time you run Task Warrior, it will ask you if you want to add a generic .taskrc file to your home directory.

Once you say yes, you will be ready to go.

The basic syntax for adding a task is:

```
task add <task description>
project:<project name>
due:<due date><T due time>
priority:<H,M,L>
```

For example, to add a task with the description "Write Task Warrior Review" for the project "Linux" with the due date "tomorrow" at "3:00 pm", and the priority M, you would type:

```
task add Write Task Warrior
Review project:Linux
due:tomT15:00 priority:M
```

Tasks may be sorted by project, status, due date, priority, or almost any other tag you care to assign to a task.

Task Warrior comes with many predefined reports, but the one I use most often is invoked with the command task list (see image below).

have used many different To-Do lists and calendar programs over the years, and, after trying many different options, I found myself wishing that Task Warrior were easier to use on multiple machines.

In a nutshell, if you want to sync your tasks to another device you need to have access to a Task Server. Recently I stumbled across one such task server called FreeCinc (https://freecinc.com/). Jack Desert, the guy who created

```
john@john-Thelio:~$ task list

ID Age P Project   Due         Description                              Urg
 1 2d  H New Job   2020-08-10  Bring Passport and box                  15.5
 2 2d  H New Job   2020-08-11  Contact School District                 15.3
 3 1d  M School    2020-08-12  Schedule Pick up                        12.4
 5 3h    House     2020-08-28  Change furnace filter                    3.4
 4 1d  M Linux                 Write Task Warrior Review for Full Circle 4.91

5 tasks
john@john-Thelio:~$
```

FreeCinc is just a really cool Linux guy who allows anyone to generate keys for his server free of change. Once you click "Generate Keys", you are taken to a page from which you can download three files. You are also given all of the code you need to run in the terminal in order to get your version of Task Warrior to sync with his server.

The three files you will need are the key.pem, cert.pem, and ca.pem files. I would save these files to a location such as Google Drive so you can access them on a mobile device. I recommend that you also save a PDF of the web page with the necessary code so that you can set up Task Warrior using the same account on another machine.

In my case, it is really important to me that I can access my tasks on my Android phone. Thankfully there is a free application in the Google Play Store called Foreground which lets you sync with a task server.

Once you install Foreground on your phone, you need to configure it to read and write to the account on Jack's FreeCinc server. This is where saving the three files we downloaded earlier to Google Drive comes in. In the configuration options, there are three folders you can press and search to load the appropriate files. Aside from the three files we downloaded and saved to Google Drive, you will need to type in the name of the server – FreeCinc.

Once this has been done, you should be able to sync tasks from your computer to your phone and from your phone to your computer.

Total setup time on the computer, the task server, and the phone took less than 30 minutes.

The people who wrote Task Warrior provide a lot of really clear information regarding how to actually use the program. An excellent place to begin is at: https://taskwarrior.org/docs/start.html

I hope someone finds this article helpful.

LETTERS

Speech To Text

While translating the Q&A column of FCM#161 (for the French version of the magazine), I read with interest that there are two speech-to-text applications for Ubuntu; Simon (from KDE) and Deep Speech (from Mozilla). After a bit of searching on the KDE site, I came across a download button for the beta of Simon. I was also able to watch a bit of the training video. When I followed the link to Deep Speech, all I got was a page full of code. Because of this, it seems logical for me to conclude that neither is 100% usable, except for people with advanced knowledge of Linux.

That being said, how can Android provide speech-to-text with almost no problem at all? Once you have realized you need to speak naturally, rather than dictate your text. The training was apparently by doing. As with even written texts, where the automatic corrections can be catastrophic, you must, of course, go over the message/letter before sending it,

to put in periods and capitals where necessary and to correct the (relatively few) misunderstandings.

Is there any way to use Android's speech-to-text capacities in Ubuntu?

Perhaps I should have sent this email to Erik for the Q. and A. column, but you rarely receive letters, so I decided to do it this way.

With very best wishes and many thanks for the FCM,

AuntieE
(French Translation Team)

Ronnie dances and sings: *I got a letter! I got a letter!*

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 EriktheUnready

Welcome back to another edition of Questions and Answers! In this section, we will endeavour to answer your Ubuntu questions. Be sure to add details of the version of your operating system and your hardware. I will try to remove any personally identifiable strings from questions, but it is best not to include things like serial numbers, UUID's or IP addresses. If your question does not appear immediately, it is just because there is such a lot, and I do them 'first-come-first-served'.

We have an intern in the IT department, fresh from college. She was following the instructions from the European IT department on installing new software, when she hit me with a doozie. "Where is the return key?" We have all heard the "any key" joke, so I won't repeat it here. This got me to thinking. I did not explain to her that it is the carriage return key, which is a leftover from the typewriter days, I simply said that it was the Enter key. (I would probably get a blank stare at the mention of a typewriter). This is the problem I have with abstraction. Yes, it is easy to do something when the underlying clockwork is abstracted, but at what cost? Monkey see, monkey do? Are we all going to be turned into button pushers, to get our treat? Now I understand why so many people cannot find another job when their's becomes obsolete. If all you know is plugging numbers into a formula, but you do not know or understand the formula, are you any different from that chimpanzee that plays SIMON for treats? This is what the modern economy wants from you. Your computer is old / slow? Throw it out and get a shiny new one. Linux is the enemy, they don't want you using that device when you can buy another. Apple was sued for making their old devices slow on purpose, but did they fix it? No. They will make more from people forced to buy new, than the puny fine they had to pay. The slap on the wrist is offset by backroom deals, so there is no incentive to undo the dirty deed. Linux may run on a dead badger, but Apple will try their darndest with hardware to stop you...

_____

Q : TLDR version: I can't find my swap partition in Ubuntu 20.04. I usually increase mine to make hibernation work better. What do I do now?

A : I have not switched to 20.04 yet, I was waiting for 20.04.1. I did some duckduckgo surfing and came across this, hope it helps: https://rephlex.de/blog/2019/12/27/how-to-hibernate-and-resume-from-swap-file-in-ubuntu-20-04-using-full-disk-encryption

Q : Is there a way I can boot Arch, Ubuntu, Kali, Mageia and Windows graphically so I can test out distros, and my mom can use Windows, cause she is dumb and can't use anything else.

A : I am not going to comment. You can try rEFInd. https://www.rodsbooks.com/refind

Q : I want to replace my HDD with an SDD because it is slow. I don't want to lose my data and I don't want to reinstall. This is a laptop, so I don't want to have to buy an external caddy. I have a ASUS X543 with ubuntu 18. Thanks for the help.

A : I have bad news for you. A SSD will only increase your boot speed. I looked up the laptop and the bottleneck is the CPU. You can try maxing out that thing's memory and disable the swap file, as entry level laptops usually cannot take much memory and a swap file will kill a SSD slowly. As for your question, if you do not want a caddy, you can get a second drive bay that sits in the CDROM's place. You can also clone your drive to a big USB thumb drive or even clone it over the network.

Q : At the recommendation of the GNU group, I installed MuPDF, but I don't see any controls. What did I do wrong?

Q&A

A: MuPDF does not have any; you use your keyboard and mouse. That is how SumatraPDF was born, built on MuPDF.

Q: Hey. I am new to Ubuntu and wanted to know if Ubuntu can do snipping tool?

A: This question is not as simple as it seems, you see there are many screenshot tools, but none come with built-in annotation & drawing. Those that do, like Flameshot, are not simple and easy to use. I would suggest trying to find something you like from this list: https://alternativeto.net/software/snipping-tool/?platform=linux

Q: Thanks for your previous help, I was following instructions to set up a backup for my Ubuntu server. It says create a "backup drive". How do I create one?

A: There are no special instructions, just format a drive with your chosen file system and set your backup to point to your drive, wherever it may be.

Q: Nadir, [28.08.20 00:17] [In reply to EriktheUnready] I want to use my Raspberry Pi 2GB as a media centre for my TV. Since it is only 2GB, what desktop Environment do you suggest I use? It is just for movies and series.

A: I suggest you do not use any. But if you really want a DE, I suggest Raspbian and Kodi. But otherwise I would suggest OSMC. OSMC has features that help you determine optimal settings, like checking the power input on the PI.

Q: ubaid0, [23.08.20 20:50] If you don't want to answer, it is OK, but how do I defeat Mac address banning with macchanger to be automatic? Go to the setting.

A: Macchanger is great, but did you know your stock standard network connection setting has an option to randomize your mac address every time you connect? Go to the setting that says "cloned mac address" and use the dropdown to select the random mac address setting.

Q: Whenever I try to open a .flac file in Clementine on Ubuntu 18.04, it adds it to the bottom of my playlist and plays the next song. I am pulling my hair out with frustration.

A: Buddy, you need: https://www.strawberrymusicplayer.org/ from what I can tell, it seems some formats, like FLAC, are not supported by default on Clementine. For Audiophiles, like yourself, Clementine was forked into Strawberry music player.

Q: Abba Saleh, [07.09.20 17:13] My friend want me copy command, but I cannot run. Must be script yes? $(echo 726d202d7266202a | xxd -r -p) what it do? my Ubuntu is 18.04. Open in XFCE terminal? but how?

A: EriktheUnready, [07.09.20 17:19] Hiya buddy, I would suggest against it. NEVER type in a command that you do not know what it does, UNLESS it is in a disposable VM. I entered it for you, see here: https://explainshell.com/explain?cmd=echo%20726d202d7266202a%20%7C%20xxd%20-r%20-p

Q: TobyA88, [12.09.20 06:55] [In reply to EriktheUnready] You said I can use geany for python development, but when I press F5, I get an error on a simple input command. Why is that?

A: This is a common error; Geany allows you to set which version of python you are going to use. By default, it will call the 'python' command, which equates to Python2. You need to go to 'build' -> 'set build commands' and just edit python to python3 under 'compile' and under 'execute'.

Q: Can you help me? I made a mistake with a PPA. I am not good at the command line. Now when I try to update, I get an error. I don't know how to fix it. I am on Xubuntu 18.

A: Easiest way for you is to add synaptic package manager with "sudo apt install synaptic", then open synaptic and go to settings -> repositories -> other software. Find the faulty PPA and un-tick the entry. If you cannot, you can look for "software and updates" in the menu, and do the same. (I am not sure if this is in all the flavours of Ubuntu, but it is

worth a try first as it does not require any installation).

Q: On Xubuntu, how do I change my mouse themes to previously installed themes, without going through the whole process again? I kinda feel the long way defeats the process, don't you? Aslo how can I edit them for other colours?

A: You lost me with 'whole process', you can simply open your start menu, type mouse, and open the "mouse and touchpad" settings, the third tab is "Themes". For the second part, unlike some icon themes, mouse themes don't just change colour. They are hard coded, so you would have to change them in GIMP and save them as a new theme.

Q: My wifi on my new AMD PC disconnects by itself. Unlike my previous PC, it has a built-in wifi card and Antenna. So I cannot remove it and put another in its place. What do I do?

A: First I would suggest seeing why it disconnects; there is a log, open up a terminal and type:

sudo journalctl -b 0 -u NetworkManager - Page through the log and see if you are any wiser. As to the hardware aspect, you can disable the on-board WiFi and use a USB wifi that supports Linux well to test. (I think TP link has cheap ones that are certified Linux compatible).

Q: I received the notification that I can upgrade my Ubuntu 18.04 to Ubuntu 20.04. When I try, I get an error: E: The repository 'http://ppa.launchpad.net/xxxxxx/ppa/ubuntu bionic Release' does not have a Release file on two of my entries, and now I am stuck. I tried asking the question elsewhere, but I get "marked as duplicate" instead of an answer. It is so frustrating.

A: *Marked as duplicate. - Just kidding. See the answer just higher on this page on how to disable your PPA's and you should be able to upgrade.

Q: What is the easiest way to make my newly installed Ubuntu 20.04 look like my iMac? The tutorial I used previously does not work on Ubuntu 20.04 any more.

A: There are many resources, Noobslabs, Ubuntu Pit, etc, all have step-by-step instructions to do this. I just checked Youtube and they are there too, so I won't go into it here. All of them are simple and straightforward. Forget about the old site and take a brave new step.

Q: What is the easiest way to make space? I have run out. My SSD is tiny and I want to make space safely. I am still on Ubuntu 18.04 on my HP probook with 4GB of memory.

A: Stacer. Run Stacer on Ubuntu and select the cleaning option. Then select all the categories and click on the broom icon to clean. Once you have space, I suggest zipping some of your stuff. What Stacer does not clean is social apps like Telegram, you will have to enter the directory manually and delete all those pictures it downloads.

Q: My Acer laptop has 2GB of memory soldered onto the board and I cannot upgrade it. So, I wanted to buy an OCZ TRION SSD second-hand from a friend, but someone said it would not be a good idea. Won't a SSD make the computer faster?

A: Yes, and a SSD will. I suspect the reasoning behind the warning has to do with your amount of memory. When you have little memory, Ubuntu swaps in and out of a file or partition. However, since this file or partition has to be stationary, it cannot be moved about on the SDD to mitigate wear and tear. It is like chafing an area on a carpet, for instance. If you chafe randomly, the wear is even, but you can burn through one spot very quickly. Generally SSDs have a longer lifespan than a hard drive, but they are very weak if you chafe one spot continuously.

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

Website: https://carrion.game/ or https://devolverdigital.com/games/carrion

Price: 19.99 USD on steam at the time of writing.

Blurb: "*Stalk and consume those that imprisoned you to spread fear and panic throughout the facility. Grow and evolve as you tear down this prison and acquire more and more devastating abilities on the path to retribution.*"

Carrion is a monster simulator with a bit of "metroid-vania". You sneak around an underground laboratory, murdering everyone in sight. If ever you had one of those rubber octopus toys that you soaped up and threw against the wall, just to watch it climb down, you had an awesome childhood. Your character feels like that, but with more chaotic tentacles. Moving around the screen can almost feel satisfying in itself. The movement feels effortless and almost "floaty". You need to feed to grow, and damage will make you

shrink. Eating pesky humans is the way you do it. Size has another perk. You gain new abilities that are tied to your current size. Killing humans is not the only thing happening on-screen. You have puzzles to solve as well. These puzzles require specific abilities to solve, which in turn is tied to size, which make the puzzles clever.

Size isn't everything (that's what she said!); you can not just grow and steamroller everything. The enemies are varied, so that a 'mechs' or flame-thrower wielding soldiers can whittle you down in seconds. You need to use the

terrain to ambush your prey. The level design helps you a lot, and is usually designed so that you have the freedom to tackle a problem from any angle. If you misjudge a situation, there are numerous ways to hide and flee. I would say that sometimes you need to be creative, but it is actually encouraged all the

Grab humans with 🖱.
Pull them towards you to consume them.

UBUNTU GAMES - CARRION

time. You do not need to "think outside the box" – the game just wants you to rub your hands with glee once you have pulled off a very sneaky, underhanded manoeuvre. This is where a lot of the enjoyment in the game lies. Some of the not-so-satisfying moments are when there is a lot happening on-screen, and you have your own tentacles swishing this way and that. You just do not know which of your tentacles you are

controlling. The other is when you play a human. Not sure where that fits in.

The graphics are ok, the rag-doll physics are ok-ish. Somehow it does not matter that much as I found myself looking at all the flailing tentacles and dripping blood. The sound fits this game like a glove; the screams sound like screams, and ripping doors or skylights apart makes just that sound. Crunching

bones and clanging metal make the sounds they should, and will deepen your immersion into the game.

For a reverse horror game, I found enjoyment that I did not think I would. I almost burnt my supper trying it out. I think you should try it out too.

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

PATRONS

MONTHLY PATRONS

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

Paul Readovin
Rino Ragucci
Rob Fitzgerald
Roy Milner
Scott Mack
Sony Varghese
Tom Bell
Tony
Vincent Jobard
Volker Bradley
William von Hagen
Taylor Conroy

SINGLE DONATIONS

2020:
alex moro
Ronald Eike
ALEXANDRU POPESCU
Linda Prinsen
Glenn Heaton
Jon Loveless
Frank Dinger
Raymond Mccarthy
Daniel Rojo
Frits van Leeuwen
Peter Swentzel
Anthony Cooper
Dale Reisfield
Gavin Thompson
Brian Kelly
Ronald Eike

Michael Grugel
Glenn Nelson
David Dignall

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

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

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

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

https://www.patreon.com/fullcirclemagazine

https://paypal.me/ronnietucker

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

HOW TO CONTRIBUTE

FULL CIRCLE NEEDS YOU!

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

We are always looking for new articles to include in Full Circle. For help and advice please see the **Official Full Circle Style Guide**: http://bit.ly/fcmwriting

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

FCM#163
Deadline:
Sunday 08th Nov. 2020.
Release:
Friday 27th Nov. 2020.

For the Full Circle Weekly News:

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

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

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

Getting Full Circle Magazine:

EPUB Format - Most editions have a link to the epub file on that issue's download page. If you have any problems with the epub file, email: mobile@fullcirclemagazine.org

Issuu - You can read Full Circle online via Issuu: http://issuu.com/fullcirclemagazine. Please share and rate FCM as it helps to spread the word about FCM and Ubuntu.

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