



Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

ISSUE #177 - January 2022



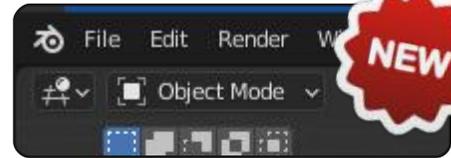
GRAPHICAL USER INTERFACE THE ORIGINS OF THE GUI



HowTo



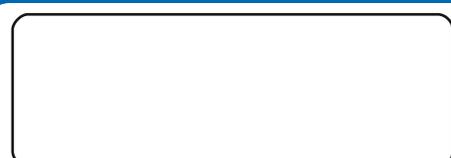
Python p.22



Blender p.26



... p.XX



... p.XX



Inkscape p.29

Graphics



Full Circle

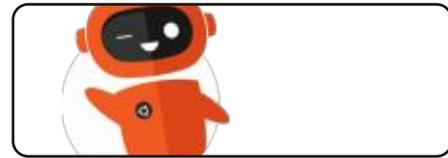
THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

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

Command & Conquer p.20



... p.XX



Ubuntu Devices p.40



The Daily Waddle p.33



My Opinion - GUI p.45



Letters p.XX



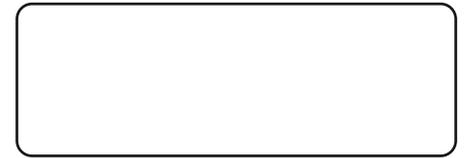
Q&A p.54



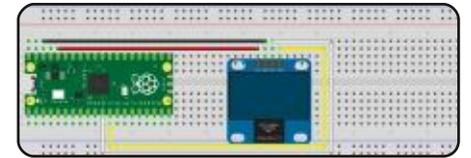
Review p.XX



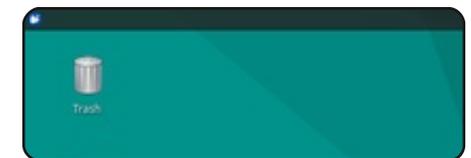
Linux News p.04



Everyday Ubuntu p.XX



Micro This Micro That p.35



Review p.50



Ubuntu Games p.57



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

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



WELCOME TO THE LATEST ISSUE OF FULL CIRCLE

Welcome to the first issue of 2022. I hope you're staying safe and well. As ever, we have Python, Inkscape, and Micro articles. Continuing last month's new series we have more Blender. Starting this month we have a series on the GUI. The Graphical User Interface that you probably use on a daily basis. Where did it come from? Who's idea was it? Find out in the coming month. Filling the third HowTo place we have; nothing. Probably for the first time in FCM history we only have enough for two HowTo articles. Yes, we really are that desperate for new stuff. As stated last month: **we DESPERATELY need articles**. I've got nothing spare. So, **PLEASE**, spend a few moments to write SOMETHING about what you know. It can be ANYTHING as long as it has something to do with Linux; hardware/software reviews, a how-to on something, even your story of how you found Ubuntu/Linux (of any flavor). Send whatever to: ronnie@fullcirclemagazine.org

Don't forget: if you're looking for some help, advice, or just a chit chat: remember, we have a **Telegram** group: <https://t.me/joinchat/24ec1oMFO1ZjZDc0>. I hope to see you there. Come and say hello.

Anyway, stay safe!

Ronnie

ronnie@fullcirclemagazine.org



This magazine was created using :



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/>

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



SLACKEL 7.5 DISTRIBUTION

RELEASED:

01.01.2022

Slackel 7.5 is out, built on the Slackware and Salix projects and fully compatible with the repositories they offer. A key feature of Slackel is its use of the constantly updated Slackware-Current branch. The graphical environment is based on the Openbox window manager. The size of the bootable image capable of running in Live mode is 2.4 GB (i386 and x86_64). The distribution can be used on systems with 512 MB of RAM.

The new release is in sync with the current Slackware branch and ships with the Linux 5.15 kernel. Updated versions of programs, including firefox 95.0.2, thunderbird 91.4.1, libreoffice 7.2.0, filezilla 3.56.0, smplayer 21.10.0, gimp 2.10.30. The fbpanel panel and a set of auxiliary graphical applications for configuring the system are used. Implemented full support for

installing the distribution kit on external USB-drives or SSD-drives to get a portable working environment. It supports the ability to update the environment installed on external media, and encrypt user data.

<http://www.slackel.gr/forum/viewtopic.php?f=3&t=688>

PINTA 2.0 GRAPHICS EDITOR

RELEASED:

02.01.2022

Pinta 2.0 is out, which is an attempt to rewrite the Paint.NET program using GTK. The editor provides a basic set of drawing and image processing capabilities, targeting novice users. The interface is simplified as much as possible, the editor supports an unlimited rollback buffer, allows you to work with multiple layers, is equipped with a set of tools for applying various effects and adjusting images. The Pinta code is licensed under the MIT license. The project is written in C # using Mono and the Gtk # binding. Binaries are

prepared for Linux (Flatpak , Snap), macOS and Windows.

<https://www.pinta-project.com/releases/2-0>

SOLUS PROJECT DRAMA:

02.01.2022

Strobl Joshua, a key developer of the Budgie desktop, announced his resignation of the Managing Board (Core Team) Solus the project and the abdication of leadership authority, responsible for interaction with the developers and the development of the user interface. Myers Beatrice, responsible for the technical part of the Solus, assured that the development of the distribution will continue in the near future, An announcement will be made about changing the project structure and the restructuring of the composition of the team. As the reason for leaving, Joshua cites the conflict that arose against the background of attempts to voice and solve problems that hinder the promotion of changes in Solus, both from direct project

participants and from stakeholders from the community. Joshua does not disclose the details of the conflict, as not to wash dirty linen in public. It is only mentioned that all his attempts to change the situation and improve work with the community were rejected and none of the problems voiced were resolved.

<https://twitter.com/JoshStrobl/status/1477391260340137988>

SPEED UP BUILDING THE LINUX KERNEL BY 50-80%:

03.01.2022

Ingo Molnar, a well-known Linux kernel developer and author of CFS Task Scheduler (Completely Fair Scheduler), proposed for discussion of Linux kernel development mailing list a series of patches, affecting more than half of all the files in the kernel source and provides an increase in the total rebuilding core speed 50-80% depending on the settings. The implemented optimization is remarkable in that it is associated

with the addition of the largest set of changes in the history of kernel development - 2297 patches were proposed for inclusion at once, changing more than 25 thousand files (10 thousand header files in the directories "include/" and "arch/*/include/" and 15 thousand source files).

The performance gain is achieved by changing the method of handling header files. It is noted that over thirty years of kernel development, the state of header files has taken on a depressing form due to the presence of a large number of cross-dependencies between files. The restructuring of the header files took over a year and required a significant redesign of the hierarchy and dependencies. During the restructuring, work was done to separate the type definitions and APIs for different kernel subsystems.

<https://lkml.org/lkml/2022/1/2/187>

KDE PLANS TO ACHIEVE PARITY BETWEEN WAYLAND AND X11 SESSIONS IN 2022:

03.01.2022

Nate Graham, QA developer for the KDE project, shared his thoughts on the direction the KDE project will take in 2022. Among other things, Nate believes that in the coming year, a KDE Plasma session based on the Wayland protocol will become ready for everyday use by a significant proportion of users and can completely replace the session based on X11. There are currently about 20 known issues being tracked when using Wayland in KDE, and the issues being added to the list have become less significant lately. The most important recent Wayland change is the addition of GBM (Generic Buffer Manager) support to the proprietary NVIDIA driver, which can be used by KWin.

<https://pointieststick.com/2022/01/03/kde-roadmap-for-2022/>

UBUNTUDE 21.10

RELEASE:

01/04/2022

UbuntuDDE 21.10 (Remix) distribution is out, based on Ubuntu 21.10 codebase and supplied with the DDE (Deepin Desktop Environment) graphical environment. The project is an unofficial Ubuntu edition, but the developers are trying to get UbuntuDDE included in the official Ubuntu editions. The size of the iso image is 3 GB.

In UbuntuDDE there is the Deepin 5.5 desktop and a set of specialized applications, developed by the project, including the file manager Deepin File Manager, DMusic music player, video player DMovie and messaging DTalk. Of the differences from Deepin Linux, there is a redesign and delivery of the Ubuntu Software Center with support for Snap and DEB format instead of the Deepin application store catalog. Kwin, developed by the KDE project, is used as the window manager.

<https://ubuntudde.com/blog/ubuntudde-remix-21-10-impish-release-note/>

REMNANTS OF THE PRECURSORS GAME

RELEASED:

01/04/2022

After six years of development, the release of "Remnants of the Precursors", a Java modernization of the 1993 game "Master of Orion", has been released. During the game, players compete with each other in the field of space colonization, developing their own civilizations, producing resources, constructing spaceships, conducting research and increasing their technological level. The program allows you to play one of 10 races, trying to lead it either to galactic domination or, in case of failure, to humiliating defeat. The micromanagement is minimal and the gameplay is smooth and streamlined. Source code is distributed under the GPLv3 license.

<https://rayfowler.itch.io/remnants-of-the-precursors/devlog/330806/remnants-of-the-precursors-is-released>

RELEASE OF LAZARUS 2.2.0:

06.01.2022

After three years of development, the integrated development environment Lazarus 2.2 is published, based on the FreePascal compiler and performing tasks similar to Delphi. The environment is designed to work with the FreePascal 3.2.2 compiler. Ready-made installation packages with Lazarus are prepared for Linux, macOS and Windows.

<https://forum.lazarus.freepascal.org/index.php/topic,57752.0.html>

TWENTY-FIRST FIRMWARE

UPDATE FOR UBUNTU

TOUCH:

06.01.2022

The UBports project, which took over the development of the Ubuntu Touch mobile platform after Canonical pulled out of it, has published an OTA-21 (over-the-air) firmware update. The project is also developing an experimental port of the Unity 8 desktop, which has been renamed to Lomiri.

Ubuntu Touch OTA-21 update has been generated for the following smartphones BQ E4.5 / E5 / M10 / U Plus, Cosmo Communicator, F (x) tec Pro1, Fairphone 2/3, Google Pixel 2XL / 3a, Huawei Nexus 6P, LG Nexus 4 / 5, Meizu MX4 / Pro 5, Nexus 7 2013, OnePlus 2/3/5/6 / One, Samsung

Galaxy Note 4 / S3 Neo +, Sony Xperia X / XZ / Z4, Vollaphone, Xiaomi Mi A2 / A3, Xiaomi Poco F1, Xiaomi Redmi 3s / 3x / 3sp / 4X / 7, Xiaomi Redmi Note 7/7 Pro. Separately, without the "OTA-21" tag, updates will be prepared for Pine64 PinePhone and PineTab devices.

<https://ubports.com/blog/ubports-news-1/post/ubuntu-touch-ota-21-release-3798>

QBITORRENT 4.4 RELEASED WITH BITTORRENT V2

SUPPORT:

01/07/2022

More than a year after the publication of the last major branch, the release of Bittorrent 4.4.0 is out, written using the Qt toolkit and developed as an open

alternative to µTorrent, close to it in terms of interface and functionality. QBittorrent features include an integrated search engine, the ability to subscribe to RSS, support for many BEP extensions, remote control via a web interface, sequential download mode in a given order, advanced settings for torrents, peers and trackers, bandwidth scheduler and IP filter, interface for creating torrents, support for UPnP and NAT-PMP.

<https://www.qbittorrent.org/news.php>

NEW VERSIONS OF GECKOLINUX:

07.01.2022

Now available, there is a new version of GeckoLinux - 999.220105 (rolling) and 153.220104 (static), based on openSUSE and paying more attention to optimizing the desktop and small things, such as high-quality rendering fonts. Version 999.220105 (1.6 GB) is developed as part of the rolling model for preparing updates, built on the Tumbleweed repository and Packman's own repository. Branch



DistroWatch.com

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

153.220104 (1.4 GB) is based on openSUSE 15.3 release. Additionally, there is a GeckoLinux NEXT branch with KDE Plasma 5.23.5 and Pantheon desktops (from elementary OS 5), which is built on openSUSE Leap 15.3, but with newer versions of user environments from separate openSUSE Build Service repositories.

<https://github.com/geckolinux/geckolinux-project/releases/>

RELEASE OF LINUX MINT 20.3:

07.01.2022

The Linux Mint 20.3 distribution is out, continuing the development of the branch based on Ubuntu 20.04 LTS. The distribution is fully compatible with Ubuntu, but significantly differs in the approach to organizing the user interface and the selection of default applications. The Linux Mint developers provide a desktop environment that matches the classic canons of desktop organization, which is more familiar to users who do not accept the new methods of building the GNOME 3 interface. DVD images based on

MATE 1.26 (2.1 GB), Cinnamon 5.2 (2.1 GB) and Xfce 4.16 (2 GB) are available for download. Linux Mint 20 is classified as a Long Term Support (LTS) release, with updates to be rolled out until 2025.

<http://blog.linuxmint.com/>

CANONICAL ANNOUNCES SNAPCRAFT TOOLKIT

REWORK:

08/01/2022

Canonical has revealed plans for a major overhaul of the Snapcraft toolkit for generating, distributing and updating self-contained Snap packages. Note that the current Snapcraft codebase has been declared legacy and will be applied when it is necessary to use old technologies. The drastic changes underway will not affect the current usage model - projects related to Ubuntu Core 18 and 20 will continue to use the old monolithic Snapcraft, and the new modular Snapcraft will begin to be applied starting with the Ubuntu Core 22 branch.

The old Snapcraft will be replaced by a new, more compact and modular option that will make

it easier for developers to create snap packages and eliminate the hassle of creating portable packages that work across distributions. The base for the new Snapcraft is the Craft Parts mechanism, allows, when assembling packages, it to receive data from different sources, process it in different ways and form a directory hierarchy in the FS, suitable for deploying packages. Craft Parts implies the use of portable components in the project, which can be independently downloaded, assembled and installed.

The choice of a new or old implementation of Snapcraft will be carried out through a special fallback mechanism integrated into the build process. Thus, existing projects will be able to build snap-packages without changes and will require modification only when the packages are transferred to a new version of the base of the Ubuntu Core system.

<https://snapcraft.io/blog/the-future-of-snapcraft>

MESSOR, A DECENTRALIZED INTRUSION DETECTION

SYSTEM RELEASED:

08/01/2022

After two years of development, the first release of the Messor project is available, a free, independent and decentralized software for securing networks and transparently collecting attack and scan data. Messor.Network launched and published a plugin for the OpenCart3 e-commerce platform. The plugin code is written in PHP and is distributed under the Apache 2.0 license. A module for nginx / apache2 (C++), a plugin for Magento (php) and a plugin for Wordpress (php) are in development.

The project provides a combination of IPS, HoneyPot and a hybrid P2P client that implements scan protection, regardless of the purpose, be it exploiting vulnerabilities, bots, search engines or other applications. The main difference between Messor and other IPSs is its network structure. The linked sites form a single P2P Messor-Network, each member collects data about attackers, sends information to other network members, and receives daily

database updates. Each member of the Messor network is responsible for distributing the current database to other network members and sending the collected attack data to the central servers of the network.

<https://messor.network/ru/>

RELEASE OF EASYOS 3.2: 09/01/2022

Barry Kauler, founder of the Puppy Linux project, has published the experimental EasyOS 3.2 distribution that attempts to combine Puppy Linux technologies using container isolation to run system components. Each application, as well as the desktop itself, can be launched in separate containers, which are isolated using their own Easy Containers mechanism. The distribution package is managed through a GUI developed by the project. The size of the boot image is 580MB.

Among other features of the distribution, we can note the default operation with root privileges with resetting privileges at startup of each application, since EasyOS is positioned as a Live-

system for one user (it is also possible to work under the unprivileged 'spot' user). The distribution kit is installed in a separate subdirectory (the system is located in /releases/easy-3.2, user data is stored in the /home directory, and additional containers with applications in the /containers directory) and can coexist with other data on the drive. It is possible to encrypt individual subdirectories (for example, /home) and install SFS meta-packages, which are Squashfs mount images that combine several regular packages.

<https://bkhome.org/news/202201/easyos-version-32-released.html>

CLONEZILLA LIVE 2.8.1 RELEASED: 09/01/2022

Clonezilla Live 2.8.1, intended for fast disk cloning (only used blocks are copied), is out. The tasks performed by the distribution are similar to those of the proprietary Norton Ghost product. The size of the iso-image is 339 MB (i686, amd64).

The distribution is based on

Debian and uses the code of projects such as DRBL, Partition Image, ntfsclose, partclone, udpcast. Bootable from CD / DVD, USB Flash and network (PXE). Supported LVM2 and FS ext2, ext3, ext4, reiserfs, reiser4, xfs, jfs, btrfs, f2fs, nilfs2, FAT12, FAT16, FAT32, NTFS, HFS +, UFS, minix, VMFS3 and VMFS5 (VMWare ESX). There is a mode of mass cloning over the network, including the transmission of traffic in multicast mode, which makes it possible to simultaneously clone the source disk to a large number of client machines. Cloning from one disk to another is possible, as well as creating backups by saving a disk image to a file. (Cloning entire disks or individual partitions is possible.)

<https://sourceforge.net/p/clonezilla/news/2022/01/stable-clonezilla-live-281-12-released/>

A COMPLETELY FREE VERSION OF THE LINUX-LIBRE 5.16 KERNEL IS AVAILABLE: 10/01/2022

With a short delay, the Latin American Free Software Foundation published a completely free version of the Linux 5.16

kernel - Linux-libre 5.16-gnu, stripped of firmware elements and drivers containing non-free components or code fragments whose scope is limited by the manufacturer. In addition, Linux-libre disables the kernel's functionality to load non-free components that are not shipped with the kernel, and removes the mention of using non-free components from the documentation.

<https://www.fsfla.org/pipermail/linux-libre/2022-January/003466.html>

SDL 2.0.20 MEDIA LIBRARY RELEASE: 01/12/2022

The SDL 2.0.20 (Simple DirectMedia Layer) library, aimed at simplifying writing games and multimedia applications, has been released. The SDL library provides facilities such as hardware-accelerated 2D and 3D graphics output, input processing, audio playback, 3D output via OpenGL/OpenGL ES/Vulkan, and many other related operations. The library is written in C language and distributed under the Zlib license.

Bindings are provided for using SDL capabilities in projects in various programming languages.

<https://libsdl.org/>

THIRD RELEASE CANDIDATE FOR SLACKWARE LINUX 15.0:

13.01.2022

Patrick Volkerding has announced the third and final release candidate of the Slackware 15.0 distribution, which has reached a freeze of 99% of packages before release. An installation image has been prepared for download, 3.4 GB in size (x86_64), as well as a reduced image for running in Live mode.

In the final changes before freezing, the Linux kernel is updated to version 5.15.14 (the possibility of including release 5.15.15 in the release is being considered), KDE Plasma 5.23.5, KDE Gear 21.12.1, KDE Frameworks 5.90, eudev 3.2.11, vala 0.54.6, iproute2 5.16.0, firefox 91.5, thunderbird 91.5.0, sqlite 3.37.2, mercurial 6.0.1, pipewire 0.3.43, pulseaudio 15.0, mdadm 4.2, mesa 21.3.3, wpa_supplicant 2.9, xorg-

server 1.20.14, gimp 2.10.30, gtk 3.24, freetype 2.11.1.

Slackware has been in development since 1993 and is the oldest distribution in existence. The distro features a lack of complexity and a simple initialization system in the style of classic BSD systems, which makes Slackware an interesting solution for learning how Unix-like systems work, experimenting and getting to know Linux.

http://www.slackware.com/changelog/current.php?cpu=x86_64

ARTI 0.0.3, A RUST IMPLEMENTATION OF THE TOR CLIENT:

01/15/2022

The developers of the Tor anonymous network have released Arti 0.0.3, a Tor client written in Rust. The project has the status of experimental development, it lags behind the functionality of the main Tor client in C and is not yet ready to fully replace it. In March, the release of 0.1.0 is expected, which is positioned as the first beta release of the project, and in the fall, the

release of 1.0 with the stabilization of the API, CLI and settings, which will be suitable for initial use by ordinary users. In the more distant future, when the Rust code reaches a level capable of completely replacing the C version, the developers intend to make Arti the main implementation of Tor and stop maintaining the C implementation.

Unlike the C implementation, which was originally designed as a SOCKS proxy and then tailored to other needs, Arti is initially developed as a modular plug-in library that can be used by various applications. In addition, developing a new project, all past Tor development experience is taken into account, which will avoid known architectural problems, make the project more modular and efficient. The code is distributed under the Apache 2.0 and MIT licenses.

<https://gitlab.torproject.org/tpo/core/arti/>

SYSTEMRESCUE 9.0.0

RELEASED:

17.01.2022

SystemRescue 9.0.0 is available, a specialized Live distribution based on Arch Linux designed for system recovery after a crash. Xfce is used as the graphical environment. ISO image size is 771 MB (amd64, i686).

In the new version, the translation of the system initialization script from Bash to Python is highlighted, as well as the implementation of initial support for setting system parameters and autorun (autorun) using files in YAML format. The core package includes the aq, libisoburn, patch, python-llfuse, python-yaml, and rdiff-backup packages, as well as a selection of documentation from the site. The Linux kernel has been updated to the 5.15 branch, which introduces a new NTFS driver with write support.

The sysrescue-customize script has been implemented to build custom versions of ISO images with SystemRescue. The full Mesa package has been replaced with a stripped-down version, saving 52

MB of disk space. They removed the xf86-video-qxl driver due to stability issues. The package inetutils (telnet, ftp, hostname) was returned, which was previously excluded from the base composition by mistake.

<https://www.system-rescue.org/>

RELEASE OF GHOSTBSD 01/22/12:

17.01.2022

The 01/22/12 release of GhostBSD based on FreeBSD 13-STABLE and offering MATE user environment is out. By default, GhostBSD uses the ZFS file system. Both Live mode and installation to a hard drive are supported (using its own ginstall installer written in Python). Boot images are generated for the x86_64 architecture (2.58 GB).

The new version removes components from the base system that provide optional support for the OpenRC init system. The dhcpd package has also been removed from the distribution in favor of a regular DHCP client from FreeBSD. VLC media player was rebuilt with UPNP support. It

provides distribution identification in the /etc/os-release file (instead of FreeBSD 13.0-STABLE it is now GhostBSD 01/22/12) and GhostBSD in the output of the uname command. The initgfx package is used to automatically tune AMD Radeon HD 7000 and older GPUs. Extraction of data about security problems from the vuxml.freebsd.org database and marking of packages with unpatched vulnerabilities is provided. They removed p7zip from base distribution due to vulnerabilities and maintenance issues.

<https://ghostbsd.org/ghostbsd-22.01.12-iso-is-now-available>

RELEASE OF MUMBLE 1.4:

17.01.2022

After more than two years of development, the Mumble 1.4 platform is released, focused on creating voice chats that provide low latency and high voice quality. The project code is written in C++ and distributed under the BSD license. Builds prepared for Linux, Windows and macOS.

The project consists of two modules - the mumble client and the murmur server. The graphical interface is based on Qt. The Opus audio codec is used to transmit audio information. A flexible system of access control is provided, for example, it is possible to create voice chats for several isolated groups with the possibility of separate communication of leaders in all groups. Data is transmitted only over an encrypted communication channel. By default, authentication based on public keys are used.

Unlike centralized services, Mumble allows you to store user data on your own servers and fully control the operation of the infrastructure, if necessary, connecting additional handler scripts, for writing which a special API based on the Ice and GRPC protocols is available. You can use existing users for authentication or connect sound bots, which, for example, can play music. It is possible to manage the server through a web interface. Users can search for friends on different servers too.

<https://www.mumble.info/>

ESSENCE - AN ORIGINAL OPERATING SYSTEM WITH ITS OWN KERNEL AND GRAPHICAL SHELL:

17.01.2022

Available for initial testing is a new Essence operating system. It comes with its own kernel and graphical user interface. The project has been developed by one enthusiast since 2017, created from scratch and is notable for its original approach to building a desktop and graphics stack. The most notable feature is the ability to divide windows into tabs, which make it possible to work in one window with several programs open at once and group applications into windows depending on the tasks being solved. The project code is written in C++ and distributed under the MIT license.

The kernel includes a task scheduler with support for multiple priority levels, a memory management subsystem with support for shared memory, mmap and a multi-threaded memory pager, a network stack (TCP / IP), an audio subsystem for mixing sounds,

VFS and an EssenceFS file system with a separate layer for data caching. In addition to its own FS, drivers for Ext2, FAT, NTFS and ISO9660 are provided. It supports the removal of functionality into modules with the ability to load modules as needed. Drivers prepared for ACPI with ACPIA, IDE, AHCI, NVMe, BGA, SVGA, HD Audio, Ethernet 8254x and USB XHCI (storage and HID)

The system can run on legacy hardware with less than 64 MB of RAM and takes up approximately 30 MB of disk space. To conserve resources, only the active application runs and all background programs are suspended. Loading and shutting down takes less than a second. The project publishes daily new ready-made builds suitable for testing in QEMU.

<https://nakst.gitlab.io/essence>

FFMPEG 5.0 RELEASED:

17.01.2022

After ten months of development, the FFmpeg 5.0 multimedia package is available, which includes a set of applications and a collection of libraries for

operations on various multimedia formats (recording, converting and decoding audio and video formats). The package is distributed under the LGPL and GPL licenses, the development of FFmpeg is carried out adjacent to the MPlayer project. The major change in the version number is due to significant changes in the API and the transition to a new release generation scheme, in accordance with which new major releases will be formed once a year, and releases with extended support time - once every two years. FFmpeg 5.0 will be the first LTS release of the project.

<https://ffmpeg.org/index.html#news>

RELEASE OF GNU RADIO

3.10.0:

18.01.2022

After a year of development, a new major release of the free digital signal processing platform GNU Radio 3.10 is out. The platform includes a set of programs and libraries that allow you to create arbitrary radio systems, modulation schemes in the form of received and sent signals in which are set programmatically, and simple

hardware devices are used to capture and generate signals. The project is distributed under the GPLv3 license. The code for most of the GNU Radio components is written in Python, the parts critical to performance and latency are written in C++, which allows the package to be used when solving problems in real time.

In combination with universal programmable transceivers that are not tied to the frequency band and type of signal modulation, the platform can be used to create devices such as base stations for GSM networks, devices for remote reading RFID tags (electronic certificates and passes, smart cards), GPS receivers, WiFi, FM radio receivers and transmitters, TV decoders, passive radars, spectrum analyzers, etc. In addition to USRP, the package can use other hardware components for signal input and output, for example, drivers are available for sound cards, TV tuners, BladeRF devices, Myriad-RF, HackRF, UmTRX, Softrock, Comedi, Funcube, FCOMMS, USRP and S-Mini.

The set also includes a collection of filters, channel codecs, synchronization modules,

demodulators, equalizers, voice codecs, decoders and other elements necessary to create radio systems. These elements can be used as building blocks for the finished system, which, combined with the ability to determine data flows between blocks, allows you to design radio systems even without programming skills!

<https://www.gnuradio.org/news/2022-01-17-gnuradio-v3.10.0.0-release/>

RELEASE OF THE DEEPIN 20.4:

18.01.2022

Based on Debian 10, but developing its own Deepin Desktop Environment (DDE) and about 40 user applications, including the DMusic music player, the DMovie video player, the DTalk messaging system, the Deepin installer and installation center, the Deepin 20.4 distribution was released. The project was founded by a group of developers from China, but has been transformed into an international project. It is distributed under the GPLv3 license. The size of the bootable iso image is 3 GB (amd64).

Desktop components and applications are developed using C/C++ (Qt5) and Go languages. The key feature of the Deepin desktop is the panel, which supports multiple modes of operation. The efficient mode is somewhat reminiscent of Unity, mixing indicators of running programs, favorite applications and control applets (volume / brightness settings, connected drives, clock, network status, etc.). The program launcher interface is displayed in full screen and provides two modes - viewing favorite applications and navigating through the catalog of installed programs.

<http://www.deepin.org/?language=en>

ONLYOFFICE Docs 7.0

OFFICE SUITE:

18.01.2022

The release of ONLYOFFICE DocumentServer 7.0 has been published with the implementation of a server for ONLYOFFICE online editors and collaboration. Editors can be used to work with text documents, spreadsheets and presentations. The project code is

distributed under the free license AGPLv3. At the same time, ONLYOFFICE DesktopEditors 7.0 was released, built on a single code base with online editors. Desktop editors are designed as desktop applications that are written in JavaScript using web technologies, but combine client and server components in one set, designed for self-sufficient use on the user's local system, without resorting to an external service. For collaboration at your own facilities, you can also use the Nextcloud Hub platform, which provides full integration with ONLYOFFICE.

ONLYOFFICE claims full compatibility with MS Office and OpenDocument formats. Supported formats include: DOC, DOCX, ODT, RTF, TXT, PDF, HTML, EPUB, XPS, DjVu, XLS, XLSX, ODS, CSV, PPT, PPTX, ODP. It is possible to extend the functionality of the editors through plugins, for example, plugins are available for creating templates and adding YouTube videos. Ready builds are generated for Windows and Linux (deb and rpm packages).

<https://www.onlyoffice.com/blog/2022/01/onlyoffice-docs-7-0/>

STABLE RELEASE OF WINE 7.0:

19.01.2022

After a year of development and 30 experimental versions, a stable release of the open implementation of the Win32 API - Wine 7.0 is out, which has absorbed more than 9100 changes. Among the key achievements of the new version are the translation of most Wine modules in PE format, support for themes, the expansion of the stack for joysticks and input devices with the HID interface, the implementation of the WoW64 architecture for running 32-bit programs in a 64-bit environment.

Full-fledged work of 5156 (5049 a year ago) programs for Windows was confirmed in Wine, another 4312 (4227 a year ago) programs work fine with additional settings and external DLLs. 3813 programs (3703 years ago) have minor problems that do not interfere with the use of the main functions of applications.

<https://www.winehq.org/announce/7.0>

RELEASE OF ALPHAPLOT, A SCIENTIFIC PLOTTING

PROGRAM:

20.01.2022

The AlphaPlot 1.02 release is out. It provides a graphical interface for the analysis and visualization of scientific data. Development of the project began in 2016 as a fork of SciDAVis 1.D009, which in turn is a fork of QtiPlot 0.9rc-2. During the development process, a migration from the QWT library to QCustomplot was carried out. The code is written in C++, uses the Qt library and is distributed under the GPLv2 license.

AlphaPlot aims to be a data analysis and graphical representation tool providing powerful mathematical processing and visualization (2D and 3D). There is support for various methods of approaching given points using curves. Results can be saved in raster and vector formats such as PDF, SVG, PNG and TIFF. It supports the creation of scripts for automating the construction of graphs in the JavaScript language. Plugins can be used to expand functionality.

The new version improves the system for managing the placement of elements on 2D charts, expanded navigation through 3D charts, also adding tools for saving and loading templates, a new dialog with settings, and also implemented support for arbitrary filling patterns, cloning charts, saving and printing 3D graphs, vertical and horizontal grouping of panels.

<https://github.com/narunlifescience/AlphaPlot/releases/tag/1.02>

LIGHTTPD HTTP SERVER RELEASE 1.4.64:

20.01.2022

The lighttpd 1.4.64 lightweight http server has been released. The new version introduces 95 changes, including applying previously planned changes to defaults and cleaning up deprecated functionality. Lighttpd 1.4.64 also fixes a vulnerability (CVE-2022-22707) in the mod_extforward module that causes a 4-byte buffer overflow when processing data in the Forwarded HTTP header. According to the developers, the problem is

limited to a denial of service and allows you to remotely initiate an abnormal termination of a background process. Operation is only possible when the Forwarded header handler is enabled and does not appear in the default configuration.

<https://blog.lighttpd.net/articles/2022/01/19/lighttpd-1.4.64-released/>

OPENSUSE WEB INTERFACE FOR THE YAST INSTALLER:

21.01.2022

Following the announcement of the move to the web interface of the Anaconda installer used in Fedora and RHEL, the developers of the YaST installer revealed plans to develop the D-Installer project and create a front-end to manage the installation of openSUSE and SUSE Linux distributions through a web interface.

It is noted that the project has been developing the WebYaST web interface for a long time, but it is limited by the possibilities of remote administration and system configuration, is not designed for use as an installer, and is rigidly tied

to the YaST code. D-Installer is seen as a platform that provides several installation frontends (Qt GUI, CLI and Web) on top of YaST. Associated plans include work to shorten the installation process, separate the user interface from YaST internals, and add a web interface.

Technically, D-Installer is an abstraction layer implemented on top of the YaST libraries and provides a unified interface for accessing functions such as package installation, hardware verification, and disk partitioning via D-Bus. The graphical and console installers will be migrated to the specified D-Bus API, as well as a browser-based installer that interacts with D-Installer through a proxy service that provides access to D-Bus calls over HTTP.

<https://yast.opensuse.org/blog/2022-01-18/announcing-the-d-installer-project>

GNU OCRAD 0.28 OCR RELEASED:

21.01.2022

Three years since the last release, Ocrad 0.28 (Optical

Character Recognition), a text recognition system developed under the auspices of the GNU project, is out. Ocrad can be used both as a library for integrating OCR functions into other applications, and as a standalone utility that outputs text in UTF-8 or 8-bit encodings based on the input image.

For OCR, Ocrad uses the feature extraction method. It includes a page layout analyzer that allows you to correctly separate columns and blocks of text in printed documents. Recognition is supported only for characters from "ascii", "iso-8859-9" and "iso-8859-15" encodings.

It is noted that the new release includes a large portion of minor fixes and improvements. The most significant change was the support for the PNG image format, implemented using the libpng library, which greatly simplified the program, since previously only images in PNM formats could be input.

<https://www.mail-archive.com/info-gnu@gnu.org/msg02997.html>

NEW VERSION OF MONITORIX 3.14.0:

01/21/2022

Monitorix 3.14.0 is presented, designed for visual monitoring of the operation of various services, like monitoring CPU temperature, system load, network activity and responsiveness of network services. The system is controlled via a web interface, the data is presented in the form of graphs.

The system is written in Perl, RRDTool is used to generate graphs and store data, the code is distributed under the GPLv2 license. The program is quite compact and self-sufficient (there is a built-in http-server), which makes it possible to use it even on embedded systems. A fairly wide range of monitoring parameters is supported, from monitoring the task scheduler, I/O, memory allocation and OS kernel parameters to visualizing data on network interfaces and specific applications (mail servers, DBMS, Apache, nginx).

<http://www.monitorix.org/news.html?n=20220118>

ARCHLABS RELEASE 2022.01.18:

01/21/2022

ArchLabs 2022.01.18 distribution was released, based on Arch Linux and delivered with a lightweight user environment based on the Openbox window manager (i3, Bspwm, Awesome, JWM, dk, Fluxbox, Xfce, Deepin, GNOME, Cinnamon, Sway). For a stationary installation, the ABIF installer is offered. The base package includes applications such as Thunar, Termite, Geany, Firefox, Audacious, MPV and Skippy-XD. The size of the installation iso-image is 959 MB.

The new version adds support for the tiled window manager dk and Sway, which uses Wayland. Improved session implementation based on Xfce. They updated packages in the repository and updated the theme. A new version of the BAPH utility for working with AUR repositories has been used, which has a new option to check for updates.

<https://forum.archlabslinux.com/t/archlabs-2022-01-18-available-for-download/6136>

LINUX KERNEL VFS VULNERABILITY ALLOWING PRIVILEGE ESCALATION:

21.01.2022

A vulnerability (CVE-2022-0185) has been identified in the Filesystem Context API provided by the Linux kernel that could allow a local user to gain root privileges on the system. The researcher who identified the problem has published a demo of an exploit that allows code to be run as root on Ubuntu 20.04 in the default configuration. The exploit code is planned to be published on GitHub within a week after distributions release an update fixing the vulnerability.

The problem has been manifesting since the Linux 5.1 kernel and was fixed in yesterday's updates 5.16.2 , 5.15.16 , 5.10.93 , 5.4.173 . Vulnerability package updates have already been released for RHEL , Debian , Fedora , Ubuntu . The fix is not yet available on Arch Linux , Gentoo , SUSE and openSUSE . As a security workaround for systems not using container isolation, you can set the

value of sysctl
"user.max_user_namespaces" to 0:

<https://twitter.com/ryaagard/status/1483592308352294917>

FRAMEWORK COMPUTER HAS OPENED ITS FIRMWARE CODE FOR LAPTOPS:

22.01.2022

Notebook manufacturer Framework Computer, which is a supporter of giving users the right to self-repair and trying to make their products as convenient as possible for disassembling, upgrading and replacing components, announced the publication of the source code for the Embedded Controller (EC) firmware used in the Framework Laptop. The code is open under the BSD license.

The main idea is to provide the ability to build a laptop from modules, similar to how a user can assemble a desktop computer from separate components that are not imposed by a specific manufacturer. The Framework Laptop can be ordered as parts and assembled into a final device by the user. Each component in the device is clearly

labeled and easily removed. If necessary, the user can quickly replace any module, and in the event of a breakdown, try to repair his device on his own, using the instructions and videos provided by the manufacturer with information on assembly / disassembly and replacement of components.

In addition to replacing memory and storage, it is possible to replace the motherboard, case (different colors are offered), keyboard (different layouts) and wireless adapter. Through Expansion Card slots without dismantling the case, you can connect up to 4 additional modules with USB-C, USB-A, HDMI, DisplayPort, MicroSD and a second drive to the laptop. This feature allows the user to select the required set of ports and replace them at any time (for example, if there is not enough USB ports, you can replace the HDMI module with USB). In case of a breakdown or for upgrading, you can separately buy components such as a screen (13.5" 2256x1504), battery, touchpad, webcam, keyboard, sound card, case, fingerprint sensor board, screen hinges and speakers.

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

SWAY 1.7 RELEASE USING WAYLAND:

23.01.2022

Sway 1.7 compositing manager is out, built using the Wayland protocol and fully compatible with the i3 tiled window manager and the i3bar panel. The project code is written in C and distributed under the MIT license. The project is intended for use on Linux and FreeBSD.

Compatibility with i3 is provided at command level, configuration files and IPC, which allows Sway to be used as a transparent replacement for i3, using Wayland instead of X11. Sway allows you to place windows on the screen not spatially, but logically. Windows are laid out in a grid that makes optimal use of screen space and allows you to quickly manipulate windows using only the keyboard.

To create a complete user environment, related components are offered: swayidle (background process with the implementation of the KDE idle protocol), swaylock (screen saver), mako (notification manager), grim (creating

screenshots), slurp (selecting an area on the screen), wf-recorder (video capture), waybar (application bar), virtboard (on-screen keyboard), wl-clipboard (clipboard management), wallutils (desktop wallpaper management).

Sway is developed as a modular project built on top of the wlroots library, which contains all the basic primitives for a composite manager.

<https://github.com/swaywm/sway/releases/tag/1.7>

JINGOS 1.2, TABLET PC DISTRIBUTION RELEASED:

23.01.2022

JingOS 1.2 distribution is available for download. It provides an environment specially optimized for installation on tablet PCs and laptops with a touch screen. The project code is distributed under the GPLv3 license. Release 1.2 is only available for tablets based on ARM architecture (previously there were releases for the x86_64 architecture as well, but the focus has recently shifted to the ARM architecture).

The distribution is based on Ubuntu 20.04 and the user environment is based on KDE Plasma Mobile. To create the application interface, Qt, a set of Mauikit components and the Kirigami framework from the KDE Frameworks are used, allowing you to create universal interfaces that automatically scale to different screen sizes. On-screen gestures are actively used for control on touch screens and touchpads, such as zooming by pinching and changing pages by swiping.

OTA update delivery is supported to keep the software up to date. Installation of programs can be done both from the Ubuntu repositories and the Snap directory, and from a separate application store. The distribution also includes a JAAS layer (JingPad Android App Support), which allows, in addition to stationary Linux desktop applications, you to run applications created for the Android platform (you can run programs for Ubuntu and Android side by side).

<https://en.jingos.com/>

RELEASE OF SANE 1.1 WITH SUPPORT FOR NEW SCANNER MODELS:

24.01.2022

The sane-backends 1.1.1 package is out. It includes a set of drivers, the scanimage command-line utility, a daemon for organizing network scanning saned and libraries with the SANE-API implementation. The project code is distributed under the GPLv2 license.

The package supports 1747 (in the previous version 1652) scanner models, of which 815 (737) have the status of full support for all functions, for 780 (766) the level of support is rated as good, for 129 (126) - acceptable, and for 23 (23) - minimal. Additionally, for 462 (464) devices, there is an incompletely tested driver implementation. Support for 470 (478) scanners remains unrealized.

The release of version 1.1.1 following 1.0.32 marked a change in the model for assigning version numbers. Instead of continuous mixed numbering for major and minor releases (1.0.30, 1.0.31, 1.0.32), the change to the second

digit in the version number (1.1.0, 1.2.0, 1.3.0, etc.) will now be for major releases, and the third digit (1.1.2, 1.1.3) will change for corrective updates. Release 1.1.0 was skipped due to a bug discovered late in the release process.

<https://alioth-lists.debian.net/pipermail/sane-devel/2022-January/039340.html>

VULKAN 1.3 GRAPHICS STANDARD PUBLISHED:

25.01.2022

After two years of work, the Khronos graphics standards consortium has published the Vulkan 1.3 specification, which defines an API for accessing GPU graphics and computing capabilities. It is noted that the requirements of the Vulkan 1.3 specification are designed for OpenGL ES 3.1 class graphics hardware, which will ensure support for the new graphics API in all GPUs that support Vulkan 1.2. The Vulkan SDK toolkit is scheduled to be published in mid-February. In addition to the main specification, they plan to offer additional extensions for mid-range and high-

end mobile and desktop devices, which will be supported as part of the Vulkan Milestone edition.

At the same time, a plan is presented to implement support for the new specification and additional extensions in graphics cards and device drivers. Products with support for Vulkan 1.3 are being prepared to be released by Intel, AMD, ARM, and NVIDIA. For example, AMD has announced that Vulkan 1.3 support will soon be available on the AMD Radeon RX Vega series of graphics cards, as well as on all cards based on the AMD RDNA architecture. NVIDIA is preparing to publish drivers supporting Vulkan 1.3 for Linux and Windows. ARM will add Vulkan 1.3 support to Mali GPUs.

<https://www.khronos.org/news/press/vulkan-reduces-fragmentation-and-provides-roadmap-visibility-for-developers>

RETROARCH 1.10.0

RELEASED:

25.01.2022

After a year and a half of development RetroArch 1.10.0 released, an add-on for emulating

FULL CIRCLE WEEKLY NEWS



Join our host Moss Bliss 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>



various game consoles, allowing you to run classic games using a simple, unified graphical interface. The use of console emulators such as Atari 2600/7800/Jaguar/Lynx, Game Boy, Mega Drive, NES, Nintendo 64/DS, PCEngine, PSP, Sega 32X/CD, SuperNES, etc. is supported. Gamepads from existing game consoles can be used, including the Playstation 3, Dualshock 3, 8bitdo, Xbox 1 and Xbox360, as well as general purpose gamepads such as the Logitech F710. The emulator supports advanced features such as multiplayer games, state saving, image enhancement of older games with shaders, game rewind, hot plugging of game pads, and video streaming.

<https://www.libretro.com/index.php/retroarch-1-10-0-release/>

RELEASE OF SCRIBUS 1.5.8: 26.01.2022

Scribus 1.5.8, a free document layout package, has been released to provide professional print layout tools, including flexible PDF generation tools and support for split color profiles, CMYK, spot

colors, and ICC. The system is written using the Qt toolkit and is licensed under the GPLv2+ license. Ready binary builds are prepared for Linux (ApplImage), macOS and Windows.

The 1.5 branch is billed as experimental and includes features such as a new Qt5-based user interface, a redesigned file format, full table support, and enhanced word processing. Release 1.5.8 is marked as well tested and already quite stable for working on new documents. After the final stabilization and recognition of readiness for widespread implementation, a stable release of Scribus 1.6.0 will be based off the 1.5 branch.

<https://www.scribus.net/scribus-1-5-8-released/>

SDL MEDIA LIBRARY MOVES TO USE WAYLAND BY DEFAULT:

27.01.2022

A change has been made to the SDL (Simple DirectMedia Layer) library codebase - by default enabling work based on the Wayland protocol in environments

that provide simultaneous support for Wayland and X11. Previously, Wayland environments with the XWayland component enabled output using X11 by default, and to use Wayland, you had to start the application with the "SDL_VIDEODRIVER=wayland" environment variable. The change will be part of the SDL 2.0.22 release scheduled for March. It is noted that for full operation of SDL-based applications in Wayland, the presence of the libdecor library for decorating windows on the client side is required.

<https://github.com/libsdl-org/SDL/commit/8ceba27d6291f1195e13608033ec439aec621fc6>

BIND DNS SERVER 9.18.0 RELEASED WITH DNS-OVER- TLS AND DNS-OVER- HTTPS SUPPORT: 27.01.2022

After two years of development, the ISC has released the first stable release of a major new branch of the BIND 9.18 DNS server. Branch 9.18 will be supported for three years until Q2

2025 as part of an extended maintenance cycle. Support for the 9.11 branch will end in March, and the 9.16 branch in mid-2023. An experimental branch of BIND 9.19.0 has been built to develop the functionality of the next stable version of BIND.

The release of BIND 9.18.0 is notable for the implementation of support for DNS over HTTPS (DoH, DNS over HTTPS) and DNS over TLS (DoT, DNS over TLS) technologies, as well as the XoT mechanism (XFR-over-TLS) for secure transmission of DNS-over-TLS content. zones between servers (both sending and receiving zones via XoT is supported). With appropriate settings, a single named process can now serve not only traditional DNS queries, but also queries sent using DNS-over-HTTPS and DNS-over-TLS. DNS-over-TLS client support is built into the dig utility, which can be used to send queries over TLS when the "+tls" flag is specified.

The implementation of the HTTP/2 protocol used in DoH is based on the ngtcp2 library, which is included in the optional build dependencies. Certificates for DoH and DoT can be provided by the

user or generated automatically during startup.

<https://www.mail-archive.com/bind-announce@lists.isc.org/msg00622.html>

RELEASE OF OPNSENSE 22.1:

27.01.2022

The OPNsense 22.1 firewall distribution was released, which is a fork of the pfSense project, created with the goal of creating a completely open distribution that could have functionality at the level of commercial solutions for deploying firewalls and network gateways. Unlike pfSense, the project is not controlled by one company, it is developed with the direct participation of the community and has a completely transparent development process, as well as providing the opportunity to use any of its code in third-party products, including commercial ones. The source code of the distribution components, as well as the tools used for building, are distributed under the BSD license. They offer a LiveCD and a system image for writing to Flash drives (339 MB).

The core of the distribution is based on the FreeBSD code. Among the features of OPNsense, one can single out a completely open assembly toolkit, the ability to install over regular FreeBSD, load balancing tools, a web interface for organizing user connection to the internet (Captive portal), the presence of mechanisms for tracking connection states (stateful firewall based on pf), setting bandwidth limits, filtering traffic, creating a VPN based on IPsec, OpenVPN and PPTP, integration with LDAP and RADIUS, support for DDNS (Dynamic DNS), a system of visual reports and graphs.

The distribution provides tools for creating fault-tolerant configurations based on the use of the CARP protocol and allowing you to launch a spare node in addition to the main firewall, which will be automatically synchronized at the configuration level and take over the load in case of failure of the primary node. The administrator is offered a modern and simple interface for configuring the firewall, built using the Bootstrap web framework.

<https://forum.opnsense.org/index.php?topic=26536.0>

FULL CIRCLE WEEKLY NEWS



Join our host Moss Bliss 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>



THE

VIRTUALBOX NETWORKING PRIMER

Connecting and Configuring
Virtual Machines



Robin Catling

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](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=16002536
99&s=books&sr=1-1](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://www.kobo.com/us/en/ebook/the-virtualbox-networking-primer)



COMMAND & CONQUER

Written by Erik

Most of us don't think twice when typing at a terminal, it's just there. What is actually happening is that you are typing commands into a shell that is interpreting your commands. By default, Ubuntu ships with bash, but you can install another if you like. I will focus on bash. The reason being, that if you ever get to use servers or containers on the internet, chances are good that it is bash.

A lot of time I hear, it is difficult to learn shell scripting. Sure, bash is not the friendliest of syntaxes, but sometimes you just need a question answered and a small shell script will do. The question may be something arbitrary like: Do I have any money? If you are a wage serf like me, the only time you have money is when you have been paid. So to simplify the question, I might ask, have I been paid? Now that is a yes or no question, the kind the computer likes. Yes or no questions translate easily into true or false. There is a dilemma in this question, as the computer has no access to my bank account and therefore has

no idea if I have been paid. However, I can do the next best thing, check if it is 'payday'. The computer has no idea what 'payday' is, but we do, it is on the 27th of each month. We have a target and a yes or no question, so we can write a simple script.

Open those terminals and let's have some silly fun. Type:

```
nano gotmoney.sh
```

and press enter.

As you know, all shell scripts require a path to the shell. So we start off by adding: `#!/bin/bash` -to the header of our script. We can use bash's if-statement to ask the question; "If today is the 27th ... TRUE or FALSE" then we got paid or



did not get paid" The if statement in bash does this nifty trick of opening with the word "if" and closing with "fi" (the reverse). The way it works is, IF TRUE, then do something, else do nothing.

So this is our outline, we now need to formulate the question. Here is where you need to think about what you want. Remember, garbage in, garbage out.

Let's open another terminal or tab and test out our 'question'. This is important, if your command does not work in bash, it will not work in your script.

First, let's see what the switches are, type: `man date` and press enter. Going down a bit you will see - "FORMAT controls the output."

After all, I am only interested in the day, as I get paid on the 27th , every month and therefore the month and the hours and seconds are of no use to me. If you don't understand what I am talking about, simply type: `date` and press enter.

According to the man page, "%d" is what I am looking for as it represents the day. Let's test that out, type: `date %d` and press enter. What happened? You did not pay attention to the Synopsis part; it says, "+FORMAT". Now try: `date +%d` and press enter. Great, now we have something that works. Last issue we learned about the \$ variable. Now it's time to use it. So... if some condition is true -do something. (The technical description of an if statement, hehe)

Clear this tab or second terminal and type: `if [$(date +%d) -eq 27] ; then echo "Quick, spend some before the wife sees it!!"; fi` and press enter

COMMAND & CONQUER

What happened? Well unless today was the 27th, you would get no output. Remember, if the condition is FALSE, nothing happens. (Else is not part of this discussion)

Let us dissect that query: IF [TRUE]; then blah...

The [TRUE] part is then broken into a question to determine true or false. If our variable "date + %d" (it is in brackets as it is not one word) -eq (equal to) 27. We have -eq as our comparison operator like -gt for 'greater than' and -lt for 'less than'. You learned these in school, so you should know them by now. If you are unsure, see:

<https://tldp.org/LDP/abs/html/>

[comparison-ops.html](#)

Also remember the spaces with the square brackets as they are commands. "[" and "[[" (we will get to these much further down the line when we get to test)

Before anyone comments, I am not a code monkey, this reads easier for me than all the unnecessary white space.

Save it and set the script executable. (You should know that much by now)

Just to test the validity, try changing the 27 to today's date and see if the echo fires.

Congratulations! You wrote a script! You can put this into cron and have it fire once a day. That way when it is 'payday', you will get a friendly reminder!

Cron is beyond the scope of this article, but we will visit it in the future.

Did I make a mistake?
misc@fulcirclemagazine.org



```
gewgaw@rooikopstiefkind: ~/Documents
GNU nano 4.8 gotmoney.sh
#!/bin/bash
if [ $(date +%d) -eq 27 ] ; then echo "It is PAYDAY!"; fi
```



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.



Happy New Year! Let's all hope that 2022 will bring great things to us all.

As most of you know, I work very closely with Don Rozenberg who created PAGE which is a GUI designer for Python. On 24 December, 2021, Don released version 7 of PAGE. While the latest version looks almost exactly like the previous version when you just look at the interface. The majority of the changes are under the hood. Let's take a quick look at the parts of the interface that changed...

The first thing you will probably notice is that there is a new option provided at the top of the toolbox.



The Toplevel widget gives you the ability to add a new toplevel form to an existing project. In previous versions of PAGE you had

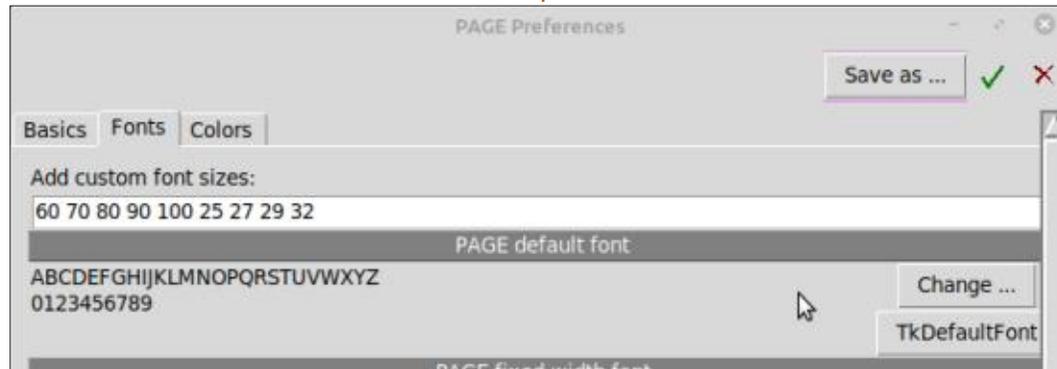
to create a multiform project by creating a project for each form, which meant that there were three files (one .tcl and two .py) for each form. Now, since all forms can be part of the same project, there are only three files. More about this in a moment.

The next change you should notice is in the Widget Tree form.



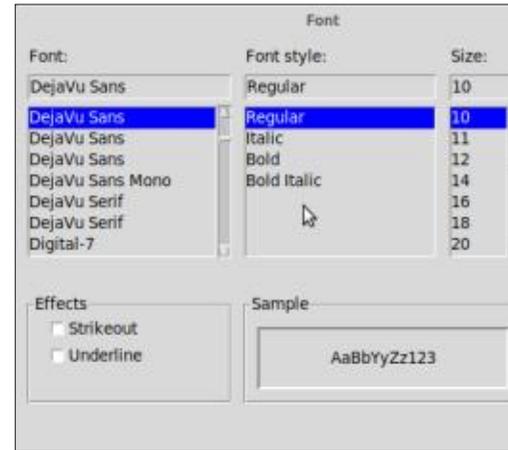
It is much more comprehensive and functional.

One of the other (out of many others) new things won't be seen



unless you open the preferences menu (shown bottom middle).

On the Fonts tab is a new entry box that allows you to add more font sizes to the redesigned Font dialog. While there is a way to pick a one-time font size for a special purpose, you can use the Add Custom Font Sizes feature to pick font sizes that suit your needs.



There are a wealth of changes to the User Interface that we won't touch on right now since we are going to discuss the biggest thing that has changed in PAGE 7.0. I'll save them for you to find and I might touch on some of them in future articles. In my mind, the biggest change is the code creation engine, which takes the native .tcl information from the GUI designer and turns it into Python code. It has changed significantly. Don has pretty much rewritten it from scratch.

When you created a GUI with PAGE Version 6, any Tk variables were exposed to your program through a special function called `set_Tk_var`. This would allow you to easily create dynamic labels or access and/or change various attributes of widgets like the TCombobox through the textvariable attribute (next page, top right).

This function was created by PAGE in the support module. As you can see from the code snippet

HOWTO - PYTHON

above, if you wanted to change the text of the label, you could simply use

```
LabelDisplay.set()
```

to change the value to whatever you needed to display.

In the support module, PAGE 6 had a startup function called `init` which was the very last thing that was run before the user would see the form (bottom left).

Every form had this same function. To reference a widget directly we had to use `w.widgetname` to access that widget. This didn't apply to the Tk variables, just the widgets.

Since PAGE 7 supports the ability to have multiple forms in one code module, changes to the paradigm had to change, since the possibility of variable conflicts became a very real concern. Assume that you have two forms both having a Label widget that has a textvariable called

"LabelDisplay". If this were to happen, you couldn't be sure which Label widget would have its text changed. In order to keep this from happening, Don came up with a very creative way to keep this from happening. Here is a simple two form program that shows how easily it is taken care of (shown bottom right).

Both of the Label widgets have the same TK variable name `DynLabel` which, while not very good programming practice, is easy to do without thinking and under PAGE version 6 and below was completely reasonable. PAGE has had a feature called "borrow" for a long time and in version 7 has been enhanced to allow you to "borrow" forms and widgets from an earlier project that had multiple forms easily. All widgets and their variables can be imported into the new project. In order to keep this from being a problem, PAGE separates each form into its own class within the GUI module. Shown right is a quick snippet of what the GUI contains in version 7.

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

```
def set_Tk_var():
    global LabelDisplay
    LabelDisplay = tk.StringVar()
    LabelDisplay.set('Label')
```

```
class Toplevel1:
    def __init__(self, top=None):
        ...

        top.geometry("360x225+2875+173")
        top.minsize(1, 1)
        top.maxsize(4225, 1410)
        top.resizable(1, 1)
        top.title("Toplevel 0")
        # The Tk variables are defined here...
        self.top = top
        self.DynLabel = tk.StringVar()

        self.Button1 = tk.Button(self.top)
        ...

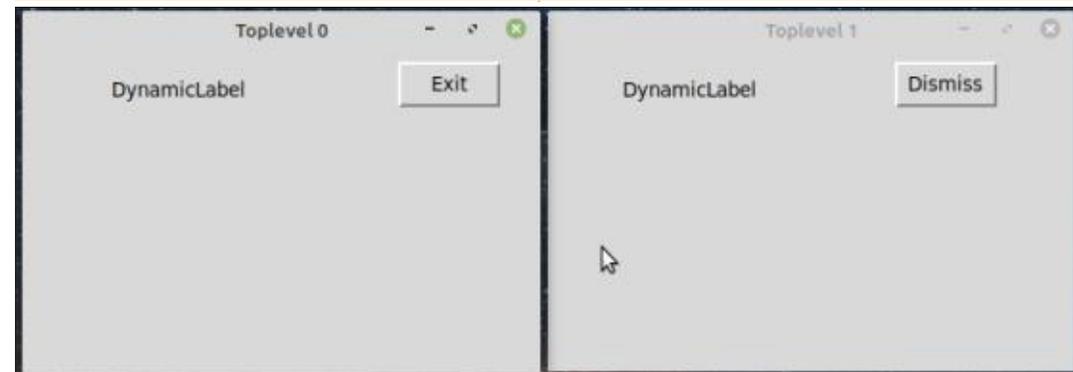
class Toplevel2:
    def __init__(self, top=None):
        ...

        top.geometry("366x224+3240+173")
        top.minsize(1, 1)
        top.maxsize(4225, 1410)
        top.resizable(1, 1)
        top.title("Toplevel 1")

        self.top = top

        self.DynLabel = tk.StringVar()

        self.Button1 = tk.Button(self.top)
        ...
```



Since each form is now defined as a separate class, the variable names can easily be the same between forms. I'll show you (top right) how they are accessed in a moment. For now let's take a look at the support module changes. There is no longer an `init` function, it is now called `main`.

The `V7` in the snippet above is the actual project name. Each form now has an alias that starts with `_w` and then a number to reference the form, so `_w1` is the alias for the widgets on form1 and `_w2` refers to those on form2.

To change the text for the label on the first form, we would simply use the following code.

```
_w1.DynLabel.set("Some Text")
```

And to change the label text on the second form we would code it like this.

```
_w2.DynLabel.set("Other Text")
```

In order to reference the form itself in version 6, we used `root`, which allowed us to do things like move the form on the screen or to terminate the program. In version

7, you can see in the above code, `root` has been aliased to `_top1` for the first form and `_top2` for the second. It's important to remember to use this new alias, since just using `root` will again cause some very unexpected (and often unwanted) issues.

To end the program in version 6, you would use the PAGE generated function `destroy_window()`. In PAGE 7, you simply call `root.destroy()`. Calling `root` from your code can be problematic, since all forms have a `root`. However, calling `root.destroy()` to terminate the program is fine, since it will close all active windows.

When you run a multiform project created with version 7, you might be surprised to see that all the forms show right away. If you look at the code again, you can see why. All of the forms get created at startup and when the `root.mainloop` line gets run, it shows everything that has been defined. Sometimes, that is the desired action and that's ok. However, if you have an application that you want one form to be the main form and the others to only show on demand, that isn't acceptable. It's really easy to

```
def main(*args):
    '''Main entry point for the application.'''
    global root
    root = tk.Tk()
    root.protocol('WM_DELETE_WINDOW', root.destroy)
    # Creates a toplevel widget.
    global _top1, _w1
    _top1 = root
    _w1 = V7.Toplevel1(_top1)
    # Creates a toplevel widget.
    global _top2, _w2
    _top2 = tk.Toplevel(root)
    _w2 = V7.Toplevel2(_top2)
    root.mainloop()
```

address this though. In version 6, I created two small functions in the support module that I called `show_me` and `hide_me`. When my main form needed to show another form, I would call the `hide_me` function to minimize the main form and then I would call the `show_me` function to show the second form. When the use of the second form was done, I would call the `hide_me` function for the second form and then the `show_me` function for the main form. It was a logical set of steps for the purpose. Here is what they looked like for version 6.

```
def show_me():
    global root
    root.deiconify()
```

```
def hide_me():
    global root
    root.withdraw()
```

We can't use these functions directly in version 7, since they both

call the `root` object directly. However, it's simple to modify them to work properly. Simply create a set for each form.

```
def show_me1():
    global _top1
    _top1.deiconify()
```

```
def hide_me1():
    global _top1
    _top1.withdraw()
```

```
def show_me2():
    global _w2
    _w2.deiconify()
```

```
def hide_me2():
    global _w2
    _w2.withdraw()
```

You could even make a single set of functions that handles any number of forms by passing the object as a parameter to the function that defines the form you want to work with.

```
def show_form(which):
    which.deiconify()
```

```
def hide_form(which):  
    which.withdraw()
```

To use this, we need to create a startup function that runs just before the forms are shown and we need to call it just before the last line of the main function.

You can use the startup function (or whatever you would like to call it) to do things like startup databases or initialize various settings and libraries. Our startup function, in this case, needs to have this at a minimum.

```
def startup():  
    hide_form(_top2)
```

This will minimize or hide the second form, leaving the first (main) form visible. In order to show the second form and hide the first, we need a button on the main form to trigger the functions. We'll call the button `btnShow2` and define its callback function as

```
def on_btnDismiss(*args):  
    print('V7_support.on_btnDismiss')  
    for arg in args:  
        print('another arg:', arg)  
    sys.stdout.flush()  
    hide_form(_top2)  
    show_form(_top1)
```

`on_btnShow2`. In its callback, we would code it like this (bottom right).

(The lines that are not bold are part of the new callback function skeleton that PAGE creates.) This hides the first form and shows the second. Then for the callback to restore the first form and hide the second form, we do the following (bottom left).

Super easy and clean. All you have to remember is which form alias you need to use to manipulate that particular form.

I have to admit, it took a little while before I became completely comfortable with the changes, but they aren't as drastic as it might seem.

There are a number of other changes in PAGE version 7 that really enhance the capabilities, but I won't go through all of them right now. If you want to have a quick

```
def main(*args):  
    '''Main entry point for the application.'''  
    global root  
    root = tk.Tk()  
    root.protocol('WM_DELETE_WINDOW', root.destroy)  
    # Creates a toplevel widget.  
    global _top1, _w1  
    _top1 = root  
    _w1 = V7.Toplevel1(_top1)  
    # Creates a toplevel widget.  
    global _top2, _w2  
    _top2 = tk.Toplevel(root)  
    _w2 = V7.Toplevel2(_top2)  
    startup()  
    root.mainloop()
```

```
def on_btnShow2(*args):  
    print('V7_support.on_btnShow2')  
    for arg in args:  
        print('another arg:', arg)  
    sys.stdout.flush()  
    hide_form(_top1)  
    show_form(_top2)
```

and (somewhat) easy way to create Tkinter GUIs for your Python programs, you really need to get the latest version of PAGE. You can download it from <https://sourceforge.net/projects/page/>

Due to all the changes that PAGE 7 brings to the table, I wrote a new tutorial for new users and a document that shows an easy set of migration tips for those who are familiar with previous versions of PAGE. Both are included in every distribution of PAGE in the docs folder.

Until next time, as always; stay safe, healthy, positive and creative!



Greg Walters is a retired programmer living in Central Texas, USA. He has been a programmer since 1972 and in his spare time, he is an author, amateur photographer, luthier, fair musician and a pretty darn good cook. He still is the owner of RainyDaySolutions a consulting company and he spends most of his time writing articles for FCM and tutorials. His website is www.thedesignatedgeek.xyz.



HOW-TO

Written by Erik



Blender Pt.2

Since I misplaced parts 2 and 3, I was recreating them, when blender made the jump to 3.0. So if anything does not look 100% the same, just go with the flow. The first few articles are going to try and draw your attention to different parts of the interface so you can orientate yourself. By no means is this series a deep dive, I do not work with Blender every day, I got it to make isometric landscapes and threw up my hands in frustration more than once, so I know how it feels. Picking up a book, I found that they used words for things I had no idea of. On top of that, English is not my first language and I know there are many more people like me. Here, I attempt to get you up to speed, the fastest way I know how. Later on we will make a penguin and maybe a low poly landscape.

When you fire up blender now, you have the choice of keeping the old settings or loading the new settings for blender 3.0. The choice is yours. I suggest going for the newer settings as things only move forward.

Last issue we looked at the mouse buttons, ctrl, alt, and we identified the lights, camera, action part. If you missed this, quickly start there.

Our focus today is to get around with your mouse and keyboard. If you have a trackball, bully for you! If you have a mouse and keyboard, pay attention.



Like our lights, camera, action bit, we will be looking at our red, green and blue bit. If you are colour blind, it is the X,Y,Z component, representing the three axes. As soon as you click on any of these points, the demo cube in the centre view screen immediately rights itself to that axis. With a cube, this is not easy to see, as all sides are the same. Trust me on this for now. Here is where mouse users need to follow along please. Clicking and dragging inside the invisible sphere that these red, green and blue points occupy, will rotate your cube. Try it a few times, you may not get it the first time. If you don't get it at all, right click and hold X,Y,Z and drag it in any direction. The nice thing about Blender is, there are many ways to do things. If you have a numpad, press the numpad key, activate and press the numbers 1-9 to see what each one does. Move your mouse cursor into the main viewport and holding down the middle mouse button, move the mouse for the same effect. You will notice that your view moves around "action". That point in the middle of the

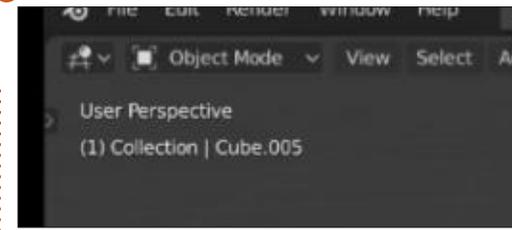
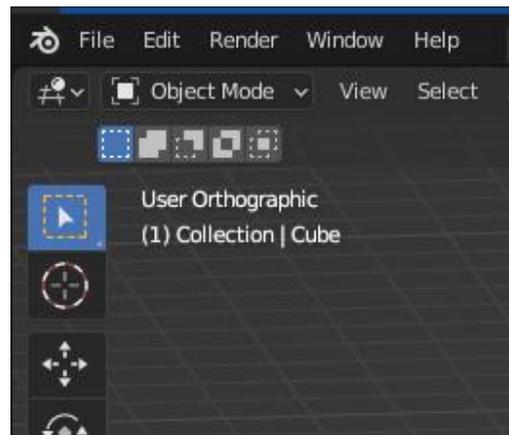


HOWTO - BLENDER

viewport. What if you need to move the “action”, but not the generic grey cube in the center? You apply the same principle as you did earlier, you click on the hand icon, below the red, green and blue, and while holding your primary mouse button down, move the mouse. You can also hold down the shift key and middle mouse button drag. We talked about these modifiers in the first issue. Do that again, but without holding down the shift key. The mouse zoom may seem intuitive, but just like the previous examples, you can use the magnifying icon above the hand to get you close or further away. Left click and hold the magnifying glass and move your mouse forward and back. Again we have another option, the plus and minus keys on your numpad. You may have noticed that I skipped 0 on the numpad. That is because it gives you a view from your camera. (See what your camera sees) So left clicking the camera icon, will give you the same view. The last one is an oddball, that grid icon. No, it does not bring up your brain, when you click it, it actually switches views. Look at the top left of the screen when you click on it. You will notice that the word orthographic changes to perspective. And that is

actually what it is, it gives you perspective on your objects, currently the nondescript grey cube. This only becomes apparent when you have more than one object. You will also notice that the icons along the right side disappeared. We will talk about different modes in an upcoming issue.

Sometimes your try at home may not gel with what I am telling you here. That is because blender acts on what is selected, or not selected. Remember we said that selected objects go orange? (Default) Now a lot of times, things will work the same, but not always, please remember that. If something behaves differently than expected, click the object, or if the object is selected, click the background and try your operation again.



Above red, green and blue you will see four spheres. This is the level of viewport shading. The first one will give you outlines, or wireframe. This becomes important with complex shapes. Go ahead and click each one to see what they do to your cube. If you have a lower end CPU, there may be a slight delay in clicking the sphere and anything actually happening. I tried this on a Celeron laptop with 2Gb memory and it actually took a few seconds before it reacted, so if you have something in that line, take a deep breath. Next to the four spheres, is a drop down arrow, the contents here change in relation to the sphere selected. We won't cover that now, but please be aware of it.

Well done! You now know more than most people about blender. Please don't stop playing with it just because the next issue has not arrived yet. There are many great tutorials on the web, but I find that they tend to start explaining and

then continue like you know everything they are talking about. We are taking it slow and steady, and if anything was unclear, please drop us a message at: 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.



HOW-TO

Written by ????????

RUNNING OUT OF HOW-TO ARTICLES.

Like I keep saying: we need your help to keep going. I'm almost out of How-To articles. The only thing keeping the How-To section going is Mark, Greg and Erik.



HOW-TO

Written by Mark Crutch

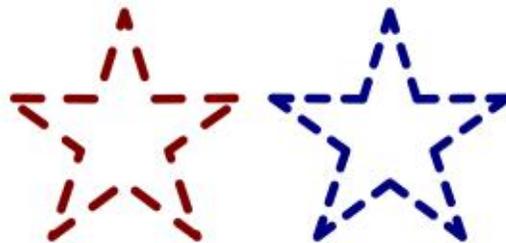
Inkscape - Part 117

Over the next few instalments I'll be looking at the new Live Path Effects (LPEs) that have been added with Inkscape 1.0 and 1.1. Earlier articles in this series provide a general introduction to LPEs and what they are (part 42), as well as a deeper dive into the LPEs that arrived with earlier Inkscape releases (parts 42 to 47, 65 to 69). The previous instalment detailed the more general changes to the LPE dialog that took place with version 1.0, but this month I'm going to look at the first of the new LPEs, plus an old one that has had something of a rebirth.

DASHED STROKE

At first glance there doesn't appear to be an obvious need for an LPE that renders the stroke of an object as dashes. After all, the Fill & Stroke dialog already offers various dash patterns which form part of the native SVG format (remember, LPEs are an Inkscape-specific extension). But although the standard SVG dashes are often sufficient, they do lack some nuance in the way they're

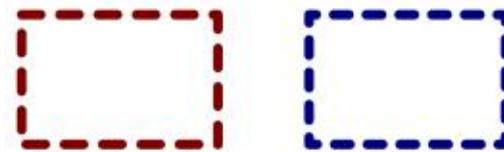
distributed along a path which can give a less than aesthetically pleasing result. This is the niche that this new LPE aims to address. As a quick example, look at these stroked stars, the red one on the left using standard SVG dashes, and the blue one on the right using the Dashed Stroke LPE.



Pay close attention to the way that the corners – both concave and convex – are rendered. The SVG version is mirror-symmetric along the vertical axis, but only because I adjusted the dash offset to give that effect. Without that manual intervention there was no symmetry to the dashes at all. Even with that change, however, the corners differ as you move around the star: for the tips of the shape we would probably like them all to look like the point at the top, and not like the remaining four. None of

the inner corners are really what we would like in most cases. Notice that the LPE version achieves exactly the right look, with the corners all appearing pleasingly similar and symmetric.

Let's look at another example. Dashed outlines are commonly used around simple rectangular boxes in flow charts and other diagrams. Which do you think looks better: the red SVG version (left) or the blue LPE version (right)? Once again, focus on the corners.



The reason for this difference is that the SVG stroking spec uses a very simple algorithm to determine how to draw the lines and spaces. It simply starts at the beginning of the path and draws a series of alternating dashes and spaces, based on the pattern described in the stroke-dasharray attribute or CSS property. It doesn't care about curves or corners, it just plods

along from start to finish rendering a repeated series of strokes and dashes, regardless of the underlying shape. You can control the position of the first dash, using the stroke-dashoffset attribute or property (which is exposed via the Fill & Stroke dialog in Inkscape), but all that does is shift the entire pattern along. It doesn't affect the length of each stroke or space, so you are still likely to end up with unbalanced dashes as they wrap around the corners of your shape.

The LPE, on the other hand, works a little differently. The biggest change is that it can work on each segment of a path individually, rather than treating the entire path as a single stretch to be dashed as one. This is the secret of those better looking corners – drawing half a dash at each end of a segment results in dashes that are pleasingly symmetric as the path turns a corner. Let's look at the options available in this LPE, and the settings I used for the blue rectangle.

Dashed Stroke

Number of dashes - +

Hole factor - +

- Use segments
- Half start/end
- Equalize dashes

Note

Add "Fill Between Many LPE" to add fill.

I'm going to describe these parameters out of order, as this is the best way to explain what each option does. I'll start with the "Use segments" parameter: when unchecked this results in the other parameters applying to the entire path (much like the native SVG dashes). When checked, each segment of a path is treated separately. In most cases you'll probably want this checked.

Going back to the top, the "Number of dashes" parameter defines the number of dashes that will be rendered along the length of the whole path, or along each individual segment. But the actual count will also depend on the "Equalize dashes" option, as we'll see shortly. This parameter is at the heart of the fundamental difference with the LPE dashes, though: SVG dashes don't have a count or limit, they'll simply keep rendering as long as there is any path left to fill; the LPE dashes, on

the other hand, aim to fit a specific number of dashes into each path or segment, subdividing the available length according to this parameter and then distributing the dashes and spaces evenly within.

The relative lengths of the dashes and spaces can be adjusted using the "Hole factor". Leave it at zero to have them the same size, increase it (up to +0.99999) to increase the size of the dashes and reduce the spaces, or decrease it (down to -0.99999) to adjust the balance in the opposite direction. Reducing it to its lowest value results in each dash appearing as nothing more than a pair of line caps, as set in the Fill & Stroke dialog: a circle (for round caps) or a square (for square caps). Watch out if you use the "Butt cap" style, however, as that effectively causes the dashes to disappear completely at the lowest hole factor. Note, however, that using a single ratio like this means that the LPE can't produce the sort of dash and dot combinations that the stroke-dasharray allows for in normal SVG strokes.

The "Half start/end" parameter determines whether to only draw half a dash as the start and end

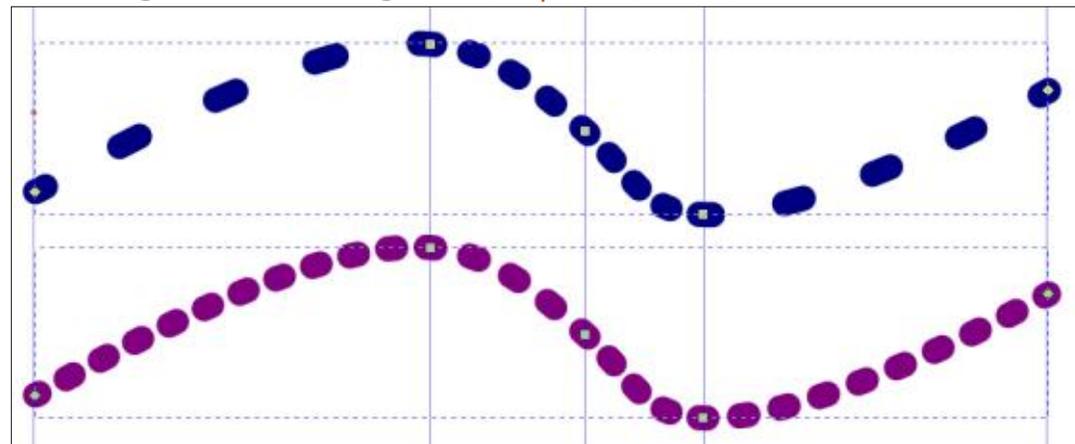
shapes (checked), or to draw a full dash at the start and end if possible (unchecked). Usually this is best left checked in order to gain the aesthetic benefits of symmetry and even spacing. Each half dash still contributes to the "Number of dashes" count, so a count of 5 with this parameter enabled actually means 3 whole dashes and two half dashes, rather than the 4 whole dashes (plus two halves) that you might expect if you were adding the parts up numerically.

Finally, the "Equalize dashes" parameter has the potential to upend the "Number of dashes" count entirely. When this is checked, the algorithm first creates the desired number of dashes for the shortest segment in the path. The length of each dash in that segment is then used when rendering all the other segments,

fitting more than the actual count in if there is space. A demonstration might make this a little clearer.

In the image below, the two paths are the same, but the top one has "Equalize dashes" unchecked, whereas the lower one has it checked. I've positioned some vertical guides to make it clearer where the nodes of the path are – i.e. where each segment starts and ends. The top path honours the "Number of dashes" count completely: each segment has 5 dashes (3 whole, 2 half). But this leads to different spacing between the dashes across the segments, and even differently sized dashes in the middle two segments where they've had to be squeezed into a smaller space.

The lower path, on the other hand, clearly shows that all the



HOWTO - INKSCAPE

dashes and spaces are even across all the segments. But it does this at the expense of the “Number of dashes” value. That parameter is used when calculating the smallest segment (the third one), but then the resultant dash and space size is simply used for all the other segments, regardless of the count. As you can see, the end result looks better, and is probably what you are likely to want, but the first and last segments have way more than 5 dashes each.

There’s one additional part of the UI in the screenshot from the LPE dialog: not another parameter, but a note in a box, which says ‘Add “Fill Between Many LPE” to add fill’. What on earth does that mean, and why is it necessary?

FILL BETWEEN MANY

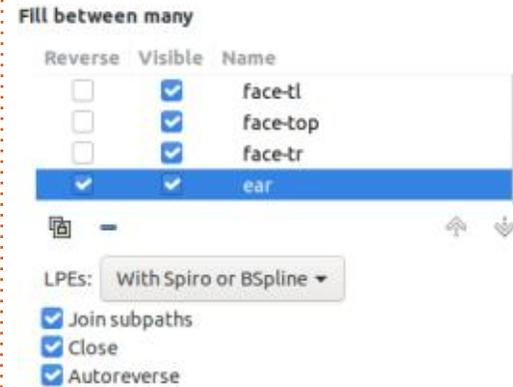
Remember that the output from an LPE is just an SVG path, so all the clever things that LPEs can do must ultimately be rendered using normal SVG capabilities. As we’ve already seen, raw SVG can’t produce the sort of dashes that we’re getting from the Dashed Stroke LPE, so what actually are we seeing in our rendered output?

The result is actually a new complex path, made up of lots of individual subpaths, one for each visible dash. Trying to add a fill to this will actually just fill the subpaths, not the whole shape. Because most of the subpaths only have two nodes, even that fill isn’t generally visible. The exception is the corners, where three nodes are used in a triangular configuration. Sure enough, adding a fill to a Dashed Stroke path does result in a web of colour in the corners, but not the filled shape we’re looking for. As an example here’s our star from earlier, but with the stroke width reduced for clarity and an orange fill applied.



This has long been an issue for many LPEs, not just the Dashed Stroke, so the Inkscape developers addressed it head-on a long time

ago, by adding the “Fill Between Many” LPE back in version 0.92. I covered this LPE in some detail back in part 67 (FCM issue #127), though the UI has expanded a little since then. In older versions you only had the ability to add paths to the LPE, flagging some of them as needing to be reversed. The new UI, when used with the same “Frankie” image I used in part 67, looks like this.



The basic functionality remains the same: you have to create a sacrificial path on which to apply this LPE, then add each of your source paths by copying each one to the clipboard and adding it to the list in the LPE, as described in the earlier article. It can be a time-consuming and difficult process when dealing with lots of paths, though it’s not too bad for adding a fill to a shape with the Dashed Stroke LPE as there’s only one path

to add in that case. These are the steps needed to add a fill to our rectangle, for example:

- Draw a sacrificial path (usually just a simple two-node line)
- Add the Fill Between Many LPE to the sacrificial path
- Select the path which has the Dashed Stroke LPE applied (the rectangle) and copy it to the clipboard
- Re-select the sacrificial path in order to bring up the UI for the Fill Between Many LPE
- Click the “Link to path in clipboard” button to add the Dashed Stroke path to the list
- Adjust the fill and stroke values to suit your needs

With luck you’ll now find that your rectangle has a fill, but things don’t always go so smoothly. In my own experiments, trying those steps with a star rather than a rectangle results in either no fill, or an oversized fill object that is wrongly positioned and can’t be moved. There are definitely some bugs in this LPE that have yet to be ironed out.

Compared with v0.92, the newer version of this LPE also provides some additional parameters to tweak. There is a “Visible” checkbox

HOWTO - INKSCAPE

for each path, allowing you to temporarily remove it from the filled shape, perhaps to test whether or not it is contributing anything useful prior to completely removing it from the list. The “Join subpaths” checkbox lets you fill each subpath individually (unchecked), or use the older behaviour of joining the subpaths to create a single shape to fill (checked). The latter is almost always going to be what you want. Another checkbox (“Close”) now lets you leave the new path unclosed between the first and last paths in the list – probably more useful if you are using this LPE to add an extra stroke rather than a fill and, again, usually something you would want to leave checked. Finally the “Autoreverse” option overrides the individual “Reverse” checkboxes for each path: with this checked the algorithm will try to join paths based on the proximity of their endpoints, rather than strictly following the direction of each path. Usually this does a good job, and is best left checked, but you do have the option to turn this off and manage path reversal on a per-entry basis as before, should you wish to.

The pop-up menu is also a new

addition, choosing how the source paths should be interpreted. Usually leaving this as “With Spiro or BSpline” is a good option: this will essentially use the shape you originally drew, whether it was created using simple SVG paths, or you used the Spiro or BSpline options that Inkscape exposes in some drawing tools. In practice these are implemented as LPEs, so this option tells Inkscape to use the output from those LPEs as the source, if they exist, or to use the plain path data otherwise. Alternatively you can select “Without LPEs” to only use the original path data, regardless of any LPEs applied. Conversely the “With all LPEs” option will use the path data that comes out of whatever series of LPEs has been applied to the shape. Be aware that this can quickly lead to very complex shapes if you’re not careful, so isn’t often the choice you want.

Looking back at the number of steps needed to add a fill to a path with the Dashed Stroke LPE applied, you may feel it’s not worth the extra effort and confusion, preferring to stick to SVG dashes or to draw the fill as a separate object. The “Fill Between Many” LPE can certainly be a tricky feature to get

your head around, and in other use cases where you need to add multiple paths to the dialog it can be a time consuming pain. Luckily the Inkscape developers have realised that this complexity gets in the way of an otherwise useful feature, so with version 1.1 they’ve added a new menu entry, Path > Fill between paths, which will silently create a sacrificial path and add the “Fill Between Many” LPE to it, already populated with any paths from your drawing that were selected at the time. This makes it trivial to use this LPE in most cases: just select the path or paths that need to be filled and select the menu option. You can then select the newly added fill in order to access the LPE parameters if you need to (e.g. to reverse specific paths).

Note that the sacrificial path added by Inkscape is of zero length: its “inkscape:original-d” attribute just consists of an “M 0,0” command, which doesn’t actually draw anything. As such, be careful not to either hide the LPEs visibility, or that of all its listed paths, otherwise you won’t be able to select it on the canvas. In that case you’ll have to find it in the XML editor (look for a path with that “M

0,0” value) in order to select it for further editing or deletion.

This new menu entry is a great addition for working with LPEs, as it helps to get around one of the most fundamental problems most users will come across as they begin to use them. For this reason alone it may be worth upgrading to version 1.1.x if you haven’t already done so. It’s a shame, however, that the “Fill Between Many” LPE, even when added using this menu entry, can still be rather buggy, even for simple examples. Hopefully future releases will make it more robust, which will help to make LPEs in general a far more useful tool than they already are.



Mark uses Inkscape to create comics for the web (www.peppertop.com/) as well as for print. You can follow him on Twitter for more comic and Inkscape content: [@PeppertopComics](https://twitter.com/PeppertopComics)

THE DAILY WADDLE

LOOKING BACK, WHAT WAS THE
MAIN SYMPTOM OF COVID-19?

PARANOIA





BACK NEXT MONTH



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. Feedback welcome at: acer11kubuntu@gmail.com



Well, it's been almost a year since I've started this column. I hope that this has been helpful to those of you who are interested in microcontrollers.

Once again, we will deal with the SSD1306 display, but this time we will be talking about creating custom fonts. This time, we are going to (as the subtitle suggests) create special fonts for use on our OLED display. For the most part, the code should work fine on the Raspberry Pi Pico, ESP32 and ESP8266. The only concern might be the memory usage due to the size of the font files.

Back in issue Full Circle Magazine issue #172 Part 6, I introduced the Micropython Font To Py library from Peter Hinch, but only a small part of that package. We used a couple of his fonts and his writer library. If you didn't get it then, here is where you can get it. <https://github.com/peterhinch/micropython-font-to-py>

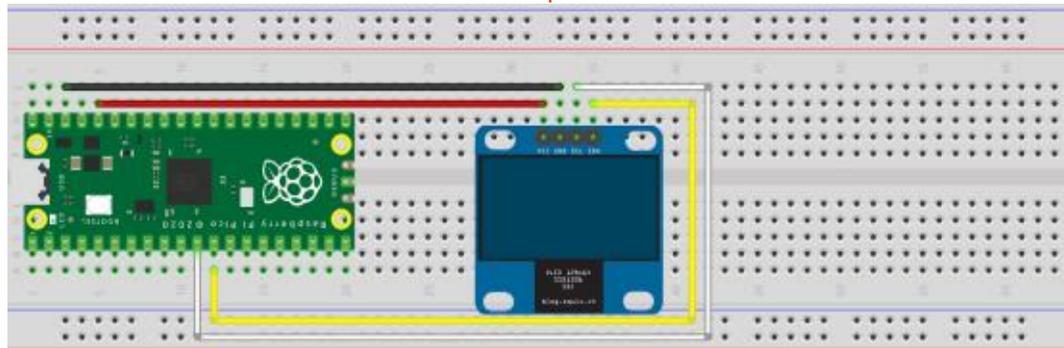
This time, we'll be creating our own fonts for the OLED display

using our normal desktop computers and regular Python. Once we have the fonts created and tested, we will transfer them to the microcontroller and use a simple test program there.

For this project, we'll use the Raspberry Pi Pico microcontroller and a SSD1306 128x64 OLED Display. Below is a picture showing the wiring just in case you need it.

And here is the wiring grid.

	Pico	Display
3V3	5	Vcc
Ground	3	Gnd
SCL	33 (GP8)	SCL
SDA	32 (GP9)	SDA



fritzing

Next, you will want to download and unpack the micropython-font-to-py package into a folder and then create a folder somewhere to serve as your work folder. You will want to copy (at a minimum) the font_to_py.py, font_test.py, writer.py and the ssd1306_setup.py files into your working folder. If you followed along back in part 6, you probably have the writer.py file and the ssd1306_setup.py file.

You need to modify the ssd1306_setup.py file to make it support our Pico I2C setup. There are only a couple of places that you need to change, though. Near the top of the file, after the comments, you will find the following lines of code.

```
import machine
from ssd1306 import
SSD1306_SPI, SSD1306_I2C
WIDTH = const(128)
HEIGHT = const(64)
```

You need to change the WIDTH and HEIGHT to match your display. While you CAN use a smaller OLED display (128x32 pixels), you will have to modify the code later on since it assumes you have 64 rows of pixels to make the text fit.

Now at the bottom of the file, you will see the following lines of code. Make sure you change the i2c = machine.I2C(2) to i2c = machine.I2C(0)

```
else:
    # i2c = machine.I2C(2)
    i2c = machine.I2C(0)
    ssd = SSD1306_I2C(WIDTH,
HEIGHT, i2c)
return ssd
```

Save your file. Copy the Writer.py and ssd1306_setup.py file

to the Pico, if they are not already there.

Now we can create our special font files. The idea is to convert a font from your system to a file that Micropython can use. In this article we'll use the Ubuntu-Regular ttf file. You should have it on your normal desktop/laptop system. If not, you can find it on the Internet.

In order to use the conversion utility, you need to have the freetype library installed on your desktop or laptop (not on your Microcontroller). You can install it easily by using

```
pip install freetype-py
```

Or

```
pip3 install freetype-py
```

Now that we have all our dependencies taken care of, we can start to work with the font_to_py program. You need to be working in a terminal session for the next steps.

The command line should be

```
font_to_py.py {full path to font file} {size} {-x} {file name to save to}
```

You MUST include a fully qualified path to the font you want to use. On my system there are two font locations. The first in /usr/share/fonts/truetype (for the truetype fonts) and then each font has its own folder. The second is a hidden folder in my home directory and is named .fonts. I keep speciality fonts in this one. Either path is seen by my system as a font folder. You can use pretty much any .ttf font, but I don't really see any reason to fool with fonts too fancy or too big or small, so that's why I chose Ubuntu-Regular.

You also need to be sure to include the -x parameter when you create the font file. This makes sure that the font is horizontally mapped. This information is not included in the readme directly and it took me the better part of an hour to break everything down to understand why my fonts were not working. Here is the terminal session that I did when I created my font. We will be creating a 12 point font, so I saved it as ubuntu12.py.

```
$ python font_to_py.py /usr/share/fonts/truetype/ubuntu/Ubuntu-R.ttf 12 -x ubuntu12.py
```

Writing Python font file.

Height set in 1 passes.
Actual height 12 pixels.

Max character width 11 pixels.

ubuntu12.py written successfully.

Now we can see that the characters are 12 pixels in height and the width is 11 pixels. This will help decide how many characters you can fit on a line and how many lines you can fit on the 128x64 display. Now that we have the font created, let's do a quick test. Peter has provided a utility to do just that. It's called font_test.py. You provide the font name as the first parameter and what you want to view as the second. Here is the terminal output that I got.

```
$ python font_test.py ubuntu12 Hello
```

Horizontal map

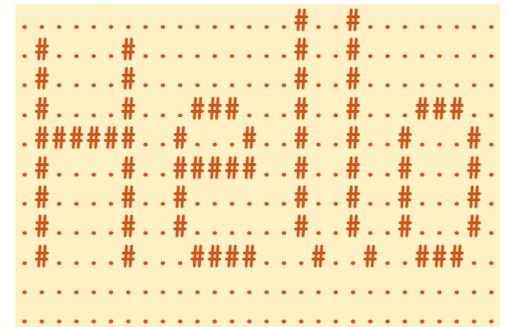
Normal bit order

Proportional spacing

```
Dimensions height*max_width  
12 * 11
```

```
Start char " " (ord 32) end  
char "~" (ord 126)
```

```
.....
```



Now we need to do the same thing but this time make the font a 16 point font. We'll keep the font as Ubuntu-Regular.

```
$ python font_to_py.py /usr/share/fonts/truetype/ubuntu/Ubuntu-R.ttf 16 -x ubuntu16.py
```

Writing Python font file.

Height set in 1 passes.
Actual height 16 pixels.

Max character width 15 pixels.

ubuntu16.py written successfully.

Now let's test the 16 point font file.

```
$ python font_test.py ubuntu16 HelloHorizontal map
```

Normal bit order

Proportional spacing

```
Dimensions height*max_width  
16 * 15
```


We finish up with making a call to `ssd.show()` in order to actually display the data we just sent and then we'll sleep for 3 seconds and after the 3 second delay, we'll clear the display with `ssd.fill(0)` which turns off all the pixels.

Next, we'll display the same data in the 16 point font. Because the font is 16 point, we can't fit all of the text in one screen like we did with the 12 point font. We'll break it into two parts, the first two lines of alphabet characters, the numbers 1 to 0 and then some shifted characters. Notice we also have to shift some of the characters from one line to another to get it all to fit. We'll push the data to the display to show it, sleep for another 3 seconds and then clear the display screen again.

```
wri = Writer(ssd, ubuntu16)
Writer.set_textpos(ssd,0,0)
wri.printstring('ABCDEFGHIJKLM')
Writer.set_textpos(ssd,16,0)
wri.printstring('NOPQRSTUVWXYZ')
Writer.set_textpos(ssd,32,0)
```

```
wri.printstring('Z -
1234567890')
Writer.set_textpos(ssd,48,0)
wri.printstring('!@#$
%^&*()_+==')
ssd.show()
time.sleep(3)
ssd.fill(0)
```

Finally we'll send out the lower case characters in the 16 point font as we did above. We sleep again for 3 seconds, then clear the display to avoid burnin.

```
Writer.set_textpos(ssd,0,0)
wri.printstring('abcdefghijkl
m')
Writer.set_textpos(ssd,18,0)
wri.printstring('nopqrstuvwxyz
z')
ssd.show()
time.sleep(3)
ssd.fill(0)
ssd.show()
```

Finally, we call the function from code to get it all to start.

```
test()
```

That's it. Save the program as `ubuntutest.py`.

Now when we run the program you can see the result of the hard work.

I've put the font files and the



code files on my repository at https://github.com/gregwa1953/FCM-177_MicroThisMicroThat

At this point, you should be able to pick a font of your choice and convert it to run on your microcontroller to support your own programs.

Until next time, as always; stay safe, healthy, positive and creative!



Greg Walters is a retired programmer living in Central Texas, USA. He has been a programmer since 1972 and in his spare time, he is an author, amateur photographer, luthier, fair musician and a pretty darn good cook. He still is the owner of RainyDaySolutions a consulting company and he spends most of his time writing articles for FCM and tutorials. His website is www.thedesignedgeek.xyz.



Linux on Your iPad

For as low as \$4.95, you can have your own personal Linux cloud computer in minutes on any device.





Ubuntu Touch is the privacy and freedom-respecting mobile operating system by UBports.

Today we are happy to announce the release of Ubuntu Touch OTA-21, the very latest update to the system! OTA-21 will become available for the following supported Ubuntu Touch devices over the next week:

- BQ E4.5 Ubuntu Edition
- BQ E5 HD Ubuntu Edition
- BQ M10 (F)HD Ubuntu Edition
- BQ U Plus
- Cosmo Communicator
- F(x)tec Pro1
- Fairphone 2
- Fairphone 3
- Google Pixel 2XL
- Google Pixel 3a
- Huawei Nexus 6P
- LG Nexus 4
- LG Nexus 5
- Meizu MX4 Ubuntu Edition
- Meizu Pro 5 Ubuntu Edition
- Nexus 7 2013 (Wi-Fi and LTE models)
- OnePlus 2
- OnePlus 3 and 3T
- Oneplus 5 and 5T
- OnePlus 6 and 6T

- OnePlus One
- Samsung Galaxy Note 4 (910F, 910P, 910T)
- Samsung Galaxy S3 Neo+ (GT-I9301I)
- Sony Xperia X
- Sony Xperia X Compact
- Sony Xperia X Performance
- Sony Xperia XZ
- Sony Xperia Z4 Tablet (LTE or Wi-fi only)
- Vollaphone
- Vollaphone X
- Xiaomi Mi A2
- Xiaomi Mi A3
- Xiaomi Mi MIX 3
- Xiaomi Poco F1
- Xiaomi Redmi 3s/3x/3sp (land)
- Xiaomi Redmi 4X
- Xiaomi Redmi 7
- Xiaomi Redmi Note 7 and 7 Pro
- Xiaomi Redmi 9 and 9 Prime
- Xiaomi Redmi Note 9, 9 Pro, 9 Pro Max and 9S

The Pine64 PinePhone and PineTab are updated independently of the rest of these devices. The stable channel for the PinePhone and PineTab will not receive an update labelled "OTA-21".

WHAT'S NEW?

This release of Ubuntu Touch is still based on Ubuntu 16.04.

The storage statistics (free / occupied space) in the system settings got an overhaul, more categories are being shown, and the calculation of space occupied by the system partition layout and Ubuntu Touch core files has been made more precise.

Greeter redesign: Most notably the style of the Greeter (the thing that is shown when the screen is about to be unlocked) has improved significantly to give a more modern and slick appearance. It shows a different style depending on if PIN or password unlock has been selected. Try it out!

Preparations for Helium 10 ongoing: We made some fixes that will allow automatic installation and updating of Helium 10 devices, to not further block that road. Note that actual devices with Helium 10 will still take some time

Tamil language font was added

Helium 9 devices got magnetometer and compass plugins enabled in the sensor framework middle-ware. Unfortunately compass readings might be terribly off due to necessary pre-filtering of the sensor values, but it's a step towards having a compass available.

Media-Hub refactoring: After the rewrite of the media-hub service, the media-hub client library has also been rewritten using Qt classes. This allows to avoid depending on a few unmaintained libraries (though for the time being these libraries are still used in other projects, so they still end up being in our images — at least for now), and switch to a code base which is more easily maintainable by our developers; there are a few things that we have in the plans for the media-hub, like improving playback of remote media, support changing volume (from within the application) and playback speed, and finally supporting playback of multiple concurrent videos in the

same application.

This last feature is going to be essential for our TELEports application. As an immediate benefit, this rewrite brought us a small saving in the on-disk and RAM footprint of qt ubuntu-media, the plugin which exposes the media-hub functionality to QtMultimedia: since both the plugin and the client library are now written in Qt, a lot of glue code was removed. Thanks Alberto for that!

KNOWN ISSUES

On Pixel 3a (sargo) but maybe also on other Helium 9 devices there might appear to be an issue with video recording. This is due to incomplete cleanup of the gstreamer cache on startup due to a race condition. We decided to classify this as “known issue” since its not critical and can be fixed either by:

- one or two more reboots
- manual cleanup of the gstreamer cache by issuing `rm -rf $HOME/.cache/gstreamer-1.0`

BUGS AND ISSUES FIXED

- Adding a new APN entry in a

Helium 9 device will now be persisted correctly

- MMS can be retrieved also when in 2G network mode only
- MMS that could not be downloaded will be now displayed with a red warning text and the user has the option to try to redownload them. There is a small issue with that: Make sure network connection is established before pressing the download button, as otherwise the connection might be tried forever and no real timeout to reset the UI model happens. This will be fixed in the next version
- User metrics (the statistics info on the lock screen’s circle interior) was not showing anything in Helium 9 devices
- Contact selection in the messaging app now works correctly when sharing an item (like a contact) into it: Contacts can be added with the Plus sign, before this would eventually display the recipient’s details and not add them to the To: field.
- Also a frequent crash in messaging app on startup has been fixed, that seems to have affected only certain devices though
- In the clock app the clock hands would do a weird animation when setting up a timer and moving them across the 12 o'clock position.

- Custom message and ringtone sound selector had a layout issue where part of the screen was cropped
- Microphone access in the Morph Browser has been fixed. This is an important prerequisite to getting calls in the browser going. Certain conferencing / communication services should already benefit from that.
- Google account setup is now finally fixed for good!

HOW TO GET OTA-21

EXISTING UBUNTU TOUCH USERS

Existing users of Ubuntu Touch on Stable channel (which is selected by default in the UBports Installer) will receive the OTA-21 upgrade by using the Updates screen of System Settings. Devices will randomly be allocated the update from today and through the next week. This spread is to give us a breathing space to stop the rollout of any bad update (should that ever become necessary), and is not to accommodate any bandwidth restrictions.

If however you would like to receive the update immediately, turn on ADB access and issue the following command over `adb shell`:

```
sudo system-image-cli -v -p 0 --progress dots
```

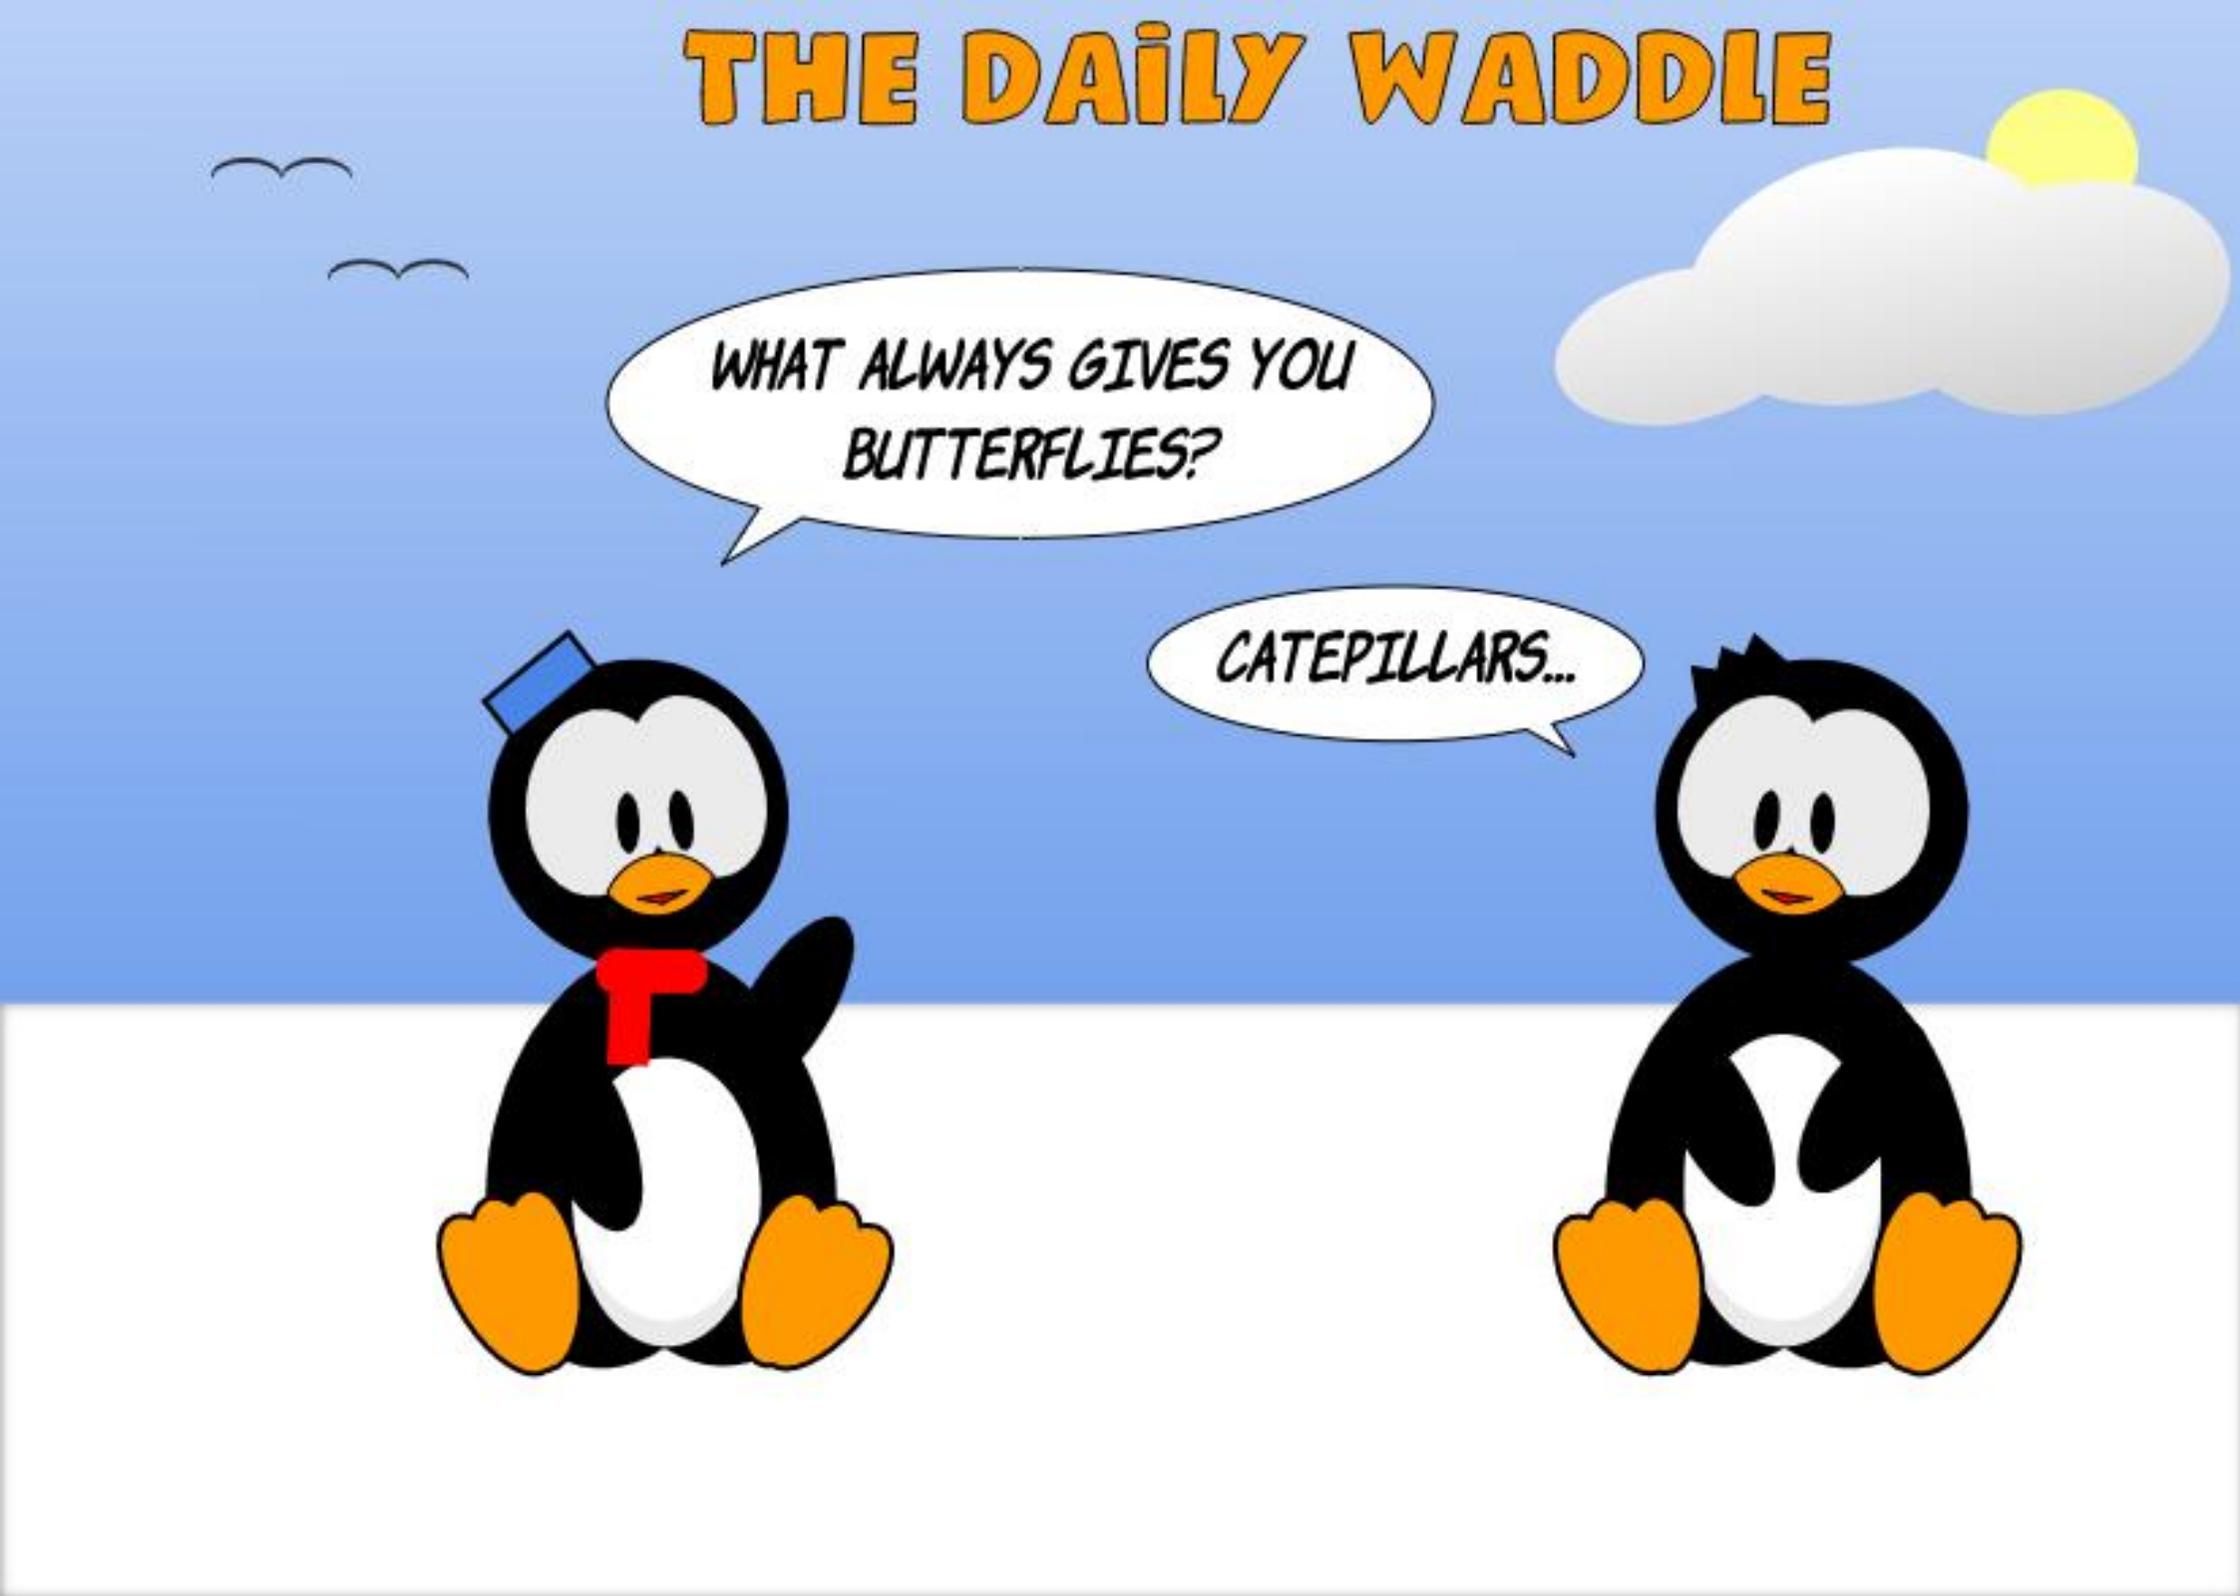
Your device should then download the update and install it. This process may take a while, depending on your download speed.

NEW UBUNTU TOUCH USERS

You will find instructions for installing Ubuntu Touch on your device at devices.ubuntu-touch.io.



THE DAILY WADDLE



WHAT ALWAYS GIVES YOU
BUTTERFLIES?

CATEPILLARS...



MY STORY

Written by Chris Burmajster

I have a pine phone. To those of you who have not heard of pine phone look at pine64.com. It's a Linux phone which can run about 20 different Linux OS's, but it's not sold as a complete phone, it was sold as a development phone. Until the software is perfect. I will explain why I bought one.

I do not like the two main OS's that phones come with. There's Android, which is Google and iOS which is an Apple product. Both very lacking in privacy! About 9 years ago Ubuntu reckoned that the world was ready for a Linux phone, and I was very excited by this and ordered one. It was a Spanish company called BQ who made it. I ordered it straight away. It was WONDERFUL to have an Ubuntu phone! The privacy it offered was absolutely staggering. To date, I still have this phone and am using it.

I purchased my pine phone last year when I bought my new car for my retirement. It had a feature whereby it had a wireless phone charging app on the car. The



PROCESSOR
Allwinner A64 @ 1.2GHz
(Quad-Core A53)

REPLACEABLE BATTERY
Samsung J7 3000mAh

RAM & STORAGE
2GB 16GB
3GB 32GB

5.9" DISPLAY
IPS LCD (1440x720)

KILL SWITCHES
Wi-Fi • Cellular
Camera • Speaker
Microphone

CAMERAS
2MP 5MP

GRAPHICS
Mali 400 MP2 GPU

Data / Charging / Video Out
USB TYPE-C PORT

AFFORDABLE & AWESOME
PINEPHONE BY PINE64

salesman chatted about it and put my 9 year old phone on it.

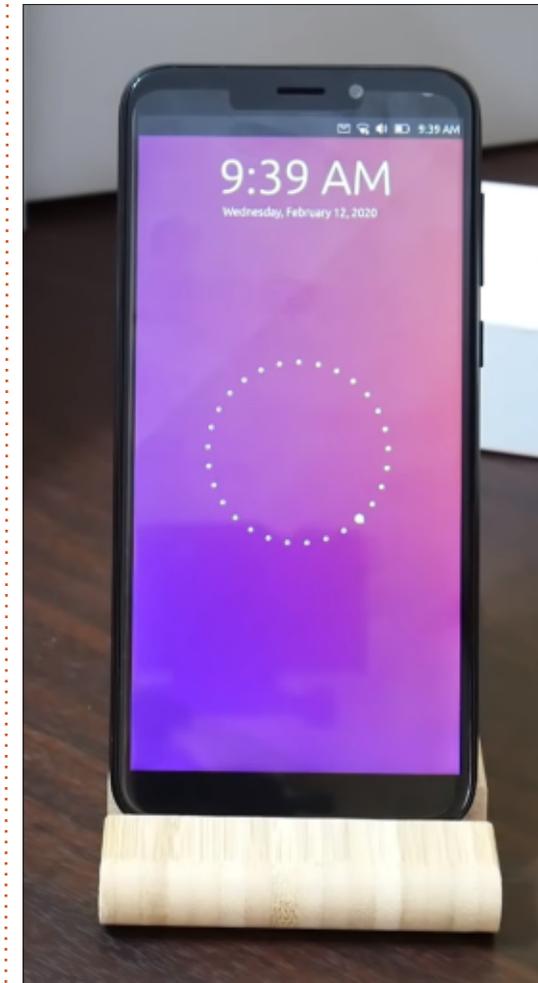
Now those of you with 9 year old phones will know that they don't have wireless! But I didn't know that at the time. Anyway, my attention was elsewhere! I drove the car back home with the phone on the wireless charging on (I didn't know how to switch it off). When I got home I took the phone off and it was dead. I was puzzled by this, because everything about the phone was solid.

So I ordered a pine phone, which

came with Manjaro OS, but I installed UBports, which was Ubuntu in all but name. Ubuntu had ceased to develop the OS for the mobile phone and UBports had stepped in to develop the phone. Those of you who have seen UBports will have (hopefully!) a good view of it.

Anyway, the pine phone was not a daily phone even with UBports on it. So I looked at my old phone and discovered that it was just the battery that had gone. So I ordered another battery. When it came I fitted it and, hey presto, it all

The Pinephone

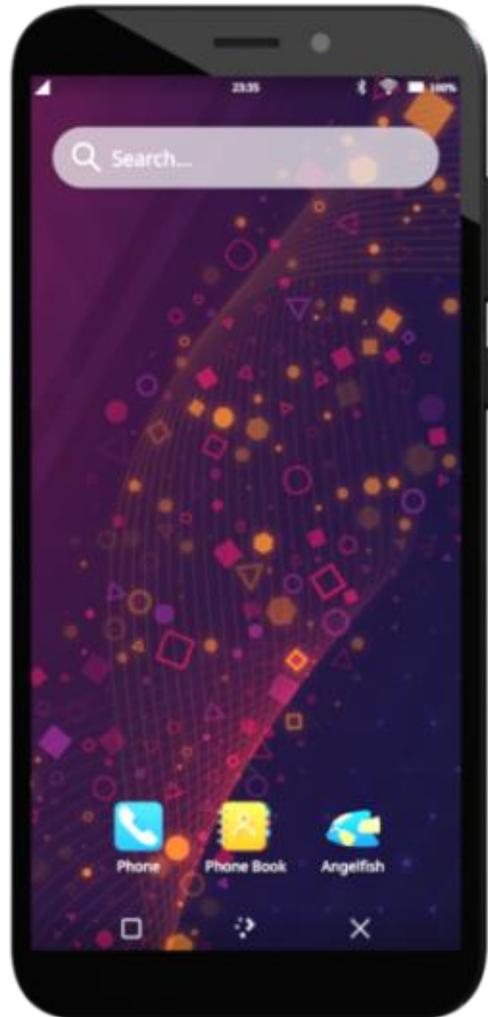


worked! So I was back to an Ubuntu phone.

So I have kept my pine phone all this time looking at UBports website and the pine phone, but so far they have not yet cracked the software. So I decided to put KDE

Plasma Mobile on it just to try it and see what had been achieved. This had been the mobile OS of choice for Pine64 and it was interesting to see it.

However, one of my computers had given up the ghost, and so I bought a new one just a few months ago. It was a Windows



computer to go with the Ubuntu computer that was my main computer. Then came Windows 11 which I put on it and it runs OK.

Now, the last time I used my Ubuntu computer it was rather complex to put an OS on the pine phone, because it came with multiple drives. I didn't want to go



through that again. So I used Windows 11. I downloaded the KDE Plasma Mobile, also downloaded Balena Etcher for the flashing process. My Windows 11 computer is set up according to Rob Braxman's new computer setup, which means that I have two users, one a local account, which I use regularly, and the admin account which I hardly ever use. I don't use any anti virus, except for Microsoft's own. That was going to be the problem.

I did everything except turn off the anti virus, and this made Balena Etcher twitchy. It didn't flash properly. I tried it again and once again, it didn't work. But then something was wrong with my pine phone. It didn't work. It was dead! So I got onto Pine64's forums where I asked why this had happened and was it going to be fixable. I'll cut the time spent (5 days!) down.

One thing that worked was Jumpdrive. This was a utility that allowed the eemc to be directly exposed to the flashing process. I had an SD card which ran Jumpdrive and the phone worked with it. But it didn't offer the computer an emmc drive to which the computer

wanted to flash it. Eventually, I got the SD card to be flashed by a mobile OS, KDE Mobile to be precise. This worked because the phone booted from it, but it was very slow. Then I discovered with the back off of the phone, that there was a reset button. I asked what it did and was told that it reset the hardware to factory settings. Oh wow! I immediately used the reset button and found that I could flash UB ports to the phone on eemc!

So I did and now all is well. It's so good that the pine phone has all these features and I'm very pleased with it.



In this series of articles, I will be going through the history of the Graphical User Interface (GUI). The questions addressed will include:

- What is a GUI?
- How old is it?
- Who created the first one?
- Why did they create it?
- How have they evolved over time?

I will try to cover as much of this as I can.

WHAT IS A GUI?

A GUI (pronounced gooey) is how you interact with a computer or other electronic device without needing to type in commands using a keyboard. The most common ways of interacting are by way of a stylus, mouse/trackball, touch pad or just your finger. You tap or click on icons that represent functions you wish to perform. When the early computer pioneers created text input via keyboards, the computer became more broadly used. The early adoption was slow and had a niche following. They knew a better way was needed.

How old is it? The idea started longer ago than you might think. It all started with a device called the light pen. The light pen was created around 1955 at the Massachusetts Institute of Technology (MIT), a private land-grant research university in Cambridge, Massachusetts. This was part of a larger project. The U.S. Navy's Naval Research Lab approached MIT about creating a computer to act as a flight simulator to train bomber crews. After some initial talks they decided to fund the development under the name Project Whirlwind.

The light pen was a light-sensitive wand that was connected to a computer terminal's cathode-ray tube (CRT) display. The light pen detects changes in brightness of screen pixels when updated by the CRT's electron beam and sends that timing information to the computer. Since the update of pixels was done one at a time by the electron beam and at a known refresh rate, the computer terminal could figure out where the light pen was touching the screen. Over time, once it was determined to be reliable, it was more widely used on the Situation Display consoles of the AN/FSQ-7, a large computer

system for military airspace surveillance built by IBM (though MIT still aided in the development). After IBM took over development, they made use of this technology on many of their other terminals.

A few things about the terminology and the technology that was used in the 1950's and 1960's: These computer systems were multiple cabinets that would fill an entire floor of a building. When they say it was a graphics terminal, by today's standards it was text. These Graphics Display Units (terminals) were only input and output devices, as they needed to connect to a Central Processing Unit to be of any use. The graphics terminals themselves processed all the input and output for the Central Processing Unit. A graphics terminal consisted of a CRT, light pen and a keyboard. To shown left is an example of an IBM 2250 Graphics Display Unit.

The Central Processing Unit is what we would refer to as the computer. It contains the core





memory, disk storage, keyboard, control panel and a line printer.

Shown top left is an example of an IBM 1131 Central Processing Unit and below is an example of a IBM 1130 Central Processing Unit.

Bottom right shows them connected together. (Images courtesy of Columbia University in NYC, IBM, The Museum of Applied Arts & Sciences and By Martin Skøtt – Flickr: IBM 1130, CC BY-SA 2.0 via Wikimedia)



Countless people at many companies and universities continued further development throughout the 1960's. The US Government continued their involvement by funding many of these projects. Here are some notable creations:

The Sketchpad was a computer program by Ivan Sutherland in 1963. It is considered to be the ancestor of both Computer Aided Design (CAD) and a major breakthrough in computer graphics in general. It was the first program to use a complete graphical user interface. It ran on the Lincoln TX-2, which was an advanced transistor-based computer built in 1958 by

MIT's Lincoln Laboratory. Users would draw on the screen using a light pen. They could create an initial drawing and then further modify it using geometric constraints. Geometric constraints were a major invention. In short, this allowed for precise drawing as opposed to free hand drawing. It would correct crooked lines to an exact length, and angles could be maintained while the drawing is modified. Here is a 7 minute 16mm film from 1963, it is a demonstration of how Sketchpad works: <https://youtu.be/57wj8diYpgY>

The RAND Tablet, developed by the Rand Corporation in 1964, was



an input device that used a stylus attached to a 10 by 10 (254 mm by 254 mm) printed-circuit screen using capacitive sensors. It looks and functions similar to today's drawing tablets, like a Wacom tablet. The Rand Tablet had 100 lines per inch (25.4 mm) resolution, capable of digitizing 1 million locations, and used a handwriting recognition program called GRAIL (Graphical Input Language). GRAIL could identify 53 hand drawn numbers, letters, symbols and geometric shapes. It was also able to use gestures to manipulate text and etc on the screen. This product was the inspiration for the later creation of devices like Wacom's drawing tablet, Apple's Newton and the Palm Pilot by Palm Inc..

RAND continued to develop computing devices and computer systems. They also partnered with IBM to create the VGS (Videographic System). It was a very interesting computer system, but a little too much to explain here. It combined video content (viewable on modified TV's) and the RAND Tablet to allow a way of annotating the content. Cartographers found this system an amazing tool to create and annotate maps. I will link some PDF files below.

RAND Videographic System. This is a technical overview of the complete system. Shows all the components that were used: <https://www.rand.org/pubs/reports/R0753.html>

RAND and the Information Evolution. For information specifically on the Videographic System, start at page 89: https://www.rand.org/pubs/corporate_pubs/CP537.html

Light pen usage continued into the 1980's, as early personal computers had the ability to use them. They fell out of use later in the decade, due the fact that you needed to keep your arm extended for long periods of time. A similar technology was used on game consoles, where the light pen was made to look like a gun, such as was used in a popular arcade game called Duck Hunt.

Linked here are two YouTube videos on how the RAND Tablet functioned. The first one is a 2 minute excerpt of a 97 minute presentation done by Alan Kay called "Doing With Images Makes Symbols: Communicating With Computers". He is one of the many

computer scientists that were responsible for the graphical interfaces that we use today, and is best known for his work on object-oriented-programming and windowing graphical user interface design. The second video is the original video Alan was commenting on in his presentation. Alan started it at the 7 minute and 30 second mark. The entire video is 14 minutes long and was produced by the RAND Corporation from the mid 1960's. It is a very well produced video and the computer screen is very clear and legible. If you don't want to take the time watching it; I would suggest at least skipping through it, as I think you will be surprised at its functionality. <https://youtu.be/QQhVQ1UG6aM>

Alan Kay "Doing With Images Makes Symbols: Communicating With Computers". <https://youtu.be/2Cq8S3jzJiQ>

The last two projects of the 1960's I am going to mention were also revolutionary in how modern graphical user interfaces operate. The following was the work of Douglas Engelbart, he was an engineer and inventor. He joined the Stanford Research Institute (now known as SRI International) in

1957. His interests centered around how people interact with objects they use in day to day activities, with a particular interest in how children do the same. He recruited others in the 1960s to form the Augmentation Research Center at the Stanford Research Institute. Out of the many contributions he and his group made to computing, I will mention three of them that I consider ground breaking.

He created the idea of hypertext. It is the foundation of how point and click actually works. The icon is a link that refers to another object. Professor Sir Tim Berners-Lee, a Computer Scientist at CERN (European Nuclear Research), used this idea to create HTML in 1989, and released it in 1990. To quote him "I just had to take the hypertext idea and connect it to the Transmission Control Protocol and domain name-system ideas and—ta-DA!—the World Wide Web".

Continuing on with Douglas, he and his group (he is named on the patent) created the computer mouse. It was given that as a nickname because the cord resembled a mouse tail. They intended to create an official name

but the nickname caught on. The first prototype was in 1964 with its public demonstration in 1968. The patent was filed in 1967 and granted in 1970. You might find it interesting that the trackball predated his invention. It was invented by two separate people in separate countries a few years apart. Going by the earliest date first, an English engineer named Ralph Benjamin invented it in 1946 and was granted the patent in 1947, he called it roller ball. A few years later in 1952 a British electrical engineer in Canada named Kenyon Taylor and his team at the Royal Canadian Navy's Digital Automated Tracking and Resolving project (DATAR) created their version of the trackball.

I saved the best for last. One of the three aforementioned projects Douglas worked on was the Augmentation of Human Intellect project. That project created the oN-Lin System (NLS). Douglas was inspired by some papers written by Vannevar Bush in 1945. He was an early engineer, inventor and science administrator for the US Government. His concept was called Memex. It outlines how our computers work today, really shocking how he described the

functionality 76 years ago.

The NLS was software that allowed many functions. To name a few, there was document creation/removal and editing. You could select text and link it to other portions of the document or even another document (a form of hypertext linking). Edit a document with multiple people remotely, which makes it the first groupware application. (Think of it as a Wiki.) This had the ability to move and resize multiple windows. It was also the first time anyone had ever seen a mouse before. The early development was using the CDC 3100 in 1965 and later using the Scientific Data Systems (SDS) 940 in 1968. These were smaller yet more powerful computers compared to those used in the 1950's. Instead of filling a floor or multiple floors of a building, these computers (fully connected) were about the size of a bus. Each unit was about 4 ft wide (1.14m) and 6 ft tall (1.87m). The amount of these connected together determined the computation ability and the amount of memory available, for a lack of a more technical explanation.

Developments of design concepts began and were

supported by the US Air Force from 1959 to 1960, and Douglas later published a framework in 1962. In 1968, development moved to the above-mentioned SDS 940 using the Berkeley Timesharing System, which was an operating system that allowed multiple people to use it. Non-time-sharing operating systems would need to process the request of one person before it could process another.

Douglas and his team gave a demonstration at the Association for Computing Machinery / Institute of Electrical and Electronics Engineers (ACM/IEEE)—Computer Society's Fall Joint Computer Conference in San Francisco, California on December 9th, 1968. This demonstration is referred to as The Mother of All Demos. Douglas's terminal was connected via a home built 1200 baud modem back to their lab in Menlo Park. Additionally a mouse, keyboard and a 5 key chorded keyset was connected to the terminal. They also had two microwave transmitters, a video switcher and cameras on loan from NASA. Which provides video between the conference and their lab in Menlo Park. Output of the terminal was projected on a large screen for the audience. After the

90 minute long demonstration the audience gave a standing ovation. I will link below the 90 minute video and a highlights playlist of 10 excerpts plus 2 additional videos from the 90 minute video. It is about 24 minutes if you watch all 12. I found it quite amazing at what this software could do. There is an over the shoulder view of the keyset, mouse and keyboard. The keyset is on the left, it looks like a 5 key mini piano: <https://youtu.be/2nm47PFALc8>

THE MOTHER OF ALL DEMOS

Excerpts from The Mother of All Demos: <https://youtu.be/VScVgXM7lQQ?list=PLCGFadV4FqU2yAqCzKaxnKKXqnJBURKTE>

To see the list of videos click once on the = in the top right of the video window, if it does show them click again.

In the next article, we will enter into the 1970's and the beginning of the GUI as you know it today.



HOW-TO

Written by Ronnie Tucker

Write For Full Circle Magazine

GUIDELINES

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

RULES

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

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

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

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

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

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

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

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

TRANSLATIONS

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

REVIEWS

GAMES/APPLICATIONS

When reviewing games/applications please state clearly:

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

HARDWARE

When reviewing hardware please state clearly:

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

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



Xubuntu 21.10 is the 32nd release of this Xfce-based distribution and came out on 14 October, 2021. It is a “standard” release and will be supported for nine months, until July 2022.

This particular Xubuntu development cycle has been a bit unusual. Normally, the first standard release after the last Long Term Support (LTS) version brings the largest number of changes and these are then polished and developed over the next two standard releases, leading to the LTS release at the end of the cycle. The Xubuntu developers decided to do things differently this cycle. The first release, Xubuntu 20.10, introduced no changes at all over the previous LTS, as the developers concentrated on moving their development platform from Launchpad to GitHub. The next release, Xubuntu 21.04, brought only a few small changes, so I was keen to see if the last release before the next LTS would bring small or large changes. The answer is “small”.

BOOTING IT UP

I downloaded the ISO file from the Xubuntu website via BitTorrent and carried out a SHA256 check to make sure that the downloaded file was good. I then used UNetbootin to write it to a USB stick, as it conveniently leaves the stick in FAT32 format. For testing, I just ran it from the USB stick.

I tested Xubuntu 21.10 on my new, high-end hardware, with a 4.7 Ghz quad core processor and 32 GB of RAM and, as expected, it loaded and ran very fast.

SYSTEM REQUIREMENTS

The recommended minimum system requirements for Xubuntu 21.10 have not changed since 21.04 and remain:

- 1.5 Ghz Dual Core processor
- 2 GiB RAM
- 20 GB of hard-drive space

These are probably realistic minimum system requirements for basic performance and represent a computer about ten years old.

NEW

Xubuntu 21.10 includes a new point version of the desktop from Xubuntu 21.04, in this case Xfce 4.16.2, based upon the GTK 3.24.30 toolkit. This brings some improvements to themes, including Greybird, the default theme, as well as some small tweaks to core applications. For instance, the file manager, Thunar, now displays the full directory path to each file in the header which is useful.

Most of the changes in Xubuntu 21.10 are in the form of the inclusion of some new applications. These are the Gnome Disk Analyzer and the Gnome Disk Utility, both to help manage storage media and the Rhythmbox music player. All of these are taken from Ubuntu and the Gnome desktop.

Pipewire has also been added. This is a “server and user space API to deal with multimedia pipelines” and promises improved audio playback quality, especially for



REVIEW

Bluetooth audio.

The Pidgin IRC client has been removed since it was replaced by Hexchat. Pidgin is still in the repositories so can be installed if needed.

SETTINGS

Greybird remains the default window theme in Xubuntu 21.10 although the color scheme has been slightly improved in this version of Xfce. There are six window schemes included along with six icon themes. The default icon theme is Elementary Xfce Darker.

Xubuntu 21.10 has a new wallpaper for this release which is distinctly green in color. Most recent default wallpapers have been blue, the last green one being for Xubuntu 18.04 LTS. Most recent Xubuntu wallpapers have been abstract designs based on circles but this one is more angular. If it is not to your taste then there are 18 other wallpapers included: a mix of abstract designs and landscape photographs, or you can use your own.

The Whisker Menu has been a feature of Xubuntu since 14.04 LTS and remains the main menu system in 21.10. It is a nicely laid-out menu and is highly customizable, being one of the few Linux desktop

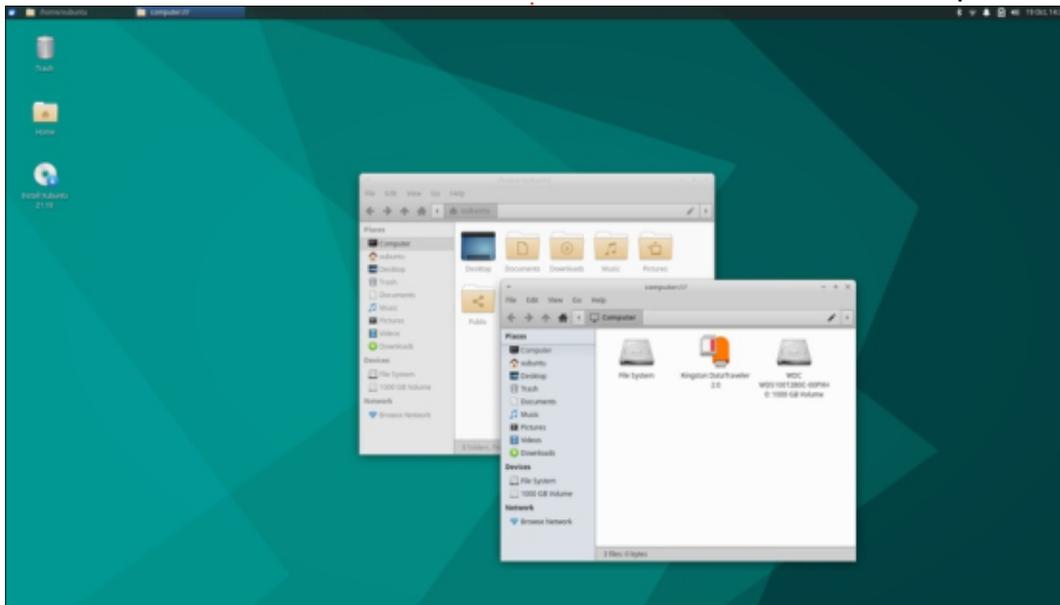
menus that can be resized. Whisker remains the main feature that differentiates Xubuntu from other distributions.

APPLICATIONS

Some of the applications included with Xubuntu 21.10 are:

- Atril 1.26.0 PDF viewer
- CUPS 2.3.3 printing system*
- Catfish 4.16.2 desktop search
- Firefox 93.0 web browser
- GIMP 2.10.24 graphics editor
- Gnome Disk Utility 41.0 disk space and health monitor
- Gnome Disk Usage Analyzer 41.0 disk display
- Gnome Software 40.4 package

- management system
- Gparted 1.2.0 partition editor
- Hexchat 2.14.3 IRC client*
- LibreOffice 7.2.1 office suite
- Mousepad 0.5.4 text editor
- Parole 4.16.0 media player*
- PulseAudio 15.0 audio controller
- Ristretto 0.11.0 image viewer
- Rhythmbox 3.4.4 music player
- Simple Scan 40.5 scanning utility (re-badged by Gnome as "Document Scanner")
- Software Updater 21.10.4 (update-manager) software update manager
- Synaptic 0.90.2 package management system*
- Thunar 4.16.8 file manager
- Thunderbird 91.1.2 email client
- Transmission 3.00 bittorrent client*
- Wget 1.21 command line webpage



REVIEW

downloader*
Xfburn 0.6.2 CD/DVD burner*
Xfce4 Panel 4.16.3 desktop panel
Xfce4 Power Manager 4.16.0
system power manager*

* indicates same application version as used in Xubuntu 21.04.

It seems that, over time, Xubuntu is relying more and more on Gnome applications, especially for utilities.

As in recent Xubuntu releases, there is no default webcam or video editing application although there are several of each in the repositories that can be installed.

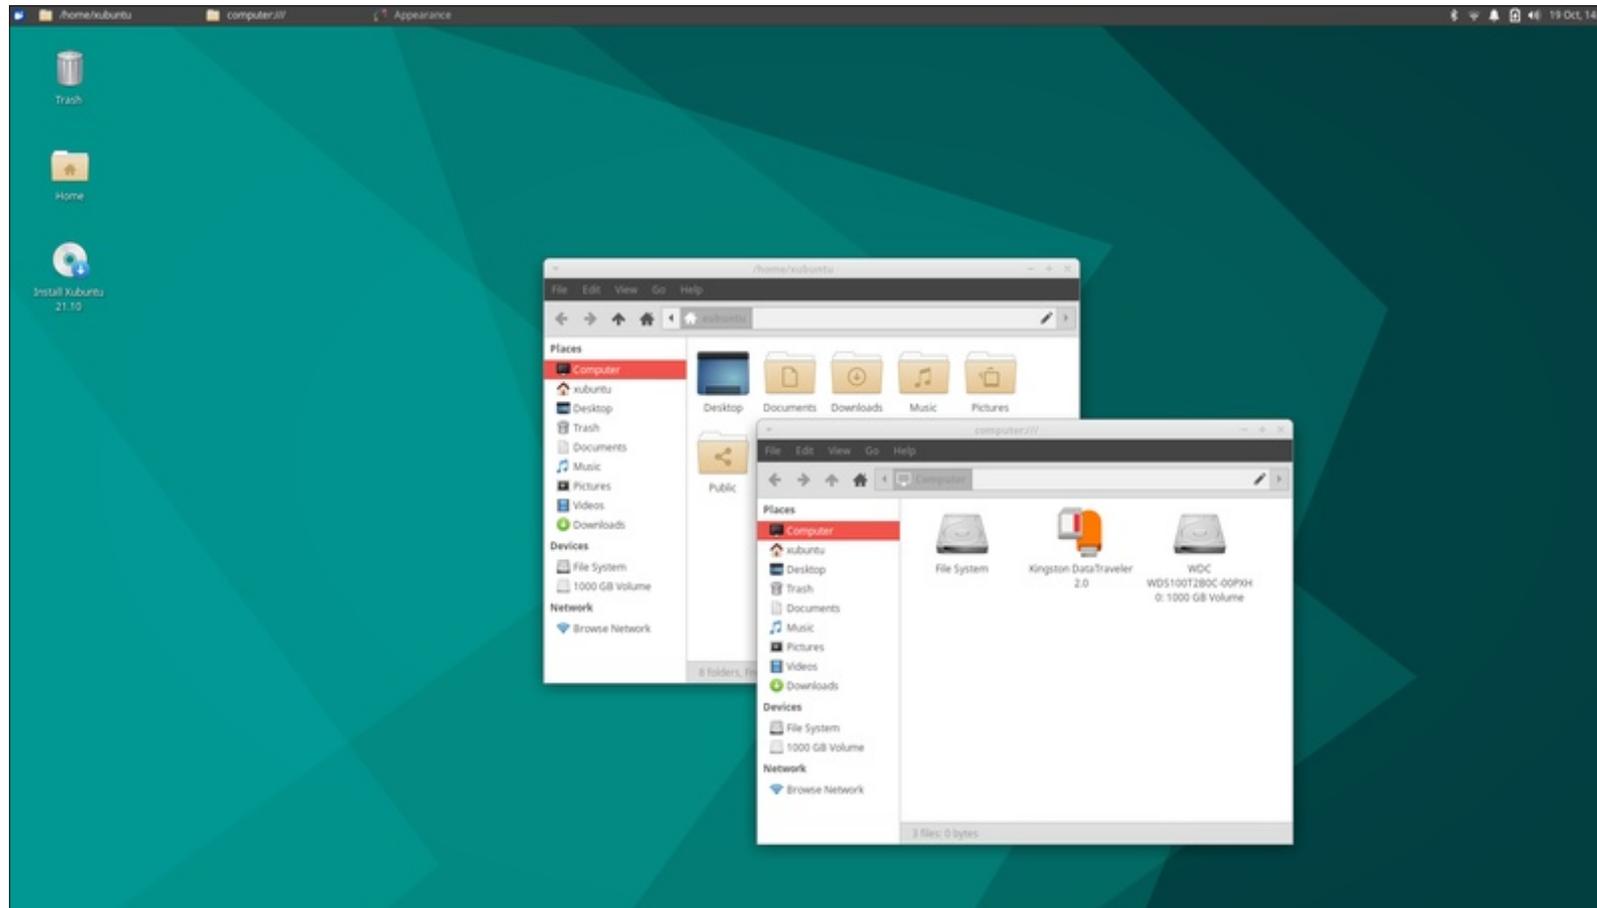
Xubuntu 21.10 LTS includes LibreOffice 7.2.1 which is, as usual, lacking only LibreOffice Base, the database application. It can easily be installed if needed.

This release still comes with a CD/DVD burning application, Xfburn. Since it is now almost ten years since new computers came with CD/DVD players, it is starting to make Xubuntu feel dated to include it by default, rather than let users install it on older computers, if needed. It will be interesting to see if Xfburn gets dropped in the upcoming LTS release.

Xubuntu 21.10 is a good, solid release with no major notable issues to report.

This Xubuntu development cycle has introduced only some minor changes which is not unexpected in a distribution that has been around for 15 years. The Xubuntu developers and its users all seem to agree that they like how it works and no one wants to see any big changes. Expect the LTS, Xubuntu 22.04, to be very similar to 21.10 when it comes out on 21 April, 2022.

CONCLUSIONS



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.



LETTERS

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

Join us on:



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



twitter.com/#!/fullcirclemag



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



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

FULL CIRCLE NEEDS YOU!



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

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

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



Q&A

Compiled by EriktheUnready

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

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

As I was unpacking, I came across my old ZX Spectrum 48k computer. Then another... then another... and another and a ZX81!!? Try as I may, I have no recollection of having more than one, ever. (And a ZX81, never) Clive Sinclair died 16 September 2021 (aged 81) and I can admit that I was sad. If he were born 40 years later, I am sure he would have been the UK's answer to Jeff Bezos, as he also sold electronics via mail order at an age most of you were still at school. It wasn't until 1982, however, things

went 'viral' for him. In 1982, Sinclair Research released the ZX Spectrum. *fade to memory sequence:* I had a friend who got a 16K ZX Spectrum as it launched. This was the first computer I had seen that didn't have blocky graphics and one of the few I knew of that came with a whole 16K. That was like 14,000 four letter words! My mind was blown. I had to have one! The ZX Spectrum sold over 5 million units. (and could probably have sold more, but it was always sold out where I lived) This little rascal ushered in a generation of bedroom developers who are still shaping the software we use and the games we play today. Even Linus Torvalds cut his teeth on a Sinclair QL. (the not-so-successful computer) Dumpster diving for 1 litre soda bottles to turn in for money, it only took me about six months to afford a shiny 48K ZX Spectrum! To put it into perspective, I would have to dumpster dive for about six years now, to afford a ZX Spectrum NEXT. So it was VERY affordable. Also my other choice was a Commodore VIC-20 at double the price running

at only 1Mhz (Speccy ran at 3.5Mhz) and had fat graphics like an Atari 2600. (yes, 1977 called, wanting their graphics back) ... Anyway, we had no books or magazines on computers, but I did stop by the library every week, to check. One friend (or I should say his mom) would print out some homebrew magazine in dot matrix every two months, or else we were dead in the water. Maybe that is why I have so much respect for FCM, bringing Ubuntu to the masses, for free – helping beginners get going, Yes, now we have Youtube etc, but even when FCM launched, the internet was so expensive where I lived, one guy would download FCM and share it with everyone.

Q: I have traded in my first gen i3 laptop for a 10th gen i7 after many years. What is the best way for me to get all my stuff on the new one? By this I mean it must include my emails and my passwords, firewall rules etc, as I have it the way I like it now. I use Ubuntu Budgie.

A: Copy your home folder including the hidden files, that is the important bit, onto a backup drive, and copy them in on the new one. You will obviously have to import a few things again, but it should be identical.

Q: Is there a quicker way? I have a master image of ubuntu for Virtual box. Once I have messed up the OS, I have to delete the image, then copy the master and set it up again for testing. I am currently working for home so I don't have and abundance of resources.

A: You didn't give me hardware and software details, but may I suggest Vagrant or Docker? As I don't know what you are doing, and what you are trying to achieve other than create and destroy VM's.

Q: My laptop came with a 256GB SSD and a 1TB hard drive. I installed steam and started downloading my games, but it got too full. I can move my steam folder to the 1TB, but I am a bit confused

reading about hard and soft links. Can you clarify please?

A: Actually, no. I need to keep QnA short and concise to give everyone a turn. Suffice to say you can not only move your Steam folder from within Steam settings, but set it so that all future games install there as well, freeing your SSD, requiring nothing more from your part.

Q: I bought one of the HP servers from the office sale, and it came with Ubuntu server. The catch is that the BIOS is locked and I can't boot from USB. IT says the password they have does not work. Can I turn this into a usable machine or does it need to go to a specialist?

A: If you have a root user, you can simply install a desktop, I don't know how old that server is, but you can try XFCE as it is lightweight: `sudo apt install xubuntu-desktop` -and once installed test it with: `startx` (see: <https://www.tecmint.com/install-xfce-desktop-in-ubuntu-fedora/-for-v4/>)

Q: I have Telegram desktop installed on my Xubuntu machine, but the other day it updated and now I have 2 entries in my menu. How do I know which one is the right one?

A: They are probably the same thing, edit your menu to be sure, and once confirmed just remove one.

Q: I inherit laptop my uncle: <https://www.notebookcheck.net/Review-HP-ZBook-14-Workstation.114091.0.html> Very powerful laptop. I wanted to make Ubuntu. Is no install.

A: The HP Zbooks are notoriously difficult to install Linux on as they use some proprietary caching to a mini SSD. You will need to play with the BIOS to get just the right settings. I suggest you trawl some forums for this as I don't have one to test the settings for you.

Q: Somehow FireFox has gone out of bounds on my screen. The top part you use to move the window, has gone under my top panel and even out of the display.

Any ideas on how to get it back? It started when I clicked the min/max button and it flickered and just went back to filling the screen, every time I clicked it, now I can't reach it.

A: Try the F11 key or use your system manager to kill it. (They differ from desktop to desktop) You can also try resizing your screen resolution smaller then bigger again. Or, even use the windows key (Sometimes called super) plus the arrow keys, that should tile the current window.

Q: I am feeling dumb at the moment. I was working through orientation, when it said to open a new tab in the terminal to continue. WTH?

A: Not all terminal emulators support tabbed terminals, a quick way to see is right click inside the terminal emulator. If you are working on TTY, it is not going to happen, you need to open another TTY.

Q: I am using a static address on my Ubuntu, when I go and change it, it still shows the previous address. Instead of me opening the command line every time, is there no other way to make it refresh to the new address.

A: That you know to open the terminal and stop /start the network card shows you know what to do. Simply click your Wifi icon and tick off the box enable WiFi and then repeat enabling the Wifi from your desktop. (Or see if you don't have a short-cut key combination) It has the same effect.

Q: My once powerhouse XP PC has been relegated to the minors. I went with XFCE, and an old square screen for the kids, 1024x768 I think. The menu bar is a bit fat. Right click properties, only gives transparency. It needs to shrink. Settings-Panel profiles has me scratching my head. Did I mention I was new to XFCE?

A: Right place, wrong spot. Right click, panel -> panel preferences. Panel profiles is usually a theme thing.

Q: You can rag on the new guy, but updates give me a GPG error. I Googled it and it says goes on about pretty good privacy an email. I fail to see what that has to do with my update. I went here, <https://knowledgebase.progress.com/articles/Article/ws-ftp-professional-gpg-errors> as it is supposed to be a knowledge base, but is is useless. I surfed around for at least 2 hours to find mostly junk.

A: GPG is Gnu Privacy Guard and PGP is Pretty Good Privacy. It's a finger error Mr butterfingers. Do you need help with the error, or are you good? You seem to be a go-getter.

Q: I have taken an online course and I am stuck before even starting. I am still on Ubuntu 20.04, but I can install 21.04 if I need to. The first error is: Error: No such command '/usr/local/bin/celery' then it snowballs. I don't really see anything in dmesg as that is filled with [83558.304803] audit: type=1400 audit(1641483245.375:1330): apparmor="ALLOWED" operation="open"

profile="libreoffice-soffice" name="/home/Michael/Documents/lu5492592raxuh.tmp" pid=569209 comm="soffice.bin" requested_mask="wrc" denied_mask="wrc" fsuid=1000 ouid=1000 and I don't know what that is.

A: Come on dude, we removed the Celery disclaimer last year. (like a month or two ago) Type \$PATH (Capitals) and see if it is in there. Does that file actually exist when you go to /usr/local/bin? Did you follow the official installation method?

Q: I have installed Xubuntu on my dead Mac. I didn't like where it went after Snow leopard, so now it cannot get updates, hence Xubuntu. Though there are similarities, -- I even themed it to look like the latest OSX -- There are a lot of differences. Where things were connected to my profile, on OSX, in Xubuntu it is not. I clicked on open my windows when I log in again and it worked well, until it didn't. Where do I clear it?

A: In the menu go to Settings -> Session and Startup and in there is a tab labelled "Saved

Sessions". At the bottom you can clear it.

Q: Let's be blunt, I am computer illiterate. With that out of the way, I can safely say I did not edit /etc/update-manager/release-upgrades in any way. Yes, I was on a few forums and people tell me I need to edit this file. Fine. What is foremost on my mind is- why did it change? I only want LTS upgrades. I have never needed to and I have been using Ubuntu sine 12.04. What changed?

A: There is another... As master Yoda says. If you open your "Software and Updates", you will notice that the third tab says "Updates" right at the bottom of that tab, you will notice "Notify me of new Ubuntu version", that is most probably where it changed with the slip of a mouse, or something. Change that to LTS.

Q: When I start my computer in recovery mode, I see this, failed to start Load apparmor profiles. Ubuntu 21.04, i7, 32GB RAM and then it doesn't load the recovery. Failed to start default target: Transaction for graphical.target is

destructive.

A: I would suggest booting from an Ubuntu CD or USB and running fsck on that drive. If all else fails, copy your home folder and reinstall.

Q: My laptop is Ubuntu 18.04.6 LTS Kernel 4.15.0.166 Mesa 3.0 v20.0.8 and then when I use Steam proton, I get an error saying Dx11 Feature Level 10.0 Is Required to Run the Engine.

A: I am afraid your graphics card driver does not seem to be loaded, Mesa is a generic driver. You need to go to "drivers" and install the proprietary driver for your graphics card. That said, the card needs to support DX11. If it does not, you will not be able to play DX11 games.



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: <http://www.soldak.com/Drox-Operative/Overview.html>

Price: 19.99 USD

Blurb: "Drox Operative is a spaceship action RPG with warring alien races, fierce space battles, a dynamic, evolving galaxy, and co-op multiplayer for Windows, Mac, and Linux."

After the longest description from a website so far, let us dive into the actual game. This is a top down Elite, or Diablo in space with 4x elements, but, instead of you colonizing planets and harvesting resources, you influence, terrorize, coerce, spy and so forth. Mmm... on second thought, maybe you are the CIA in space. You start off picking a race, however, this only determines your bonuses and nothing else really. You can 'try before you buy' – which I think is great – here: <http://www.soldak.com/Drox-Operative/demo.html>

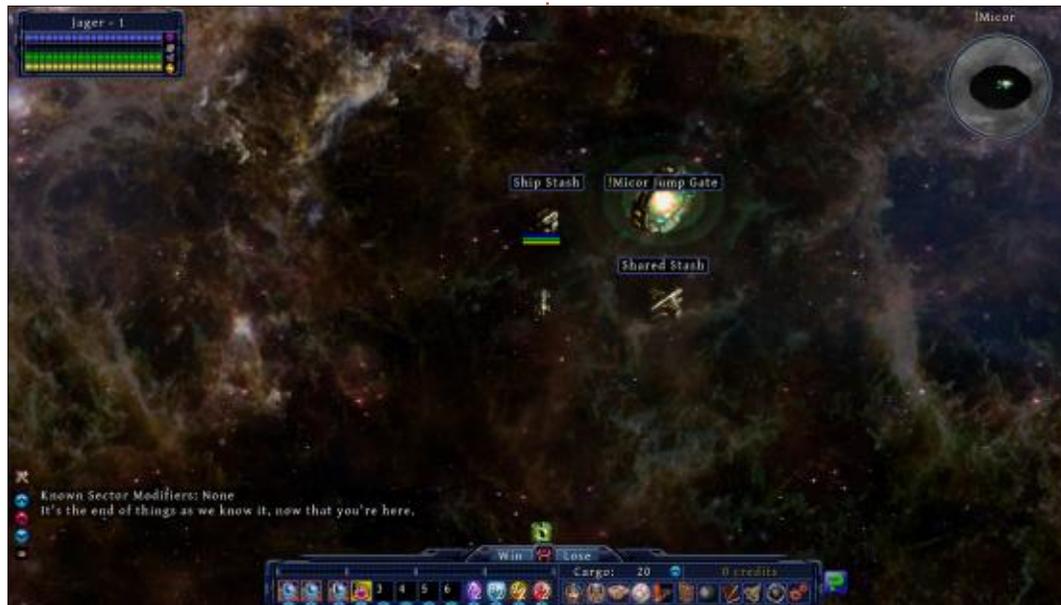
What it is not – even though I grabbed the game thinking it was another Ascendancy – it is not a 4x game. There is a 4x game being run in the background by the AI, but you are only one ship, with missions to take, kick ass and chew bubblegum, even though bubblegum went out of fashion 2 centuries ago.

Which brings me to the graphics. In terms of Soldak games, this is one of their better ones, but the artists still lack imagination. All the "ships" look alike except for one, all the races even look alike, right out

of star trek 1977, just different colour make up. Being in space I suppose the background does not interfere with the foreground as much. (but it can still get overwhelming) If you played Kivi, one of their other games, you will know what I mean, characters get lost against the background. The UI design is in line with their other games and just makes it as functional.

If you look at the screen-shot, you will see three 'health' bars, top left. Blue is shield power, green is structural integrity and yellow is

power. For combat, this is what you will need to manage. The bottom part is quick slots and cargo and top right is the mini map. All fairly standard, with a gorgeous background, but notice how my ship almost disappears. The pace of the game is also much faster than previous Soldak games, so that is great. The downside is there is a lot of grind, no real story and little brainfood in the reward tree. This game needs a grand space opera story with unlocking story elements for reward points and it would take it out of the mediocre category. Now don't get me wrong – the game is not mediocre in gameplay, it is actually fun, it is just that there is no rhyme or reason, other than to beat the game six ways. Yes, there are six ways to win. However I can't see myself being bothered. Ten hours in and I am already bored. I see the same old-same old, wherever I go. This is the only issue I have with all Soldak's games, Din included. It is like the company belongs to a family and the son is the sales manager and the daughter the creative manager and so forth. They can't do any wrong



and the end product suffers for it. Let me quickly point one out, all the races have crew +2 perk.(so why even bother?) You need to have constraints and balances. The weapons are all the same, they all do the same damage, they even all look the same and sound the same. I can see there was a lot of effort put into the game, but it is almost like the budget dried up. There is so much potential here. On the one side I am worried these guys plan another Hearts of Iron, where there are endless DLC to milk the players.

If you were hoping to win the game by diplomacy, know that there is no real diplomacy, you have to do missions for a race to increase diplomacy, and that's it. So currently I have a love-hate relationship with this game.

The mechanics also need work, you cannot play "smartly", I got caught in an ambush and high tailed it out, crossing a minefield hoping it would affect the other ships too, but no, they only affect you. The minefield was next to a planet of one faction – presumably placed by that faction to protect their planet, yet every other faction can fly through it, even the ones they are

at war with, but me. It reminded me of the new Bond movie, where he has this terribly powerful EMP in his watch, that will fry a guy's bionic eye, but ignores the cellphone in Bond's pocket and the transceiver in his ear and the "smart blood" in his veins. (all that, to make a terrible pun). Also do not expect great trading or intricate combat. Even your overlord missions will sometimes be in another star system you don't even know about, never mind the faction missions. 'Protect person Y from planet A to B, but you haven't discovered planet A or B and person Y dies a minute later? Or things that have no purpose – paying or not paying your crew, the happiness stat is affected,

but does not do anything. Would you continue working if your boss did not pay you?

So if you just want to mess around and not think too much about the game, give it a spin, but I would say only at \$9, not \$19. As I said, the game has a LOT of potential, if they fix that and make it a bit more engaging, I would pay the \$19 with a smile.

As you can see from the screenshot the "help" is a bunch of big green question marks that march across the screen, and don't go away after reading. UI design, UI design, UI design...



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
 Bob C
 Brian Bogdan
 CBinMV
 Darren
 Dennis Mack
 Devin McPherson
 Doug Bruce
 Elizabeth K. Joseph
 Eric Meddleton
 Gary Campbell
 George Smith
 Henry D Mills
 Hugo Sutherland
 Jack
 Jason D. Moss
 Joao Cantinho Lopes
 John Andrews
 John Malon
 John Prigge
 Jonathan Pienaar
 JT
 Kevin O'Brien
 Lee Allen
 Leo Paesen
 Linda P
 Mark Shuttleworth

Moss Bliss
 Norman Phillips
 Oscar Rivera
 Paul Anderson
 Paul Readovin
 Rino Ragucci
 Rob Fitzgerald
 Robin Woodburn
 Roy Milner
 Scott Mack
 Sony Varghese
 Taylor Conroy
 Tom Bell
 Tony
 Vincent Jobard
 Volker Bradley
 William von Hagen

János Horváth
 Ronald Eike
 John Porubek
 Hans van Eekelen
 Kees Moerman
 Jon Loveless
 Jim Hibbard

SINGLE DONATIONS

2021:
 Floyd Smith
 Dale Reisfield
 Jan Ågren
 Linda Prinsen
 melvyn smith
 Frits van Leeuwen
 Raymond Mccarthy
 Robert Kaspar
 Frank Dinger
 Ken Maunder
 Brian Kelly

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#178

Deadline:
Sunday 06th Feb. 2022.
Release:
Friday 25th Feb. 2022.



Full Circle Team



Editor - Ronnie Tucker
ronnie@fullcirclemagazine.org

Webmaster -
admin@fullcirclemagazine.org

Editing & Proofreading

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

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

Getting Full Circle Magazine:

For the Full Circle Weekly News:



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



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



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



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



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

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