



Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

ISSUE #204 - April 2024




UBUNTU BUDGIE AND WUBUNTU REVIEWED

HowTo



Python p.26



Stable Diffusion p.30

Table 1: Demonstration Table

Municipality	Jan	Apr	July	Oct
Prince Rupert	5.6C	10.2C	16.2C	11.1C
Tofino	8.3C	11.9C	18.9C	13.1C
Nanaimo	6.9C	14.1C	23.9C	14.1C
Victoria	7.6C	13.6C	22.4C	13.1C

Latex p.33



... p.XX



Inkscape p.37

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.24




... p.XX



Linux News p.04




... p.XX




Ubuntu Devices p.47



The Daily Waddle p.41



Review p.59



My Story p.51



Letters p.XX



Review p.55



Q&A p.64



Micro This Micro That p.42



Ubuntu Games p.67



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

Another full house this month folks with Python, Micro This Micro That, Latex, Stable Diffusion, and Inkscape.

Our reviews this month are Ubuntu Budgie 23.10 and Wubuntu 11.4.3. Speaking of reviews, I'm sure you'll know by now that 24.04 LTS has been released. No doubt Adam is warming up his word processor and grabbing the ISO files as we speak. It's an LTS (Long-Term Support) release. So let's see if they've taken any chances or risks with 24.04. Probably not, but we'll see.

Elsewhere, we have some Q&A, a My Story about Balena Etcher, and a game review. Something to keep you busy while the weather changes. Hopefully. Maybe. Probably not where I am (UK).

More important than all of that: **this month marks our 17th birthday!** This little magazine only has another couple of years left before it's no longer a teenager. I just hope our little magazine doesn't end up buying booze, getting drunk and end up married.

Don't forget: we have a Table of Contents which lists every article from every issue of FCM. Huge thanks to **Paul Romano** for maintaining: <https://goo.gl/tpOKqm> and, if you're looking for some help, advice, or just a chinwag: remember that we have a **Telegram** group: <https://t.me/joinchat/24ec1oMFO1ZjZDc0>. I hope to see you there. Come and say hello.

All the best!

Ronnie

ronnie@fullcirclemagazine.org



This magazine was created using :



Find Full Circle on:



facebook.com/fullcirclemagazine



twitter.com/#!/fullcirclemag



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

Weekly News:



<https://fullcirclemagazine.org/podcasts/index.xml>



<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>



RELEASE OF SYSVINIT 3.09:

25/03/2024

A new release of the classic init system SysVinit 3.09 has been published, which was widely used in Linux distributions in the days before systemd and upstart, and now continues to be used in distributions such as Devuan, Debian GNU/Hurd and antiX. The versions of the insserv and startpar utilities used in conjunction with sysvinit have not changed. The insserv utility is designed to organize the boot process, taking into account dependencies between init scripts, and startpar is used to ensure the parallel launch of several scripts during the system boot process.

<https://github.com/slicer69/sysvinit/releases/tag/3.09>

RELEASE OF GNU EMACS

29.3:

25/03/2024

The GNU Project has published a new release of the GNU Emacs 29.3 text editor. Until the release of GNU Emacs 24.5, the project developed under the personal leadership of Richard Stallman, who handed over the post of project leader to John Wiegley in the fall of 2015. The project code is written in C and Lisp and is distributed under the GPLv3 license.

GNU Emacs 29.3 is being touted as an unscheduled emergency release with vulnerabilities fixed. According to preliminary data, the vulnerabilities allow code execution to occur when opening certain content or viewing emails with specially designed attachments in Gnus. The vulnerabilities in the list of changes are not yet detailed; there is only information about added protection methods:

<https://www.mail-archive.com/info-gnu@gnu.org/msg03269.html>

AFTER TRUENAS CORE 13.3, THE FREEBSD-BASED

BRANCH WILL BE PUT INTO

MAINTENANCE MODE:

25/03/2024

ixsystems has announced plans to further develop a distribution for the rapid deployment of network storage TrueNAS CORE, which continues development of the FreeNAS project. In June, they plan to release TrueNAS CORE 13.3 based on FreeBSD 13.3 (the previous release was formed in 2021 based on FreeBSD 13.0), OpenZFS 2.2.3 and Samba 4.19. Going forward, development will focus only on TrueNAS SCALE, which uses the Linux kernel and the Debian package base.

The TrueNAS CORE and TrueNAS SCALE projects will continue to coexist, but the TrueNAS CORE branch will be placed in maintenance mode, where they plan to fix bugs and security issues for many years to come. New

features and new versions of components will be developed in the TrueNAS SCALE branch. There are no plans to release TrueNAS CORE based on FreeBSD 14.

<https://www.truenas.com/blog/truenas-core-13-3-plans/>

STARTING WITH UBUNTU 14.04, SUPPORT FOR LTS

RELEASES WILL BE INCREASED TO 12 YEARS:

25/03/2024

Canonical has announced the introduction of a 12-year update period for LTS releases of Ubuntu. The extension will affect all current LTS releases, starting with Ubuntu 14.04, whose support was supposed to expire in April of this year, but has now been extended until April 2026. Updates for Ubuntu 24.04 will be generated until 2036, Ubuntu 22.04 - 2034, Ubuntu 20.04 - 2032, etc. Updates, as before, will be publicly available for 5 years after release, and the

remaining 7 years will be provided only to users of the Ubuntu Pro service.

Ubuntu Pro is a paid service for commercial use. For personal use, you can connect to Ubuntu Pro for free to receive extended updates for up to 5 machines. Official members of the Ubuntu Community can receive extended updates for up to 50 machines for free.

<https://canonical.com/blog/canonical-expands-long-term-support-to-12-years-starting-with-ubuntu-14-04-lts>

FEDORA 41 PLANS MOVE TO DNF5 PACKAGE MANAGER:

25/03/2024

In the Fedora 41 release, they propose moving the default distribution to the DNF5 package manager. Initially, the migration to DNF5 was planned for the Fedora 39 release, but the migration was delayed due to the unavailability of the tooling. Despite the fact that full parity in functionality with the old tools has still not been achieved, the developers believe

that the distribution is ready for migration, and the missing features can be implemented later. The proposal has not yet been considered by the FESCo (Fedora Engineering Steering Committee), which is responsible for the technical part of the development of the Fedora distribution. If the proposal is accepted, the dnf, libdnf and dnf-automatic packages will be replaced with the DNF5 toolkit and the new libdnf5 library, and the /usr/bin/dnf symlink will begin to point to the dnf5 executable.

The rework also made it possible to implement some improvements that improve the usability of the package manager. For example, the new DNF provides a more visual indication of the progress of operations; also added support for using local RPM packages for transactions; added the ability to display in reports on completed transactions information produced by scripts built into packages; as well as a more advanced input completion system for bash has been proposed.

<https://www.mail-archive.com/devel-announce@lists.fedoraproject.org/msg03249.html>

EXT2 CHANGES:

26/03/2024

A change has been made to the code base on which the Linux 6.9 kernel is based, moving the driver with the implementation of the Ext2 file system from the supported category to the deprecated category. The reason given is that the driver supports only 32-bit inode time counters, which will overflow on January 19, 2038.

To work with existing partitions, instead of the ext2 driver, dev's suggest you use the ext4 driver, which supports working with the Ext2 file system and is fully compatible with it, but can use timestamps in ext2 partitions that are not subject to the 2038 problem if the file system is created with an inode, larger than 255 bytes (in the ext2 driver, 32-bit time counters were used regardless of the inode size).

Ext2 users are advised to check the inode size being used. The mkfs.ext2 utility was moved to use 256-byte inodes by default as of

version 1.46.5 (shipped with Ubuntu 22.04), prior to which the command "mkfs.ext2 -I 256" could be used to increase the inode size. If 128-byte inodes are used, they should be converted to 256-byte ones:

```
$ sudo tune2fs -l /dev/sda1 | grep "Inode size"
```

```
Inode size: 128
```

```
$ sudo e2fsck -f /dev/sda1
```

```
$ sudo tune2fs -I 256 /dev/sda1
```

According to Theodore Ts'o, creator of the Ext4 file system, if someone misses the ext2 driver due to its simplicity (for example, the driver may be in demand to save resources on embedded systems), it is not too late to move support to it 64-bit time counters, which is not difficult.

<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id%3Db960e8093e7a57de98724931d17b2fa86ff1105f>

FEDORA LINUX 40

DISTRIBUTION ENTERS BETA

TESTING:

03/26/2024 23:58

Testing of the beta version of the Fedora Linux 40 has begun. The beta release marked the transition to the final stage of testing, where only critical bugs are fixed. The release is scheduled for April 23. The release covers Fedora Workstation, Fedora Server, Fedora Silverblue, Fedora IoT, Fedora CoreOS, Fedora Cloud Base, Fedora Onyx and Live builds, delivered in the form of spins with user environments KDE Plasma 5, Xfce, MATE, Cinnamon, LXDE, Phosh, LXQt, Budgie and Sway. Builds are generated for x86_64, Power64 and ARM64 (AArch64) architectures.

<https://fedoramagazine.org/announcing-fedora-linux-40-beta/>

RELEASE OF BUBBLEWRAP

0.9:

27/03/2024

After a year of development, a new version of Bubblewrap, v0.9, used to restrict individual applications of unprivileged users, has been published. In practice, Bubblewrap is used by the Flatpak project as a layer to isolate applications launched from packages. The project code is written in C and is distributed under the LGPLv2+ license.

For isolation, traditional Linux container virtualization technologies are used, based on the use of cgroups, namespaces, Seccomp and SELinux. To perform privileged operations to configure a container, Bubblewrap is launched

with root rights (an executable file with a suid flag) and then resets privileges after the container is initialized.

The key difference from the similar Firejail project, which also uses a setuid launch model, is that in Bubblewrap the container creation layer includes only the necessary minimum capabilities. All the advanced functions necessary for running graphical applications, interacting with the desktop and filtering requests to Pulseaudio, transferred to the Flatpak side and executed after the privileges have been reset. Firejail, on the other hand, combines all the related functions in one executable file, which makes it difficult to audit and maintain security at the proper level.

<https://github.com/containers/bubblewrap/releases/tag/v0.9.0>

UPDATE TO QUBES OS

4.2.1:

27/03/2024

Qubes 4.2.1, which implements the idea of using a hypervisor to strictly isolate applications and OS components (each class of applications and system services runs in separate virtual machines), is out. Recommend system has 16 GB of RAM (minimum 6 GB) and a 64-bit Intel or AMD CPU with support for VT-x with EPT/AMD-v with RVI and VT-d/AMD IOMMU technologies, preferably an Intel GPU (GPU NVIDIA and AMD are not well tested). The installation image size is 6 GB (x86_64).

The new release marks the update of versions of programs that form the basic system environment (dom0). A template has been prepared for creating virtual environments based on Fedora 39. By default, Linux 6.6 kernel is used, which has reduced the number of situations where it is necessary to install the latest kernel package on systems with new hardware.



DistroWatch.com

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

<https://www.qubes-os.org/news/2024/03/26/qubes-os-4-2-1-has-been-released/>

RELEASE OF TAILS 6.1:

27/03/2024

A release of a specialized distribution, Tails 6.1 (The Amnesic Incognito Live System), based on Debian 12, supplied with the GNOME 43 desktop and designed for anonymous access to a network, has been created. Anonymous access to Tails is provided by the Tor system. All connections other than traffic through the Tor network are blocked by the packet filter by default. Encryption is used to store user data in the save user data between runs mode. An iso image capable of working in Live mode, 1 GB in size, has been prepared for downloading .

There are new version updates of Tor Browser 13.0.13 (based on Firefox 115.9.1) and Thunderbird 115.9. A new microcode version has been added for Intel CPUs with the changes necessary to block the RFDS vulnerability. The persistent storage configurator ensures that

all enabled additional features are displayed.

<https://tails.net/>

SAMBA 4.20.0 RELEASE:

28/03/2024

After 6 months of development, a new release of Samba 4.20.0 was presented, that continues the development of the Samba 4 branch with a full implementation of a domain controller and Active Directory service, compatible with the Windows 2008. It is capable of serving all Microsoft-supported versions of Windows clients, including Windows 11. Samba 4 is a multifunctional server product that also provides implementation of a file server, print service and identification server (winbind).

<https://lists.samba.org/archive/samba-announce/2024/000663.html>

VPN LANEMU 0.11.6

RELEASED:

28/03/2024

Lanemu P2P VPN 0.11.6 was released - an implementation of a decentralized virtual private network that works on the Peer-To-Peer principle, in which participants are connected to each other, and not through a central server. Network participants can find each other through a BitTorrent tracker or BitTorrent DHT, or through other network participants (peer exchange). The application is a free and open analogue of Hamachi VPN, written in Java (with individual components in C) and distributed under the GNU LGPL 3.0 license.

<https://gitlab.com/Monsterovich/lanemu/-/releases/0.11.6>

A FORK OF THE REDIS:

28/03/2024

The Linux Foundation has announced the creation of the Valkey project, which will continue the development of the open source code base of the Redis DBMS, distributed under the BSD license. The project will be

developed under the auspices of the Linux Foundation on an independent platform with the involvement of a community of developers and companies interested in continuing to preserve the open source code base of Redis. Companies such as Amazon Web Services (AWS), Google Cloud, Oracle, Ericsson and Snap Inc. have joined the project.

<https://www.linuxfoundation.org/press/linux-foundation-launches-open-source-valkey-community>

RELEASE OF GNU

COREUTILS 9.5:

28/03/2024

A stable version of the set of basic system utilities GNU Coreutils 9.5 has been published , which includes programs such as sort, cat, chmod, chown, chroot, cp, date, dd, echo, hostname, id, ln, ls, etc.

<https://www.mail-archive.com/info-gnu@gnu.org/msg03270.html>

THE FURYGPU PROJECT DEVELOPS FPGA-BASED GPUs:

29/03/2024

A working prototype of the FuryGpu project was presented, a homemade GPU based on the Xilinx Zynq UltraScale+ FPGA, designed as a separate board connected to a PC via the PCIe interface. Descriptions of hardware units are implemented in the SystemVerilog language, and the board design is prepared in the free computer-aided design system for printed circuit boards KiCAD. In its current form, the FuryGpu GPU can already run the Quake game with a rendering performance of 60 FPS at a resolution of 1280 x 720.

The project has been in development since 2021 as a hobby by a programmer specializing in computer game development and interested in hardware. Board diagrams, Verilog descriptions and drivers are not yet publicly available, but the author of the project has announced his intention to open all developments related to FuryGpu after agreeing on some legal issues with his employer (approval is necessary to insure

oneself against possible claims from the employer).

At the current stage of development, the capabilities of FuryGpu are compared with an advanced video card from the mid-1990s. GPU functionality is defined by a set of fixed functions, software shader execution is not yet supported, and the focus is on rendering, rasterization, and texture processing components. To perform graphics operations, it provides its own graphics API, FuryGL, which is similar to the Vulkan API, but is not compatible with it due to the lack of shader support.

<https://www.furygpu.com/blog/hello>

A BACKDOOR WAS DISCOVERED IN THE XZ/LIBLZMA LIBRARY THAT ALLOWS ENTRY VIA SSHD:

29/03/2024

The XZ Utils package, which includes the liblzma library and utilities for working with compressed data in the ".xz"

format, has a backdoor (CVE-2024-3094) that allows interception and modification of data processed by applications associated with the liblzma library. The main target of the backdoor is the OpenSSH server, which in some distributions is bundled with the libsystemd library, which in turn uses liblzma. Linking sshd with a vulnerable library allows attackers to gain access to the SSH server without authentication.

The backdoor activation code was hidden in m4 macros from the build-to-host.m4 file used by the automake toolkit when building. During assembly, during the execution of intricate obfuscated operations based on archives (bad-3-corrupt_lzma2.xz, good-large_compressed.lzma), an object file with malicious code was generated, which was included in the liblzma library and changed the operation logic some of its functions. The m4 macros that activate the backdoor were included in the release tarballs, but were not in the Git repository. At the same time, malicious test archives were present in the repository, i.e. the person who implemented the backdoor had access to both the repository and

the release generation processes.

<https://www.openwall.com/lists/oss-security/2024/03/29/4>

DEBIAN 10 "BUSTER" MOVED TO ARCHIVE:

30/03/2024

Debian 10 "Buster" repositories are moved to archive.debian.org, after which the distribution will soon cease to be available through the main network of mirrors. Debian 10 packages are scheduled for deletion from mirrors or architectures for which there is no LTS support should be mid-April.

Debian 10 was released on July 7, 2019 and was maintained regularly until September 2022. As part of the LTS cycle, the release of updates for Debian 10 will be available for the i386, amd64, armhf and arm64 architectures until June 30, 2024. The extension of the Extended LTS program will be applied to the distribution, implying the publication of updates with the elimination of vulnerabilities in a limited set of Debian 10 packages for the amd64

architecture. Support for the Debian 10 Extended LTS will be supported in a separate repository until June 30, 2029.

<https://lists.debian.org/debian-devel-announce/2024/03/msg00003.html>

RELEASE OF NETBSD 10.0:

31/03/2024

A year and a half after the last update, a new release of the NetBSD 10 operating system was announced. Installation images of 630 MB, available in builds for 57 system architectures and 16 different CPU families, ready for download. The new branch includes a number of significant improvements, such as support for access control lists in the FFS file system, significant performance optimization, disk encryption using the Adiantum algorithm, and implementation in the driver core for WireGuard.

The project supports 8 primary ports that make up the core of the NetBSD development strategy: amd64, i386, ev.barm, evbmips, evbpmc, hpcarm, sparc64 and xen,

as well as 49 secondary ports associated with CPUs such as alpha, hppa, m68010, m68k, sh3, sparc and vax. Secondary ports are still supported, but have already lost relevance or do not have a sufficient number of developers interested in their development. One port (acorn26) is included in the third category, in which inoperable ports are placed, claiming to be removed, unless enthusiasts are interested in the development.

https://blog.netbsd.org/tnf/entry/netbsd_10_0_available

RELEASE OF WIDELANDS 1.2:

01/04/2024

After a year and a half of development, the free-to-play game Widelands 1.2 has been released, incorporating elements of a real-time strategy (RTS) and city-building simulator. The game is inspired by Settlers II and has similar gameplay. Both multiplayer mode and single-player campaigns with computer players are supported. The project code is written in C++ and Lua using the

SDL library and is distributed under the GPLv2 license. Builds are for Linux (ApplImage, Flatpak, PPA), Windows and macOS.

<https://www.widelands.org/news/2024/Mar/30/widelands-1-2-released/>

NITRUX 3.4.0 IS AVAILABLE:

01/04/2024

A new release of NitruX 3.4.0, built on Debian, KDE technologies and the OpenRC init system, has been published. The project offers its own desktop, NX Desktop, which is an add-on to KDE Plasma. Based on the Maui library, a set of standard user applications is being developed for the distribution that can be used on both desktop systems and mobile devices. To install additional applications, a system of self-contained ApplImages packages is being promoted. The full boot image is 3.4 GB in size. The project's developments are distributed under free licenses.

<https://nxos.org/changelog/release-announcement-nitruX-3-4-0/>

RELEASE OF THE ROSA FRESH 12.5:

03/04/2024

The company STC IT ROSA has announced a corrective release of the freely distributed and community-developed distribution kit ROSA Fresh 12.5, built on the rosa2021.1 platform. Builds designed for the x86_64 platform in versions with KDE Plasma 5, GNOME, Xfce, LXQt and without GUI have been prepared for free download. Users who already have the ROSA Fresh 12 distribution installed will receive the update automatically.

<https://rosa.ru/rosa-anonsiruet-novuyu-versiyu-besplatnoj-os-dlya-domashnego-ispolzovaniya-rosa-fresh-12-5/>

OPENING THE MOJO LANGUAGE STANDARD LIBRARY:

03/04/2024

The developers of the Mojo programming language have

begun converting the project's code into open source software. The first to open source was the standard library code, which is now available under the Apache 2.0 license with exceptions from the LLVM project that allow mixing with code under the GPLv2 license. In addition to publishing code, the development process has shifted towards openness and the ability to communicate third-party changes through submitting pull requests to GitHub. The source code of the compiler is planned to be opened after the design of the internal architecture is completed.

<https://www.modular.com/blog/whats-new-in-mojo-24-2-mojo-nightly-enhanced-python-interop-oss-stdlib-and-more>

LIBSYSTEMD DEPENDENCY REDUCTION INITIATIVE:

03/04/2024

Some developers of the systemd system manager, are discussing reducing the dependencies of the libsystemd library, which links not only to systemd components, but also to many external applications. For example, in Fedora, more than

150 packages use libsystemd in their dependencies. The initiator of the discussion believes that adding additional third-party libraries to libsystemd that are not controlled by systemd developers significantly increases the attack surface in the event of third-party libraries being compromised, as happened with the liblzma library.

In addition to liblzma and glibc, libsystemd also loads libzstd, liblz4 and libgcrypt, the security of which becomes critical. libsystemd provides access to 12 basic APIs (sd-bus, sd-daemon, sd-device, sd-event, sd-hwdb, sd-id128, sd-journal, sd-login, sd-netlink, sd-network, sd-path and sd-resolve) and a situation arises where an application, for example, using libsystemd only for the sake of calling the sd_notify function to inform systemd about a state change or sd_journal to write data to the log, links with all other libraries and API handlers. As a way out, they propose to split libsystemd into several separate libraries responsible for separate APIs, which will allow third-party dependencies to be loaded only where they are needed.

<https://github.com/systemd/systemd/issues/32028>

RELEASE OF REDICT 7.3.0:

03/04/2024

Drew DeVault, author of the Sway user environment, the Hare programming language, the Aerc email client and the SourceHut collaborative development platform, presented the release of the Redict 7.3.0 project, which is a fork of the Redis DBMS. Redict is forked from Redis 7.2.4, the latest version distributed under the BSD license. The changes developed by the Redict project are distributed under the LGPLv3.0 license, but the code borrowed from Redis remains under the BSD license.

Redict 7.3.0 is positioned as the first separate release of the project, the changes, so far are limited to rebranding and other minor changes, for example, the redis-cli utility has been renamed to redict-cli, and the configuration file /etc/redis.conf to /etc/redict.conf. Redict 7.3.0 is fully compatible with the Redis 7.2.4 release and can be used to replace it, but is not related to or compatible with the currently

in development Redis 7.3.0 branch.

<https://redict.io/posts/2024-04-03-redict-7.3.0-released/>

CENO 2.1 WEB BROWSER, USES A P2P NETWORK TO

ACCESS SITES:
03/04/2024

After a year of development, the release of the mobile web browser CENO 2.1.0 has been published. It is designed to allow access to information in conditions where direct access to said information is impossible, for example, when Internet segments are disconnected from the global network as a result of failures or actions of attackers. The browser is built on the GeckoView engine (used in Firefox for Android), enhanced by the ability to exchange data through a decentralized P2P network, where users participate in redirecting traffic to external gateways that provide access to information. The project's code is distributed under the MIT license. Ready-made builds are available on Google Play.

P2P functionality is included in a

separate library, Ouinet, which can be used in arbitrary applications. The CENO browser and the Ouinet library allow you to access information in situations up to a complete Internet shutdown, for example, as a result of global outages. In this case, content can be distributed from the cache or local storage.

<https://github.com/censorship-no/ceno-browser/releases/tag/v2.1.0>

GOOGLE INTRODUCED THE JPEGLI LIBRARY:

04/04/2024

Google has introduced a new open source jpegli library with an implementation of a JPEG image encoder and decoder. The library includes additional optimizations to improve encoding efficiency, allowing for up to 35% higher compression rates for high-quality images than traditional JPEG codecs. Compared to libjpeg-turbo, the jpegli library allows you to achieve a similar level of quality while reducing the bitrate by 32%. At the API and ABI level, the library is fully compatible with libjpeg62 and can be used to transparently

replace it. The library code is written in C++ and distributed under the BSD license.

<https://opensource.googleblog.com/2024/04/introducing-jpegli-new-jpeg-coding-library.html>

X.ORG SERVER 21.1.12 UPDATE WITH 4

VULNERABILITIES FIXED:
04/04/2024

Corrective releases of X.Org Server 21.1.12 and the DDX component (Device-Dependent X) xwayland 23.2.5 have been published, which ensures the launch of X.Org Server for the execution of X11 applications in Wayland-based environments. The new version of X.Org Server fixes 4 vulnerabilities. One vulnerability can be exploited for privilege escalation on systems where the X server is running as root, as well as for remote code execution in configurations that use X11 session redirection via SSH for access.

The remaining three vulnerabilities can lead to server memory leaks or crashes. These

three vulnerabilities are exploited by using a different byte order on the client side than on the server. In this regard, the new release offers the ability to block the connection of clients from systems with a different byte order, which is activated through the configuration parameter "AllowByteSwappedClients" or the command line option "+byteswappedclients".

<https://www.mail-archive.com/xorg-announce@lists.x.org/msg01706.html>

SCHLESWIG-HOLSTEIN AND OPEN SOURCE:

05/04/2024

The government of Schleswig-Holstein, a region in northern Germany, has approved the migration from Windows to Linux and from MS Office to LibreOffice on 30 thousand computers in various government agencies. To organize collaboration in the new infrastructure, Nextcloud, Open Xchange and Thunderbird will be used instead of Microsoft Sharepoint and Microsoft Exchange/Outlook, and instead of Active

Directory, a directory service based on open source software will be proposed. The transition to open source software will affect all computers of government employees, including workstations of teachers in schools. The decision was made after the successful completion of a pilot project that lasted two and a half years.

<https://blog.documentfoundation.org/blog/2024/04/04/german-state-moving-30000-pcs-to-libreoffice/>

RELEASE OF THE FFMPEG 7.0:

05/04/2024

After five months of development, the FFmpeg 7.0 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 LGPL and GPL licenses, FFmpeg development is carried out adjacent to the MPlayer project.

<http://ffmpeg.org/download.html%23releases>

OPENBSD 7.5 RELEASE: 05/04/2024

The release of the free UNIX-like operating system OpenBSD 7.5 is presented. The OpenBSD project was founded by Theo de Raadt in 1995. The full installation ISO image of the OpenBSD 7.5 base system is 630 MB. In addition to the operating system itself, the OpenBSD project is known for its components, which have become widespread in other systems and have proven themselves to be one of the most secure and high-quality solutions. Among them: LibreSSL (a fork of OpenSSL), OpenSSH, PF packet filter, OpenBGPD and OpenOSPF routing daemons, OpenNTPD NTP server, OpenSMTPD mail server, text terminal multiplexer (analogous to GNU screen) tmux, identd daemon with an implementation of the IDENT protocol, BSD alternative GNU groff package - mandoc, protocol for organizing fault-tolerant systems CARP (Common Address Redundancy Protocol), lightweight http server,

OpenRSYNC file synchronization utility.

<https://www.mail-archive.com/announce@openbsd.org/msg00522.html>

NETPLAN 1.0 IS AVAILABLE: 05/04/2024

After seven years of development, the release of the Netplan 1.0 toolkit has been published. Netplan provides parameter storage in YAML format and provides backends that abstract configuration access for NetworkManager and systemd-networkd. The project code is written in Python and C, and is distributed under the GPLv3 license.

In addition to employees from Canonical, employees from Microsoft and Deutsche Telekom also take part in the development of the project. Netplan has already been used in Ubuntu since the 17.10 release, is included in Ubuntu 23.10 Desktop as the default backend in NetworkManager, and is used by default to manage network settings in official Debian 12 cloud

builds. Netplan support is also implemented for cloud-init and the Calamares installer. The new release of Netplan 1.0 will be included in the Ubuntu 24.04 distribution and will ship with Debian 13.

Netplan 1.0 marks the first significant stable release, eliminating some long-deprecated features. Removing technical debt made the codebase easier to maintain and removed barriers that were preventing some new ideas from being implemented.

<https://canonical.com//blog/introducing-netplan-v1>

RELEASE OF KODI 21.0 OPEN MEDIA CENTER: 06/04/2024

After more than a year of development, the open media center Kodi 21.0, previously developed under the name of XBMC, released an update. The media center provides an interface for viewing Live TV and managing your collection of photos, movies and music. It supports navigation on TV shows, has an electronic cart

guide and does video recordings on schedule. Ready-made installation packages are available for Linux, FreeBSD, Raspberry Pi, Android, Windows, macOS, tvOS and iOS. The project code is distributed under the GPLv2+ license.

Initially, the project was aimed at creating an open multimedia player for the Xbox game console, but in the process of development it was transformed into a cross-platform media center running on modern software platforms. Based off Kodi, about a dozen commercial set-top boxes and several open projects (Boxee, GeeXboX, 9x9 Player, MediaPortal, Plex) are developing their thing.

<http://kodi.tv/>

UPDATES FROM THE GNOME PROJECT: 06/04/2024

The developers of GNOME summed up the work carried out in the project over the past week. The film manager of Nautilus (GNOME Files) has published plans to create a file selection interface (Nautilus.org.freedesktop.implport

al.FileChooser), which can be used in applications instead of the dialogs to open the files provided by GTK (GtFillChileChooserDialog). Compared to the implementation of the GTK in the new interface, the new interface will provide better integrated behavior and visual design, as well as elements and advanced features of the GNOME platform, such as libadwaita and wildcard selection.

<https://thisweek.gnome.org/posts/2024/04/twig-142/>

AV LINUX BASED ON MX-23.2:

07/04/2024

The new release of the AV Linux 23.2 distribution containing a selection of applications for creating/processing multimedia content, is out. The distribution is based on a package database of MX Linux and the KXStudio repository with a collection of applications for sound processing and additional packages of their own build (Polyphone, Shuriken, Simple Screen Recorder, etc.). The distribution can function in Live mode and is available for x86_64

(5.4 GB).

The Linux kernel in AV Linux comes with a set of Liquorix patches to increase the responsiveness of the system during work related to sound processing. A custom environment based on Enlightenment is supplied. The delivery includes Ardour, ArdourVST and Harrison Mixbus, Blender 3D design systems, video editors Cinelerra, Openshot, LIVES and tools for converting multimedia file formats. When installing, the option to use sysvinit or systemd is available. The distribution provided a detailed illustrated manual too.(PDF, 72 pages)

<https://www.bandshed.net/2024/04/05/av-linux-mx-edition-23-2-iso-update/>

PRESUMABLY 14 MILLION PEOPLE USE THE XFCE DESKTOP ENVIRONMENT:

07/04/2024

Alexander Schwinn, who is involved in the development of the Xfce desktop environment and Thunar file manager, tried to

calculate an estimated number of Xfce users. Assessing the popularity of the main Linux distributions, it was concluded that about 14 million users work in Xfce!!

https://alexcons.github.io/blogpost_9.html

BCACHEFS PATCHES TO FIX THE FS:

08/04/24

Kent Overstreet, developer of Bcachefs FS, offered patches that allow the Linux kernel to work with the Bcachefs FS even after damage to a significant amount of metadata, if necessary, rebuilding spoiled b-trees on metadata from inode and dirent structures. The changes were adopted by Linus Torvalds and included in today's test update of the 6.9-rc3 core.

The changes provide mounts of damaged FS and provide access to all or almost all data in the FS affected by the recent error that led to the destruction of the FS. The error occurred when, after the update, the toolkit initialised, but did not end, an attempt to update the structures of the FS, after which

the core with the old implementation of Bcachefs made an attempt to downgrade the version of the structures, ignoring the unfinished update. This situation led to the infamous "split brain mode," when none of the operations can be completed, and the structures of b-tree FS are no longer consistent with the real contents.

The mentioned patch allows you to continue working even in such cases.

<https://lore.kernel.org/lkml/wwkq7ugdewzde6gdej5bi6kb3bsv0qzqkexxejcl64d5r3pow@46qmmgq5wx4y/>

RELEASE OF GNU STOW 2.4:

08/04/24

Almost 5 years after the last release, GNU Stow 2.4 symbolic link management system has been presented, using symbolic links to separate directories from the contents of programs and related data. Stow may be useful for installing and integrating programs installed from source code or for

managing configuration files. The Stow code is written in Perl and distributed under the GPLv3 license.

Stow uses a simple approach to the placement of applications: all the components associated with the program are installed by the standard tools in a separate subdirectory in the `/usr/local/stow` directory (for example, `/usr/local/stow/emacs`). The Stow's task is to set up symbolic links in the hierarchy of system directories `/usr/local` and the control of their integrity.

When the package is deleted, all associated symbolic links are also deleted inside the `/usr/local/stow` directory, the files of the installed programs always remain unchanged, Stow operates only with symbolic links. In addition to managing system packages, it is possible to manage programs and configuration files (`/.bashrc`, `/.zshrc`, etc.) in the user's home directory.

<https://www.mail-archive.com/info-gnu@gnu.org/msg03275.html>

THE PUMPKINOS PROJECT DEVELOPS PALMOS

REINCARNATION:

08/04/24

The PumpkinOS project attempted to recreate the PalmOS operating system used in Palm communicators. PumpkinOS allows you to directly run applications created for PalmOS, without the use of the PalmOS emulator and without requiring the original firmware from PalmOS. Applications built for the m68K architecture can run on systems with x86 and ARM processors. The project code is written in C and is licensed by GPLv3.

The environment for running applications work in Linux and Windows, but can also be built as a self-sufficient downloadable OS (the system environment is based on the Linux kernel and the BusyBox package). A graphic shell is provided to navigate the available applications and switch between them. The delivery includes four basic applications, the code, transferred from the PalmOS SDK: address book, Todo-notque system, text editor MemoPad and calendar.

<https://github.com/migueletto/PumpkinOS>

RELEASE OF APT-MIRROR2 4:

08/04/24

The release of the 'apt-mirror2' 4 toolkit, designed to create local mirrors of apt-repositories of distributions based on Debian and Ubuntu, has been published. Apt-mirror2 can be used as a transparent replacement instead of the apt-mirror utility, which it has not been updated since 2017. The main difference in apt-mirror2 is the use of the Python language with the asyncio library (the original apt-mirror code was written in Perl), as well as the use of integrity checks at all stages of the mirror to prevent a mirror failure. The code is distributed under the GPLv3 license.

<https://gitlab.com/apt-mirror2/apt-mirror2/-/releases/v4>

PIVPN PROJECT

DISCONTINUED:

08/04/24

The developer of the PiVPN tool, designed to quickly set up a VPN server based on the Raspberry Pi, announced the publication of the final version 4.6, which summed up the 8 years of the project's existence. After the release of the repositories, it was transferred to archive mode, and the author announced the complete termination of the support of the project.

As a reason, a loss of interest in development is mentioned when the project has fulfilled its mission and has lost relevance in modern realities, as other tools that solve the problem are better. PiVPN also said that it does not intend to transfer the rights to those wishing to pick up the development due to the lack of credible candidates and the feeling that it has no right to decide who to transfer the project to. Those wishing to continue the development of PiVPN can create a fork and develop it under a different name.

It was implied that with the help

of PiVPN, any user without technical skills would be able to quickly create a VPN server with optimal security settings, running one command, and subsequently use the pivpn command line utility to add, delete and view clients.

<https://github.com/pivpn/pivpn/releases/tag/v4.6.0>

FIRST EDITION OF PINGORA:

08/04/24

Cloudflare has published the first release of the Pingora framework, designed to develop secure high-performance network services in the Rust language. A proxy was used for about a year in the content delivery network Cloudflare instead of nginx and processes more than 40 million requests per second, built with Pingora. The code is written in Rust and published under the Apache 2.0 license.

<https://github.com/cloudflare/pingora/releases>

FEDORA 41 APPROVED

DNF5:

09/04/24

The FESCo (Fedora Engineering Steering Committee), responsible for the technical part of the development of the Fedora Linux distribution, approved the default use of the DNF5 packet manager in the autumn release of Fedora 41. The dnf, libdnf and dnf-automatic packages will be replaced in Fedora 41 with the DNF5 toolkit and the new libdnf5 library, and the symbolic link /usr/bin/dnf will start pointing to the dnf5 executable.

Despite the fact that full parity in functionality, with the old tools is still not achieved, the developers believe that the distribution is ready for migration, and the missing opportunities can be implemented later. For example, the functionality associated with managing the transaction history underlying the "dnf history" command is not yet available. Also, work has not yet been completed to integrate support for the dnf5daemon background process in GNOME Software application manager.

<https://pagure.io/fesco/issue/3191#comment-904697>

RIVENDELL 4.2 IS AVAILABLE:

09/04/24

The Rivendell 4.2 project, that develops a digital sound processing platform, which can be used to automate the work of multimedia studios and professional radio broadcasting stations is out. Rivendell includes a set of components that solve tasks, such as ether grid planning, ordering and rotation of compositions, automatic inclusion of voice, export of data for the announcement of the current composition and the announcement of the next one, bundles. The code is written in C++ and is distributed under GPLv2 and LGPLv2 licenses.

<https://lists.linuxaudio.org/hyperkitty/list/linux-audio-announce@lists.linuxaudio.org/thread/WJFB554MUA2BHOZMYIANTDAG7WCESIC2/>

RELEASE OF NXS-BACKUP

3.4.0:

09/04/24

Nxs-backup 3.4.0, which allows you to create backups, rotate them and store them on local or external storages, has a new release. In addition to backup files, it can support creation backups of various DBMSs. The project code is distributed under the GPLv3 license. Starting with the nxs-backup 3.x branch, the project code is completely rewritten in Go (previously used by Python) to improve performance and functionality.

<https://github.com/nixys/nxs-backup/releases/tag/v3.4.0>

RELEASE OF OPENSSL

3.3.0:

09/04/24

After five months of development, the release of the OpenSSL 3.3.0 library with the implementation of SSL/TLS protocols and various encryption algorithms was formed. Support for OpenSSL 3.3 will be available until

April 2026. Support for the past OpenSSL 3.2, 3.1 and 3.0 LTS will last until November 2025, March 2025 and September 2026, respectively. The project code is distributed under the Apache 2.0 license

<https://www.mail-archive.com/openssl-announce@openssl.org/msg00445.html>

KUBUNTU PROJECT PRESENTED AN UPDATED LOGO AND BRANDING ELEMENTS:

10/04/24

The results of the competition among graphic designers to update the elements of branding of the distribution are summed up. As part of the competition, an attempt was made to get a recognizable and modern design, reflecting the specifics of Kubuntu, positively perceived by beginners and old users, and harmoniously combined with the style of KDE and Ubuntu.

As a result of the competition, recommendations were developed for the modernization of the

project logo, desktop screensaver, color palette and fonts. They also published vector files with a new logo, proposed in various colors. Based on new recommendations, the execution will happen in the Kubuntu 24.04 release.

<https://kubuntu.org/news/celebrating-creativity-announcing-the-winners-of-the-kubuntu-contests/>

RELEASE OF I2P 2.5.0 ANONYMOUS NETWORK:

10/04/24

Anonymous network I2P 2.5.0 and C++ client i2pd 2.51.0 was released. I2P is a multi-layered anonymous distributed network running on top of the conventional Internet, actively using end-to-end encryption that guarantees anonymity and isolation. The network is built in P2P mode and is formed thanks to the resources (passage) provided by network users, which allows you to do without the use of centrally managed servers (communications within the network are based on the use of encrypted unidirectional tunnels between the participant

and peer-s).

On the I2P network, you can anonymously create websites and blogs, send instant messages and emails, share files and organize P2P networks. For the construction and use of anonymous networks for client-server (sites, chats) and P2P (files, cryptocurrencies) applications, I2P clients are used. The basic I2P client is written in Java and can run on a wide range of platforms such as Windows, Linux, macOS, Solaris, etc. I2pd is an independent implementation of the I2P client in the C++ language and is distributed under a modified BSD license.

https://geti2p.net/en/blog/post/2024/04/08/new_release_i2p_2.5.0

PROJECT GENTOO AND SPI:

10/04/24

The developers of the Gentoo distribution announced the transition of the project to under that of the patronage of the non-profit organization SPI (Software in the Public Interest), which delegated solutions to non-technical tasks, which will allow

them to do without maintaining their legal entity, get rid of the administrative routine and increase the attractiveness for the transfer of donations by commercial companies, since such donations will give the opportunity to get a tax deduction in the United States.

Among other things, SPI will supervise the admission of donations, resolve legal issues, manage assets and trademarks, ensure the storage of the project, pay costs, conclude contracts, conduct audits and maintain accounting. Currently, there are 44 open projects under the auspices of SPI, including Debian, Arch Linux, LibreOffice, X.Org, systemd, O.A.D, PostgreSQL, FFmpeg, freedesktop.org, OpenWrt, OpenZFS, Jenkins and OpenEmbedded.

<https://www.gentoo.org/news/2024/04/10/SPI-associated-project.html>

OPENZITI 1.0:

11/04/24

The 1.0 release of the OpenZiti toolkit is presented, which

allows you to integrate tools for distributed network interaction and access to services in an isolated network environment. OpenZiti makes it possible to deploy a specialized network for the developed applications, running on top of the usual Internet and using mesh-routing, at which each point of the network is connected through neighboring points. The project code is written in Go and distributed under the Apache 2.0 license. API's are provided for the languages C, Go, Python, Swift, Java, Node.js and C#(NET), as well as tools for integration with applications for the Android platform.

The OpenZiti system is originally designed to build networks that guarantee the protection and isolation of traffic in unprotected environments or nodes in which can be compromised (architecture with zero trust). The technology can be used to communicate with applications instead of a VPN. The operation of the network is provided through the use of the controller process, which manages the configuration of the network, authentication and configuration of services, as well as nodes that are allowed to act as routers that form

a mesh network and perform transit traffic transfer through themselves. You can add services to balance the load and ensure fault tolerance.

<https://blog.openziti.io/announcing-openziti-v1>

OS ZONE INITIATIVE REVOKES 54 ANTI-OPEN SOFTWARE PATENTS:

11/04/24

The Open Invention Network (OIN) and Unified Patents announced the successful development of a joint initiative, the Open Source Zone (OS Zone), created in 2019 following the patent aggression of the Rothschild Patent Imaging company against the GNOME project. Since then, in addition to the patent that the GNOME developers were accused of violating, OS Zone participants were able to have 53 more patents invalidated, one way or another aimed at open source software.

Some of the open projects whose interests were protected by having the patents associated with them invalidated, mention is made

of Apache Cocoon, Apache mod_evasive, Argo, Bluez, Fail2Ban, FreeMesh, LibreMesh, iFolder, LibVLC, Linkerd, Linux kernel, Automotive Linux, Kubernetes, KVM, Magento, Mastodon, Firefox, Mycroft, OpenACH, OpenSwan, QEMU and WebM.

<https://openinventionnetwork.com/founding-open-source-zone-members-oin-the-linux-foundation-and-microsoft-mark-five-years-of-successes-in-protecting-open-source-software-from-patent-assertion-entities/>

ORACLE HAS PUBLISHED DTRACE 2.0.0-1.14 FOR LINUX

11/04/24

An experimental release of the DTrace dynamic debugging toolkit for Linux 2.0.0-1.14 is out, implemented as a user-space process that uses the eBPF subsystem and the standard tracing mechanisms provided by the Linux kernel. In terms of functionality, the eBPF-based DTrace implementation is close to the first DTrace implementation for Linux, implemented in the form of a

kernel module. The project code is distributed under the GPLv2 license.

The toolkit can be used with stock Linux kernels that support BPF. To use it, you need the libctf library with support for the CTF (Compact Type Format) debug format, which is included in the binutils package starting with release 2.40, or the libdtrace-ctf library, ported from Solaris. Optionally, two patches are offered for the 6.7 kernel, which allow you to use advanced features to obtain additional data about the modules and the kernel.

<https://lore.kernel.org/all/ZhBRSM2j0v7cOLn%252F@oracle.com/T/%23u>

RELEASE OF ARDOR 8.5:

12/04/24

A new release of Ardor 8.5 has been published, designed for multi-channel recording, processing and mixing of sound. Ardor provides a multi-track timeline, an unlimited level of rollback of changes throughout the entire process of working with a file (even

after closing the program), and support for a variety of hardware interfaces. The program is positioned as a free analogue of professional tools ProTools, Nuendo, Pyramix and Sequoia. The code is distributed under the GPLv2 license. In the near future, unofficial builds for Linux will be generated in Flatpak format .

<https://ardour.org/whatsnew.html>

UBUNTU 24.04 BETA

RELEASE:

12/04/24

The beta release of the Ubuntu 24.04 "Noble Numbat" distribution was presented, after the formation of which the package database was completely frozen, and the developers moved on to final testing and bug fixes. The release is scheduled for April 25. The 24.04 branch is classified as a long-term support (LTS) release, where updates are generated over a period of 12 years (5 years - generally available, plus another 7 years for users of the Ubuntu Pro service). Ready-made test images are created for Ubuntu, Ubuntu Server , Lubuntu , Kubuntu , Ubuntu

Mate , Ubuntu Budgie , Ubuntu Studio , Xubuntu , UbuntuKylin (China edition), Ubuntu Unity , Edubuntu and Ubuntu Cinnamon .

<https://fridge.ubuntu.com/2024/04/12/ubuntu-24-04-lts-noble-numbat-beta-released/>

RELEASE OF KDE FRAMEWORKS 6.1.0:

13/04/24

A month and a half after the release of KDE 6.0, the release of the KDE Frameworks 6.1.0 platform was published, providing a restructured and ported to Qt 6 basic set of libraries and runtime components that underlie KDE. The framework includes 72 libraries, some of which can work as self-contained add-ons to Qt, and some of which form the KDE software stack. The release follows an initiative to deliver predictable monthly updates to the KDE Frameworks.

<https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://kde.org/announcements/frameworks/6/6.1.0/>

LAKKA 5.0:

14/04/24

Lakka 5.0, which allows you to turn computers, TV consoles or single-board computers into a full-fledged game consoles for the launch of retrogames, is out. The project is a modification of the LibreELEC distribution, originally designed to create home theaters. Lakka builds are for i386, x86_64 platforms (GPU Intel, NVIDIA or AMD), Raspberry Pi, Orange Pi, Banana Pi, Hummingboard, Cubox-i, Odroid C1/C1+/XU3/XU4, etc. To install, simply write the distribution to an SD card or USB drive, connect the gamepad and boot the system.

Lakka is based on the emulator, RetroArch, providing emulation of a wide range of devices and supporting advanced capabilities such as multiplayer games, save states, improving the image quality of old games with shaders, rewinding the game back, hot connection of gamepads and video streaming. Among the emulated consoles are Atari 2600/7800/ Jaguar/Lynx, Game Boy, Mega Drive, NES, Nintendo 64/DS,

PCEngine, PSP, Sega 32X/CD, SuperNES, etc. Gamepads from existing game consoles are supported, including Playstation 3, Dualshock 3, 8bitdo, Nintendo Switch, Xbox 1 and Xbox360.

<https://lakka.tv/articles/2024/04/13/lakka-5.0/>

RELEASE OF LUTRIS 0.5.17:

14/04/24

Lutris 0.5.17, which provides tools to simplify the installation, configuration and management of games in Linux, has been published. The project code is written in Python and distributed under the GPLv3 license. Ready-made builds are available in flatpak format.

The project supports a directory for quick search and installation of games, allowing you to run Linux games through a single interface with a single click, without worrying about installing dependencies and settings. Runtime components for launching games are supplied by the project and are not linked to the distribution used. The runtime is a distribution-based library suite that includes components from

SteamOS and Ubuntu, as well as various additional libraries.

It is possible to install games distributed through GOG, Steam, Epic Games Store, Battle.net, Amazon Games, Origin and Uplay services. At the same time, Lutris itself acts only as an intermediary and does not sell games, so for commercial games, the user must independently purchase the game in the corresponding service (free games can be run with one click from the Lutris graphical interface).

Each game in Lutris is tied to the download script and the handler describing the environment to run the game. Including ready-made profiles with optimal settings for launching games under Wine. In addition to Wine, the games can be run using emulators of game consoles, such as RetroArch, Dosbox, FS-UAE, ScummVM, MESS/MAME and Dolphin.

<https://github.com/lutris/lutris/releases/tag/v0.5.17>

THE SERVO BROWSER ENGINE & MOZILLA SPIDERMONKEY:

15/04/2024

The developers of the Servo browser engine presented an initiative to strengthen modularity and move away from the use of low-level APIs of the SpiderMonkey JavaScript engine, developed by Mozilla and used in Servo to provide support for JavaScript and WebAssembly. In the future, Servo plans to move to a higher-level API and an abstraction layer for interacting with JavaScript engines, which will make it possible to get rid of unsafe direct calls to SpiderMonkey code performed in unsafe blocks. "In the long term, the change will allow us to move away from being strictly tied to SpiderMonkey and provide support for other JavaScript and WebAssembly engines, such as the V8 engine used in Chrome".

The Servo project is written in Rust and is distinguished by its support for multi-threaded rendering of web pages, parallelization of operations with the DOM (Document Object Model) and the use of safe programming

mechanisms provided by the Rust language. The Servo project was originally developed by Mozilla, but then came under the wing of the Linux Foundation. Servo is natively designed to support breaking DOM and rendering code into smaller subtasks that can run in parallel and make more efficient use of multi-core CPU resources. Firefox already integrates some parts of Servo, such as the multi-threaded CSS engine and the WebRender rendering system.

<https://servo.org/blog/2024/04/15/spidermonkey/>

DESCENT 3 OPEN SOURCE:

16/04/2024

Kevin Bentley, one of the developers of the game 'Descent 3', got the management of Outrage Entertainment to open the source code of the project. Kevin, who has taken charge of the support of the new project, is recruiting a team of enthusiasts to revive and continue the development of the game. The code is written in C++ and is open under the MIT license.

Descent 3 was released in 1999. The game combines a first-person shooter and a flight sim in open and closed spaces. Single-player mission mode and multiplayer battles are supported. The latest Descent 3 "1.5 Patch" codebase has been published, including changes prepared several years ago to update code written in the 1990s.

The delivery does not include music and libraries for processing ACM and MVE audio and video formats, as Outrage Entertainment does not own the rights. To ensure a correct build of the project, missing libraries will be replaced with stubs for now. In the future, the videos used in the intro and final scenes are planned to be converted into a format for which open libraries are available. The tasks that they intend to solve first include providing the ability to build for different platforms, creating an infrastructure for testing in a continuous integration system, cleaning the code and removing components of the old version control system.

<https://translate.google.com/website?sl=auto&tl=en&hl=en-US&u=https://github.com/kevinbentley/Descent3>

RELEASE OF GNU TALER 0.10:

16/04/2024

After a year and a half of development, the GNU Project has released GNU Taler 0.10, a free electronic payment system that provides anonymity for buyers, but retains the ability to identify sellers for transparent tax reporting. The system does not allow tracking of information about where the user spends money, but provides tools for tracking the receipt of funds (the sender remains anonymous), which solves the problems inherent in BitCoin with tax audits. The code is written in Python and distributed under the AGPLv3 and LGPLv3 licenses.

The development is funded by grants from the European Commission, the Swiss State Secretariat for Education and the Swiss State Secretariat for Research and Innovation (SERI). Within the framework of the NGI TALER project, work is underway to create a product based on GNU Taler, ready for use in the European Union.

<https://taler.net/en/news/2024-06.html>

THE GENTOO PROJECT HAS BANNED THE ADOPTION OF CHANGES PREPARED USING AI TOOLS:

17/04/2024

The Gentoo Linux distribution's governing board has approved rules that prohibit Gentoo from accepting any content created using AI tools that process natural language queries, such as ChatGPT, Bard, and GitHub Copilot. Such tools should not be used when writing Gentoo component code, creating ebuilds, preparing documentation, or submitting bug reports.

<https://www.mail-archive.com/gentoo-dev@lists.gentoo.org/msg99327.html>

VIRTUALBOX 7.0.16

RELEASED:

17/04/2024

Oracle has published a bug-fix release of VirtualBox 7.0.16, which contains 15 fixes. In addition to these changes, the new version eliminates 13 vulnerabilities, 7 of which are marked as dangerous (four problems have a danger level of 8.8 out of 10, and three have a danger level of 7.8 out of 10).

Details about the vulnerabilities are not disclosed, but judging by the severity level set, they allow access to the host environment from guest systems. Two vulnerabilities appear only on Linux hosts and two only on Windows hosts. One of the vulnerabilities allows a remote attack via HTTP without authentication, but the severity level for this problem is set at 5.9 out of 10 due to the complexity of exploitation

<https://www.mail-archive.com/vbox-announce@virtualbox.org/msg00230.html>

LXQT 2.0.0 DESKTOP

ENVIRONMENT AVAILABLE:

17/04/2024

A new release of the LXQt 2.0.0 desktop environment (Qt Lightweight Desktop Environment), which continues the development of the LXDE and Razor-qt projects, has been presented. The LXQt interface follows the ideas of the classic desktop layout, but introduces modern design and techniques that increase ease of use. LXQt is positioned as a lightweight, modular, fast and convenient environment that incorporates the best features of LXDE and Razor-qt. The code is hosted on GitHub and is licensed under GPL 2.0+ and LGPL 2.1+. Ready builds are expected for Ubuntu (LXQt is offered by default in Lubuntu), Arch Linux, Fedora, openSUSE, Mageia, FreeBSD, ROSA and ALT Linux.

<https://lxqt-project.org/release/2024/04/15/release-lxqt-2-0-0/>

RELEASE OF VIDEO PLAYER MPV 0.38:

18/04/2024

The open source video player MPV 0.38 was originally released in 2013, a fork from the code base of the MPlayer2 project. MPV focuses on developing new features without worrying about maintaining compatibility with MPlayer. The MPV code is licensed under LGPLv2.1+, some parts remain under GPLv2, but the transition to LGPL is almost complete and the "--enable-lgpl" option can be used to disable the remaining GPL code.

<https://github.com/mpv-player/mpv/releases/tag/v0.38.0>

LINUS TORVALDS SPOKE OUT AGAINST KCONFIG PARSERS THAT DO NOT SUPPORT TABS:

18/04/2024

Linus Torvalds refused to accept the kernel change to replace the tab character with a space in the delimiter of the FTRACE_RECORD_RECURSION_SIZE parameter in the Kconfig kernel

configuration. The change was proposed by the Fedora project developer with the note that using tabs causes the configuration parser to break. Instead of the proposed change, Linus included his own patch in the kernel, which specifically adds tab characters to the definition of the PAGE_SHIFT setting, which sets the offset for various sizes of memory pages.

Adding tabs to settings greater than FTRACE_RECORD_RECURSION_SIZE is explained by the fact that the kernel settings file allows the use of both spaces and tabs, so if the parser cannot parse a string with tabs normally, this is a parser problem that should be fixed in it. The "make defconfig" command accepts tabs correctly, so external parsers should process them too.

<https://lore.kernel.org/lkml/CAHk-%3Dwqg2NW5tar-Xew614JZPKfyTdet5fC0mqwK%2B2sUsZ0EkW@mail.gmail.com/>

MIRACLE-WM 0.2:

19/04/2024

Matthew Kosarek from Canonical presented the second release of the miracle-wm composite manager, which uses the Wayland protocol and components for building Mir composite managers. Miracle-wm supports tiling windows in the style of the i3 and Sway window managers. Waybar can be used as a panel. The project code is written in C++ and distributed under the GPLv3 license. Current builds are generated in snap format.

The goal of the project is to create a composite server that uses tiled windowing, but is more functional and stylish than projects like Swayfx. Miracle-wm is expected to be useful for users who prefer a tiled layout but want visual effects and a more vibrant graphic design with smooth transitions and colors. To install miracle-wm you can use the command "sudo snap install miracle-wm --classic".

<https://discourse.ubuntu.com/t/release-v0-2-0-of-miracle-wm-a-wayland-compositor-built-on-mir/44322>

NEW DEBIAN PROJECT LEADER ELECTED:

20/04/2024

The results of the annual Debian project leader election have been announced. The title went to Andreas Tille, who was elected for the first time. This year, two participants competed for the position of leader. Jonathan Carter, who served as leader for the past four years, did not participate in the election. 362 developers took part in the voting, which is 36% of all participants with voting rights (last year the turnout was 28%, the year before 34%, the historical maximum in 2000 was 62.25%, the minimum in 2016 was 27.56%).

Andreas Tille has been involved in maintaining packages for Debian for over 25 years and is the author of the Debian Med project, aimed at optimizing the distribution for solving problems related to medicine and biological research. In addition to development, Andreas also provides mentoring, teaching newcomers who want to join development. Andreas maintains 1,591 packages and ranks first in

the project in terms of the number of packages downloaded (since 1998, he has downloaded 17,254 packages), and also ranks 8th in the number of bugs fixed (5,870 fixes).

The second contender for the leadership post this year was Sruthi Chandran, a Debian developer from India who joined development in 2016 and is involved in maintaining 198 packages, including packages with libraries for Ruby, Node.js and Go, and also font packs.

<https://lists.debian.org/debian-vote/2024/04/msg00063.html>

THUNDERBIRD & MICROSOFT EXCHANGE:

20/04/2024

The developers of the Thunderbird email client announced the start of integration of components written in Rust into the code base. The next major release of Thunderbird, expected to be released in July of this year, will include code implementing Microsoft Exchange Web Services (EWS) mail protocol, implemented

in Rust. Support for accessing the Microsoft Exchange calendar and address book will be added at a later date. The built-in implementation will allow you to do without installing third-party add-ons, through which support for Microsoft Exchange has so far been provided.

Disadvantages of integrating Rust support include, an increase in the code base, the need to create missing bindings and the adaptation of some asynchronous handlers that are not compatible with the model of asynchronous operations in Rust to work with Rust code.

<https://blog.thunderbird.net/2024/04/adventures-in-rust-bringing-exchange-support-to-thunderbird/>

RELEASE OF THE NIRI 0.1.5:

21/04/2024

There was a release of the composite server Niri 0.1.5, developed under the GNOME-extension PaperWM and implements the method of mosaic (tiling) layout, in which the windows

are grouped into an infinitely scrolled tape on the screen. Opening a new window leads to an extension of the tape, and previously added windows never change their size. The project code is distributed under the GPLv3 license. Packages are collected for Fedora, NixOS, Arch Linux and FreeBSD.

Niri is based on the Wayland protocol, but allows you to run X11 applications using the Xwayland DDX server. There is a built-in interface for creating screenshots and recording screencasts, a feature of which has the ability to exclude individual windows with confidential information, from recording.

<https://github.com/YaLTeR/niri/releases/tag/v0.1.5>

RELEASE OF KATA CONTAINERS 3.4:

21/04/2024

A new release of the Kata Containers project, v3.4, has been published. They develop a stack for the organization of containers using insulation, based on full-

fledged virtualization mechanisms. The project was created by Intel and Hyper by combining Clear Containers and runV technologies. The project code is written in Go and Rust, and is distributed under the Apache 2.0 license. The development of the project is supervised by a working group created under the independent organization, OpenInfra Foundation, which includes companies such as Canonical, China Mobile, Dell/EMC, EasyStack, Google, Huawei, NetApp, Red Hat, SUSE and ZTE.

<https://github.com/kata-containers/kata-containers/releases/tag/3.4.0>

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)



This issue I'd like to talk more about firewalls. This all stems from a few Telegram chats I've had with some of you. Let me start with how to see existing firewall rules. Usually, when working with firewall rules, you need to be root. So remember to use sudo (see image bottom left).

`sudo iptables -L`

You should see an INPUT, a FORWARD and an OUTPUT. What you are seeing is the default filter, so it would be the same as typing: `sudo iptables -t filter -L` (see image bottom right).

I am not showing you guys the whole thing, as I have docker

installed and that adds another section and I do not want to confuse you, so you do yours and we work from there.

Because I'm in a virtual machine here, I use NAT, network address translation. We can look at that table too. It differs in that you should see the standard INPUT, FORWARD and OUTPUT, as well as PREROUTING and POSTROUTING. I've shown you how to check a table name (we checked the table named filter), so please go ahead and check the table named nat. Well done if you did it! In my book you are no longer a total newbie. If you did not manage to figure it out, take the previous command and

```
ed@eddv: ~$ sudo iptables -t nat -L
Chain PREROUTING (policy ACCEPT)
target     prot opt source                destination
DOCKER    all  --  anywhere              anywhere
E match dst-type LOCAL

Chain INPUT (policy ACCEPT)
target     prot opt source                destination
DOCKER    all  --  anywhere              !localhost/8
E match dst-type LOCAL

Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination
DOCKER    all  --  anywhere              anywhere
E match dst-type LOCAL
```

substitute the word "filter" for "nat". If you had set up port forwarding, you would see PREROUTING containing some address and port. I do not have anything like that set up. If you are a gamer, chances are you do (see image above).

We are not doing an in-depth look here, I just want newbies to recognise the patterns. When you look up the man page for iptables, you may notice a few more tables listed. Go ahead and try out these tables and see what you end up

```
ed@eddv: ~$ sudo iptables -L
[sudo] password for edd:
Chain INPUT (policy ACCEPT)
target     prot opt source                destination
ACCEPT    tcp  --  anywhere              anywhere
:domain /* generated for Multipass network mpqemubr0 */
ACCEPT    udp  --  anywhere              anywhere
:domain /* generated for Multipass network mpqemubr0 */
ACCEPT    udp  --  anywhere              anywhere
:bootps /* generated for Multipass network mpqemubr0 */

Chain FORWARD (policy DROP)
target     prot opt source                destination
```

```
ed@eddv: ~$ sudo iptables -t filter -L
Chain INPUT (policy ACCEPT)
target     prot opt source                destination
ACCEPT    tcp  --  anywhere              anywhere
:domain /* generated for Multipass network mpqemubr0 */
ACCEPT    udp  --  anywhere              anywhere
:domain /* generated for Multipass network mpqemubr0 */
ACCEPT    udp  --  anywhere              anywhere
:bootps /* generated for Multipass network mpqemubr0 */

Chain FORWARD (policy DROP)
target     prot opt source                destination
ACCEPT    all  --  anywhere              anywhere
```


with and if you can make any sense of it. Trust me on this, doing is best. Use your VM so you don't break anything by accident.

I trust that you noticed that each of these has (policy ACCEPT) in parenthesis. This means that if a rule is not in our list, this is the elif. In other words, a fallback. Now a lot of the time, when newbies do not understand firewall rules, they start with (policy DROP). Don't do that unless you are an advanced Linux user. Something I often see if I scroll through the command history is: `sudo iptables -A INPUT -p tcp --dport 22 j DROP`

What this usually means, is that someone new to firewalling, has read up on Fail2ban or some such and copied commands without understanding what they mean.

Let's quickly dissect that line, shall we? The `sudo iptables` is a given, the `-A` is for adding a rule, the `-p` is for the protocol, in this case, `tcp` and `dashdash` (Office packages and two dashes are not friends) `dport` is for destination port, in this case `22`, then the `j` is for jump to and the action is `DROP`. This means that all `tcp` traffic on port `22` will now be dropped, but

not `udp` traffic. This means that you may now notice that you can no longer transfer files via `sctp`. It may all be good, until 5 months down the line, you need to transfer a file and now you cannot. (Now comes the swearing like you saw in the previous issue. This is why companies have change requests that need to be signed off and documented, as you can see how this can become a huge problem, when multiple people are involved.)

As we have demonstrated and you hopefully mimicked, you have realised that it is not hard to understand the syntax, and it is not hard to add rules, but you also need to understand what the rule you just implemented does or what it affects.

On the other side of the coin, if you set your default policy to `DROP`, and you work on allow lists, all it takes, is someone to accidentally flush your rules. It is as simple as: `sudo iptables -F`. Though this clears all the rules, it does not reset the defaults back to the original defaults. This is another issue I have come across with novices, hence why I said not to start with `DENY` unless you understand 100% what you are

doing. This is really, really important when working with remote machines, that you do not have physical access to. I have had to drive like an hour to the data centre enough times to know not to do this. Making an appointment to see your own server with blood and urine samples is not fun at all, specially if someone else did this. Victor, if you ever read this, know I'm referring to you... LOL. If all of this was not scary enough as it is, realise that there are dynamic ports and you will have to check RELATED traffic as well, but that is a topic for another day. Internalise this chunk we talked about, it isn't much, but without training your eyes and your brain, you will miss the small things and firewall configurations are all about the small things.

TIP: If you are learning, my advice to you is, play on a virtual machine on your system first and do whatever you need to and test if you can access it remotely, before mirroring your changes on a remote server. Set your history to 1000 lines, if you accidentally flush your rules, you can add them from history. Train your eye to look for `DENY` default rules. Remember that a server serves, it cannot do

that if you deny all traffic. Sure, the intruders cannot get in, but neither can your clients. (if you would like to see this in action, look at the Digital Ocean forums, where devs have locked themselves out of their VPS's, every other day!)

If you think we helped you wrong, misc@fullcirclemagazine.org

Ref: man iptables

SUGGESTED READING

<https://www.geeksforgeeks.org/50-common-ports-you-should-know/>



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



THE MYSTERIES OF THE SPHINX

Greetings fellow Beings. Beaming from somewhere in time and space, I come again to, hopefully, provide some well needed information.

This month, I will present something that I've wanted to do for the last 8 months. As always, something would always get in the way and require my attention. This month, I purposefully refused to let anything take precedent over this topic.

As you can see from this month's title, it has to do with Sphinx. Not the massive limestone structure in Giza, but the Python software package used to create documentation.

Sphinx uses reStructured Text (rst) files to create HTML and, if desired, EPUB files, that can be easily included in your project.

So let's jump right in and see

how to get started.

INSTALLING SPHINX

First, the people at Sphinx suggest you use a virtual environment to install into. So, here are the steps. First, create the virtual environment and activate it. I usually create a directory in my project folder called "sphinx" and start from there. Open a terminal and use these steps.

```
$ python -m venv .venv
```

```
$ source .venv/bin/activate
```

Next, use pip to install it into the virtual Python instance.

```
(.venv)$ pip install sphinx
Collecting sphinx...
```

This takes only about a minute on my machine. Now check the version.

```
(.venv) $ sphinx-build --version
sphinx-build 7.2.6
```

Ok. Now, run the quickstart script which will ask questions

about your project, then create the folders and the base files for you.

```
(.venv)$ sphinx-quickstart docs
Welcome to the Sphinx 7.2.6 quickstart utility.
```

Please enter values for the following settings (just press Enter to accept a default value, if one is given in brackets).

```
Selected root path: docs
```

You have two options for placing the build directory for Sphinx output. Either, you use a directory "_build" within the root path, or you separate "source" and "build" directories within the root path.

```
> Separate source and build directories (y/n) [n]: y
```


```
The project name will occur in several places in the built documentation.
> Project name: sandbox
> Author name(s): Greg Walters
> Project release []: 0.1
```

If the documents are to be written in a language other than English, you can select a language here by its language code. Sphinx will then

translate text that it generates into that language.

```
For a list of supported codes, see
https://www.sphinx-doc.org/en/master/usage/configuration.html#confval-language.
> Project language [en]:
```

Here is what the directory tree looks like.



Now, change to the /docs/source/ folder. There are two files there that we will eventually need to edit. First is the index.rst file and the second is the conf.py file. If you plan to use any graphics files like screenshots of your applications, create an "images" folder in the source folder as well. So let's change to the /docs/source folder.

```
(.venv)$ cd docs/source
```

Now, let's look and see what the index.rst file contains.

```
(.venv)$ cat index.rst
```

```
.. sandbox documentation
master file, created by
sphinx-quickstart on Sat
Dec 16 10:04:46 2023.
You can adapt this file
completely to your liking,
but it should at least
contain the root `toctree`
directive.
```

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

```
.. toctree::
: maxdepth: 2
: caption: Contents:
```

It's basic but that's all we need to get started.

Notice at the top of the file is a line that starts with two dots. That is a commented line, so it will be ignored by Sphinx when creating the HTML file.

Then at the bottom of the file are three lines that will be used to create the table of contents.

The two dots that start each of these lines is called a directive.

Anytime you want to do something special, outside of bold or italic markings, which are the same commands as markdown (.md), so to set a word or phrase as bold you would wrap the phrase with two asterisks and italics would be just one.

We'll look at more directives in a little bit.

Now, to create an HTML version of the documentation for our project, use the sphinx-build script, at least the first time. This should be done in the folder that you originally set up, which, in my case, was /sphinx. This tells Sphinx where the source folder is and the build folder. The build folder will hold our finished HTML files. So change back to the sphinx (parent of everything) folder.

```
(.venv)$ sphinx-build -M html
docs/source/ docs/build/
```

```
Running Sphinx v7.2.6
making output directory...
done
building [mo]: targets for 0
po files that are out of date
writing output...
building [html]: targets for
1 source files that are out
of date
updating environment: [new
config] 1 added, 0 changed, 0
removed
```

```
reading sources... [100%]
index
looking for now-outdated
files... none found
pickling environment... done
checking consistency... done
preparing documents... done
copying assets... copying
static files... done
copying extra files... done
done
writing output... [100%]
index
generating indices...
genindex done
writing additional pages...
search done
dumping search index in
English (code: en)... done
dumping object inventory...
done
build succeeded.
```

```
The HTML pages are in docs/
build/html.
```

That's it. Now let's look at getting it to actually create something. We'll have to modify the index.rst file for that.

EDITING THE SPHINX DOCUMENT

After trying many markdown editors, my favorite one to use is formiko. Formiko is included in most Linux distributions, but you can also find it at <https://github.com/ondratu/formiko>.

This editor ticks all the boxes for me to be my 'go to' .rst editor. Automatic spell checking, syntax highlighting, it's all there. It also handles .md files.

THE CONFIG.PY FILE

For the most part, the only thing that I've ever had to modify here is the next to last line, which sets the theme file for the output (we'll discuss themes below). This line is by default...

```
html_theme = 'alabaster'
```

If you want a different built-in theme file, just change the word 'alabaster' to the theme of your choice. For many of the other themes, it is also this simple.

IMAGES IN YOUR SPHINX DOCUMENT

In order to use images in your Sphinx document, you will need to create a sub-folder to hold them in the source folder. Do NOT use spaces in your image filenames.

Adding an image is pretty simple. You use the directive `.. image::` a space and then the

HOWTO - PYTHON

filename with the path. So it would be something like this...

```
.. image:: /images/  
startup.png
```

There are options that include height, width, scale, and others. Here is an example of the image directive with the scale option.

```
.. image:: /images/  
sclbDemo20.png
```

```
:scale: 50%
```

Notice the spacing. Spacing is VERY important in rst, so if there is an error in the build, that would be the first place to look. For the most part, any directive options are indented 4 spaces. Also note that when continuing your text after an image directive (or just about any directive), it is important to have at least one blank line between the directive and your next text.

THEMING YOUR DOCUMENT

You will need to edit the config.py file in order to change the theme for your document. There are MANY themes available to suit your style. There is a web page that provides a gallery of the themes at

<https://sphinx-themes.org/#theme-mozilla-sphinx-theme>.

My favorite theme is the bizstyle, which is built in. There are 10 different themes that are already built into Sphinx. Any other themes that you might want to use, will have to be downloaded and installed. Some themes have special steps you need to do, so follow the instructions for that theme.

As I said above, you will have to edit the config.py file to set the theme that is used.

SOME OTHER USEFUL DIRECTIVES

There are WAY too many directives to try to cover them all here, and there are plenty of web sites that cover rst directives, so I'll cover only a few of the ones that I commonly use.

Remember that how the directives look in the finished product depends on the theme you are using.

Note

The actual entry method depends on your Operating system. Some of the options might not be available on your system.

ADMONITIONS

I use admonitions often to highlight notes or warning statements. According to one of the rst directive web pages "Admonitions" ("safety messages" or "hazard statements") can appear anywhere an ordinary body element can. They contain arbitrary body elements. Typically, an admonition is rendered as an offset block in a document, sometimes outlined or shaded."

Here is a simple note in code...

```
.. Note::
```

```
The actual entry method  
depends on your Operating  
system. Some of the options  
might not be available on  
your system.
```

Remember the spacing is VERY important both before and after the directive.

Shown top right is what it looks like rendered using the bizstyle theme (but can look different in other themes).

CODE BLOCKS

You can use code blocks to highlight console commands or code snippets. Here is the directive for highlighting a terminal command.

```
.. code:: console  
  
python colours1.py
```

And below is a screenshot of how it gets rendered.

LANGUAGE BASED CODE BLOCKS

Like terminal code blocks, you can use language specific code blocks as well.

Create an instance of the widget

```
python colours1.py
```


HOWTO - PYTHON

giving it all of the information you might want use to customize the look of the widget (top right).

And how it gets rendered is shown below.

BULLETED LISTS

Bulleted lists are REALLY easy in rst. Just place a dash in front of each item.

From top to bottom on the left side of the image are:

- Hex value
- RGB Value
- Calculated Closest Colour name

Create an instance of the widget giving it all of the information you might want use to customize the look of the widget.

```
sclb = ScrolledCheckedListBox.ScrolledCheckedListBox(
    _w1.FrameSCLB,
    width=350,
    height=490,
    font="-family {DejaVu Sans} -size 12 -weight bold -s
    sorted=True,
    # background="cornsilk3",
    background="antiquewhite2",
    foreground=fg,
    activebackground="antiquewhite3",
)
```

Then you can use the Tkinter **place** geometry manager to place the widget into the parent.

- Calculated Best Colour for text on this colour.

Renders out like this...

From top to bottom on the left side of the image are:

- Hex value
- RGB Value
- Calculated Closest Colour name
- Calculated Best Colour for text on this colour.

As I said, there are many other directives you can use (and probably will) so in the next section I've provided a few of my favorite reStructured Text "go to" web pages.

HELPFUL LINKS

Create an instance of the widget giving it all of the information you might want use to customize the look of the widget.

```
.. code-block:: python

sclb = ScrolledCheckedListBox.ScrolledCheckedListBox(
    _w1.FrameSCLB,
    width=350,
    height=490,
    font="-family {DejaVu Sans} -size 12 -weight bold
    -slant roman",
    sorted=True,
    # background="cornsilk3",
    background="antiquewhite2",
    foreground=fg,
    activebackground="antiquewhite3",
)
```

Then you can use the Tkinter **place** geometry manager to place the widget into the parent.

<https://www.sphinx-doc.org/en/master/tutorial/first-steps.html>

<https://www.sphinx-doc.org/en/master/usage/restructuredtext/basics.html#sections>

<https://docutils.sourceforge.io/docs/user/rst/quickref.html>

<https://docutils.sourceforge.io/docs/ref/rst/directives.html#top>

Looking at my page count, it's about time for me to close for this month.

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 Jon Hoskin

Stable Diffusion 13

I have mentioned several times about using prompts that are available without really defining how you would do it. So, let me give an example. Go to: <https://civitai.com/> which is a location you likely visit to get Models, and look at the Images. I selected: <https://civitai.com/images/7287702> and copied the positive (**Photo of artistic stone cup with 3D carvings, little mouse theme with forest background, decorated with amber accents, masterpiece of art, visually stunning, intricate details, sharp focus, 55mm f/ 1.8 lens, depth of field, natural daylight**) and negative (**blurry, painting, drawing, sketch, cartoon,**



anime, manga, render, CG, 3d, watermark, signature, label, (worst quality, low quality, normal quality:2),) prompts. I changed the Positive prompt to include possum and not mouse (mouse). I used both the sd_xl_base_1.0.safetensors and childrensStories_v1SemiReal.safetensors models and got two different example images:

Both are impressive enough but do not accurately represent possums. General models are unlikely to have much of a possum photographic base and you might have to create your own if needed. This simple technique of borrowing prompt information will help newer

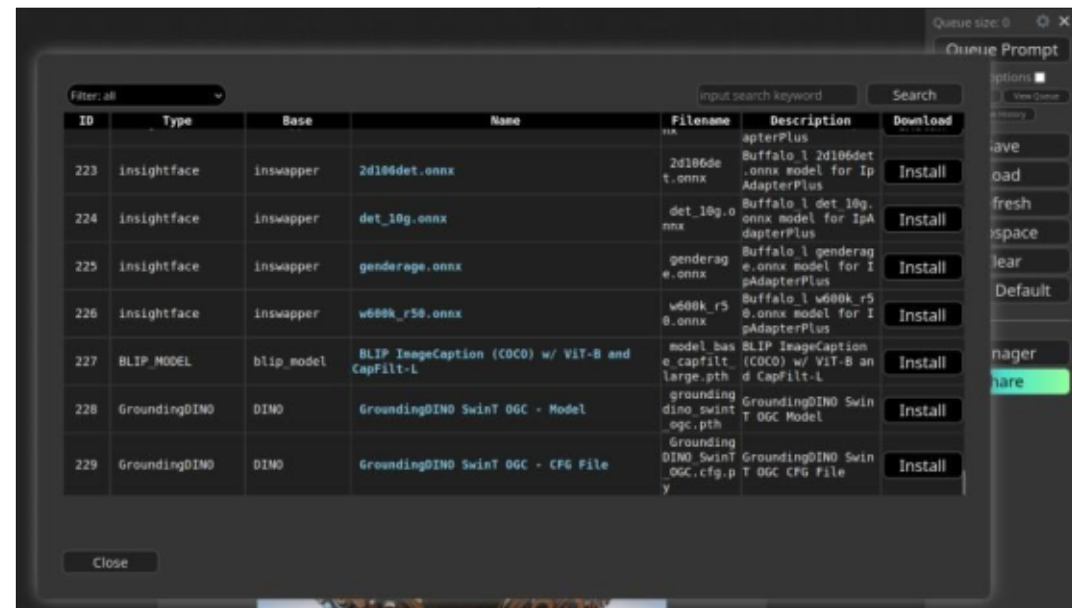


and even seasoned Stable Diffusion users to go in the image direction needed.

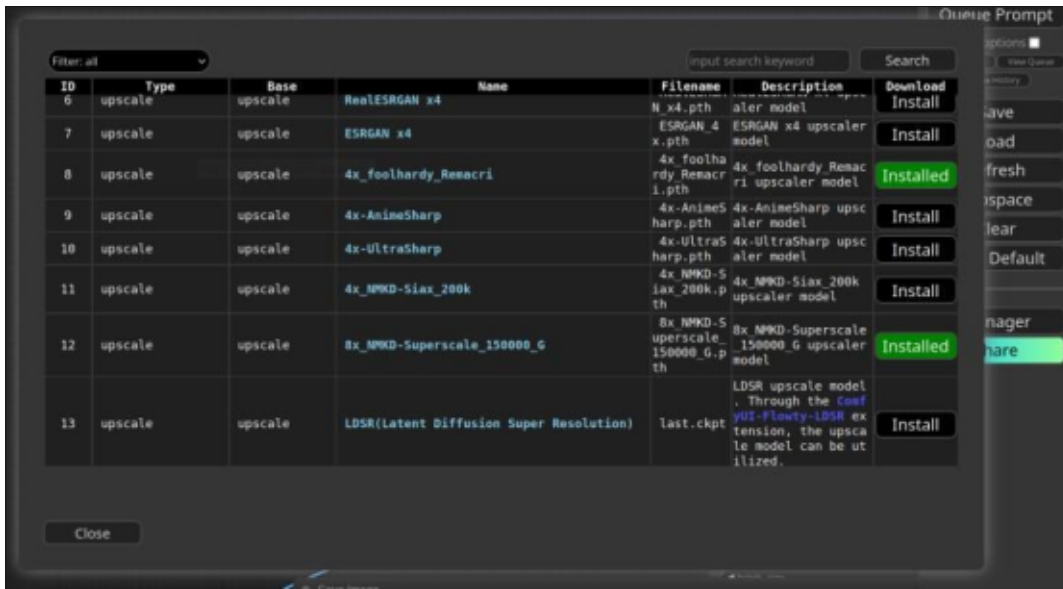
One other time saver and space saving option with ComfyUI is to have it point to the models you have already loaded in Automatic1111 for example (unneded if you started SD using the ComfyUI version). In your ComfyUI folder you should find a file named: extra_model_paths.yaml.example. Remove the.example text from the file name making it a yaml file and edit it to reflect your actual

Automatic1111 folder and ComfyUI should then look to that folder for models. The base_path: path/to/stable-diffusion-webui/ line is what needs to be changed appropriately and then the new file saved to effect this change.

But getting back to the purpose of this article, we want to start using ComfyUI. We might first want to add another model or two. You can do this by opening the ComfyUI Menu Manager which we added previously. After selecting the Install Models button, a list of over 300 options becomes available as



HOWTO - STABLE DIFFUSION



ID	Type	Base	Name	Filename	Description	Download
6	upscale	upscale	RealESRGAN x4	RealESRGAN_x4.pth	RealESRGAN x4 upscaler model	Install
7	upscale	upscale	ESRGAN x4	ESRGAN_x4.pth	ESRGAN x4 upscaler model	Install
8	upscale	upscale	4x-foolhardy_Remacri	4x-foolhardy_Remacri.pth	4x-foolhardy Remacri upscaler model	Installed
9	upscale	upscale	4x-AnimeSharp	4x-AnimeSharp.pth	4x-AnimeSharp upscaler model	Install
10	upscale	upscale	4x-UltraSharp	4x-UltraSharp.pth	4x-UltraSharp upscaler model	Install
11	upscale	upscale	4x_NMKD-Siax_200k	4x_NMKD-Siax_200k.pth	4x_NMKD-Siax_200k upscaler model	Install
12	upscale	upscale	8x_NMKD-Superscale_150000_6	8x_NMKD-Superscale_150000_6.pth	8x_NMKD-Superscale_150000_6 upscaler model	Installed
13	upscale	upscale	LDSR(Latent Diffusion Super Resolution)	last.ckpt	LDSR upscale model. Through the ComfyUI2-Flowty-LDSR extension, the upscale model can be utilized.	Install

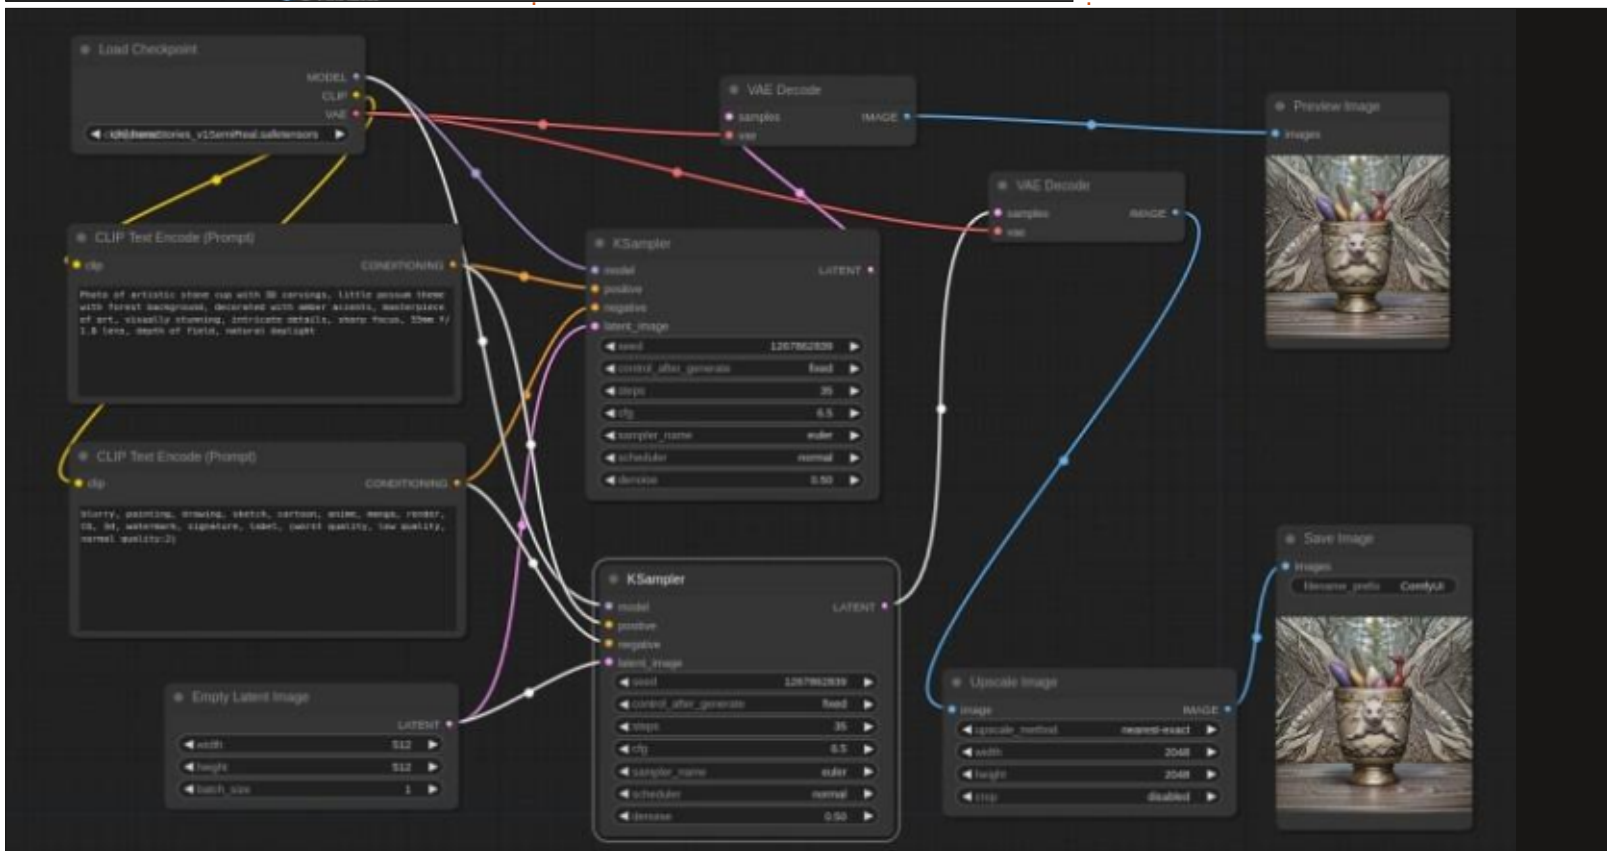
shown. You can use the filter to show all, installed or not-installed. In my experience, you may need to wait after selecting a model and then restart the system and look again to make sure it's installed.

Many up-scaling options also exist as optional models. When installed they may be found in the Models/Upscale models in the ComfyUI models folders depending on file type (pth or safesensor) as noted in the type column. You may

want to install one or more of those available. You can use the search option for upscale related models.

Next we will learn how to build our own workflow. We could start from scratch with the Queue Prompt and clear the current workflow using the Clear button, however you should always consider if you may want to re-use the present workflow. If so, you can select the Save button and use an appropriate name for the workflow which is saved as a json file. In our case we will start with the default workflow or whichever you may have that works with the ultimate goal of making something that looks like the provided workflow.

Initially (shown left) it looks much more complex than what we have done previously but it's essentially two workflows connected. To get the feel of it, right click somewhere on the workflow and select Add Node, then Image and Upscale for example. You will find over a hundred Node options, a bit mind boggling at first. If you started with a blank page you could add each node as found in your default but



HOWTO - STABLE DIFFUSION

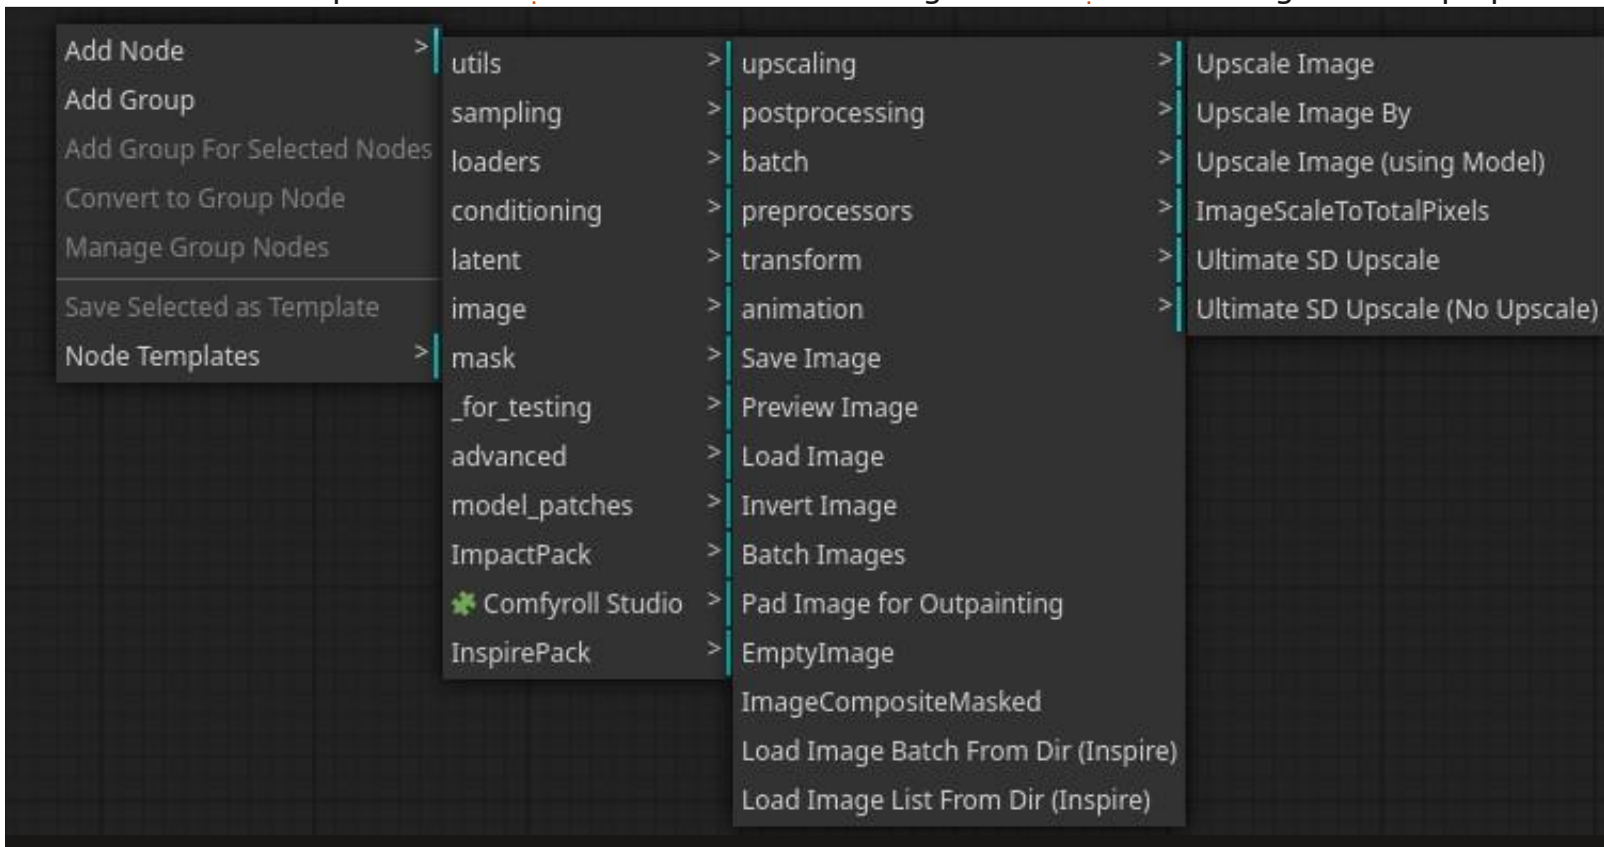
it's easier to begin with something and simply modify it. Each node can be moved around by selecting it and using your mouse. You can also copy each (Ctrl+C) and (Ctrl+V) to paste or more importantly (Ctrl+Shift+V) will not only paste a copy of the node but keep all the same options and connections. (Make sure you don't move the mouse as you paste as you might get a hundred copies pasted.)

You will want to duplicate the

KSampler node in that manner. Be careful to move each node so you can see it's inputs and outputs, sometimes holding the Ctrl button down and using the scroll wheel on your mouse to adjust the size makes this placement easier. The original Latent Image output from the KSampler Node went to the VAE Decode node to the Preview Image node or Save Image node. (The lower right corner of image nodes can be pulled down to make space for the image.) That converts or decodes the latent image which

is not a viewable image into a standard image. Likewise the Latent Image output from the second KSampler node does that same thing except in between it goes to an Upscale Image node. What's happening is the initial Prompt information is going to two KSamplers with one processing at 512 x 512 and the other 2048 x 2048. The problem is that while the larger image is larger, it contains no more information than the smaller image. But this ComfyUI process is the same regardless of purpose.

Next time we will produce a larger image with more information which gives more detail.





HOW-TO

Written by Robert Boardman

Latex - Tables

This time, I am going to come back to a topic I did quickly before: tables. I'll start with what was used in documents to present data in columns before word processing.

Before computers, there were typewriters, mechanical devices. On a typewriter, to make columns, the user sets up tabulation positions, usually called tab stops. Each tab stop was a specific physical distance from the edge of the sheet of paper inserted into the machine. There are tab stops in word processing software, but they are rarely used except by very experienced operators. Tab stops in word processing have more options and more flexibility than on typewriters. Tables are even more flexible but require at least as much thinking as tab stops.

Tabs can be used in Latex in the tab environment. It is relatively easy to set tab stops at specific distances from the left margin in either the tabbing or the tabto environments. However, if you want information arranged in columns

```
\begin{tabular}{|l|c|c|c|c|}
\hline\hline Municipality & Jan & Apr & July & Oct\\ \hline
Prince Rupert & 5.6C & 10.2C & 16.2C & 11.1C \\ \hline
Tofino & 8.3C & 11.9C & 18.9C & 13.6C\\ \hline
Nanaimo & 6.9C & 14.1C & 23.9C & 14.6C\\ \hline
Victoria & 7.6C & 13.6C & 22.4C & 14.2C \\ \hline \hline
\end{tabular}
```

and rows, the tabular environment offers more flexibility and more options.

This is not a simple topic. In The Latex Companion (Third Edition) there's about seventy pages of detailed information discussing the various table environments and their various extensions. If that seems like more than you want to read, there are eleven pages dedicated to a simplified discussion of the table environment in LATEX2e: An unofficial reference manual available at <https://latexref.xyz>. These pages cover a selection of the possible extensions and packages for making tables. Remember, Latex is a modular tool. Instead of loading every possible tool and its variations when starting, Latex allows you to use the packages you need for your document when you need them,

and it will generate a typeset document (PDF) to your specifications.

Ideally tables are used to organize and present information so the contents of the table are clear to the reader. Tables organize information in vertical columns and horizontal rows. A table is a framework, a skeleton into which the writer puts information. A table is described by indicating how many rows and columns it has. Often a table has the same number of columns in each row and the same number of rows in each column. However table cells can be joined in various ways. Table cells are often separated from each other with vertical and / or horizontal lines. All of this is familiar to you, I am sure.

In Latex, a table is built inside a table environment. As Erik showed

in one of his last columns here is the basic code for a table with some contents (shown above).

{|l|c|c|c|c|} means set up a table with five columns, the left hand column is left-aligned, all the others are centred. The single vertical line between each alignment option indicates a vertical line will be placed between each column. The double line on the left-hand and right-hand ends indicated the outside lines will be doubled. The \hline instruction places a horizontal line below each row, the repeated \hline instruction makes a double line at the top and bottom of the table. The double back-slash marks the end of a row. As you can probably guess, the & ampersand symbol separates the contents of one cell from its neighbours in its row (see Table 1 in Figure 1 next page).

In Latex tables are one of the items which are floated when typeset. That has two consequences for page designers and authors. The first consequence is the table will appear on the page where there is room for it. If there is not room on the current page then it is "floated" to the next page. The second consequence is a table in the tabular environment cannot extend past the end of a page, that is one table has a maximum size of one page. If you need more space than one page for your table, investigate the use of the supertabular and the longtabular environments.

The following example shows how to connect two cells in the same row together. There is a similar instruction for connecting two cells in two different rows. The & after {Ene} is necessary to build the third cell in the first row.

```
\begin{tabular}{|c|c|c|}
\hline
\multicolumn{2}{|c|}{Ene} & \
\hline
Mene & Muh! & Curly \\
\hline
\end{tabular}
```

Similar to the parameters for

multicol, multirow has three required parameters: how many rows, the width of the resulting cell, and the contents of the spanned cell. The hline instruction between rows one and two is omitted. If it is included it will cross into the spanned cell. A difference which you cannot see from the code below is that multirow is a separate package and must be loaded (\usepackage{multirow}) in the document preamble. Multicolumn is built into the version of Latex I use and into other versions as well.

```
\begin{tabular}{|c|c|}
\hline
\multirow{2}{14mm}{Ene} &
Mene \\
& Muh! \\
\hline
\end{tabular}
```

Since I started this column by discussing tabs and tab stops, I am going to finish by showing two methods of getting and using tab stops in Latex. The first way is in the tabbing environment, the second is with the tabto package. (See Figure 2 - next page - for the results of this code.)

```
\begin{tabbing}
First \= Second \= Third
\=\kill
one \> two \> three \> four\\
```

Table 1: Demonstration Table

Municipality	Jan	Apr	July	Oct
Prince Rupert	5.6C	10.2C	16.2C	11.1C
Tofino	8.3C	11.9C	18.9C	13.6C
Nanaimo	6.9C	14.1C	23.9C	14.6C
Victoria	7.6C	13.6C	22.4C	14.2C

Table 2: Multicolumn demo

Ene		
Mene	Muh!	Curly

Table 3: Multirow demo

Ene	Multirow Muh!
One	Two

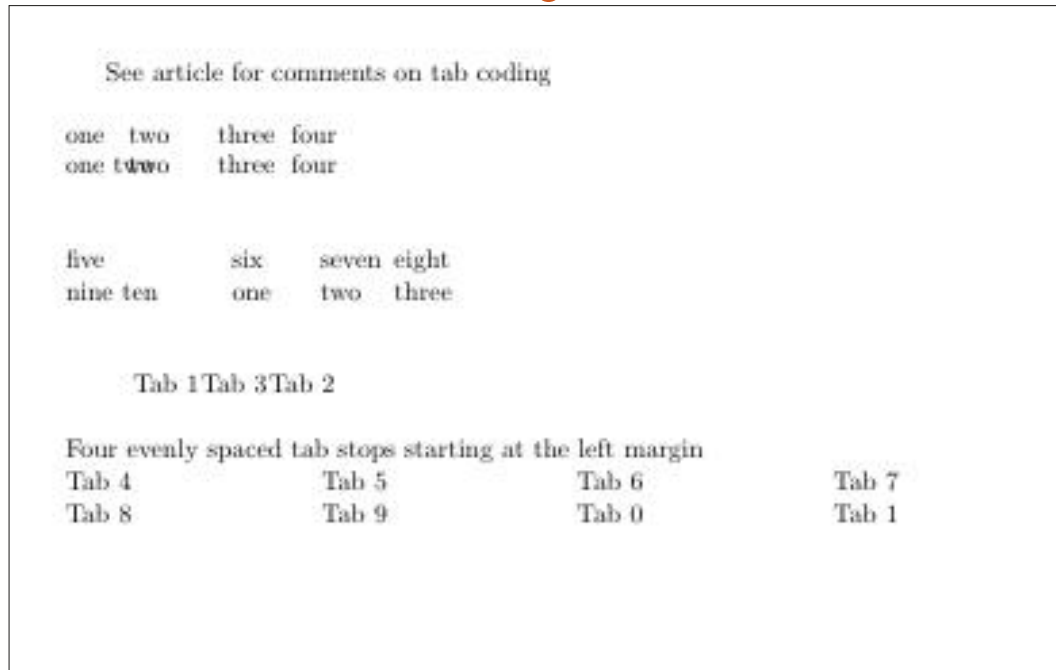
HOWTO - LATEX

```
one two \> two \> three \>
four\\
\end{tabbing}
```

The first line after the `begin{tabbing}` instruction sets the tab stops, the distance from the left margin for the start of the text (or other contents). Notice there are five characters plus one space between the margin and the first tab stop. The first item in the first line will display with no problems, it contains three characters and one space. The first item in the second line has six characters and two spaces. It will display as usual. However the second item in the second line starts with five characters and one space from the left margin, and will overwrite part of the first item. The next code example shows how to fix this problem.

```
\begin{tabbing}
First tab stop \= Second \=
Third \=\kill
one \> two \> three \> four\\
one two \> two \> three \>
four\\
\end{tabbing}
```

As stated above, `tabto` is a package but it is not an environment, unlike `tabbing`. An author needs to load it in the preamble but can use the package



without beginning and ending an environment.

```
\tabto*{1cm} Tab 1
\tabto*{3cm} Tab 2
\tabto*{2cm} Tab 3 % in
document
```

The distances are distances from the left margin. Notice the text `Tab 3` will be printed between `Tab 1` and `Tab 2`. (The asterisk suppresses line breaks.)

The `tabto` package will also allow a user to set a specific number of evenly spaced tabs or specify exactly where each tab stop should be (just as can be done with a typewriter). Both parts of the

code below go into the document.

```
\TabPositions{0,0.25\linewidth
h,0.5\linewidth,0.75\linewidth
h} % in preamble
```

Then use `\tab` in the document. (This is equivalent to `\NumTabs{4}`)

```
\NumTabs{4} % in document
%Four evenly spaced tab stops
starting at the left margin
\tab Tab 1 \tab Tab 2 \tab
Tab 3 \tab Tab 4 \\\
```

There is a brief one-page document for the `Tabto` package available on ctan.org. A little experimenting will be helpful if you are not used to using tab stops.

That is enough Latex fun for one issue. I will spend more time working with tables in the next column. In the meantime, if there is anything in this column that is new to you, I encourage you to try it out, play with the code, experiment with the choices, perhaps even read the documentation. Learning something new takes some effort but the rewards with Latex are worth the effort.

KILOBYTE MAGAZINE

Kilobyte Magazine is a fanzine for 8bit enthusiasts. It covers consoles, computers, handhelds and more, as well as new games for old systems. If you grew up with Commodore, Atari, Sinclair or Amstrad, this magazine is for you.

<https://retro.wtf/kilobytemagazine/>



HOW-TO

Written by Mark Crutch

Inkscape - Part 144

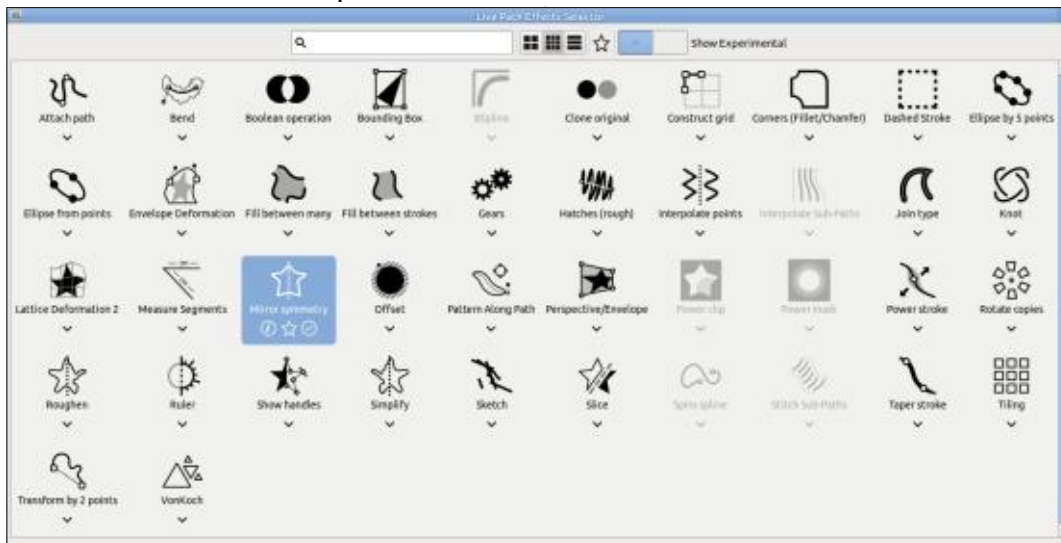
This month we're continuing to look at the changes to dialogs in Inkscape version 1.3. It's almost become a cliché to see substantial changes to the Live Path Effects dialog with each significant Inkscape release, but here we are again with yet another round of UI updates. As a reminder, in version 1.2 the main LPE dialog looked like this (once a couple of effects had been added):



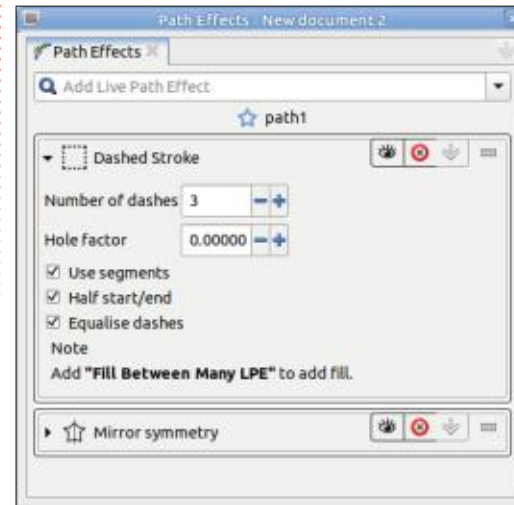
Clicking the '+' button would open a large dialog from which you could select the LPE to add, as well

as search, view the entries as a list, or as a more/less compact grid, set LPEs as favorites (and filter to show only those), and toggle the visibility of experimental effects. In fact there's so much functionality baked into this UI that I expected it to hang around for years to come. How wrong I was.

Fast forward to version 1.3. The main LPE dialog no longer shows a list of the applied effects, with a single panel of settings for the currently selected one. Instead, the configuration panels are shown directly, in an 'accordion' UI pattern, such that by expanding one panel, all the others will collapse. You can



collapse the current panel, too, if you just want to see a shorter list of effects, which makes it easier to reorder them by dragging the handle at the top-right of each panel.



With this UI, there's no longer a '+' button, and no longer a large dialog that opens. Instead you have two options for selecting the new LPE to add:

- Type something into the search box to be presented with an alphabetically ordered list of LPEs that match the search term. (Hint: if you just want to see a list of all of them, putting a single space in this field seems to do the trick).
- Click the drop-down button next to the search field to be presented with a pop-up that contains a categorised grid of effects.



There's no way to change the size or shape of this pop-up, nor the spacing between items, so I'm not sure how well it will work on small screens – especially if you enable the experimental LPEs. Speaking of which, the option for that is now hidden away in the preferences dialog: Edit > Preferences > Behaviour > Live Path Effects (LPE). Enabling the 'Show experimental effects' checkbox in there will add an extra section to the bottom of the pop-up. In that same settings panel, you'll find an option to 'Show deprecated LPE gallery' (i.e. the older UI), if you really don't like the new design. Personally I advise against switching back, as this option – and the whole gallery dialog – will certainly be removed in a later release. You may as well get used to the new approach now, and hope that this one finally sticks around for a while.


I think hiding the option for showing the experimental LPEs was a good move. This isn't something that should be temptingly presented to novice users every time they add an effect. The experimental LPEs are marked as such for good reason. They should only be used with full knowledge of

the potential problems – not only in terms of Inkscape's stability, but also considering possible issues loading any files that use them in future releases.

Any LPEs that are set as favorites, appear in a section at the top of the pop-up. Unfortunately they no longer appear in their categorised sections, and the Favorites section is shown in alphabetical order, so any sense of grouping is lost for these. You can either have an LPE categorised or flagged as a favorite, not both. The list view that drops down when you use the search field is purely alphabetical, not only lacking the categories, but also lacking any indication of which effects have been set as favorites. It would be nice to see favorites displayed with a small star or other badge, and I'd rather see the list view also grouped in the same way as the icon pop-up, for consistency.

Without the gallery view, there's a question of how you set (or unset) the favorite flag on an effect. There's no context menu available in the pop-up – right-clicking on an LPE will behave the same as a left-click, and add it to the list of applied effects in the main body of

the dialog. Instead, you have to add the effect to a path (perhaps a temporary, sacrificial one), then use the new pop-up menu in the effect's controls pane to toggle the favorite flag. On my system, there's some oddness with the outline style for this pop-up's button, but it still behaves correctly. In this screenshot you can see that there's a group of three buttons which are visible whether the pane is collapsed or open: the first button toggles the visibility of the LPE, the second deletes it from the list, and the third opens the pop-up menu.



The pop-up contains a few useful options for working with effects:

'Duplicate' is pretty self-explanatory, but note that the duplicated effect will appear directly below the original, and you may need to move it up or down in the list by hand. This can be done by dragging with the handle at the

right, as previously mentioned, or you can use the 'Move up' and 'Move down' entries in this menu. Only valid options are presented, so you won't see 'Move down' in the menu for the last LPE in the list, for example (as in this screenshot).


'Flatten' is an interesting new addition. It's always been possible to 'fix' the output of an LPE chain by selecting the Path > Object to Path menu option. This changes the drawing such that the original path is replaced by the result of the entire LPE chain. The Flatten option does something similar, but it stops at the effect on which the option is selected, fixing all the effects up to and including that one, but leaving subsequent parts still as editable LPEs. If you use this on the last effect in the chain, it's the same as using Object to Path, but this could be very useful where you're sure that the first few effects in a chain are correct and won't need to be modified, but still want to leave yourself with the flexibility of fully editable effects for the rest of the chain.

'Set custom defaults' will use a snapshot of the current settings as the default values for this particular effect in future. If there's already a

snapshot set, this changes to 'Forget custom defaults'. Finally we have the 'Set favorite' option (shown with the British English spelling of 'favourite' in this screenshot, as I'm based in the UK). If this LPE is already set as a favorite then, as you've likely guessed, this will display 'Unset favorite' instead.

It's a little unfortunate that the only way to change favorites is to first add the effect to a path, perhaps only to throw it away again afterwards. But changing favorites is likely to be a rare enough occurrence that this additional step probably won't cause any real problems for most users.

I'm going to take a very quick look at a second dialog this month: the Symbols dialog (Object > Symbols). This has seen a bit of UI polish with 1.3, with the confusing array of buttons for adjusting the layout now rationalised into a simple pop-up menu. My tip here is to leave 'Zoom to fit' enabled and just tweak the 'Tile size' slider to adjust the size of the symbols in the dialog, which seems to work well with the predefined symbols, at least.



On that point, version 1.3 ships with a vastly increased number of symbols (596 of them, to be precise), through the addition of the 'SJB map icons' set. Of course, more symbols are always welcome – if they're high enough quality – so it would be nice to see some Open Source icon sets, of the sort often used in web pages, make their way into future releases.

The plus and minus buttons in the bottom-left replace the similarly-placed buttons from previous versions, which were always somewhat confusing due to

less-than-obvious icon designs. These buttons only become active when the 'Current document' set is selected in the dialog's drop-down list. Their tooltips have been reworded to make their purpose far more understandable now:

"Convert selected object(s) to symbol" – this will create a symbol from the currently selected objects, replacing those objects with the symbol, in-place. This may be useful if you need to use the same icon or other design repeatedly within your document, though I'd usually just stick with clones myself.

"Convert selected symbol to object, and all symbols of this type to clones of this object" – this one is perhaps more useful, allowing you to select a symbol in the dialog and have the first instance of it on the page converted to a real, editable object. Any further instances will be converted to clones of the object. I can see this coming in handy when one of the predefined symbols is close to your requirements, but needs a little tweaking.

Of course these can be used in series to first convert an existing symbol to an editable object; then edit it on-canvas; then convert the

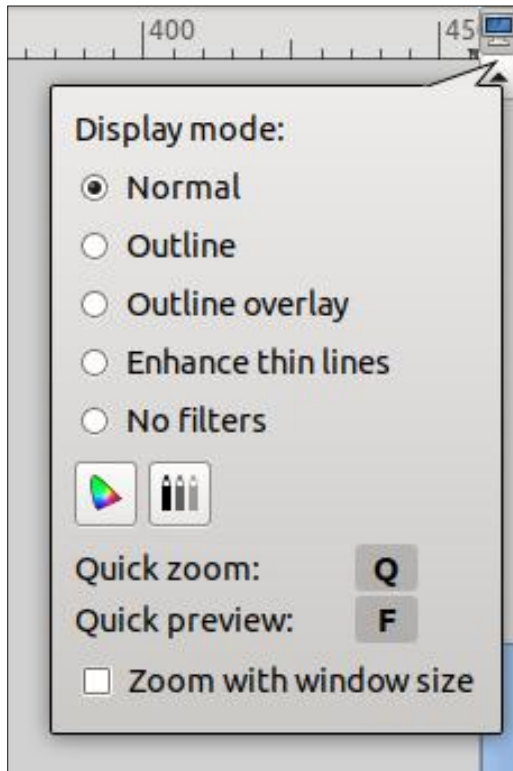
edited version back to a document-specific symbol.

Finally I'll squeeze in one more UI addition – though a pop-up rather than a dialog. At the top-right of the canvas, Inkscape has always had a rarely used button to toggle its behaviour when the window size is changed. The default is to not do anything much, other than to show or hide the relevant parts of the canvas to suit the new window size. In other words, the document is not zoomed in any way. The other mode will scale the document to suit the new window size, which may suit some users who always want to see the same amount of content, regardless of the window size. This feels to me like a feature that novice users might think they want before they discover the myriad ways to zoom and pan in Inkscape, and realise they'd rather control the zoom themselves than have it automatically change when they resize the window.

This button has now been replaced with a pop-up that still contains the previous functionality via an option at the bottom, but which also exposes radio buttons and toggles for the many display

HOWTO - INKSCAPE

modes Inkscape now has.



Note, however, that the 'Quick zoom' and 'Quick preview' entries can't be clicked – they're just there to remind you of the keyboard shortcuts. These both work as long as the key is held down – though you can also trick them into operating as on/off toggles by focusing another window, or right-clicking to bring up the context menu, while holding the key. In that case the 'temporary' view will remain on releasing the key, and a quick jab of it subsequently will

return you to the previous state.

I've covered 'Quick zoom' before – it zooms the display just enough so that the selected object(s) fill the available space. 'Quick preview' is a new addition with 1.3, and temporarily hides grids, guidelines, selection boxes, and so on. It also clips the visible display of objects to the page boundary. Essentially it is intended to give you a rapid preview of how your document might look when printed or exported, without all the Inkscape-specific parts that won't make it to the final document.



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

WE WATCHED THE SAFETY
FILM AT WORK, IT DIDN'T
HAVE A HAPPY ENDING





MORE FUN WITH DATES AND TIME

Greetings fellow Beings. Yet again, I come to you from an undisclosed location somewhere in time and space. Yes, I am a fan of all things Doctor Who, Star Trek, Star Wars, and other things.

Anyway, last month we looked at the subject of time change from “standard” time and “daylight saving” time and some of the issues that come with that on a Microcontroller. I promised that we would continue that subject. So here we go with a reminder of some of the issues.

First, the `ntptime` module, used to get the network time, has a small problem when you make too many calls to it. The server simply tells you to go away after a while and you have to start all over again. I also gave you a short snippet to use the portable version of `ntptime` from Peter Hinch.

I also talked about the fact that there are two functions that

provide the current time (once it’s set) named `gmtime` and `localtime`, and the fact that they both give GMT time. There is no timezone information handled by either of these functions.

So let’s talk about some of the possible remedies available to us.

In my search for an easy solution, I came across the fact that the MicroPython group has been working on creating a large library of modules that are useful for programmers writing MicroPython programs. Many of them attempt to mirror standard libraries in the Python world. One of these modules is the `datetime` module. MicroPython has suffered from the lack of this being in the builds. Well, now, the wait is over (at least for the most part).

By using the `micropython-lib` `datetime` module in our programs, we can, with a few modifications, use “normal” python code in our MicroPython programs.

Of course you have to have the

`datetime` module “installed” on your device. There are few ways of doing this. One is my preferred way of getting the entire source code package and copying the necessary modules onto my device. By doing this, I can make sure that any of the dependencies are available and I can examine the code to see how that programmer (or group of programmers) got their code to work. Also if there are any errors being thrown, I can track it down much quicker than relying on a library where I can’t see the source code. You can find the package at <https://github.com/micropython/micropython-lib/tree/master> and download the entire zip file. Unzip the package into a folder (I use my Downloads folder and unzip it there, since it has the folder already there for me). Once there, you can find the file in the `/python-stdlib/datetime/` folder (named `datetime.py`)

Copy the `datetime.py` to your Microcontroller device and put it into the `/lib` folder or, if you don’t have it, either create the folder and place it there, or just drop it into

your main folder on the device. (It’s better to put it into the `/lib` folder.)

The other way, which is a bit easier but only provides you the compressed `mip` file, is to use `mip`. Here is a link to using `mip`: <https://www.tomshardware.com/how-to/raspberry-pi-pico-micropython-mip>.

Anyway, I haven’t had a chance to fully explore all the functions available, but for this month, we’ll use the `datetime` library to get the `datetime.now` function. You will call it like this:

```
my_timezone =  
timezone(timedelta(hours=10))
```

```
current_time =  
datetime.datetime.now(my_time  
zone)
```

Of course, you need to import the library before you can call it, but we’ll see that in this month’s first project.

If you are a “regular” Python programmer, you will recognize the second line, which uses the `datetime.datetime.now()` method.

So, for this month's main project, we'll connect to the Internet (be sure you're using a RPi Pico-W), use Peter Hinch's portable NTPTime module, get the network time, set the Real Time clock on the Pico-W, and then start printing the time in the REPL once every 10 seconds. I've named the program 'timetestFCM1.py'. Make sure that you have the datetime.py or datetime.mpy file either in your main device's folder or in the lib folder. All of the other libraries needed are "standard" libraries, except for Peter's portable NTPtime library. I've got that in the repository listed at the end of this article.

THE CODE

First, we'll need a secret.py file. This will hold your SSID and your network password. I use this to make sure that if I ever share a project that uses networking, that my network information isn't exposed in the code. (It's in the repository, so all you need to do is to make sure you edit it once it's on your device). It looks like this...

```
# secret.py
# =====
SSID = {Your Network Name
```

```
Here}
PASSWORD = {Your Network
Password Here}
```

Really simple huh?

Moving on. Now we'll need Peter's portable NTPTime routine. That's also in the repository saved as port_ntptime.py. Make sure that's also on your device.

Now (top right) we'll start digging into the timetestFCM1.py project. First the imports. It will seem to be a REALLY big number of imports, but it's not much.

Next comes the connectNetwork function. This is a fairly straight forward function that I reuse often. First, I set up the onboard LED of the Pico (name it "led"), then make sure it's off. Next I start to set up the network modes, send in the SSID and Password to the wlan object and loop until we get connected. Once we are connected to our local network, we turn the onboard LED on (next page, top right).

Next we have the settime2 function. It's very similar to the one defined in the NTPTime module, but this one will return a 0 if, for

```
import machine
import port_ntptime
import utime
import network
import secret
import time
from time import sleep, localtime, gmtime
import sys
import gc

from datetime import MAXYEAR, MINYEAR, datetime, date,
timedelta, timezone, tzinfo
```

Now we can define our internal functions. First the GetWhichPico function, which is useful if you have both types of the RPi Pico.

```
# Check that the board is a Pico-W with wireless support
def GetWhichPico():
    which1 = sys.implementation
    if "Pico W" in (sys.implementation[2]):
        print("Pico W")
        return True
    else:
        print("NOT Pico W")
        return False
```

some reason, it can't set the time correctly. We get the time and date tuple by sending the number of seconds since our EPOCH (from the network time server) to utime.gmtime and then set the software real time clock with that value. If everything worked correctly we return a True, otherwise return a False (next page, top left).

Finally, we get to begin to play with the datetime module we went to "so much trouble" to get.

First we need to set a value for our timezone, which in my case is -5 hours from GMT. Once we go through the time change again, this will be a -6 until the next time change. Then we take that and use it to seed the datetime.datetime.now() function, which returns a time string like this...

```
2024-04-07 02:49:10-05:00
```

We'll get into formatting that next month.

```
def settime2(tz):
    t = time.time()
    if t < 0:
        return False
    else:
        tm = utime.gmtime(t)
        tm2= utime.localtime(t+tz)
        machine.RTC().datetime((tm[0], tm[1], tm[2], tm[6]
+ 1, tm[3], tm[4], tm[5], 0))
        return True
```

```
def showMyTime():
    import datetime
    my_timezone=timezone(timedelta(hours=-5))
    current_time = datetime.datetime.now(my_timezone)
    return current_time
```

That's it for all of our support functions. Now we have the main routine. This makes calls to the various functions that we've already discussed.

First, I define the two timezone offsets used here where I live. Cst is 6 hours behind GMT (starting in the Fall) and cdt is 5 hours behind (Starting in the Spring). I multiply the number of hours by 3600 (The number of seconds in an hour) and since we are behind GMT, I use a negative number.

```
cst=-6*3600
cdt=-5*3600
```

Next I get the Epoch which is where my device bases the number of seconds from in order to "decode" the value sent from the NTPServer. The Pico uses January 1, 1970. Other devices use January 1, 2000. But that actual date is January 1, 1900. Knowing this can help if there are any errors. After that, I get the time that the device thinks the time is.

```
EPOCH = time.gmtime(0)[0]
print(f"{EPOCH}")
rtc=machine.RTC()
dti=rtc.datetime()
print(f"{dti}")
```

Now, I threw in a quick check to make sure this program was running on a Pico-W. If not, well,

```
def connectNetwork():
    global myipconfig

    # Setup onboard LED
    led=machine.Pin("LED",machine.Pin.OUT)
    led.off()
    # Setup Network
    wlan=network.WLAN(network.STA_IF)
    wlan.active(True)
    # Provide SSID and PASSWORD from secret.py file
    ap=secret.SSID
    pwd=secret.PASSWORD
    # Try to connect to the wireless network
    wlan.connect(ap,pwd)
    # Loop until we get connected
    while not wlan.isconnected() and wlan.status() >= 0:
        print("Waiting to connect")
        time.sleep(1)
    # Print the status value and the ifconfig values
    print(wlan.status())
    print(wlan.ifconfig())
    myipconfig=wlan.ifconfig()
    # Turn on the onboard LED
    led.on()
```

there is no need to continue on, since we can't connect to a network. We simply call sys.exit() if it isn't a Pico-W.

```
# Check to see if the Pico
has network
which=GetWhichPico()
if which == False:
    sys.exit()
```

Now we call the connectNetwork function.

```
# Connect to the network
connectNetwork()
```

Now that we are connected, we

can try to get the time (in seconds) from the NTP server. This is fairly simple. Once we are in the loop, we make a call to Peter's portable ntp_time function. It will either return the time value from the NTP server or a 0. If the value is greater than zero, we do a quick garbage collection (to get rid of any unneeded stuff), and call the settime2 function with the proper timezone offset. If the call to settime2 is successful, we set the loop variable to False, otherwise we sleep for 2 seconds (next page, top left).


```
loop = True
while loop:
    secs=port_ntptime.time()
    print(secs)
    if secs > 0:
        gc.collect()
        success=settime2(cdt)
        if success:
            loop = False
    else:
        sleep(2)
```

```
# Now print gmtime and localtime
print(f"{gmtime()} - {localtime()}")
print(showMyTime())
```

Finally, we can enter our forever loop, printing the device's time every 10 seconds.

```
while True:
    print(showMyTime())
    time.sleep(10)
```

That's it. Here is the REPL output...

```
MPY: soft reboot
EPOCH=1970
dti=(2024, 4, 7, 6, 7, 48, 37, 0)
Pico W
Waiting to connect
Waiting to connect
Waiting to connect
3
('192.168.1.251', '255.255.255.0', '192.168.1.1',
'192.168.1.1')
0
1712494129
gmtime()=(2024, 4, 7, 7, 48, 50, 6, 98) -
localtime()=(2024, 4, 7, 7, 48, 50, 6, 98)
2024-04-07 02:48:50-05:00
2024-04-07 02:48:50-05:00
2024-04-07 02:49:00-05:00
2024-04-07 02:49:10-05:00
```

At this point, we can pretty much safely assume we got the Network time. Now all we need to do is print the time.gmtime and time.localtime to the REPL, then print the return from showMyTime (bottom left).

If you are old enough to remember "Y2K", you'll be either happy, or scared, to know that in 2038, there will be another one. Hopefully, we will be more ready for that issue than we were back in 1999. We'll talk about that next month.

The repository for this month is located at <https://github.com/gregwa1953/MTMT204>.

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.



Linux on Your iPad

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





UBPORTS DEVICES

Written by UBports Team

We have elections for both the Board of Directors and the Membership Committee coming up in May this year.

The question you may be asking now is how do I join the BoT, who can, and how do you stand in or vote in the elections? Firstly, you need to be a member of the UBports Foundation Board of Trustees. And how do I do that, I hear you ask. Well, if you have actively contributed to Ubuntu Touch for the last three months, and plan to do so for at least the next six, then you can apply to join the UBports Foundation Board of Trustees. Members of the BoT are then eligible to stand for and vote in the elections for the BoD and Membership Committee.

If you want to do that now, or are just interested in finding out more, then please read the blog (<https://ubports.com/r/4EH/m/180443>) to find out all the requirements.

If you are thinking about it for this year, then please note there is

limited time left to join the BoT for this round of elections, so don't hang about; read the blog and apply now.

BoT members are not there just to vote or stand in our elections, they are a vital part of UBports and Ubuntu Touch in many other ways too, so to have your say in, and help shape the future of, UBports and Ubuntu Touch directly, please do apply.

It cannot be stressed enough how important to us is the involvement of the UBports Community in Ubuntu Touch. We would really like to have as many of you as possible directly involved, and the BoT is one great way to do that.

UPDATES ON UBPORTS TOUCHDOWN BUCKETS

You will probably remember us mentioning the upcoming changes to our release model as discussed in this blog (<https://ubports.com/r/ttL/m/180443>)

Ratchanan has been working on the changes to the development branch as mentioned here (<https://ubports.com/r/Jsf/m/180443>). The CI changes are now complete, and running as expected. However there is an issue with some MR's failing to build due to dependencies. We're working on bootstrapping the APT archive to correct this.

CREATING SNAPS ON UBUNTU TOUCH

A new blog by Aaron Prisk of Canonical, and UBports Alfred E. Neumayer covers all you need to know about creating and using snaps on Ubuntu Touch.

<https://ubuntu.com/blog/creating-snaps-on-ubuntu-touch>

STANDALONE KERNEL BUILDS

Also on image builds, Luksas has created and is testing a standalone kernel method for building the image on the OP5/5T. This method saves on build resources and time

needed compared with the current full-system-image method. It needs to build only one generic_arm64 image (for every Halium version), which could then be used by many devices. There are many other benefits with this method, making the whole process easier on resources and maintenance.

To find out more of the details and how to try it on the OP5/5T, use this link to the Forum post (<https://ubports.com/r/euO/m/180443>).

Capsia has been in discussions on our current website and odoo and proposed improvement plans should be available shortly. Capsia is also working on a device testing app. More on that as soon as we can.

Then there are the upcoming changes to the devices webpage as mentioned in the blog from early February. We said it would be at least two months from the blog publication and that remains the plan, so expect the changes early to mid April.

We will have more information next time including Lomiri where a lot is in progress and the other UTD Buckets. We also have an online UTD Office space every Wednesday from 18.44 UTC to approx 20.30 UTC for catching up on the latest development work from the UTD so more on that also next time.

Thanks to @rubencarneiro there are three new devices on the UBorts Installer:

OnePlus Nord N10 5G (billie)

OnePlus Nord N100 (billie2)


Lenovo Tab M10 HD 2nd Gen Wifi (amar_row_wifi) (TB-X306F) - 4Gb RAM 64Gb ROM version.

Ubuntu Touch Q&A Blog and Audio available using the link <https://ubports.com/blog/ubports-news-1/post/ubuntu-touch-q-a-135-3922>



THE DAILY WADDLE

CONSIDER THIS YOUR FIRST
LESSON IN RISK ASSESSMENT






Clear and concise guides to help you learn with minimal time and effort!

Order direct from ineasysteps.com and get **25% discount!**

Enter **FC25** in the Coupon Code box at the checkout.



C++ Programming in easy steps, 6th edition is for anyone getting started with C++.

- Master C++ basics, including creating and employing variables, declaring constants, doing arithmetic, making conditional statements, using text strings, and storing and retrieving data.
- Grasp the key principles of Object Oriented Programming (OOP).
- Bring it all together, and create programs graphically using a modern IDE, such as Microsoft Visual Studio 2022 illustrated inside.

Learn C++ programming efficiently, effectively and economically – that's smart learning!

£11.99 / \$15.99 paperback (ebook version also available)

ISBN: 9781840789713 / 192 pages / By: Mike McGrath



MY STORY

Written by Art Schreck

balena Etcher

admit it – I’m a distro hopper.

If I’d see a Linux variant worth investigating, it was time to fire up UNetbootin (shown below) and create a bootable flash drive.

And so it went for probably a decade or so and a couple hundred Linux variants.

Then things started getting goofy – UNetbootin stopped working when Jammy was released and it doesn’t appear anybody’s making a great effort to correct the problems.

Hate to say it, but UNetbootin was a little long in the tooth when it came to recognition of recordable drives. Had no problems with USB thumb drives but I never got a flash card to work unless it was attached to a USB adapter.

And while it graciously offered to download several different Linux varieties, it rarely worked within the past year or so.

So, in essence, UNetbootin is a UNogo (sorry, had to stick that in there).

What do us distro-hoppers do then?

You could try dd (disk dump) or Ventoy, but they are a bit more complicated than most casual users would probably prefer.

How about something a bit more modern and not so confusing - balenaEtcher aka Etcher? You can find it at balena.io/etcher (it’s original name was Etcher – the title I’ll use from here on).

While you could go through the hassle of adding the Debian repository to your system and then installing Etcher from there, I opted to install the appimage directly from the website.

And what’s an appimage you ask? It can be used as a standalone program on your computer or used as a portable program on an external drive. You get a twofer.

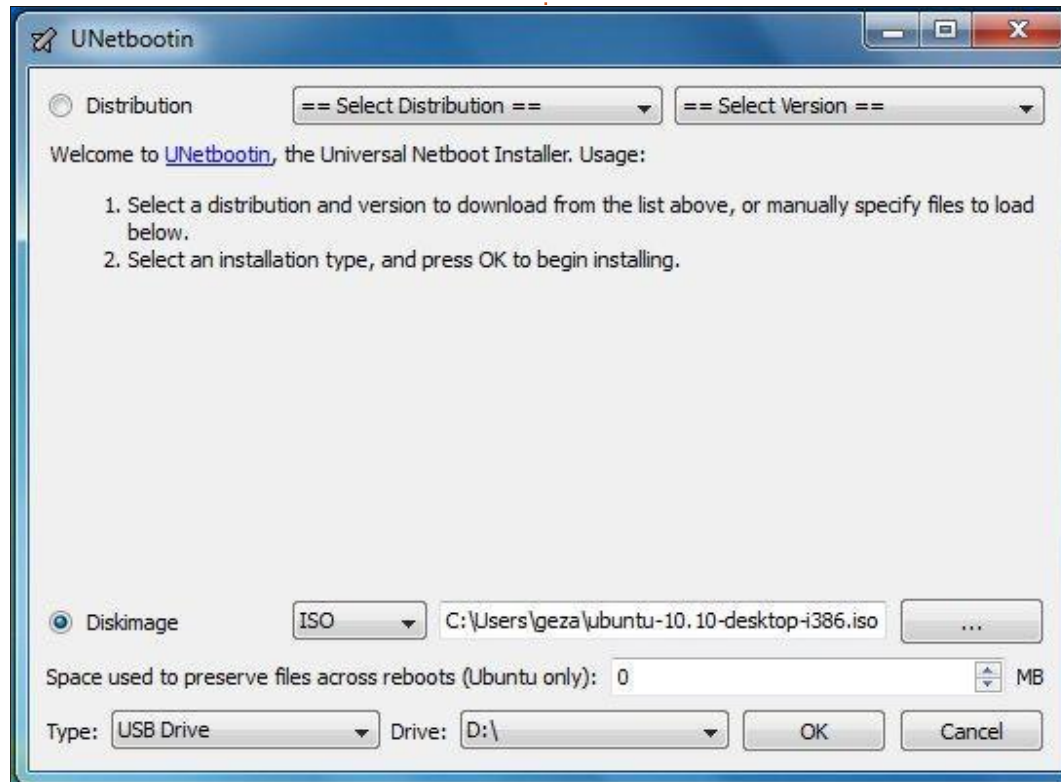
Once downloaded, you right click on the icon, go to Properties and then Permissions. Click on “Allow Executing File as a Program”.

From here it should be fairly evident what you need to do, but there are 2 options for making a bootable drive along with one for cloning your current HDD/SSD:

* **Flash from file.** In this case you download the ISO image of your choosing and then point Etcher to the location.

* **Flash from URL.** Here you provide the web address for the download and Etcher will fetch it for you.

Given those options, the first one should be your choice. Some download pages have more entries than an IHOP menu and who knows what you might get.



MY STORY

If you're looking for recommendations like UNetbootin offers, no such luck here. No downloads, either. You do the research and procurement, Etcher does the rest.

As for the Clone drive option, that's for those who already are satisfied with their current operating system and want to clone it for backup or transfer to another computer. It's incredibly simple to use but that's a discussion for another day.

Now you move on to Select target. Pick your thumb drive or flash card. As a helper it'll list the ones available. Unlike UNetbootin, Etcher has no qualms about SD cards and the like.

In fact, if you have more than one USB or flash drive mounted, you might want to remove all but the one you plan to use. Don't want to mistake one for the other.

And, finally, Flash.

Without much fanfare beyond that, you should have a bootable drive within a couple minutes.

Now for the inevitable question

– is it quicker than UNetbootin?

No, sorry. Still takes just as long, it's just prettier doing it.

Nothing is perfect and Etcher is no different. One pitfall I've found is that it stinks at housekeeping.

If you attempt to load a new OS on your flash drive, UNetbootin will overwrite anything on that drive, even previous operating systems. It may use bits and pieces of what it needs to work, but anything else is history.

Etcher doesn't bother. Old ISO images will clog up your USB drive

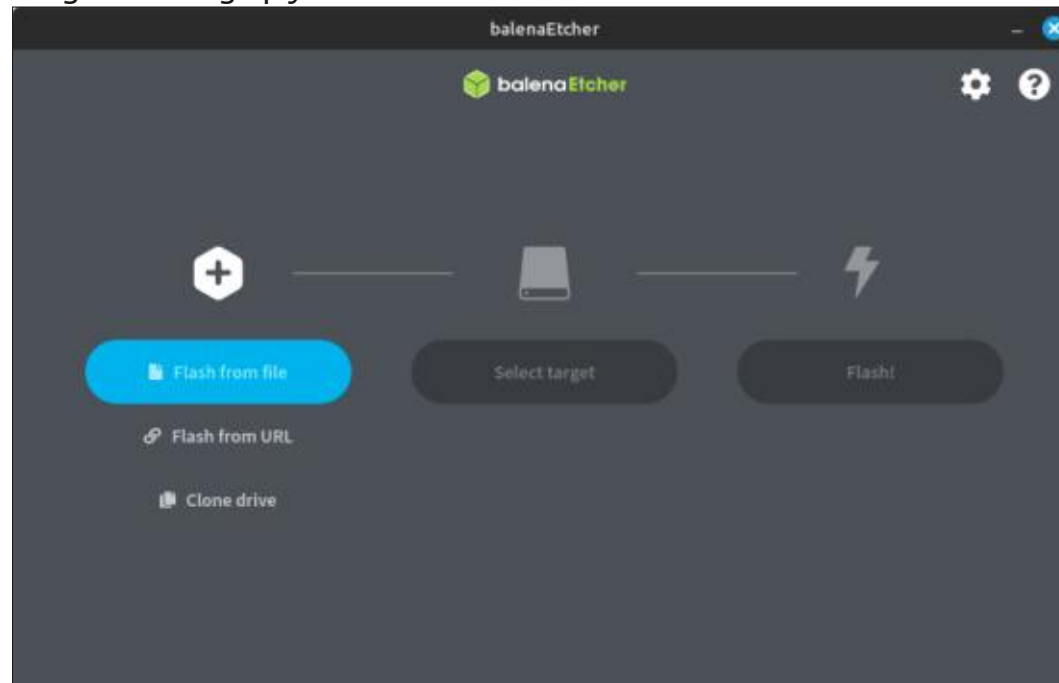
until you clear them out or reformat.

On the plus side, it will tell you that the destination drive doesn't have enough space before starting the process, but it won't help cleaning up.

So it creates a multi-boot drive then, right?

Wrong. Your computer, just like the ones I've tried it on, will only recognize the last image loaded. All the others are ghosts. They are still there, but at boot your computer will see just the last one.

But, in the long run, this is the easiest option now available to make a bootable iso image, so us distro-hoppers had better get used to it.






Tech Books Made Better




Python Crash Course, 3rd Edition
9781718502703
\$44.99 PB | 552 pages




The Recursive Book of Recursion
9781718502024
\$33.99 PB | 328 pages




Dead Simple Python
9781718500921
\$59.99 PB | 752 pages



How Linux Works, 3rd Edition
9781718500402
\$49.99 PB | 464 pages




Practical Linux Forensics
9781718501966
\$59.99 PB | 400 pages




DevOps for the Desperate
9781718502482
\$29.99 PB | 176 pages




The Rust Programming Language, 2nd Edition
9781718503106
\$49.99 PB | 560 pages



Data Structures the Fun Way
9781718502604
\$39.99 PB | 304 pages



Deep Learning
9781718500723
\$99.99 PB | 776 pages



Codebreaking
9781718502727
\$29.99 PB | 488 pages

Get 25% off your order at nostarch.com with code **FULLCIRCLE25**



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.



REVIEW

Written by Adam Hunt

Ubuntu Budgie 23.10

The Ubuntu Budgie developers have had a busy release cycle. Unlike some other Ubuntu flavor developers, the people at Ubuntu Budgie have introduced lots of new things in every release during this cycle and Ubuntu Budgie 23.10 is no exception.

Out on 12 October, 2023, Ubuntu Budgie 23.10 is the 16th version since its introduction as budgie-remix 16.04 LTS in April 2016. This is the third and last interim release of this development cycle – with the long-term-support version, Ubuntu Budgie 24.04 LTS, due on 25 April, 2024.

As an interim release, Ubuntu Budgie 23.10 is supported for just nine months, until July, 2024.

INSTALLATION

I retrieved the Ubuntu Budgie 23.10 ISO file from the official source using Transmission to get the BitTorrent. Once it had arrived, I carried out a command-line SHA256 sum check of the ISO to make sure the file was good, which

it was.

This release is 3.9 GB, which is 500 MB bigger than the last version, Ubuntu Budgie 23.04 which was 3.4 GB. There is no reason given for the sudden 13% jump in size, but most of the other flavors of Ubuntu have been similarly growing in size as well. The mainstream Ubuntu 23.10 download clocked in at 5.2 GB.

I tested Ubuntu Budgie 23.10 from a USB stick equipped with Ventoy 1.0.96. Ubuntu Budgie is officially listed as tested and

supported on Ventoy.

SYSTEM REQUIREMENTS

The recommended minimum system requirements for Ubuntu Budgie 23.10 have not changed from recent releases and remain: Intel or AMD 64-bit processor
4 GB RAM
UEFI PCs booting in CSM mode
Modern Intel-based Apple Macs


NEW

There is much new to consider in

Ubuntu Budgie 23.10 so let's jump into the list!

This release introduces the Budgie desktop 10.8, bringing improvements to the alternative Budgie Menu, and the alternative application indicator StatusNotifier which can be added to the panel. It also introduces a new window manager, Magpie 0.9.3.

There is quite a long story behind Ubuntu Budgie getting its own window manager. I promise that this is the short version! Magpie came about as a result of issues with how Gnome's own native Mutter window manager was being administered. Mutter works fine but Budgie was using a small program, budgie-wm, on top of Mutter specifically to handle some specialty Budgie features like key bindings. Last minute changes in Mutter for each release resulted in a similar requirement for last minute changes in budgie-wm which resulted in a mad rush for each release. There was also the ongoing problem of incorporating older regressions into the new



REVIEW

versions of Mutter so that Budgie could still run on X11 instead of Wayland. Over time, as Mutter moved on to Wayland support and away from X11, this got harder and harder. The resulting solution was that the Budgie team decided to “soft fork” Mutter at the point of Gnome 43, which saved a lot of regression catch-ups. Future plans for Magpie include working towards a 1.0 version with full Wayland support. This will, of course, have Budgie's own desktop needs already built in, which will greatly simplify making the Budgie desktop work right. So Magpie is not a short-term band aid solution, it should be around for a while and even if the desktop users won't notice much difference, it should

make the developers happier, which is a good thing.

In another change, the Budgie Trash Applet used to be a third party addition, developed by Buddies of Budgie team member Evan Maddock. Now it has been incorporated into the Budgie desktop and made a default feature. The applet improves file deletion functions and also file restoration, whenever needed.

More changes incorporated include PolicyKit graphical privilege escalation, battery indicator status applet selections added, and Appindicators which can be added to the panel in lieu of the Status Notifier applet. The system tray

now uses the Status Notifier specification for better presentation of the tray symbols.

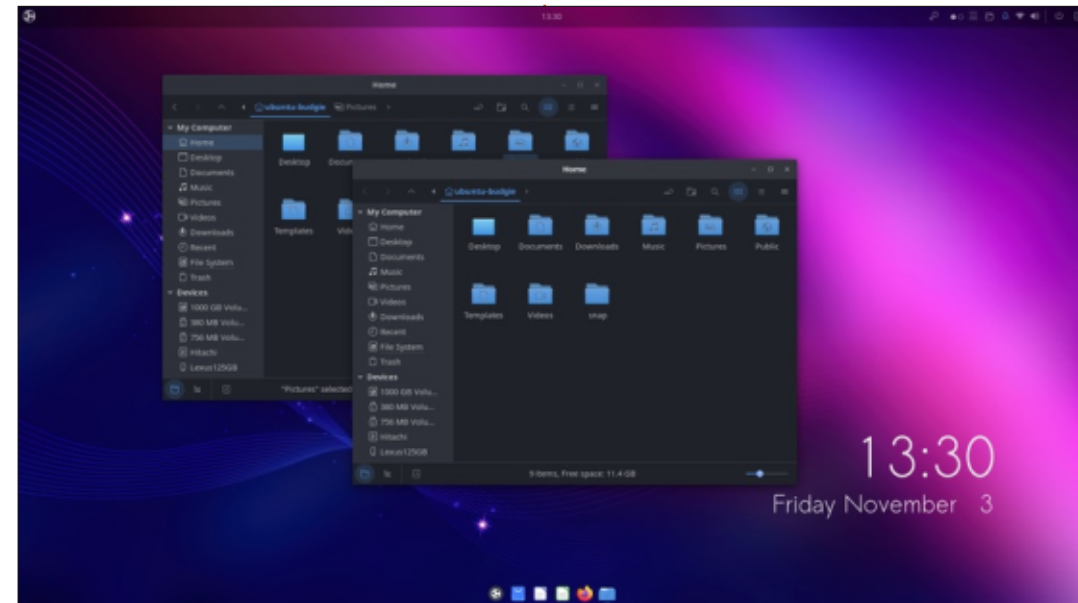
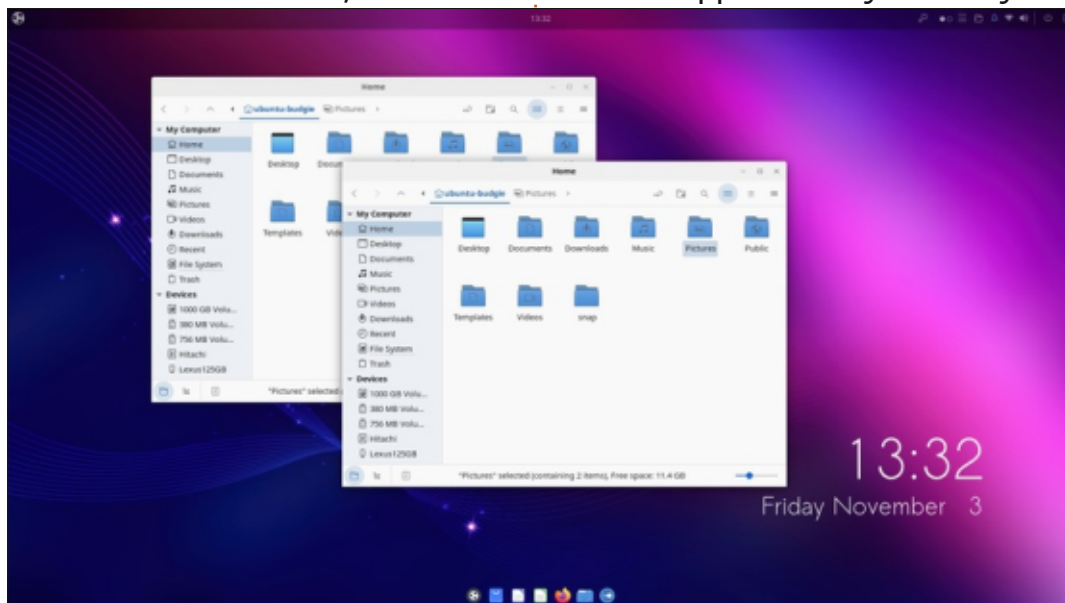
There is a new design Budgie Control Center with many upgrades and fixes.

There have been changes to the themes as well, which are now using green as the highlight color instead of blue, plus the Kvantum and Murrine themes are now hidden inside the Budgie Desktop Settings. Many of the themes are now Snap packages.

There are also many other small changes and refinements in this release.

SETTINGS

There is a sort-of-new default wallpaper, officially named "budgie-codename". This is a modified version of the standard Budgie spacey-looking wallpaper, called "ubuntu_budgie_wallpaper1" which has been in use since Ubuntu Budgie 19.04. For Ubuntu Budgie 23.04, that standard wallpaper added a lobster in honor of that release's code name, Lunar Lobster. This time, for Mantic Minotaur, the Ubuntu 23.10 Minotaur and labyrinth design has been added to the spacey wallpaper. Perhaps this is now a Budgie trend that we will see with each new release? If you are not that keen on the design



REVIEW

then Ubuntu Budgie 23.10 has 34 other wallpapers to choose from, none of which have Minotaurs on them.

Budgie Desktop Settings is where you find the window color themes under Style: Widgets. There are 14 window themes provided, with the default one remaining Pocillo-dark. Ubuntu Budgie now has some nice light themes provided that look quite good, notably Pocillo-light. Budgie Desktop Settings also has styling preferences (light or dark); ten icon styles, with Pocillo as the default; two cursor styles and four user notification positions to choose from, one in each screen corner, with top right as the default.

Budgie Makeovers & Layouts is the location for complete one-button wallpaper, window theme and icon packages, this time with nine of them to choose from. Only three of these are installed, while the remaining six have to be downloaded for use.

Budgie Extras is where the desktop applets are hidden away. On Kubuntu, these would be called "widgets" and add things like clocks, calendars, weather reports and other functionality to the desktop. In this Ubuntu Budgie release, there are 35 applets included, one fewer than in the last release.

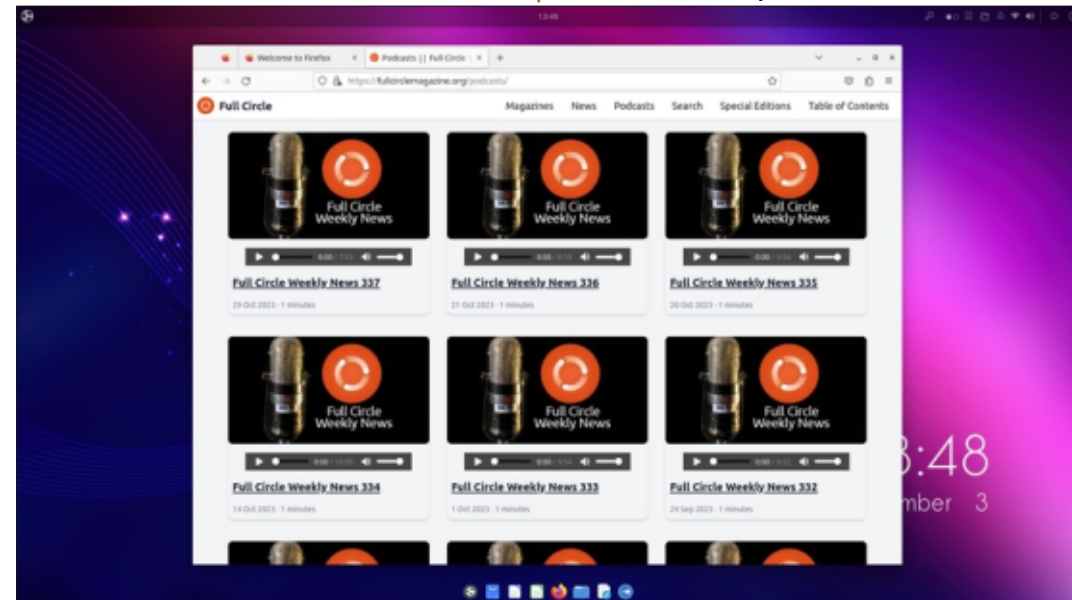
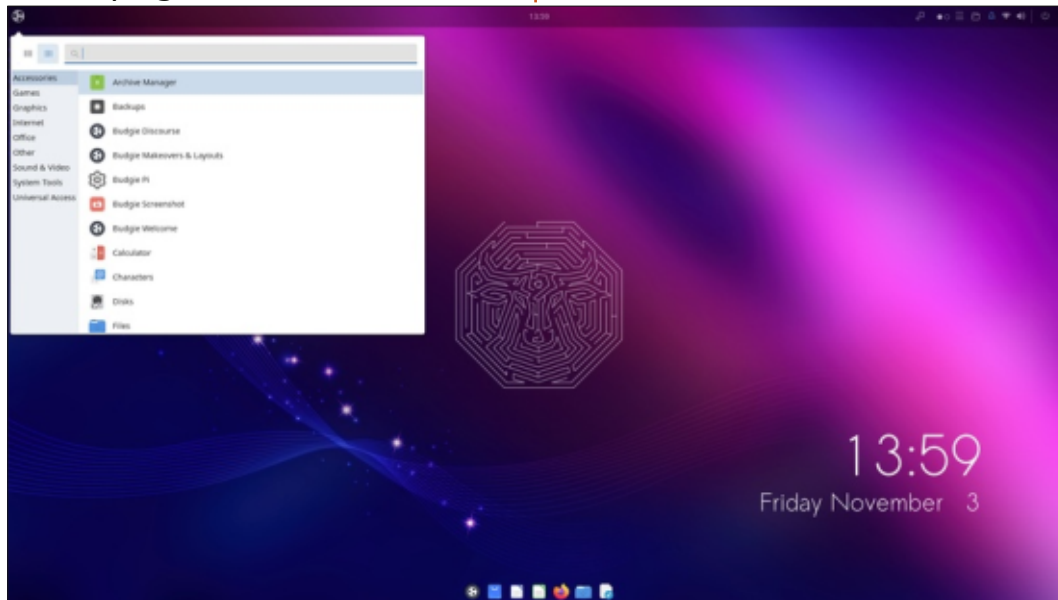
There is also the Budgie Control Center which is a modified version of Gnome Settings. Here you can configure such features as WiFi, wallpaper, sound and power settings.

Everything needed related to settings are all here in Ubuntu Budgie, but they are widely scattered. This is the one area where Ubuntu Budgie could really benefit from an effort to locate all the settings in one place. That would make life a lot easier for the users, especially for new Budgie users, trying to figure it all out.

As in past releases, there is a small Mac-like dock at the bottom of the screen. This is Plank and it

works fine, if you are the sort of user who likes docks. In normal use, it will hide if a window touches it. Settings for Plank can be accessed by right-clicking on the dock itself, as unfortunately its settings do not appear in any of the myriad of Budgie settings menus. One of the right-click choices is "quit" which closes it, but then there is no way to turn it on again other than call it up from a terminal or do a reboot. Closing Plank also means you have no list of open applications or access to minimized applications other than using alt-tab to find them.

Despite the Plank dock, Ubuntu Budgie is really a classical menu-driven desktop. The main menu can



REVIEW

be set to either show applications as tiles in alphanumerical order or as icons by category. It works well in either mode. Budgie's menu is comparable to those included with Kubuntu and Xubuntu. It is worth noting that Xubuntu used to have a dock, too, up until Xubuntu 11.10, after which it was dropped.

APPLICATIONS

Some of the applications included with Ubuntu Budgie 23.10 are:

- Archive Manager (File Roller) 43.0 archiver*
- Atril 1.26.0 PDF viewer*
- Budgie Screenshot Applet screenshot tool
- CUPS 2.4.6 printing system
- Deja Dup 44.2 back-up tool
- Document Scanner (Simple Scan) 44.0 optical scanner*
- Drawing 1.0.2 image editor*
- Firefox 118.0.1 web browser**
- Goodvibes 0.7.6 internet radio*
- Gnome Disks 45.0 disk manager
- Gnome Software 45.0 package management system
- Gparted 1.5.0 partition editor
- Gpodder 3.11.2 podcast player
- gThumb 3.12.2 image viewer*
- Guvview 2.0.8 webcam application*

- LibreOffice 7.6.2 office suite
- Lollypop 1.4.37 music player*
- Magpie 0.9.3 window manager
- Mate System Monitor 1.26.0 system resource monitor*
- Nemo 5.8.4 file manager
- Parole 4.18.0 movie player*
- Pipewire 0.3.79 audio controller
- Plank 0.11.89 desktop dock*
- Text Editor (gedit) 44.2 text editor*
- Thunderbird 115.3.1 email client
- Tilix 1.9.5 terminal emulator*
- Transmission 4.0.2 bittorrent client

* indicates same application version as used in Ubuntu Budgie 23.04
** supplied as a snap, so version depends on the upstream package manager

As can be seen, the list of default applications is very complete. Perhaps the only thing missing is a video editor but the repositories have several good ones, if desired.

After a big swap-out of the default applications in 22.10 due to libadwaita issues, there were no changes to the list in 23.04 or in this release, just some updated versions. The applications included are a mix from Gnome 43, 44 and 45, MATE, Cinnamon, plus independent projects such as

Guvview, Thunderbird and Firefox.

Ubuntu Budgie 23.10 uses the Cinnamon desktop's Nemo file manager which both looks good and works well. Its only missing feature is still bulk file renaming, which is a fairly important feature for a modern file manager to still be lacking today. There are some workarounds available, however, including standalone bulk file renamers, like GPRename, which can be installed from the repositories.

CONCLUSIONS

Ubuntu Budgie 23.10 is another good release of this popular distribution. Like the last few iterations, this version brings some useful updates that, while mostly "behind the scenes", will make it work better. Its only real downside is its widely scattered user settings. The list of applications, which saw many changes in the 22.10 release, has now been stabilized.

It will be interesting to see if the upcoming long term support release, Ubuntu Budgie 24.04 LTS due out on 25 April 2024, will continue the fast pace of changes

seen in the interim releases this cycle, or rest on its laurels.

EXTERNAL LINKS

Official website:
<https://ubuntubudgie.org/>



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.



REVIEW

Written by Adam Hunt

Wubuntu 11.4.3 LTS

Over the years, there have been several attempts to create Linux distributions that look and work like Microsoft Windows, often with the aim of luring away Windows users. Many readers may recall Lindows from 2001, which is still around today as Linspire, now a paid commercial distribution. Several others have come and gone over time, too.

A recent addition to this market is Wubuntu (Windows Ubuntu) which had its first release, numbered 11.3.22.04.2 LTS, on 25 June, 2023. Based on Ubuntu and using a modified KDE desktop, it does look very much like Windows 11.

As part of our series about Ubuntu derivatives that use the Ubuntu name, this month I thought we would have a look at this unusual distribution's second release, 11.4.3, out on 6 December, 2023.

BACKGROUND

Wubuntu comes from Brazil and

is based on Linuxfx. Both distributions are the product of the same anonymous developer and based on Ubuntu 22.04 LTS, although the KDE versions draw mostly from Kubuntu.

Wubuntu takes the Linuxfx base and adds new icons, themes, and Microsoft software to complete the Windows 11 look and feel.

The first Wubuntu release offered a choice of KDE or Cinnamon desktops, but more recent releases have stuck with a modified KDE Plasma 5 interface.



Further Cinnamon releases have been promised but not arrived. KDE has a lot of built-in user customization, but Wubuntu's developer has taken it a step further, creating a Windows-like menu and other features that would fool most seasoned Windows users at first or even second glance. Some serious effort has obviously been expended here.

GETTING WUBUNTU

I downloaded the rather heavyweight 5.1 GB ISO file from SourceForge via HTTP as there is no

BitTorrent available. There are also no SHA256 or even MD5 checksums provided, an oversight which prevents confirming that the download is valid. The first release actually did provide an MD5 sum which was better than nothing.

INSTALLING

Once downloaded, I dropped the Wubuntu ISO file onto a USB stick equipped with Ventoy 1.0.96. Wubuntu is not on the list of Ventoy tested and supported distributions but, being Ubuntu-based, it booted up smoothly.

SYSTEM REQUIREMENTS

The recommended system requirements are:
2 GHz dual-core processor or better
2 GB system memory (4 GB is highly recommended)
25 GB of free hard drive space (SSD is highly recommended)
Internet access is helpful
Either a DVD drive or a USB port for the installer media
Virtual Box: It is extremely

REVIEW

important to install guest add-ons for the system to run efficiently.

These minimum system specs are actually very close to those specified for Windows 11, although Windows has specific requirements for processor types and additional items such as Trusted Platform Module (TPM) 2.0.

TRYING OUT WUBUNTU

Wubuntu looks very much like Windows 11 when booted up. Being Linux-based though, it is more secure and stable, won't run all that Windows malware, and also doesn't come with the crapware and bloatware that the tech media has

taken Windows 11 to task for ever since its release on 5 October, 2021. Wubuntu will also run on somewhat older hardware that doesn't make the grade for Windows 11.

While Wubuntu runs well, looks good, and is easy to use, as I tried it out I was wondering "just who is this aimed at?" It clearly is not intended for seasoned Linux aficionados as they are not going to be attracted by a Linux distribution that looks like Windows and runs a considerable amount of Microsoft proprietary software. Windows is just not a "prestige brand" in the Linux world. It also is probably not going to appeal to dedicated Windows fans who are already running Windows 11 on modern


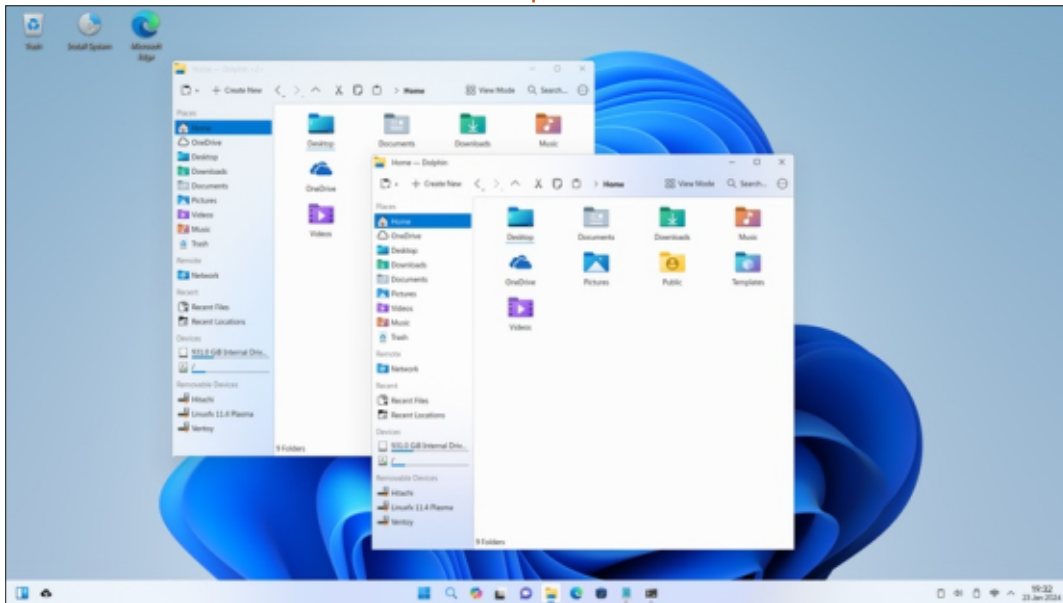
hardware, unless they are looking for better security and stability. I do think that there is a potential niche here made up of Windows users who would like to be running Windows 11 but without shelling out money to buy new hardware. These people are probably using Windows 10 for now but its mainstream support runs out on 14 October, 2025.

Wubuntu does not actually come complete as a free download. There is a package of extras included called PowerToys that can only be unlocked by purchasing the PowerToys Professional Key from the developer for a fee of \$35. The official website describes this: "PowerToys is a powerful toolset

that implements a software layer based on Microsoft Windows applications. Tools like Control Panel, Network Settings, One Drive, Android Support and many others are available through PowerToys".

So, in effect, Wubuntu is "freemium" software where the free version gives you the basics but upgrades for full functionality cost extra. The same PowerToys package is also available for Linuxfx by the same developer and is obviously being used as a source of income.

A freemium model is very unusual in the Linux desktop world and so this reinforces the notion that Wubuntu is aimed at Windows



REVIEW

users who are more used to paying for software.

On the subject of this paid extra package, the Wubuntu forum users have much to say about it. One user noted, "anytime you try to change the desktop it Nags for you to purchase Powertoys!! It's actually INFURIATING it's so often! IF I can actually get this OS setup on my PC, the way I want, I'll be happy to support your work, but right now, NO WAY! Right now, this OS is NOT ready for Prime Time!". Another user noted it was "worse than running an unactivated copy of Windows 11. Had to uninstall and move on. If you don't plan on buying the \$35 key I can't recommend Wubuntu." Another wrote, "When does the nagging stop? It is relentless, every single time." There are some indications that this issue was at least dialed back via some system updates, though.

Overall Wubuntu looks good and works well in a limited way but, unless you buy the extra PowerToys package though it is going to be short of full capabilities. That said, Wubuntu has access to the Ubuntu repositories collection of application software via APT from

the command line or from the included Discover software store, so there are lots of options for adding free software for functionality instead, much like any Linux distribution.

SETTINGS

Wubuntu comes with a built-in range of user customization, similar to Windows 11. It includes 98 wallpapers, the most of any Linux distribution I have ever seen. Many of the wallpapers are from Kubuntu releases as well as several Windows-like ones.

It also has three window color themes: light, dark, and one that is

halfway in-between, with light windows and a dark bottom panel.


From Kubuntu, it also offers 73 widgets, which are small programs that can be installed on the desktop to provide functions such as a calendar, clock or weather.

APPLICATIONS

Some of the applications included with Wubuntu 11.4.3 are:

- Android File Transfer client
- AntiMicroX 3.1.4 gamepad keyboard mapper
- Ark 23.08.3 file archiver
- ChatGPT 1.1.0 AI client
- Cheese 41.1 webcam
- Discover 2.1.2 software store

- Dolphin 23.08.3 file manager
- Google Chrome 120.0.6099.62 web browser*
- GParted 1.3.1 partition editor
- Gwenview 23.08.3 image viewer
- K3B 23.08.3 CD/DVD burner
- KCalc 23.08.3 calculator
- KMail 22.04.3 email (failed to launch)
- KTorrent 23.08.3 BitTorrent client
- KWrite 22.04.3 (failed to launch)
- Kate 23.08.3 text editor
- Konsole 23.08.3 terminal emulator
- Microsoft Copilot AI toy*
- Microsoft Edge 119.0.2151.97 web browser*
- Microsoft Office Online (Excel, OneNote, People, PowerPoint, Word)*
- Microsoft Teams chat sign-up*
- Microsoft OneDrive cloud service sign-up*
- Okular 23.08.3 PDF viewer
- OnlyOffice 7.5.1.23 online office suite (Calc, Word, PowerPoint)
- PulseAudio 15.99.1 audio controller
- Power Toys Android Support sign-up
- RemoteDesktop Client 0.23 remote desktop
- Skancelite 23.08.3 scanner
- Spectacle 23.08.3 screenshot utility
- Steam 1.0.0.74 gaming client
- Thunderbird 115.5.1 email client
- VLC 3.0.16 media player
- Webcamoid 8.8.0 webcam



REVIEW

Wine Stable 8.0.2 Windows compatibility layer
Winetricks 0.0+20210206-2 Windows compatibility layer
Wubuntu PowerToys Register

* indicates proprietary software

There are quite a lot of proprietary applications here with an emphasis on Microsoft products or launchers that give at least access to them. This will not thrill free software advocates, but is quite normal for Windows users. Some of these are Windows programs running on the included Wine compatibility layer.

It is odd these days to find a Linux distribution without the Firefox web browser and the LibreOffice office suite, but Wubuntu has neither of these. Instead, you get a choice of two proprietary browsers in their Linux versions, Microsoft Edge and Google Chrome, and two online office suites, the free software OnlyOffice and the proprietary Microsoft Office.

Aside from including two web browsers and two office suites, Wubuntu has several other redundancies such as two webcams

and two email clients.

The list of applications does show that, at its core, Wubuntu is really Kubuntu with many applications from the KDE Gear 23.08.3 collection. With the Ubuntu repositories available, it would be easy to remove the proprietary software and replace it all with free software alternatives, but instead of all that work, wouldn't it just be easier to install Kubuntu instead?

CONCLUSIONS

Wubuntu looks good and works reasonably well in its free download incarnation but, to make it work right, you need to buy a paid key for

the PowerToys package, or at least flush it out with some free software alternatives. It seems to be intended to appeal to the niche of Windows fans who have older hardware that won't run Windows 11 or perhaps for Windows users looking for better stability and security.

Most Linux fans are not going to like this distribution due to its Windows-look and reliance on proprietary software. For Linux fans, there are better options.


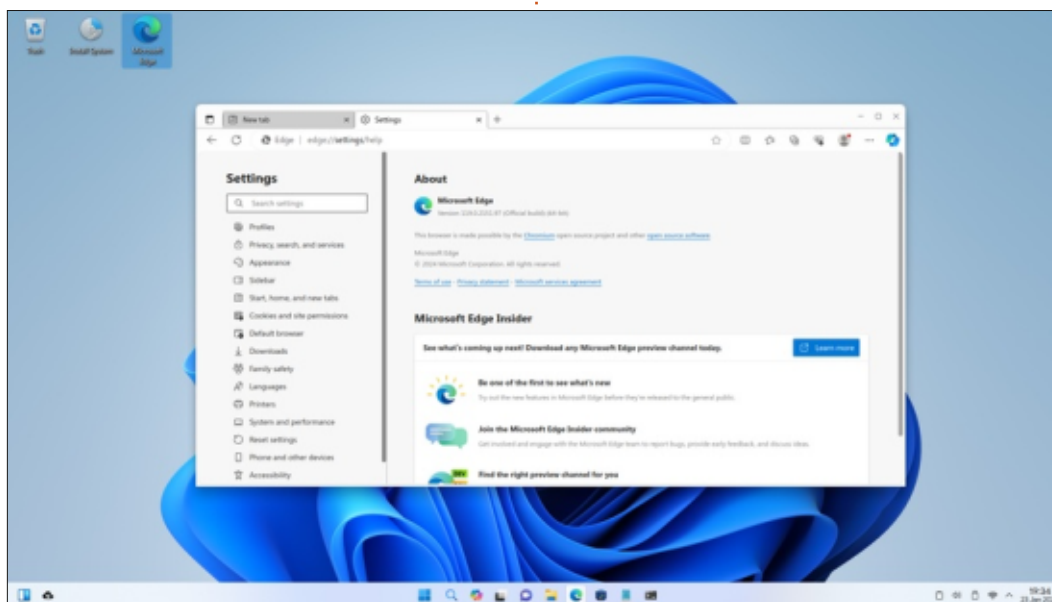
Judging by the forum traffic, the built in advertising for the paid PowerToys extras doesn't seem to be going over well among users. It will be interesting to see if

adjustments are made in this distribution, and whether or not it attracts a dedicated user base between now and the end of support for Windows 10 in October, 2025.

It is probably promising that, while I was writing this review, a new point release, 11.4.4, was issued, indicating that development is continuing.

EXTERNAL LINKS

Official website:
<https://www.wubuntu.org/>



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.

I was looking at games to review on Steam, when I came across one that seemed interesting. I had a look at the system requirements, Ubuntu 12.04, and lo and behold, the processor requirement was "potato+", I kid you not! That brought me to thinking. You see modern CPUs are a lot faster than those of fifteen or twenty years ago, and hard drives vs solid state drives offer an even bigger gap. I'm not even mentioning memory. So why is it that some things just do not get better? What gets added to

modern operating systems that make them so slow? Some people argue that it is systemd, others say it is kernel bloat, etcetera. I say, these limitations are artificially inserted. You see, some of these games play just as well on Ubuntu 6.06 as they do on 22.04. These type of games are usually written by an Indie team or someone passionate. When it is by a company, it needs to be out the door, yesterday, so they can fire all the programmers and reap the profits. So they use frameworks that are out of date as soon as the game hits the streets. This is precisely what you see with operating systems too. Microsoft and Apple. I mean that is why Linux is growing in popularity. The mainstream media always under reports Linux usage or adoption, (remember how Microsoft begged people to write to congress about "stifling innovation" during their antitrust allegations?) Anyway, modern Linux still power my old machines, my pre-core 2 CPU PC runs Bodhi Linux just fine. Probably a "potato -" processor compared to my new PC. Probably won't even

boot Windows 11 in under five minutes. You know what? I still use it to write on, I surf the interwebs and listen to audiobooks on, thanks to Linux.

Q: I can only see 3 processes when I run ps ax which is 1,10,12. On my personal laptop I can see a lot. What does that mean?

A: It means that the session you are in has a restricted namespace. My guess is probably a container rather than a full OS? Are you doing a CTF? Read up on cgroups and namespaces.

Q: I am trying to run <https://tobyfox.itch.io/deltarune> on plain Ubuntu 22.04, but it is not working. It worked on Ubuntu Mint, but that is not working with my new Nvidia 3060 any more. So now I am in a bind. My error is this --- ""libssl.so.1.1: cannot open shared object file: No such file or directory"" and I don't know where to start fixing this.

A: OK, I had a look and it seems DeltaRune has no Linux build? As it is GameMaker, I assume this is an unofficial build, as GameMaker can make Linux builds. This may help: Go to /usr/lib and make a symlink called libssl.so.1.1 to the libssl.so that you have:

```
cd /usr/lib
sudo ln -s libssl.so
libssl.so.1.1
```

or you can try this: <https://stackoverflow.com/questions/72133316/libssl-so-1-1-cannot-open-shared-object-file-no-such-file-or-directory>

Q: On Ubuntu Gnome, in the software centre, it says that the snap store must be updated, but when I do it, I get an error. Null, cannot refresh snap store, has running apps. When are they going to fix this? Snap refresh says all apps are up to date. Updaters should update, it should not be broken. I hate having to go into the terminal and type latin.

A: One cannot use a screwdriver to fix itself. You need another. If you close the Snap store, it does not close 100%, sort of like Apple apps. It is like Windows TSR so you need to go into the terminal and tell it to close completely, pkill snap-store does that for you. Now you run;

```
sudo snap refresh
```

which is the other screwdriver to fix the first. If you do not like snaps at all consider Linux Lite or Linux Mint.

Q: I have Windows on my laptop, I have downloaded an ISO to USB and installed Ubuntu on my laptop. Boot option not giving me Ubuntu. BIOS is password locked. How can I get Ubuntu now? I have a new 256gb SSD, but still I cannot boot.

A: Bioses are strange beasts, you will need to unlock it somehow, to turn off the Windows only settings. You can try some common passwords, but if you do not come right, you may have to get that BIOS chip flashed anew. It goes beyond our simple Q&A.

Q: On Kubuntu 22.04, I closed my laptop lid after an evening of setting it up just the way I like it. In the morning I opened it to find it no longer accepts my password. If I preview the password, it shows the right one, but it will not unlock? It's been 2 days now and I don't know what to do.

A: I recall seeing something about KDE themes breaking things, but I did not pay attention. My suggestion is to shut it down or try to log into another TTY. You should see from there if you are using the correct password or something is damaged. Remember you can always boot with a live boot image and change the password. You can also always: touch fsck and reboot. Password reset articles are all over the internet.

Q: I am having an issue running my VPN with Ubuntu. I open the terminal and type – rather up arrow to the string-- <removed> every time I start or open my laptop, can I not run this on a desktop icon? Ubuntu 22.04, N5105 3ghz, 4gb RAM 256GB SSD.

A: You do know that there is a graphical setting for VPN in Ubuntu 22.04? You just need to add it in the settings first, then it will show up every time you click the cluster on the top right. You can turn it on and off from there. The way you are doing it is outdated. (for quite some time actually)

Q: We use some horrible thing at work to stop people from surfing illegal things and when I am at home I still get “site blocked” for Reddit, even though I am on my own wifi. Can't Ubuntu overcome this?

A: my first answer to that is usually, use your own laptop for private things, however if it is your own device, I suggest you hardcode your DNS entries to something like GoogleDNS, 8.8.8.8 or Cloudflare 1.1.1.1 or even quad 9. You can also use Firefox's secure/encrypted DNS.

Q: Is Ubuntu becoming Windows? Every time I run an update I have to reboot or I get “gah, your

tab just crashed” in my browser.

A: That is not Ubuntu's fault, but Firefox. If you wish to continue working in your browser, when you update, simply deselect Firefox and update everything else. Once you are done with what you were doing, update Firefox and reboot.

Q: I got a mini PC for Christmas and I installed Ubuntu on it. It seems to work for the most part. The terminal is the problem. I type then like a second later, the letters appear. It is almost like it is checking what I typed before typing it in the terminal. I am using X11 and not Wayland, so it's not that. I'm kind of stumped?

A: I had a quick poke around and it seems to be a mutter issue. You can add this PPA for now until the issue is sorted. Type:

```
sudo add-apt-repository
ppa:vanvugt/mutter
```

```
sudo apt update
```

```
sudo apt upgrade
```

```
reboot.
```

Q&A

Q: I recently updated my firmware, now there is a Dell logo with 'secured by dell securebios' or some such nonsense written on it. This now shows during boot and I want it gone. I prefer not to have any writing or extra logos in my Ubuntu boot.

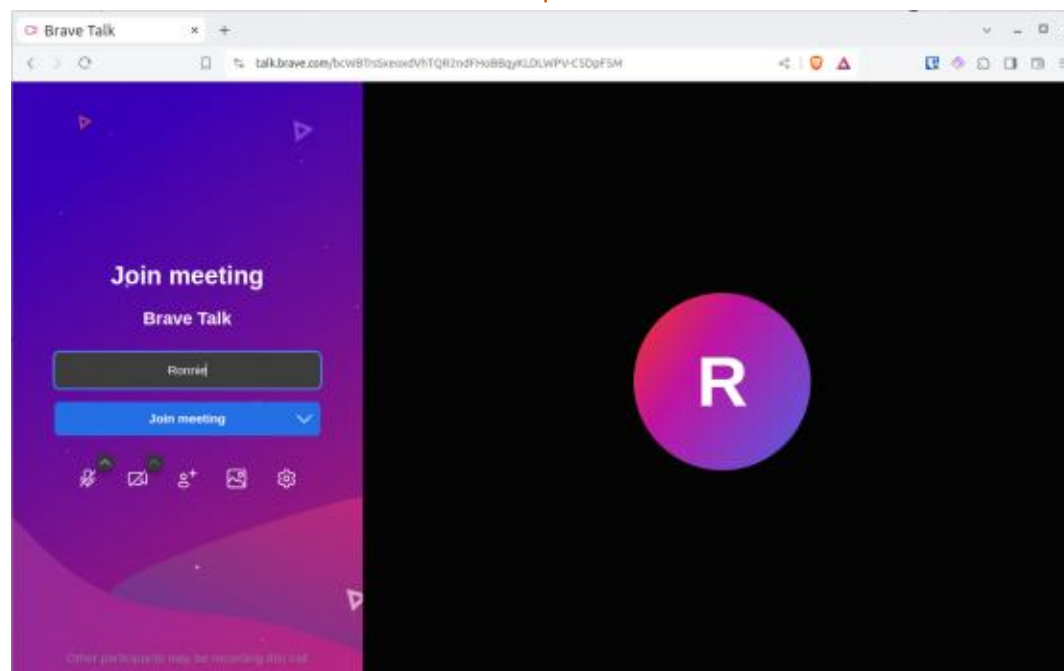
A: You may be able to turn it off in the BIOS itself, but if it is showing up during boot as well, I found this that may help: <https://askubuntu.com/questions/1238754/want-to-remove-oem-logo-during-boot-in-ubuntu-20-04>

Q: I have now lost my Skype account that I have had for many years, due to my ISP changes in the last few years. You see I created the account with my very first ISP e-mail account. That is history like 3 ISP's ago at least. Microsoft will not let me recover it without that account. Is there a Linux alternative to Skype that I can get my whole family on? Please do not say Teams.

A: There is always Jami. I almost pointed you to Jitsi. Jami is easy to install and operate and works on almost any platform. If you use Brave browser, I think they

recently launched Brave talk? I'm sure even Linphone may work for you.

Q: I have Ubuntu 23.10 installed and I noticed something different to my Ubuntu 22.04 install, in the realm of LibreOffice. On 22.04 I can double click a word, and it will highlight, allowing me to type over it. On 23.10, if I double click there, only my cursor's I-bar jumps around. If I triple click, it jumps to the end of the paragraph. Can I install the old version like 22.04 to fix it?



A: In LibreOffice, you have different types of selection. I suspect you clicked on block selection, by accident. It happens. You need to change it back to standard selection, which is double click = word, triple click = sentence. The easiest I know is to look left of the magnifier and page layout in the status bar at the bottom. It may look like your cursor's I bar or the symbol for space.

Q: How does one get those ligatures in plain bash? I want the Ubuntu icon. I see tutorials for zsh, but I prefer just plain bash. Just explain it like you are talking to a

child please, I'm just starting out with Ubuntu.

A: I'm not 100% on that, but I know if you download nerdfonts (<https://www.nerdfonts.com>), you can change your terminal default font to one of those, it is under 'appearance' and they contain what you are looking for. How you type them etcetera, is something you will need to read up on, sorry, I have never used them.



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



Website:

<https://shatteredpixel.com/>

Price: Free


Blurb: " *Shattered Pixel Dungeon is a Roguelike RPG, with pixel art graphics and lots of variety and replayability. Every game is unique, with four different playable characters, randomized levels and enemies, and over 150 items to collect and use. The game is simple to get into, but has lots of depth. Strategy is required if you want to win!*

This game is based on the freely available source code of Pixel Dungeon. It began as a project to improve on the original game's quirks, but has since evolved into a separate game, with many unique features."

You can grab Shattered Pixel Dungeon from the snap store or Flathub. I have an installer that I got from somewhere? I had a look and they all appear to be the same. Shattered Pixel Dungeon – for those of you in the know is an

Android game. If you look at various stores, there are various versions of it available too. As an android game, it works well with the one finger tap control system. However, this does not seem to translate well into Linux. Yes, we will talk about the problems first. It is December and I have four days to try out games.

UPDATE: So it turns out there was an update in January, a few days after writing this, so I grabbed the Snap, created by **Alan Pope**, and installed it to see what is what. I also grabbed the flatpak version, just to see if there is a difference.



Many people have more than one screen connected these days, I suppose I am one of the few with only one screen attached to my computer. However, I was testing a 4K external display on my laptop when I tried the game. Clicking anywhere on the second screen would immediately minimize my game. Why? This is quite annoying when you have to reply to an e-mail or instant message on the other display. The other thing that almost killed me, right out of the starting blocks, was that the music was turned up to ten when I started the game and you can only adjust it

once you can click on settings, once the game lets you. Other than that, everything seemed to work fine.

The controls were my next problem, I could use the mouse to get around, but it did not seem to register every time, so I was using the keyboard arrows, which is great, until you are faced with a trap and you need to move diagonally. Luckily the game has a special button that allows you to fight and pick up things if it does not do so automatically. If that was not there, I would have deleted the game immediately as unplayable, as clicking on enemies to attack sometimes does the weirdest things.

I want to touch on the graphics next. As you know, Shattered Pixel Dungeon is a very low-res pixel game. Some may be put off by this, but who am I to judge? I played Angband after all. To me the character is still only an "@" sign. However, that is not what I want to bring attention to. What is definitely worth mentioning is that the graphics stay crisp and fresh no

matter your screen size or zoom level. Yep, you can zoom in and out with your mouse wheel. Just like pinching the game on an Android device. The advantage of playing on a larger screen is that more of the dungeon is visible at one time.

Before you start playing Shattered Pixel Dungeon, you should know that it is huge. Not the file size, mind you, but the in-game size. Because it is procedurally generated, no play through is the same and the amount of content is delivered by the truck load. There are so many items to find and interact with, there are badges and achievements to unlock and lots of enemy types to encounter. Once you unlock the Mage, there will be crafting that you will unlock, making potions and whatnot.

You start with one character and unlock the others as you go along. There is no hand-holding, but there is an in-game manual that will pop up with helpful hints. Other than that, you are on your own discovering and exploring. Very much in the vein of the original rogue. If ever there were a “rogue-like” game, this is it.

The game is turn based, but

luckily not a one turn one movement block game. If there are no enemies, then you can move as far as you like, until you encounter an enemy. So your thinking has to be in the I go, they go, format – even if the game does not feel that way. The game relies heavily on RNG, so if you don’t make it, try again. This makes the game difficult and you need to make sure you find all the upgrades as you progress. This makes progress rewarding and frustrating. Making it to a new floor can feel like winning the game sometimes and contributes to the addictiveness of the game. Runs are usually short, unless you get really lucky, so you can play a game during

your coffee break.

The game is in constant development, I mean I have version 2.1.4 installed and version 2.2.1 is out on the Snap store and I am sure I got it in November (Editor's note: as I check this in April 2024 the dev is talking about 2.4). I don’t know if Shattered Pixel Dungeon has an end or if it is truly endless, but I know it is going to feel that way. There seems to be a resurgence in the genre lately as I was also checking out RogueCraft, a game made for the Amiga, available on itch.io: <https://badgerpunch.itch.io/roguecraft> (which looks amazing by-the-way!)

On Android and iOS you can support the developer by buying the game for a fiver. I did not see any mention on Flathub or the Snap store though. (Editor’s note: it’s also on Itch.io, Steam and GoG.)

I say give the game a go, but don’t give up after your first ten times, it is only after about ten games, that you go from frustrated to hooked.



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
 Carl Andersen
 CBinMV
 Darren
 Dennis Mack
 Devin McPherson
 Doug Bruce
 Elizabeth K. Joseph
 Eric Meddleton
 Gary Campbell
 George Smith
 Henry D Mills
 Hugo Sutherland
 Jack
 Jack Hamm
 Jason D. Moss
 Joao Cantinho Lopes
 John Andrews
 John Malon
 John Prigge
 Jonathan Pienaar
 Joseph Gulizia
 JT
 Katrina
 Kevin O'Brien

Lee Allen
 Lee Layland
 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

SINGLE DONATIONS

2024:
 Louis W. Adams, Jr.
 Sergio Arroyos
 Brian Kelly
 Linda Prinsen
 Christophe CARON
 Borso Zsolt

Ennio Quattrini
 Kimberly James Kulak

The current site was created thanks to **Arun** (from our Telegram channel) who took on the task of completely rebuilding the site, from scratch, in his own time.

The Patreon page is to help pay the domain and hosting fees. The money also helps with the new mailing list.

Several people have asked for a PayPal (single donation) option, so I've added a button below.

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#205

Deadline:
Sunday 12th May 2024.
Release:
Friday 31st May 2024.



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



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



and via TuneIn at: <https://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 - From Feb. 2023 Issuu are removing all free items with 50+ pages. So most (if not all) issues of FCM will vanish.

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