



# Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

ISSUE #189 - January 2023




Help

Part Design




Start Part

Create body



## FreeCAD NEW SERIES




## HowTo



Python p.17




Multiple Drives p.20



Latex p.27



FreeCAD p.34



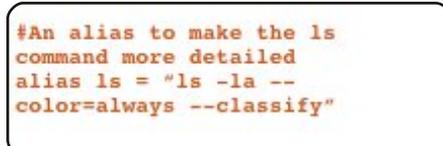
Inkscape p.30

## Graphics



# Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY



Command & Conquer p.15



... p.XX



Ubuntu Devices p.XX



The Daily Waddle p.37




My Opinion p.42



Letters p.51



Q&A p.52



Review p.47



Linux News p.04



Everyday Ubuntu p.XX



Book Review p.50



Review p.44



Ubuntu Games p.56



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

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



## WELCOME TO THE LATEST ISSUE OF FULL CIRCLE

This month we have; Python, Latex, and Inkscape for you. No Blender again this month as Erik has been a very busy panda. Filling the Blender void is an interesting article on setting up multiple hard drives. A new series starting this month is one that I threatened a while back; FreeCAD. Full disclaimer: I'm no expert with FreeCAD (not even close) so it'll be a short series. I expect about ten parts.

As the first post says: it's only taken fifteen years, but we finally have a new site! Yes, the beta is now the default site for FCM. **BIG HUGE thanks to Arun** (from the FCM Telegram group) for making the site for us. It's MUCH faster than the old site, which was an old, bloated, Wordpress site that was always under attack. New year, new site!

In other news: Issuu.com is implementing a hard limit to all files uploaded to it. Anything over 50 pages will be removed from the site. So, unfortunately, most issues of FCM will vanish from it in February. An alternative source is my MEGA folder which has every issue: <https://bit.ly/fcm-mega>

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.

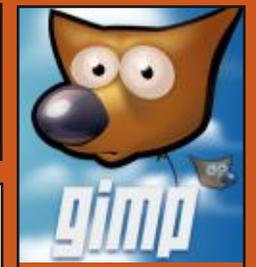
**All the best for 2023!**

*Ronnie*

[ronnie@fullcirclemagazine.org](mailto:ronnie@fullcirclemagazine.org)



This magazine was created using :



## Find Full Circle on:



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



[twitter.com/#!/fullcirclemag](https://twitter.com/#!/fullcirclemag)



<http://issuu.com/fullcirclemagazine>



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

## Weekly News:



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



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



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

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



## RELEASE OF SHOTCUT

### 22.12:

26/12/2022

Shotcut 22.12, which is being developed by the author of the MLT project and uses this framework for video editing, is available. Support for video and sound formats is implemented via FFmpeg. You can use plugins with the implementation of video and audio effects, compatible with Frei0r and LADSPA. In one of the features of Shotcut, you can do multi-track editing with the layout of video from fragments in various source formats, without the need for them to be imported for recoding. There are built-in tools for creating screencasts, processing images from a web camera and receiving streaming video. Qt5 is used to build the interface. The code is written in C++ and distributed under the GPLv3 license. Ready-made builds are available for Linux (AppImage, flatpak and snap), macOS and Windows.

<https://www.shotcut.org/blog/new-release-221221/>

## RELEASE OF SWAY 1.8:

26/12/2022

After 11 months of development, a release of the composite manager. Sway 1.8, built using the Wayland protocol and fully compatible with the mosaic window manager, i3 and the i3 panel, was published. The project code is written in C and distributed under the MIT license. The project is aimed at Linux and FreeBSD.

Compatibility with i3 is provided at the command level, configuration files and IPC, which allows you to use Sway as a transparent replacement of the i3, using Wayland instead of X11. Sway allows you to place windows on the screen not spatially, but logically. Windows are arranged, forming a grid that optimally uses screen space and allows you to quickly manipulate windows only with the help of a keyboard.

<https://swaywm.org/>

## AV LINUX MX 21.2:

26/12/2022

AV Linux MX 21.2, containing a selection of applications for creating/processing multimedia content, is available. The distribution is built from source using the toolkit used to build MX Linux as well as its own build packages (Polyphone, Shuriken, Simple Screen Recorder, etc.). AV Linux can operate in Live mode and is available for x86\_64 (3.9 GB).

The user environment is based on Xfce. It includes Ardour, Harrison Mixbus, Blender, Cinelerra, OpenShot, LiVES and tools for converting multimedia file formats. JACK Audio Connection Kit is offered for switching sound devices (JACK1/Qjackctl is used, not JACK2/Cadence). The distribution is accompanied by a detailed illustrated manual (PDF, 72 pages)

<http://www.bandshed.net/2022/12/15/avl-mxe-21-2-1-mxde-efl-21-2-2-released/>

## MXDE-EFL 21.2:

26/12/2022

The release of the MXDE-EFL 21.2 build, based on the experience of MX Linux and supplied with a desktop based on the Enlightenment environment, was announced. The project being developed by AV Linux developers is positioned as an build experimenting with the transfer of AV Linux from the Xfce desktop to Enlightenment. The build contains basic optimizations and settings of AV Linux, but it has a smaller set of specialized applications. The live-image size is 3.8 GB.

<http://www.bandshed.net/mxde-efl-build/>

## DAPHILE 22.12:

26/12/2022

Daphile 22.12, based on Gentoo Linux and designed to create a system for storing and reproducing a music collection, is out. To ensure maximum sound quality, the computer's connections from Daphile to analog amplifiers are supported through digital-analog converters with USB interface, to create multizone audio systems. The distribution can also run in an audio server, network storage (NAS, Network-Attached Storage) and wireless access points. It supports playback from internal drives, network streaming services and external USB drives. Control is through a specially created web-interface. Three builds are offered: x86\_64 (278 MB), i486 (279 MB) and x86\_64 with real-time (279 MB) components.

<https://www.daphile.com/>

## RELEASE OF APACHE NETBEANS 16:

27/12/2022

The Apache Software Foundation introduced an integrated development environment - Apache NetBeans 16, which provides support for programming languages like Java, Java EE, PHP, C/C++, JavaScript and Groovy. The finished builds are for Linux (snap, flatpak), Windows and macOS.

<https://blogs.apache.org/netbeans/entry/announce-apache-netbeans-16-released>

## A NEW MAJOR BRANCH OF THE MARIADB 11 DBMS:

28/12/2022

10 years after the founding of the 10.x branch, MariaDB 11.0.0 has been released, which offers several significant improvements and changes that break compatibility. The branch is alpha release and will become ready for work after stabilization. The next major branch of MariaDB 12, containing changes that break

compatibility, is expected no earlier than 10 years (in 2032).

A key improvement in the MariaDB 11 branch is the translation of the query optimizer to a new weight model (cost model), providing a more accurate prediction of the weights of each plan of fulfillment of the query. Despite the fact that the new model allows you to get rid of some bottlenecks in performance, it is possible that it will not be optimal in all scenarios and may slow down some requests, so users are recommended to take part in testing and notify developers in case of problems.

<https://mariadb.org/mariadb-11-0-new-optimizer/>

## RELEASE OF NTPSEC 1.2.2:

29/12/2022

After a year and a half of development, the release of the NTPsec 1.2.2 accurate time synchronization system, which is the foreline of the reference implementation of the NTPv4 protocol (NTP Classic 4.3.34), focused on the processing of the

code base in order to improve security (extinguished code is cleaned, methods of preventing attacks and protected functions for working with memory and strings) were used. The project is being developed under the direction of Eric Raymond with the participation of some developers of the original NTP Classic, engineers from Hewlett Packard and Akamai Technologies, as well as GPSD and RTEMS projects. NTPsec's source codes are distributed under the BSD, MIT and NTP licenses.

<https://blog.ntpsec.org/2022/12/29/version-1.2.2.html>

## CALCULATE LINUX 23:

29/12/2022

Calculate Linux 23, developed by the Russian-speaking community, built on Gentoo Linux, supporting the continuous cycle of release of updates and optimized for rapid deployment in the corporate environment, was announced. The new version includes the server edition of Calculate Container Manager for work with LXC, added a new cl-lxc utility, added support for the

selection of the update repository.

The following distribution editions are available: Calculate Linux Desktop with KDE (CLD), MATE (CLDM), LXQt (CLDL), Cinmon (CLDC) and Xfce (CLDX), Calculate Container Manager (CCM), Calculate Directory Server (CDS), Calculate Linux Scratch (CLS) and Calculate Scratch Server (CCM). All versions of the distribution are distributed as a bootable Live-image for x86\_64 systems with the ability to install a hard disk or USB drive.

Calculate Linux is compatible with Gentoo portage, uses the OpenRC initialization system, and uses a Rolling update model. The repository has more than 13,000 binary packages. Live USB includes both open and proprietary video drivers. Multiloader and modification

of the boot image with Calculate utilities is supported. The system supports the Calculate Directory Server domain with centralized authorization in LDAP and storage of user profiles on the server. The composition includes the Calculate project a selection of utilities for customization specially designed for it, to build and install the system. Tools are provided for the creation of specialized ISO-image adapted for the user's tasks.

<https://forum.calculate-linux.org/t/calculate-linux-23/11297>

## RELEASE OF HANDBRAKE 1.6.0:

30/12/2022

After a year of development, the release of the tool for multithreaded transcoding of video

files from one format to another - HandBrake 1.6.0, was announced. The program is available in both the command line version, and in the form of a GUI interface. The project code is written in C (for Windows GUI implemented on .NET) and is distributed under the GPL license. Binary builds are prepared for Linux (Flatpak), macOS and Windows.

The program can recode video from BluRay/DVD disks, copies of VIDEO\_TS directories and any files that are supported by libavformat and libavcodec libraries from FFmpeg. For output, files can be generated in formats such as WebM, MP4 and MKV, video codecs AV1, H.265, H.264, MPEG-2, VP9 and Theora. It can be used for video encoding, AAC, MP3, AC-3, Flac, Vorbis and Opus. Additional functions include: bitrate calculator, preview during the encoding

process, resizing and scaling the image, integrator of subtitles, a wide range of conversion profiles for specified types of mobile devices.

<http://handbrake.fr/>

## RELEASE OF SIDUCTION 2022.1:

30/12/2022

After a year of development, the release of the Siduction 2022.1 project was announced. They are developing a desktop-oriented Linux distribution, built on the Debian Sid (unstable) package. Siduction is an Aptosid fork, since July 2011. The key difference from Aptosid was the use as a custom environment of a newer version of KDE from the experimental Qt-KDE repository. For download, builds based on KDE (3.1 GB), Xfce (3 GB) and LXQt (1.8 GB), as well as a minimalistic assembly "noX" (1 GB), supplied without a graphical environment and designed for users who want to customize their system, is on tap. Unofficial assemblies with GNOME, Cinnamon and MATE are assembled separately.



# DistroWatch.com

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

<https://siduction.org/2022/12/release-notes-for-siduction-2022-1-masters-of-war/>

## VANILLA OS STABLE RELEASE:

30/12/2022

The first stable release of the user distribution Vanilla OS, based on Ubuntu, but goes beyond the usual respin, is presented. As a user environment, GNOME is used, which is offered in the form in which it is originally produced by developers, without changing the settings. Iso-image size is 1.7 GB.

<https://vanillaos.org/2022/12/29/vanilla-os-22-10-kinetic.html>

## RELEASE OF IPFIRE 2.27 CORE 172:

31/12/2022

The release of a distribution for the creation of routers and firewalls IPFire 2.27 Core 172 has been published. IPFire is distinguished by the simple installation process and the

organization of configuration through an intuitive web-interface, with graphics. The size of the installation iso-image is 373 MB (x86\_64, AArch64).

The system is modular, in addition to the basic functions of packet filtering and traffic management for IPFire, modules are available with the implementation of a system to prevent attacks based on Suricata, for the creation of a file server (Samba, FTP, NFS), mail server (Cyrus-IMAPd, Postfix, Spamassassin, ClamAV and Openmailadmin) and print server (CUPS), VoIP gateway, and Asteri, and audio streaming and video streamingIcecast, Gnump3d, VDR. A special Pakfire package manager is used to install IPFire add-ons.

<https://blog.ipfire.org/post/ipfire-2-27-core-update-172-released>

## RELEASE OF NCURSES 6.4:

01/01/2023

After two years of development, the ncurses 6.4 library, designed to create multi-platform interactive console user interfaces

and supports the emulation of the curses software interface from System V Release 4.0 was presented. The release of ncurses 6.4 is compatible at the source code level with ncurses 5.x and 6.0 branches, but extends the ABI. Popular applications built using ncurses, include aptitude, lynx, mutt, ncftp, vim, vim, vifm, minicom, screen, tmux, emacs, less.

The new version mainly notes the corrections of errors and the work to improve reliability and efficiency. The OBD terminals have added descriptions of mosh, mosh-256color, taken-16color, taken-sc, taken-vt and xgterm, as well as updated descriptions of many previously present terminals.

<https://www.mail-archive.com/info-gnu@gnu.org/msg03127.html>

## RELEASE OF CARBONOS 2022.3:

01/01/2023

The release of carbonOS 2022.3, built using the atomic layout model the system, in which the base environment is supplied as a single whole, not broken into

separate packages, was published. Additional applications are installed in Flatpak format and run in isolated containers. The size of the installation image is 2.2 GB. The project is distributed under the MIT license.

<https://carbon.sh/blog/2022-12-31-release.html>

## RELEASE OF RIGS OF RODS 2022.12:

03/01/2023

The Rigs of Rods project, that develops a realistic simulator of cars, ships, aircraft, helicopters, trains and other transportation, has been published. The project involves the OGRE graphic engine and its own physics engine of deformable bodies to simulate the movement, damage and deformation of vehicles. The code is written in C++ and distributed under the GPLv3 license.

The simulator began its development as a research project of Pierre Michel-Ricordel in the field of soft body physics. After the project attracted the attention of the community for the simulator,

thousands of different maps and vehicles were created. The user is provided with an environment with complete freedom of action and the ability to use different maps and modes of transport with a realistic model of movement and damage, right up to the simulation of complete destruction.

<https://github.com/RigsOfRods/rigs-of-rods/releases/tag/2022.12>

## RELEASE OF DRAGONFLY BSD 6.4:

03/01/2023

After a year of development, the release of DragonFlyBSD 6.4, an operating system with a hybrid core, created in 2003 for the purpose of alternative development of the FreeBSD 4.x branch, was published. From the features of DragonFly BSD, you can select a HAMMER distributed version file system, support for downloading the "virtual" cores of the system as user processes, the ability to cache data and meta-data of the FS on SSD-drives, taking into account the context of the variable symbolic links, the ability to freeze processes with saving their state on

the disk, a hybrid core using lightweight streams (LWW).

<http://www.dragonflybsd.org/release64/>

## RELEASE OF LIGHTTPD 1.4.68:

03/01/2023

The lightweight spring server lighttpd 1.4.68 was released, trying to combine high performance, security, compliance with standards and customization flexibility. Lighttpd is suitable for use on high-load systems and targets low memory and CPU resources. The project code is written in C and distributed under the BSD license.

<https://blog.lighttpd.net/articles/2023/01/03/lighttpd-1.4.68-released/>

## RELEASE OF BUSYBOX 1.36:

04/01/2023

The release of the BusyBox 1.36 package with the implementation of a set of standard UNIX utilities, designed as

a single executable file and optimized for the minimum consumption of system resources. (with a set of less than 1 MB.) The first release of the new branch 1.36 is positioned as unstable, complete stabilization will be provided in version 1.36.1, which is expected in about a month. The project code is distributed under the GPLv2 license.

<https://busybox.net/news.html>

## RELEASE OF THE NITRUX 2.6:

04/01/2023

The release of Nitrox 2.6.0, built on Debian, KDE technologies and the OpenRC initialization system, has been published. The project offers its own desktop; NX Desktop, which is an add-on over the user environment of KDE Plasma. Based on the Maui library for the distribution, a set of typical user applications is developed, which can be used on both desktop and mobile devices. AppImages is being promoted to install additional applications. The size of the full image is 3 GB. The project is distributed under free licenses.

The NX Desktop desktop offers a different style design, its own implementation of the system tray, notification output center and various plasmoids, such as a network connector and multimedia applet to control the volume and control the playback of multimedia content. Applications created using the MauiKit framework, you can note the file manager Index (you can also use Dolphin), text editor Note, Station terminal emulator, VVave music player, Clip video player, NX Software Center application center and Pix image viewer.

<https://nxos.org/changelog/release-announcement-nitrox-2-6-0/>

## PINTA 2.1:

04/01/2023

After a year of development, the release of the open raster graphics editor Pinta 2.1 was published, which is an attempt to rewrite the Paint.NET program using GTK. The editor provides a basic set of utilities for drawing and image processing, focusing on

novice users. The interface is simplified as much as possible, the editor supports an unlimited buffer of undo's and redo's, it allows you to work with several layers and is equipped with a set of tools for overlaying various effects and adjusting images. The Pinta code is distributed under the MIT license. The project is written in C, using Mono and Gtk. Binary builds are prepared for Linux (Flatpak, Snap), macOS and Windows.

<https://www.pinta-project.com/releases/2-1>

## RELEASE OF BLINK:

05/01/2023

The blink project is developing a new emulator for x86-64 processors, which allows you to run statically assembled Linux applications in a virtual machine with an emulated processor. The main purpose of the project is to enable the launch of x86-64 Linux programs collected for architecture in other operating systems (macOS, FreeBSD, NetBSD, OpenBSD) and on equipment with other hardware architecture (x86, ARM, RISC-V, MIPS, PowerPC, s390x). The project

code is written in C (ANSI C11) and is distributed under the ISC license. Dependencies only required libc (POSIX.1-2017).

<https://hiro.codes/read/emulating-an-emulator-inside-itself.-meet-blink>

## FIREWALLD 1.3 RELEASED:

06/01/2023

firewalld 1.3, is out, to prevent: "If a packet hits a pocket on a socket on a port, and the bus is interrupted as a very last resort, and the address of the memory makes your floppy disk abort, then the socket packet pocket has an error to report!" (Happy new year, Moss!) Firewalld is running as a background process that allows you to dynamically change the rules of the packet filter through D-Bus, without the need to restart the rules of the packet filter and without breaking the installed connections. The project is already used in many Linux distributions, including RHEL 7+, Fedora 18+ and SUSE/openSUSE 15+. The firewalld code is written in Python and distributed under the GPLv2 license.

To control the firewall, the firewall-cmd utility is used, which, when creating the rules, does not block IP addresses, network interfaces and port numbers, but the names of services (for example, to open access to SSH you need to perform "firewall-cmd --add --ssesh," to close SSH - "firewall-cmd --remove --servicesesh"). The firewall-config (GTK) and firewall-applet (Qt) graphical interface can also be used to change the configuration of the firewall screen. Firewalld control support via the D-BUS API is available in projects such as NetworkManager, libvirt, podman, docker and fail2 projects

<https://github.com/firewalld/firewalld/releases/tag/v1.3.0>

## RELEASE OF MYLIBRARY 2.0.1:

06/01/2023

The home library, MyLibrary 2.0.1, was released. The program code is written in the C++ and is available under the GPLv3 license. The graphical user interface is implemented using the GTK4 library. Currently, it is available for

Windows and Linux. A ready-made package is available for Arch Linux users in AUR.

MyLibrary catalogs the files of books in fb2, epub, pdf, djvu format, both available directly and packaged in zip archives, and creates its own database without changing the source files and changing their positions. Control of the integrity of the collection and its changes is carried out by creating a database of hash-totum files and archives.

<https://github.com/ProfessorNavigator/mylibrary>

## RELEASE OF OPENIPC 2.3:

07/01/2023

After a year of development, the release of the OpenIPC 2.3, Linux distribution, designed for use in CCTV cameras instead of regular firmware, most of which eventually cease to be updated by manufacturers, was announced. The project is distributed under the MIT license. Firmware images are prepared for IP cameras based on Hisilicon Hi35xx, Goke GK7205\*, Ingenic T31\*, SigmaStar SSC335 and

XiongmaiTech XM510/XM530/XM550 chips.

The proposed firmware provides features such as support for hardware motion detectors, the proprietary implementation of the RTSP protocol for the distribution of video from one camera to more than 10 customers at the same time, the ability to use hardware support for h264/h265 codecs, support for sampling speed up to 96KHz, the ability to recode JPEG images on the fly for overclocking (progressive) and support for the RAW format.

<https://openipc.org/about/>

## THE FIRST RELEASE OF THE ROLLING DISTRIBUTION OPENMANDRIVA Lx ROME: 07/01/2023

The OpenMandriva project presented the first release of the new edition of OpenMandriva Lx ROME (23.01), which uses a model of continuous delivery of updates (rolling-releases). The proposed version allows you to access new versions of packages developed for the OpenMandriva

Lx 5 branch, without waiting for the classical distribution. To download, Iso-images @ 2.8 GB with KDE and GNOME desktops that support Live mode are available.

<https://www.openmandriva.org/en/news/article/openmandriva-rome-the-rolling-release>

## RELEASE OF LIBDBX 0.12.3: 08/01/2023

The libmdbx 0.12.3 (MDBX) library was released with the implementation of a high-performance compact built-in database of the key-meaning class. The libmdbx code is distributed under the OpenLDAP Public License. All relevant operating systems and architecture are supported. libmdbx offers a developed C++ API, as well as binding bindings supported by enthusiasts to the languages Rust, Haskell, Python, NodeJS, Ruby, Go, Nim, Deno, Scala.

<https://gitflic.ru/project/erthink/libmdbx/release/35a2dcd5-7b7f-48a2-b038-74f5df96bfb8>

## RELEASE OF THE OBS STUDIO 29: 08/01/2023

OBS Studio 29, for streaming, compositing and video recording, is out. The code is written in C/C++ and is distributed under the GPLv2 license. It is available for Linux, Windows and macOS.

The purpose of the development of OBS Studio was to create a portable version of the Open Broadcaster Software (OBS Classic) application, not tied to the Windows platform, which supports OpenGL and is extensible through plugins. The difference is also the use of modular architecture, which involves the separation of the interface and the core of the program. It supports transcoding of original streams, video capture during games and streaming in PeerTube, Twitch, Facebook Gaming, YouTube, DailyMotion, Hitbox and other services. To ensure high performance, hardware acceleration mechanisms (e.g., NVENC and VA-API) can be used.

Support is provided for compositing with the construction

of the scene based on arbitrary video streams, data from web cameras, video capture maps, images, text content of applications or the entire screen. In the process of broadcasting, you can switch between several predefined scenes. The program also provides tools for mixing sound, filtering with VST plugins, leveling volume and noise suppression.

<https://github.com/obsproject/obs-studio/releases/tag/29.0.0>

## RELEASE OF DBMS LIBDBX 0.12.3: 08/01/2023

The libmdbx 0.12.3 (MDBX) library was released with the implementation of a high-performance compact built-in database of the key-meaning class. The libmdbx code is distributed under the OpenLDAP Public License. libmdbx offers a developed C++ API, as well as bindings supported by enthusiasts of Rust, Haskell, Python, NodeJS, Ruby, Go, Nim, Deno, Scala.

Technologically libmdbx offers ACID, a strict serialization of

changes and non-blocking reading with linear scaling on the cores of the CPU. It supports autocompact, automatic control of the size of the database, estimation of sample volume by range (range query estimation).

<https://gitflic.ru/project/erthink/libmdbx/release/35a2dcd5-7b7f-48a2-b038-74f5df96bfb8>

## UPDATE OPENWRT 22.03.3:

09/01/2023

There was an update of the OpenWrt 22.03.3 distribution, focused on application in various network devices, such as routers, switches and access points. OpenWrt supports many different platforms and architectures and has a build system that allows you to easily and conveniently produce cross-compiling, including various components in the build, which allows you to easily form a ready-made firmware or disk image with the desired set of preset packages.

<https://lists.infradead.org/pipermail/openwrt-devel/2023-January/040194.html>

## RELEASE OF FREECOL 1.0:

10/01/2023

After 20 years of development, the first stable release of the turn-based strategic game FreeCol 1.0 was announced. The game continues the development of the old Colonization game and offers gameplay to colonize undeveloped territories and create a new nation. The code is written in Java and distributed under the GPLv2 license.

<https://www.freecol.org/news/freecol-1.0.0-released.html>

## SOURCE CODE OF DUELYST GAME IS OPEN:

10/01/2023

Counterplay Games announced the opening of the original code of the game Duelyst, combining the properties of the card game and the turn-based strategy, in which the two players fight on the tactical battlefield, in turn pulling out cards with creatures and spells. The code is written in JavaScript and is open as

a public domain licensed CC0 1.0 (Creative Commons Zero v1.0 Universal).

Ready-made client builds are prepared for Linux, Windows and macOS. The server can be launched on your facilities or use the environment staging.duelyst.org, which also supports online game play. Further development of the game will be continued by the community within the framework of the OpenDuelyst project and anyone can join.

[https://www.youtube.com/watch?v=h3Cuz6d\\_7x0](https://www.youtube.com/watch?v=h3Cuz6d_7x0)

## HAMMER2 FILE SYSTEM PORT AVAILABLE FOR NETBSD AND FREEBSD:

11/01/2023

The initial port versions of the HAMMER2 file system for NetBSD and FreeBSD have been published. Currently, ports are still in read mode, but in the future they plan to implement NetBSD support and write support (for FreeBSD, they do not intend to add an entry support).

HAMMER2 is from DragonFly BSD developers and is notable for features such as separate snapshot mounting - available for recording snapshots, directory quotas, incremental mirroring, support for various data compression algorithms, the ability to combine multiple disk volumes, repair errors during mount (without fsck), multi-master mirroring with data distribution.

<https://github.com/kusumi/>

## INDIANS SEEK TO RENAME APACHE PROJECT:

12/01/2023

The non-profit Natives in Tech, a technology ecosystem for indigenous peoples, has called on the Apache Software Foundation to stop using the word "Apache" and Indian symbols in the logo, including renaming all projects developed under this name. To promote the initiative, a petition was published, which was signed by 40 activists. Among those who supported the petition, Bradley M. Kuhn, founder of Software Freedom Conservancy (SFC), Josh Simmons, Chairman of the Board of

Directors of OSI (Open Source Initiative), and Erin Stein, Head of Data and Tech for Good. This is what you get if you put out a code of conduct. Politics should not interfere in things that have nothing to do with them. Are they going to rename the helicopter and motor cycle too?

<https://blog.nativesintech.org/apache-appropriation/>

## RELEASE OF I2P 2.1.0:

12/01/2023

The anonymous network I2P 2.1.0 and C++ client i2pd 2.45.0 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 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 peers).

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/2023/01/09/2.1.0-Release>

## RELEASE OF IPTABLES 1.8.9:

12/01/2023

iptables 1.8.9 package filter, who has recently focused on components to maintain backward compatibility - iptables-nft and ebtables-nft, providing utilities with the same command line syntax as in iptables and ebtables, but broadcasting the rules to the nftables byte, was updated. The

original set of iptables programs, including ip6tables, arptables and ebtables, was transferred to the outdated category in 2018 and has already been replaced by nftables in most distributions.

<https://www.mail-archive.com/netfilter-announce@lists.netfilter.org/msg00253.html>

## XUBUNTU MINIMALISTIC BUILDS:

12/01/2023

The developers of the Xubuntu distribution have introduced a new version of the official builds - Xubuntu Minimal, which will complement the regular builds of Xubuntu Desktop and will be available from the release of Ubuntu 23.04. Unlike the previously offered builds, which occupy 2.9 GB, the new builds should fit on a classic CD-ROM (700 MB) with only the basic Xfce without additional applications.

It is expected that the new build will be useful for those who prefer a different set of applications than in the previously proposed basic

delivery. In Xubuntu Minimal, the user will be able to select and download the set of installed applications from the repository during the installation of the distribution. Since 2015, similar minimum Xubuntu Core builds have been informally made by enthusiasts. Unlike the Xubuntu Core, Xubuntu Minimal builds will receive official status and will be built using the main infrastructure of Ubuntu.

<https://lists.ubuntu.com/archives/ubuntu-release/2023-January/005521.html>

## RELEASE OF TOYBOX 0.8.9:

13/01/2023

The release of a set of system utilities, Toybox 0.8.9, as well as BusyBox designed as a single executable and optimized for the minimum consumption of system resources, has been published. The project is being developed by the former BusyBox maininer and is distributed under a 0BSD license. The main purpose of Toybox is to provide manufacturers with the possibility of using a minimalist set of standard utilities without

opening the original code of modified components. Toybox is still lagging behind BusyBox, but has already implemented 306 base commands of the 384 planned.

<https://github.com/landley/toybox/releases/tag/0.8.9>

## RELEASE VENTOY 1.0.88:

13/01/2023

The release of Ventoy 1.0.88 designed to create bootable USB drives, including multiple operating systems, has been published. The program is notable for the fact that it allows you to download the OS from unchanged ISO, WIM, IMG, VHD and EFI images, without requiring unpacking the image or reformatting the carrier. At any time, you can replace or add new iso-image simply by copying new files, which is convenient for testing and previewing various distributions and operating systems. The project code is written in C and distributed under the GPLv3 license.

The new version adds support for Dell PER, Porteus and Lenovo Product Recovery. They improved

the menu, added support for additional languages and implemented the ability to return to the previous menu from the secondary boot menu. In the plugin with the design theme, a parameter has been added to determine the maximum screen resolution. They also added support for files with checksums, like .md5, .sha1, .sh256 and .sha512.

<https://github.com/ventoy/Ventoy/releases>

## CHROME OS 109

AVAILABLE:

14/01/2023

Chrome OS 109, based on the Linux kernel, using the upstart system manager, ebuild/portage assembly tools, open components and Chrome 109 web browser, is available. The Chrome OS user environment is limited to a web browser, and instead of standard programs, web applications are used, however, Chrome OS includes a full-fledged multi-window interface, desktop and taskbar. The original code is distributed under the Apache 2.0 free license. The Chrome OS 109 is available for

most of the current Chromebooks. For use on conventional computers, the Chrome OS Flex is proposed.

[https://chromereleases.googleblog.com/2023/01/stable-channel-update-for-chromeos\\_13.html](https://chromereleases.googleblog.com/2023/01/stable-channel-update-for-chromeos_13.html)

## FULL CIRCLE WEEKLY NEWS



Join our host Moss Bliss as he presents you with a short podcast (<10min) with just the news. No chit-chat. No time wasting. Just the latest FOSS/Linux/Ubuntu news.

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



THE

# VIRTUALBOX NETWORKING PRIMER

Connecting and Configuring  
Virtual Machines



Robin Catling

## The VirtualBox Networking Primer

Connecting and Configuring Virtual Machines

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

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

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

Author: **Robin Catling**

Publisher: Proactivity Press

ISBN13 : 9781916119482

Amazon US link:

[https://www.amazon.com/dp/1916119484?  
ref=pe\\_3052080\\_397514860](https://www.amazon.com/dp/1916119484?ref=pe_3052080_397514860)

Amazon UK link:

[https://www.amazon.co.uk/VirtualBox-Networking-Primer-Connecting-Configuring/dp/1916119484/ref=sr\\_1\\_1?  
dchild=1&keywords=virtualbox+networking+primer&qid=1600253699&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>



# COMMAND & CONQUER

Written by Erik

Welcome back to another edition of FCM, may your year be filled with Ubuntu.

Since we have been spending a lot of time in the terminal, let's discuss some terminal tips. Bring on the disco fonts as we do some tinkering and a little deeper dive (just so you know what you are copying from some random website) for the total n00bs.

Before we start, this is a how and why, rather than a copy-paste job.

OK, just kidding – we won't be going that far, I mean lolcat and rotrflcat is fun, but you do not want to pump everything through there. Then again, Pimp my Ride was as fake as a \$9 note, so we will skip pre-made scripts.

Bash is everywhere, chances of you finding another shell when you connect to another machine are low. Yes, we know things like Fish exist, but bash is functional and usually ships with Ubuntu by default.

Bash is a collection of variables, which simply means things that can change. (Yes, sometimes it does things too, but let's go with the flow).

Here are some environment variables that an average user may encounter:

- PWD – Current working directory.
- HOME – The user's home directory location.
- SHELL – Current shell (bash, zsh, etc.).
- LOGNAME – Name of the user.
- UID – User's unique identifier.
- HOSTNAME – Device's hostname on the network.
- MAIL – User's mail directory.
- EDITOR – The Linux system default text editor.
- TEMP – Directory location for temporary files.

To view all environment variables use the "printenv" command to view all environment variables. Since there are many variables on the list, use the less command to control the view:

```
printenv | less
```

or we can taste the rainbow, but only once:

```
printenv | less | lolcat
```

```
File Edit View Terminal Tabs Help
ed@Dropped:~$ printenv | less | lolcat
SHELL=/bin/bash
SESSION_MANAGER=local/Dropped:@/tmp/
555
WINDOWID=96468995
QT_ACCESSIBILITY=1
COLORTERM=truecolor
XDG_CONFIG_DIRS=/etc/xdg/xdg-xfce:/
XDG_SESSION_PATH=/org/freedesktop/D
XDG_MENU_PREFIX=xfce-
```

If you were paying attention, you would see that they are all in caps, and because they are variables, they need to be prefixed by a \$ - dollar sign, when you use them in commands; for instance, if I wanted to see my LOGNAME, I'd type: echo \$LOGNAME.


Since I always create three users, ed, edd and eddy to test things with, it can be handy for me to see which profile I am in, you may have only one, but the idea stands.

Environmental variables can be

stored in a few places, like ~/.bashrc or .bash-profile or just .profile – so be sure to look at what "dotfiles" you have. Dotfiles are hidden configuration files, much like Windows .cfg or .ini files. If you use a fancy GPU-Accelerated terminal, you can set it in your \$TERM variable.

There are four commands you need to remember here, to manipulate your environment variables, namely: export, env, set, and unset

A lot of the settings can be set from a modern terminal, and as these vary from terminal to terminal, you may need to dig



# COMMAND & CONQUER


some. Here we will just look at a standard terminal. Usually things like, unlimited scroll back, can now be a tick-box away, though you may want to set that in the `bashrc` in your `/etc/` folder, but that is out of the scope of this article.

Fonts, colours, background images are just a setting away if you open your terminal preferences. As you can see here, usually on the first three tabs. A word of warning about fonts: you may see a font you like, that looks good on a web page, that looks terrible on the terminal. Always

test first!

Let's talk about customising your prompt:

The prompt simply says, 'you can type here'. The default prompt on most distro's out of the box is usually `USER@HOST` that you set in the `PS1` variable. You should see a `$` at the end if you are a normal user, or a `#` if you are root. It does not need to say, user at host, it can say whatever you please. I'll demonstrate at the end. (Don't judge, I am not artistic.)



**TIP:** If you use 'nerd fonts' (<https://www.nerdfonts.com/>) or powerline fonts, you can use fancy symbols to make images in your prompt.

On that note, you can peek at your selected font's glyphs: <https://www.nerdfonts.com/cheat-sheet>

that your changes are not happening once you have made them. This is because you need to reload the current session to apply the changes.

This is done by:


```
source ~/.bashrc
```

We are not going to list out all the backslash-escaped special characters you need to pimp your terminal, as it would just be reinventing the wheel. We will, however, point you to a great resource: <https://www.cyberciti.biz/tips/howto-linux-unix-bash-shell-setup-prompt.html> as this is an explanation, not a how-to.

We hope this brings a better understanding and a desire to have your own prompt rather than something you saw on the internet that every second terminal has.

Another thing to remember is that modern terminals do not only support colors, but font weights too, so you may have something in bold or dimmed.

Lastly, as a n00b, you may notice



**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.



Python 3.11.1 is now out (as of 6 December 2022) and with it comes, as usual, a number of changes and improvements. If you want to know what is new with 3.11.1, you can go to <https://docs.python.org/3/whatsnew/3.11.html> to find out all the information.

One of the new things that did come along with 3.11.1 is a new standard library module called `tomllib`, which has "support" for parsing TOML style documents. Notice that I said "support". It's not complete support, but it is support. I'll talk more about that in a few moments.

## WHAT IS TOML?

TOML stands for Tom's Obvious Minimal Language and from what I understand, was created mainly as a means of storing configuration data. Many other ways of storing configuration data don't provide a means of saving comments inline.

To get started, let's assume that I created a Python program that

used PAGE to create the GUI front end. Let's further assume that I want to give the user the ability to select the Theme for the program (since it uses `ttk Widgets`). I want the program to remember what that user decided for their theme of choice. To do this, I will use a configuration file to keep all the customization information.

Shown top right is a simplified version of the hypothetical configuration file.

At this point, it looks like a standard configuration file that you would find just about anywhere. However, if this were a "standard" configparser type config file, the first entry under the `[Themes]` section would not be possible directly, since `configparser` doesn't support lists without manipulation. In TOML, sections are called tables. The `available_themes` key has the value of an array. Once it's ported into Python, it becomes a list.

Now let's look at how to get the data into a program.

```
[Themes]
available_themes = [
    "notsodark",
    "plastik",
    "waldorf",
    "page_wheat",
    "clearlooks",
    "forest-light",
    "forest-dark",
    "default",
    "clam",
    "classic",
    "alt",
]
default_theme = "waldorf"
current_theme = "notsodark"

[Program_Info]
version = "0.7.1"
```

Of course, we have to import the `tomllib` library. Remember this is supported directly only under Python 3.11.

```
import tomllib
import pprint
```

Next, we open the configuration file and use the `load` method of the library.

```
with open("config.toml",
"rb") as f:

    data = tomllib.load(f)
```

Using `pretty print`, we can now

look at the data that was brought in from the config file (next page, top right).

You can see that it is simply just a dictionary. To access the data, we do it just like any other dictionary (next page, bottom left).

The output of our little program will look like this...

```
Available Themes:
['notsodark', 'plastik',
'waldorf', 'page_wheat',
'clearlooks', 'forest-light',
'forest-dark', 'default',
'clam', 'classic', 'alt']
```

# HOWTO - PYTHON

```
Your default Theme is:
waldorf
Your current Theme is:
notsodark
Program version 0.7.1
```

The Python tomlib library provides only two functions, `tomlib.loads` which loads a TOML string and returns a dictionary, and `tomlib.load` which reads a TOML file and returns again, a dictionary. See <https://docs.python.org/3/library/tomllib.html>.

Unfortunately, Python does not provide any way to properly write out the TOML data. The good news is that there is a third party library called `tomli_w` which will allow you to write the TOML data back to a file. So if the user decides to change his current theme from 'notsodark' to 'clam', it's a simple chore to make the change and write it back. You can install it via pip.

```
pip install tomli-w
```

Once it's installed, you can

simply write it pretty much as it's a normal, but binary file.

Remember the `tomllib` (bottom right) comes with only Python 3.11. So what do you do if you are running Python 3.8.10 and you aren't ready to upgrade quite yet? Never fear. The `tomli-w` library is currently supported in versions 3.7 all the way up to 3.11. To get the parser analog to `tomllib`, you can install `tomli`. Again, just use pip.

```
pip install tomli
```

Of course, if you are using `tomli` rather than `tomllib`, you need to do the import a bit differently.

```
#import tomllib
import tomli
import tomli_w
import pprint
```

The home pages for `tomli-w` and `tomli` can be found at <https://github.com/hukkin/tomli-w>

```
print(f"Available Themes: {data['Themes']['available_themes']}")
default_theme = data['Themes']['default_theme']
current_theme = data['Themes']['current_theme']
print(f'Your default Theme is: {default_theme}')
print(f'Your current Theme is: {current_theme}')
program_version = data['Program_Info']['version']
print(f'Program version {program_version}')
```

```
{'Program_Info': {'version': '0.7.1'},
 'Themes': {'available_themes': ['notsodark',
                                 'plastik',
                                 'waldorf',
                                 'page_wheat',
                                 'clearlooks',
                                 'forest-light',
                                 'forest-dark',
                                 'default',
                                 'clam',
                                 'classic',
                                 'alt'],
            'current_theme': 'notsodark',
            'default_theme': 'waldorf'}}
```

```
import tomli_w
```

```
Current_theme = "clam"
data['Themes']['current_theme'] = current_theme
with open("config.toml", "wb") as f:
    tomli_w.dump(data, f)
```

So an edited snippet of the config file after updating would be...

```
[Themes]
available_themes = [
    "notsodark",
    "plastik",
    ...
    "alt",
]
default_theme = "waldorf"
current_theme = "clam"

[Program_Info]
version = "0.7.1"
```

<https://github.com/hukkin/tomli>

If you want to take a look at the complete information on TOML, you can visit the home page at <https://toml.io/en/>

There is another TOML third-party package for Python at <https://github.com/sdispater/tomlkit> . It's documentation can be found at <https://github.com/sdispater/tomlkit/blob/master/docs/quickstart.rst>. I haven't had much of a chance to play with it, but it looks promising.

I've put the config.toml file and the Python file (toml1.py) in my repository at <https://github.com/gregwa1953/FCM-189>.

That's it for this month. **Happy New Year!!!**

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



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



Like most users, I do change my main working computer from time to time. When doing so, it may be a good time to start out afresh and give some thought to our hard disk setup. Even if cloud storage is widely used, limits on amounts of storage available (or its monthly fee), and network bandwidth to get our files up into the cloud, mean most users actually keep a sizable portion of their files on their local drives. Security considerations may also be a factor to consider, even with smaller files. In this article, I would like to set out several ideas that a more or less technical user could consider, strategies that may be of help both for data deduplication on your local drive, or for extending a drive with a limited capacity we are growing out of.

## WHY MIRROR DATA?

Making multiple copies of all files that are written to our computer's hard drives is a technique used in all server environments; every time a file is modified, alterations are written to more than one disk using one of

several different schemes. The thinking being that, should one of these drives stop working suddenly: the second copy would not only ensure no data is lost, but, moreover, that service could continue immediately with no interruptions. Even when making a more pedestrian use of our personal computers, at some point we need to recognise that hard drives can, and do, fail.

Perhaps things are not as critical with modern SSD technology as they were at one point with rotational drives and their moving disks. The Mean Time Between Failures (MTBF) is a commonly-used metric that seems to have gone up with the switch to solid-state technology. But still, common sense tells us it is perhaps best not to entrust all our data to a single copy on our internal drive without making any backups. Fine... but there still are obstacles. In practice, not all backups are made on a schedule and kept up-to-date. At times, large files may never even get backed up at all -- neither to the cloud because they are too large,

nor to a local external drive because we would need to think of plugging it in and performing the backup. Since, in the immortal words of Murphy's Law, anything that can go wrong will do so and at the worst possible time, it can make sense to have multiple copies of all our data, made automatically, on our daily driver computer.

Clearly, making several copies of our files on a single disk would not be very useful at all. In the event of that disk having a catastrophic physical failure, we would lose access to all data it contained and the fact that it held several copies of each file would be moot. So, our goal must include the use of more than one physical hard drive on our computer. And, since messing around with the hard drives on our daily driver is probably not a great idea, what better time to take a look at the possibilities than when acquiring a new (or new-to-us) machine?

## WHY JUST MIRROR DATA?

The term "mirrored drive"

implies the use of exactly two drives of a certain capacity. All data is written to disk in two copies, one on each disk. So, in essence, we are buying (and probably paying for) twice the disk space, and then using only the capacity of a single unit giving us a disk use efficiency of a mere 50%.

In server rooms, this would be considered wasteful, since techniques such as RAID-5, RAID-6 or others allow higher efficiencies -- given a sufficient number of disks. For instance, with an 8-disk setup and using single parity, a RAID-5 would give us a useful space of 7 times the capacity of a single disk, while buying eight times that worth of space. 7/8 or 87,5% disk usage efficiency, while retaining the ability to survive the death of any single disk without losing any data is not a bad result. However, the immediate objection is that while it may be practical to house a large number of disks in a server rack, things are a lot more complicated in a typical laptop, or even some desktop computers with slim form factors. Even the smaller grade of

## HOWTO - SET UP MULTIPLE HARD DRIVES

commercially-available Network Attached Storage (NAS) units tend to have receptacles or bays for only two hard drives, thus limiting us to mirroring or – in technical parlance – RAID-1.

On most desktop machines, the motherboard will have at least two SATA interfaces, where two rotational drives or two SSD drives with SATA interfaces may be plugged in. Even on many laptops, a second SATA hard drive may often be installed. If the laptop comes with an optical drive, a possibility would be to replace this little-used piece of equipment with a hard drive carrier that just plugs in the same space the DVD used to take. This is actually quite useful, mainly with laptops with the larger form factors.

Modern technology can also give us the possibility of using M2 interface SSD drives. These take up quite a bit less space than rotational drives with the SATA interface. Many laptop motherboards either come equipped with a disk drive with this interface, or an available M2 slot on the motherboard. However, most laptops will only have a single M2 interface, so some creativity may

be necessary to combine, for instance, one SSD on the M2 interface with a second one on a SATA interface. This setup would not be advisable from the standpoint of speed, but should be doable if data redundancy is the primary objective and maximizing throughput is not (an SSD drive should be plenty fast, even on a SATA interface).

On the other hand, desktop machines can use a PCI-Express adapter (extension card) to add an M2 interface to an existing machine. Some adapters are for a single M2 drive, while others even have space for two M2 interfaces thus allowing the user to build a two-drive system on a single adapter. What is more, the cost can be quite moderate, with some dual M2 interface adapters available for as little as 17 USD. However, maybe ensure your computer can boot off this extension card before acquiring one.

As a final thought, it may be useful to remember that using two partitions situated on the same physical disk would not be a good idea: not only will access times be lengthened since all data needs to be written twice to the very same

disk thus creating a bottleneck, but if that unit goes out physically, both partitions will be lost at once thus rendering mirroring completely useless as implemented. To sum up: mirroring requires the use of (at least) two physical disks or drives.

### WHY EXTEND AN EXISTING DISK (STRIPING) INSTEAD OF MIRRORING?

For some people, mirroring data may not even be necessary for the user's workflow. This would be the case when the user actually generates most data in the cloud – or on a company server, and the local hard drive is seen more as an archive to have a local copy of files in case the remote server is no longer available or some space needs to be freed up. Another scenario would be the “perfect” user who has one copy of each file locally, one copy on a remote server, and a third on cold storage, for instance on an archived external disk or optical media. Finally, we may at some times need a large amount of disk space to store files for a relatively short period of time, after which this large amount of data will be discarded. This is

usually the scenario found while editing video, for instance, where the final product would certainly be kept, but perhaps not most of the original material if it is no longer relevant.

In these cases, would it not make more sense to use several drives to extend the total amount of space available, while retaining just a single copy of each file. In essence, we are trading some data security for available space. In the same way, the concept of mirroring deals with combining several hard drives into a single unit in which files are stored in two copies, striping is the analogous concept of combining several drives into one unit with a larger capacity but without duplication. In technical terms, where mirroring is often spoken of as RAID-1, striping is often known as RAID-0. In very large setups, these two techniques may be combined to create RAID-10 disk arrays, but this would probably be pertinent only on larger desktop computers with many hard drive bays.

With some of the technologies used, striping may allow an existing disk drive to be extended – in some cases, even while it is mounted and

# HOWTO - SET UP MULTIPLE HARD DRIVES

working -- using a second drive.

## WHAT ABOUT THE FILE SYSTEM?

To my knowledge, there are at least four different ways of setting up either mirroring or striping on a personal computer. The first includes the use of dedicated physical hardware (thus, logically called "hardware RAID") to which the two disks are plugged in, and which manages the complete set and presents it as a single unit to the user. However, issues include difficulties in using such a system on a laptop, the fact that not all desktop motherboards include this feature, and the hassle of using proprietary software (that often is not available under Linux) to recuperate data in the event of a problem. Needless to say, I would not advise using this way of going about it in a domestic or small business setting. It should also be noted that such setups may or may not allow "hot-plugging" new drives and associating the new disks to existing arrays. Some reading may be necessary if looking into using such a system.

The second way would be to set

up a RAID array using the Linux kernel's built-in md (multiple-disk) subsystem. Utilities such as mdadm give us access to the system, which is quite mature and works well. But the interface and commands are not always user-friendly, which is a bit of a pity and also why I will not get further into this system. Interested users' are referred to Jan Mussche's "Install Mint On A Two Disk Raid 0 Set" in FullCircle issue number 104 -- switch from using RAID 0 (striping) to RAID 1 (mirroring) and the basics are about the same. In this case, "hot-plugging" new drives will be possible if your computer's hardware allows it, which is in general not the case for consumer-grade computers and their drives. Adding new drives to existing arrays is possible in software, but will require some familiarity with the commands.

## MIRRORING OR STRIPING WITH BTRFS


This leaves us with two systems that are not based on managing individual disk drives, but whole file systems. In FCM#94, I wrote about the B-tree File System (BTRFS) and some of its applications. This file

system has evolved even further since then, to the point that I consider it mature and use it on all my own computers (yes, I do eat my own dog food!). More to the point, Ubuntu's installer, and that of most other modern Linux distributions, supports installing our system on a btrfs volume out of the box. GRUB can also use btrfs volumes, so very little fiddling is needed to set up mirroring. Basically, if our computer has two hard drives, we can use both disks to set up a mirrored or striped system. The usual caveats exist, including the fact that all existing data may be erased -- so, begin by making sure you have a full backup of all your data, and also

that the computer can be safely formatted with no ill-effects (i.e. please do not try this out on a computer you need to work on within the next hour or two...).

In the following examples, I will be using two rotational hard drives connected via USB to demonstrate the principles. In all cases, a single partition of 100 GByte size will be created on each drive.

Boot from the USB live image as usual, install the operating system using ubiquity on the first hard disk. When you get to the hard drive selection, create a partition. In my example, I will be using Kubuntu



# HOWTO - SET UP MULTIPLE HARD DRIVES

20.04, and the disk to install to is /dev/sdc. Please do verify you are using the correct disk -- perhaps twice -- better than just once!

Boot from our new system, as usual. Now, using our tools of choice for partitioning (fdisk, gfdisk, KDE Partition Manager...), create a single partition on your second hard drive. It is not necessary to format it as btrfs, since whatever format we use will be overwritten.

We then add the second drive's partition to the existing partition (containing our operating system). This can be done with the computer actually running the system. Supposing our main partition on the first drive with the operating system is /dev/sda1, and the new partition on the second drive is /dev/

sdc1, we would issue commands such as these:

```
$ sudo bash
# btrfs dev add /dev/sdc1 /-f
# btrfs fil show
Label: none  uuid: aeb12e81-f5b1-4a48-b80d-d64624867456
Total devices 2 FS bytes used 8.13GiB
devid    1 size 93.13GiB used 11.02GiB path /dev/sda1
devid    2 size 97.66GiB used 0.00B path /dev/sdc1
```

At this point, the two separate partitions have been combined into a single BTRFS file system. With these two commands, we have a striped unit with a total capacity of about 200 GBytes. If striping (RAID-0) is desired, stop here.

Mirroring (or RAID-1) has not yet

been implemented, and we will proceed to do so by converting our striped unit into a mirrored one. We will be using the fact that the file system is currently mounted on / (i.e. is the root file system). Be aware this process may take some time, especially if partitions' sizes are in the TeraByte range:

```
# btrfs balance start
-dconvert=raid1
-mconvert=raid1 /
```


This converts both our data (-dconvert) and filesystem metadata (-mconvert) into RAID-1, effectively writing a second copy of each item on the other disk. Our two partitions should not report a similar amount of space being used. That's it; we have a mirrored unit and our data will from this point on be effectively duplicated with one copy on each partition.

Please note that both partitions are not quite the exact same size. With BTRFS, this is not an issue for striping. With mirroring, however, some space will be lost with unequal partitions. As a general rule of thumb, the final capacity of our mirrored drive will be about equal to that of the smallest of our two partitions.

Before continuing with our final option, let's consider recovery from a failed disk. With BTRFS, if we are booting from a mirrored filesystem, the boot process will not continue at the point at which the Linux kernel mounts the drive. To get past that, we need to boot from another medium -- an Ubuntu live image on a USB thumb drive is fine -- and fix the bad BTRFS volume. We have two choices, both of which include removing the offending bad drive. On the one hand, if we have a spare physical drive to plug in, we can physically replace the bad unit and then run the "btrfs replace" command to restore the array to operation. On the other, if we have no physical spares available at the time, we can re-balance the BTRFS volume to single-disk mode, and remove the bad drive using the "btrfs balance" and "btrfs remove" commands. At a later date -- better sooner than later -- a second, working disk may be partitioned and added to the system as described above.

## MIRRORING AND STRIPING WITH ZFS

Although ZFS (at this time, the



# HOWTO - SET UP MULTIPLE HARD DRIVES

acronym is given no precise meaning) is an old player in the server market with Sun Microsystems behind it, it took some time for this technology to come to Linux. OpenZFS is gaining traction in the Linux kernel and Ubuntu distribution worlds, though perhaps not as swiftly as some server administrators would like.

ZFS file systems share many characteristics with BTRFS, though some details are different. For example, with ZFS, a virtual device (vdev) may be a single write unit, or may implement mirroring or other forms of data duplication -- but it cannot change from one scheme to another as can BTRFS. Thus, when a vdev is created, one must specify from the onset if we want mirroring to happen, or not. A second feature is that ZFS units are usually mounted automatically by the ZFS subsystem, without needing any entries in the system configuration. In brief, file /etc/fstab may contain zero entries on a system that uses only ZFS units. Finally, in current versions of Ubuntu, ZFS drivers are not part of the basic install, but need to be added through various apt packages.

For interested readers, there is a nice writeup by Jim Salter on ArsTechnica. Though server-focused, it should give you a good initial presentation to fully understand the specific strengths (and quirks) of ZFS over other file systems: <https://arstechnica.com/information-technology/2020/05/zfs-101-understanding-zfs-storage-and-performance/>.

On Debian-based distributions such as Ubuntu, things have come on to the point that there has been talk of offering to install recent versions on a root ZFS file system, at least as a beta feature. As a matter of personal preference, for the time being I am steering away from this option, at least until ZFS tools mature slightly (they are now shown to be at version 0.8.4) and are included out-of-the-box on all Ubuntu and derivatives' live boot media. The reasoning here is that, while ZFS does have some very nifty features, the path to recovering data from a crashed boot volume -- and please note we are speaking only of a boot volume -- is a tad complex for regular users. For this reason, in the following example I will be working with:

- A root file system on /dev/sda1 (first partition of our main

hard drive), with an existing Kubuntu 20.04 installation. This is actually the very same BTRFS unit used previously. Which actual filesystem is used is not important, any familiar options such as ext4 or btrfs should work well.

```
# df -h
Filesystem      Size  Used
Avail Use% Mounted on
/dev/sda1        46G   7,7G
36G  18% /
```

- Two partitions, /dev/sda2 and /dev/sdc1. These will be combined in a ZFS pool to hold user data.

The purpose of the root file system is just to hold our operating system, and to do so separately from user data. This means it can be formatted easily without affecting user data, or even replaced by a partition on another disk (or even a live media) at a pinch. The ZFS pool will be mounted on /home, to hold user


data.


Before beginning, let us install the required packages. BTRFS is integrated into current Linux kernels, but ZFS is not. So, we will need:

```
# apt update ; apt install
zfs-dkms zfsutils
```

As previously, begin by creating the appropriate partitions. Since we will not be using the root partition for ZFS, we will leave it alone and create /dev/sda2 and /dev/sdc1, both with about the same capacity. In our case, I will be using 400 GBytes each.

Now, let us create a zpool that uses striping to combine the two partitions. It is traditional to call this zpool the tank, but feel free to choose your own name.





```
# zpool create tank /dev/
sda2 /dev/sdc1 -f
# df -h
Filesystem      Size  Used
Avail Use% Mounted on
/dev/sda1       46G   7,7G
36G 18% /
tank            756G 128K
756G 1% /tank
```

Each partition is of about 400 GBytes, but in combination they yield 756 GBytes. Note that the new zpool has automatically been mounted as /tank.

Within the zpool, we will now create a filesystem (vdev) called home, which may expand to use up all the space available.

```
# zfs create tank/home
# df -h
Filesystem      Size  Used
Avail Use% Mounted on
/dev/sda1       46G   7,7G
36G 18% /
tank            756G 128K
756G 1% /tank
tank/home       756G 128K
756G 1% /tank/home
```

The new filesystem has also been mounted automatically, on /tank/home.

We now need to transfer over the contents of our /home directory, since this is where we will be mounting our new filesystem.

```
# mv /home/* /tank/home/
```

We actually need to move the files (and not merely make a copy) since the /home directory will need to be empty. Otherwise, the ZFS filesystem will refuse to mount on a non-empty directory.

Finally, we program the new filesystem to mount automatically on /home.

```
# zfs set mountpoint=/home
tank/home
```

Now, reboot. The system should come up with tank/home mounted on /home:

```
$ mount
/dev/sda1 on / type ext4
(rw,relatime,errors=remount-
ro)
tank/home on /home type zfs
(rw,xattr,noacl)
tank on /tank type zfs
(rw,xattr,noacl)
```

```
$ df -h
Filesystem      Size  Used
Avail Use% Mounted on
/dev/sda1       46G   5,6G
38G 13% /
tank/home       756G  36M
756G 1% /home
tank            756G  36M
756G 1% /tank
```

As can be seen, the new ZFS vdev is now mounted in its place on /home, with the two striped volumes giving us the expected combined space.

With ZFS, to create a mirrored zpool instead of striped, the appropriate syntax would simply have been:

```
# zpool create tank mirror /
dev/sda2 /dev/sdc1 -f
```

Please note the “mirror” keyword. Other commands would be the same in either case, especially the creation of the ZFS filesystem (vdev) within the zpool.

## SOME FINAL THOUGHTS

A decade ago, mirroring and striping were options that would have seemed exotic, or quite frankly difficult to use, to many users. Even today, systems administrators can have cause to rely on dedicated hardware to set up such arrays. However, the spread of BTRFS and ZFS gives regular users – albeit those with a rather technical bent – the option of creating mirrored or striped units using only software that is available with the Linux kernel.

Both technical solutions have their proponents, with some very good reasons on either side – and a tendency to be quite vocal about it at times. Since our purpose is not to provoke an all-out flame war, let us merely suggest that BTRFS has perhaps a tad more flexibility as regards converting striped units into mirrored and vice-versa, can more easily extend an existing unit – even on the fly in a working system – and is also slightly better integrated into standard kernels. Setting up a mirrored root partition with BTRFS is also marginally easier than with ZFS. On the other hand, the strong points of ZFS include

## HOWTO - SET UP MULTIPLE HARD DRIVES

more flexibility setting up and mounting both zpools and vdevs in various configurations, and is more similar to the setup used in many large servers.

Additionally, both ZFS and BTRFS feature snapshots, which facilitates rolling back system changes in the event of a disaster, which is perhaps better integrated into the GRUB bootloader under OpenSUSE. They may also be used as a backup mechanism, since snapshots can be transferred between systems using the send-receive process.

All this opens interesting perspectives for the interested student of systems management -- or for the advanced user. However, we would suggest to begin by testing out either or both systems described here on a spare computer, not using them as your main system (i.e. on data you depend upon) until you feel thoroughly at home with their syntax and quirks.



**Alan** holds a PhD in Information and the Knowledge Society. He teaches engineering at Escola Andorrana de Batxillerat (high-school). He has previously given GNU/Linux courses at the University of Andorra and taught GNU/Linux systems administration at the Open University of Catalonia (UOC).



# HOW-TO

Written by Robert Boardman

Last issue, we made a one page poster advertising a lost dog. Of course the subject matter for a poster could be anything you wish. The poster got us started into an exploration of font use in Latex / Tex. Before we start, there are at least two guides you might wish to look at if this topic interests you. The first is more recent: LATEX 2ε Font Selection, compiled by the Latex Project Team (2021 March). The second is LATEX Font Encodings by Frank Mittelbach, Robin Fairbairns, and Werner Lemberg (2016 February). Because they deal with the same topic, there is some duplication of contents of course.

While pointing you to references, I should also point out The Tex Book by Donald Knuth is also available for download. It is a tex file, a large tex file, so you can open it in any text editor and examine the coding required to make this particular book. You can also compile it into a PDF (with TexStudio or whichever environment you use). If you download it, you are not allowed to

distribute it or sell it. Printed copies are available through the usual sources.

As mentioned last time, there is a series of keywords which indicate the ten default sizes for fonts. Depending on the size of the base font selected (10pt, 11pt, 12pt), the keyword sizes increase from about 5pt to almost 25pt. We used the `\selectfont` command to change the existing font to a new one. (Note: Although not technically correct, I will use font and typeface to mean the same thing in this article.) If you make a change in the code, but do not use `\selectfont` where you want the change to occur, the change will not be visible.

There are six commands which can be used to change the appearance of the text:

```
fontencoding { encoding },
fontfamily { family }, fontseries { series }, \fontshape { shape },
fontsize { size } { baselineskip },
linespread { factor }. These individual commands can be shortened to usefont { encoding } { family } { series } { shape }. In
```


this article, I will focus on family, series, and shape. You have probably made changes to these features of a typeface when using a word processor.

There are far too many font families to list. If you have done any HTML-CSS coding, you will be familiar with some font families. Word processors often use the term "font name" as a substitute for the font family. Then word processors use buttons or menus to set the series and shape. In Latex / Tex, a chosen series refers to the one of ten weights, and one of ten widths, of the font (Bold, Condensed, Light, etc.) Shape refers to normal, italic, small caps, and five others. Font size must be specified separately with the `fontsize` command if it is to be different from the default. Different measuring units can be used. If no unit is specified, then point (pt) is assumed.

```
\paragraph*{} The quick fox
jumped over the lazy yellow
dog.
\paragraph*{}\fontfamily{phv}
\selectfont The quick fox
jumped over the lazy yellow
```

```
dog.
\paragraph*{}\fontfamily{phv}
\selectfont The quick fox
jumped over the lazy yellow
dog.
```

This code produces:



The `\paragraph*{}` code gives more space between lines of text than two presses of `<enter>`. The `*` means the paragraph is not numbered, and the `{}` means the paragraph does not have a name (title).

As usual with Latex, there are other ways to vary fonts. The next method is suggested for small sections of text although it could be used for any quantity of text. There are at least nine additions to the `\text` command. All enclose the

# HOWTO - LATEX

text to be modified in curly braces. The general form of the command is `\text_{modified text here}`. When used, the underscore is replaced by two letters which indicate the desired effect. The two-letter codes try to indicate what the effect is. [Note: Not all fonts will have the appropriate letter set for the change.]

On my machine with the default font, UC did nothing and lf generated an error.

It is possible to nest the codes to get various combinations. For example, some people like to use bold and italic together. In Latex that means:


- writing the first code ending with curly braces
- writing the second code inside the first set of curly braces
- writing the text inside the inside

curly braces

```
\textbf{\textit{The quick fox
jumped over the lazy yellow dog}}
```

Here is the code and result:

```
\paragraph*{\textit{The
quick fox jumped over the
lazy yellow dog.}}
\paragraph*{\texttt{The
quick fox jumped over the
lazy yellow dog.}}
\paragraph*{\textbf{The
quick fox jumped over the
lazy yellow dog.}}
\paragraph*{
\textbf{\textit{The quick fox
jumped over the lazy yellow
dog}}}
```



Code	Effect	Code	Effect
rm	Roman – return to standard letter forms	sf	Use a sans serif typeface if available
tt	Monospaced	it	<i>Italic</i>
sl	Slanted	uc	UPPERCASE
lf	lower case	bf	<b>Bold</b>
md			

It is possible to use three of the same two-letter abbreviations with another command to achieve similar results: `\rmfamily`, `\sffamily`, `\ttfamily`. A space is required between the command and the start of the text. Do not use curly braces around the text. I assume you can handle this code on your own.

If you are eager to explore using different fonts in Latex, I suggest you start at [tug.org/FontCatalogue](http://tug.org/FontCatalogue). There are hundreds available – most Adobe Type 1 fonts, Open Type fonts and TrueType fonts can be used. The FontCatalogue page for each font will tell you if a particular font is installed by default or needs to be downloaded and installed. The page will also tell you if you need to use a specific package in the preamble of your document. Some fonts appear in the preamble as packages, others need one or more commands in order to be available to a Latex compiler. Here is an example from TexStudio.

And here is what is needed for the preamble in order to change the default typeface in this

document to Tex Gyre Adventor. (This is a sans serif typeface derived from the URW Gothic family.)

```
\usepackage{tgadventor}
\renewcommand*\familydefault{
\sfddefault}
\usepackage[T1]{fontenc}
```

Fonts are usually installed in `/usr/share/fonts/`, then under type (opentype, truetype, type1). In each type is a directory for each font name, for example my truetype directory has several subdirectories including: arphic, cm-unicode, dejavu, liberation, and others. If you want to know what fonts are available to you, without downloading anything, this is where to look. Have fun with Latex / Tex until next time.



# 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 129

Aside from the snap controls that I described last month, at first glance it may not look like much has changed in the main Inkscape window with the release of version 1.2 – apart from the welcome return of these zoom buttons to the main Control bar, of course:



Observant readers may notice that not only have the 'Zoom to fit selection', 'Zoom to fit drawing', and 'Zoom to fit page' buttons reappeared, but they're accompanied by a new sibling. Clicking this fourth button will center the page in the window, but without changing the current zoom. This can be useful when you've lost your way in a drawing, either because you've zoomed in or out too far, or have gone a bit wild with the pan option.

A less obvious addition to the main window is a change to the toolbox. Last month, I looked at the


new preferences that allow you to turn individual tool icons on and off, or to change their size. But those are not the only changes to the way these buttons are displayed. It's now possible to adjust the width of the toolbox by dragging the right-hand separator. There's no obvious visual affordance for this (i.e. a 'handle'), but if you move the mouse slowly over the dividing line between the toolbox and the ruler (or canvas, if you have rulers hidden), your mouse pointer will change to indicate where you can click-and-drag. This allows you to transition seamlessly from the classic vertical list of icons, to a squarer layout in which each group of buttons occupies its own row – or, indeed, something in-between if you prefer.

Personally I think that the one-row-per-group view is rather wasteful, as it's not possible to dock other dialogs in the empty space below the icons. But I can certainly see the benefit of changing to two columns if you're working on a screen with limited height, but still want access to all of

the icons rather than turning some off.

One thing to note is that it's possible to drag the divider the other way, to collapse your toolbox down until it's no longer visible. That's not the same as using the View > Show/Hide > Toolbox menu entry – that menu will still indicate that the toolbox is visible, even though you can't see it. To avoid confusion, I suggest not collapsing the toolbox by dragging the separator: if you want to hide it, use the menu entry instead. If that's something you do regularly, then you can assign a keyboard shortcut to the menu option via Edit > Preferences > Interface > Keyboard. It's in the 'Canvas Display' section of the keyboard shortcuts.

Once again, observant readers may notice the addition of a new icon to the end of the toolbox. This is for managing multiple pages in an Inkscape document, which is such a major new feature that I'll be covering it separately in future.




# HOWTO - INKSCAPE

There's been one other significant change to the main window but, like the resizable toolbox, it's not immediately obvious that anything is different. The color palette, at the bottom of the window, has seen a major overhaul. Most of the changes take place towards the right of the palette, where you'll find a pair of up/down buttons, and a menu button.




We'll skip the up/down buttons for now, and go straight to the menu button. Users familiar with earlier releases may know that there's been a pop-up menu hanging around in this corner for a long time, allowing you to switch between different palettes and adjust a few settings related to the display of the swatches. The old design was a text-only menu, with some submenus providing a limited selection of options (e.g. None/Solid/Wide for the border around each color swatch). The new menu is a richer UI widget, allowing for a small thin-line preview of each

palette – which makes it much easier to pick the right one. This image shows the old and new menus, for comparison.



The various options for configuring the display of the palette have now been moved into a single pop-up which is opened via the 'Configure...' option at the bottom of the menu.




The options in here are pretty self-explanatory. 'Tile size' sets the basic size for the swatches, though the actual height and width will vary based on the aspect ratio. I'll gloss over the fact that the 'Aspect' slider runs from -1.0 to 1.0 (that's not really how aspect ratios work), and just state that negative values make the swatches taller than they are wide, positive values make them short and fat, and zero makes them square. The nice thing about all the controls in this pop-up is that you can see their effect on the swatches dynamically as you change them, so it's not worth getting too caught up in the specific values – just drag the sliders until the palette looks the way you want it to.


The 'Stretch to fill' option is a little odd. Ticking it disables the Aspect slider entirely, which in 'normal' UI terms should really mean that the checkbox is put above the slider to indicate the

parent-child relationship between them. That's not what makes it odd though: the weird thing is that it really does what it suggests it will only for palettes with few entries. Let's start by looking at a case where it does work: the very limited color set of the 'Android icon palette'. Here's how that palette appears with the 'Stretch to fill' option disabled (top) and enabled (bottom).

It doesn't take a genius to see that the 'Stretch to fill' option has stretched the individual swatches to fill the available space. No big surprises there – it's exactly what you would expect an option with that name to do. But what happens if there are a larger number of colors in the palette – something like the 'Blues' palette, for example. Here's how that palette appears with the 'Stretch to fill' option disabled (top) and enabled (bottom).

If you're struggling to spot the





difference, you're not alone. In both cases, the swatches actually extend over three rows (on my screen), with those up/down buttons being used to switch between them (you can also use the scroll wheel on your mouse, if you prefer). But clearly the swatches haven't been made narrower in order to fit them all in the available space, as you might expect or want. 'Stretch to fill' very specifically means 'Stretch' and not 'Compress'.

Perhaps you're thinking that Inkscape is just being sensible, as there are too many colors in that palette to fit on one row in a usable way. Not so: unchecking that option and manually setting the tile size to 11 and the aspect to -0.7 is enough to get the entire palette to fit in a single line.

Moving on from the inconsistencies of stretching




swatches, the 'Border' option adjusts the amount of space around each swatch. If you want your palette to look like a continuous gradient of color (assuming the tones are so arranged), then set this to zero. Higher values add more space around each swatch, which may be useful to avoid mis-clicks, or simply to ensure each swatch appears as an individual item rather than blending with its neighbours.

The final option, 'Rows', exposes what I think is the most significant flaw in the new palette interface. When the number of swatches is simply too large for them all to appear in the available space, they overflow onto multiple rows. You can then scroll through them using the up/down buttons or the mouse wheel, as I described above. The 'Rows' control determines how many rows can be displayed in the palette area – basically, how tall the palette area is. This can allow you to

display an entire palette at once, rather than having to use the up/down buttons to access them all. This image shows the difference between Rows=1 and Rows=5 when using the Inkscape Default palette.

If you're the sort of person who wants to see all the colors in the palette, all the time, this might be fine. But what if you're also the sort of person who switches between different palettes, depending on your needs? Here's how the palette area looks when you've got Rows=5 but you've selected the Android Icon Palette.



That's a lot of wasted space. Enabling the 'Stretch to fill' option doesn't help much either, as that only stretches the swatches horizontally, and does nothing to fill the additional blank rows. The 'Rows' parameter sets a fixed number of rows to display, even if fewer rows would make more sense. It would be nicer, I think, to have the number of displayed rows adjust dynamically, using the 'Rows' parameter to set a maximum.

But in my opinion there's a

bigger issue to consider. In previous versions, the palette would scroll horizontally if it couldn't all fit on screen. You could switch to a vertical scrolling mode by enabling the 'Wrap' checkbox in the palette menu, which would give the same effect as the palette in version 1.2. But the new release doesn't offer an equivalent way to revert to the horizontal scrolling design of its predecessors. You're forced to use a vertically scrolling palette, whether you like it or not.

This may seem like a trivial gripe, but there are pragmatic reasons why a horizontally scrolling palette is arguably better in some cases. Let's consider that sweeping blue palette once more – here it is with Rows=3 on my machine.

Try to imagine that as one long



horizontal palette. Yes, you can only see a small 'window' of colors at a time, but scrolling the mouse wheel over the colors smoothly moves the line along, with no breaks or discontinuities. It's easy to see the relative colors of every swatch when compared with its


neighbours.

With a multi-row display, however, there are breaks artificially added to the flow of color. Your eyes have to scan from the end of one row to the start of the next in order to continue progressing along the palette. Swatches can now be completely surrounded by up to 8 immediate neighbours, not just one on each side. This affects your eye's ability to discriminate between the colors, especially where (as in this screenshot), some light colors are sandwiched between darker rows. And the position of those breaks aren't consistent, but change with the window size and the parameters chosen for the palette display. Your favourite shade of blue might be at the start of the second line one day, but buried somewhere in the middle of a block of color the next.

For now, us 'single row' advocates just have to live with this forced vertically scrolling design, but I do hope that a horizontally scrollable palette makes a return in a future release.


Moving on from the palette to a related topic: the Fill & Stroke

dialog has seen some useful tweaks and improvements. You may recall that previous versions offered tabs, and then buttons, to switch between different color pickers. The choices were RGB, HSL, HSV, CMYK, Wheel, and CMS. Those tabs/buttons have been replaced once more, with a pop-up menu that adds 'HSLuv' to the mix.



You may have noticed that there's no 'Wheel' option on the menu. Fear not! The wheel has been moved to a collapsible section in the HSL, HSV and HSLuv modes, giving you the option to use both the wheel and the sliders in combination far more easily. I don't understand why it didn't also make it into the RGB view, as I can't see any technical reason for preventing it (Inkscape uses RGB internally, so, even if you pick your colors using CMYK, what ends up in the file is actually an RGB conversion of your

color). The HSLuv view is an interesting one: I can't really make sense of it, even with the visualisation offered by the wheel view, but I daresay there are users with more specific color requirements (and knowledge) who will benefit from the addition of this mode.



The color pickers aren't the only change in the Fill & Stroke dialog... but I've reached my word count for this article, so the other new additions will be the subject of next month's instalment.



**Mark** uses Inkscape to create comics for the web ([www.peppertop.com/](http://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)



# HOW-TO

Written by Ronnie Tucker



# FreeCAD - Part 1

Well, I bought a 3D printer. Fear not, I'll review it shortly. But I realised that although I can download a myriad of things from websites, I really should learn how to make my own things.

Enter the wide world of CAD (Computer Aided Design). There are, quite literally, dozens of CAD programs out there, but I settled with FreeCAD. It runs quite happily in Linux and, as its name suggests, it's free.

**Disclaimer:** this brief series is just to show you the basics of FreeCAD. I'm far from being an expert in it.

## INSTALLATION


I'm using Ubuntu 22.04 (as of writing) and have removed the snap nonsense. I hate snaps, but that's another article I could write. Anyway, in the repositories, it shows version 0.19, but the latest version is 0.20. So, to get that version, visit the FreeCAD website: <https://www.freecadweb.org/downloads.php> and download the appimage file for Linux.

Double-click the appimage and you should have FreeCAD up and running.

If it doesn't run on a double-click, you might need to right-click on the appimage file and choose Properties. In the window that pops up, choose the Permissions tab. At the bottom of the window there's a checkbox for executable. Make sure that's ticked.


## GETTING STARTED

On the first run you'll see a page




on the right with three tabs and some examples. A lot of people get right into the settings and configure the starting page, colors, measurements, etc.. We're just going to get right into making something.


Let's start a new project. Click File > New menu. Now we have a project to work on.




The first thing we need to create within this project is a 'body'. Think of this as the thing you're going to create. But first, we need to swap workbenches. At the top of the screen is a dropdown menu that currently says 'Start'. Change that to say 'Part Design'.




There are two ways to create a body. Either through the icon that's below the dropdown, or (as we'll do) click the Tasks tab on the left panel. Think of this as a basic wizard to help you.







Click 'Create body' there. It'll then change to say 'Create sketch'.



Click 'Create sketch'.



Wow! Things have changed a bit here. Basically, it's asking us from what angle we're going to create this sketch. The sketch is going to be your drawing that we'll make into a part. Pick the XY plane. You can use either the left or right panel to choose the plane.



Think of this XY plane as looking down on top of the piece. The red dot is the midpoint. But look up top at the workbench dropdown. It said Part Design, but now it's changed to 'Sketcher'. This Sketcher workbench is where we do our technical drawing (as we called it when I was at school).

Your layout might look slightly different depending on version and screen size, but you can drag and drop the groups of icons using the little vertical dots. The pointer will change to a hand icon and you can drag them around.


We'll describe the icons as we use them, but think of the all red ones (shown above) as constraints.

These will lock your drawing in space.


The icons with the red dots and white lines (shown second from top) are ones we'll use to draw with.

Let's get something on the screen and we'll end part 1.

Click the down arrow beside the Rectangle icon and choose Centered rectangle.



Now click on the red dot, move the mouse, and click again to place the rectangle.




Beside the mouse pointer are the coordinates and a red

# HOWTO - FREECAD

rectangle. The red rectangle is to show that you're still placing rectangles. Either right-click or press Escape to come out of rectangle mode.


Now click the Close button on the left panel.



Now we're back in the Part design workbench. Have a look on the left panel again. If you're still in the Tasks tab, you'll see a list of things we can do here. For the moment pick Pad.

Boom! We have a rectangle that's 10mm thick. How do I know it's 10mm? Because the options on the left panel tell me. Feel free to tinker with that for now.

Ok, it's not the most interesting of things, but we'll add more to it next time.



**Ronnie** is the founder of Full Circle and, somehow, still editing this thing. He also paints, draws and does woodcarving in his spare time.

# The daily waddle

LOOKS LIKE A HARPOON  
TO THE HEART... WE NEED  
A COVID TEST TO BE  
SURE ...





BACK NEXT MONTH



# 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



# The Daily Waddle

I NAMED MY CAT  
"MARKETING" AS HE IS  
ALWAYS DOING HIS BUSINESS  
OUTSIDE THE BOX.

...





# MY OPINION

Written by Erik

# Linux Gaming?

I just tried to install a cross-platform file manager for a review in FCM. I downloaded the MSI file for Windows and tried to install it. The very first thing it tried to do was download Microsoft Edge components. I don't allow anything "Edge" on my system; in fact, I have removed it completely, by the roots, even in windowsSXS. Let's face it, the time of the Windows desktop is over. Now, I know gaming is the number one reason touted for Windows use, and I agree; however, there are still things that need work FOR gaming on Linux to function as it should, like fan control. I want my fans to ramp up sooner and keep my machine cooler – as I have seen too many times what heat does. (Remember MSI afterburner? I don't know now, as it does not work on my current machine, but it was pretty nifty.)

That said, Canonical was recently looking for a Linux Desktop Gaming Product Manager <https://canonical.com/careers/3776036/linux-desktop-gaming-product-manager-remote>. What is

interesting about the advert, was the second last point: "tell the story of Ubuntu for gamers". For one, I really don't think you need to sell anyone on it, simply make it work and work better, and it will sell itself. Youtubers will fall over themselves to report on Ubuntu beating Windows on the same hardware, like they do with video cards and CPUs.

This makes me excited for the Linux desktop. Not just for games, but for what will pop up around it, like the fan control I mentioned earlier as an example, or HUDs, game streaming, gaming kernels? The possibilities are mind-boggling.

Fast forward to April 2022, and one of the Ubuntu desktop team members, "Kevinvandine" posted an interesting tid-bit on the discourse: "Canonical is going all in on the gaming experience on Ubuntu and we've started building out a team dedicated to working on just that." (<https://discourse.ubuntu.com/t/introducing-early-access-to-the-steam-snap/28082>).

He exits with: "In addition, we're already looking at further future improvements for the gaming experience, such as providing easy ways to get more bleeding-edge components like Mesa drivers, and even newer kernels and proprietary drivers." I must admit, bleeding-edge drivers was not what I had in mind, but, to grab headlines and screen time from the gaming skills, you need to have those 3 extra frame rates over the Windoze version, right? I suppose it is like dominoes, get one falling, and the rest will follow? Ubuntu being the de-facto face of Linux, and probably the most well-known and talked about, it is marvelous to see them make Ubuntu attractive for PC gamers.

Kevinvandine also states that, "As part of these efforts, our first step is to simplify the process of getting your gaming environment set up without the need to add PPAs to get bleeding-edge packages or kernels. With Steam being the number 1 platform for PC gamers, it was the obvious choice

to focus our efforts here first." Steam has the most user friendly interface and is probably the most well-known, so it is natural to start there, not just because of the Steam Deck, am I right? "I bet Humble are sorry now: <https://www.pcworld.com/article/605285/the-humble-choice-subscription-is-ditching-mac-and-linux-gamers.html>. Who knows, maybe they will come around?

This will put the focus back on Canonical, I think, because we have to face facts, that they have been slacking the last few years, focussing on Snap and letting things like the Software centre die. Speaking of the software centre's UX... if they could adopt something like Solus's proposed software centre (not sure if it is only mockups), or even the Electron Apps store. I would contribute to something like that. It makes Linux SEXY.

Would you like to tell us something? [misc@fullcirclemagazine.org](mailto:misc@fullcirclemagazine.org)



# 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](mailto:articles@fullcirclemagazine.org)

## TRANSLATIONS

If you would like to translate Full Circle into your native language please send an email to [ronnie@fullcirclemagazine.org](mailto: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

# Lubuntu 22.10

Released on 20 October, 2022, Lubuntu 22.10 marked the start of a new development cycle. Like all cycles, this one will consist of three interim releases – leading to the next long-term support version, Lubuntu 24.04 LTS, due in April, 2024.

The first interim release in a new cycle is always interesting, as it tends to indicate what can be expected through the cycle and how much new will end up in the final LTS. If the release of Lubuntu 22.10 is any indication of what can be expected, the answer is “not much.”

Lubuntu 22.10 is the ninth Lubuntu version with the LXQt desktop, and the 26th release overall for Lubuntu. The very first release, Lubuntu 10.04, came out on 2 May, 2010, 12 years ago.

## INSTALLATION

I downloaded Lubuntu 22.10 from the official source via BitTorrent. This release is 2.7 GB in size compared to 2.5 GB for the last

release. In comparison, the mainstream Ubuntu 22.10 is 3.8 GB to download.

Once I had the ISO file downloaded, I did a SHA256 sum check to confirm that the file was good, and then dropped it onto my USB stick equipped with Ventoy 1.0.81 for testing.

## NEW

There is actually not a lot that is new in this release. It uses the LXQt 1.1.0 desktop, which is based on the Qt 5.15.6 toolkit. The new Linux

kernel is 5.19.

Visually, Lubuntu looks the same as other recent releases. The default theme is still Lubuntu Arc with ePapyrus icons. There is the usual wide assortment of window themes, icon sets, plus panel and menu color schemes to choose from, to customize your Lubuntu installation pretty much any way you like.


Perhaps the most notable change in this release is that new versions of both the file manager, PCManFM-Qt, and the image

viewer, LXImage-Qt, now support the WebP image format. This format for photos and drawings is becoming more common and, while web browsers like Firefox have supported it for a while, the rest of the operating system is now catching up. In general, WebP is a replacement for JPG, GIF and PNG for web photos, with smaller file sizes as well as support for transparency and animation.

This release does have new artwork in the form of several new wallpapers. At one time in the distant past, Lubuntu was noted for its beautiful, blue, wavy, abstract wallpapers, but those days are long gone. The new default “Kinetic Kudu” wallpaper was created by Lubuntu Team member Aaron Rainbolt. Fortunately, there are some alternative wallpapers provided, or you can use your own.

## APPLICATIONS

Some of the applications included with Lubuntu 22.10 are: 2048-qt 0.1.6 simple lightweight game\*



# REVIEW

Bluedevil 5.25.5 bluetooth connector  
Discover Software Center 5.25.5 package management system  
FeatherPad 1.3.0 text editor  
Firefox 105.0.3 web browser\*\*  
ImageMagick 6.9.11.60 image editor\*  
Kcalc 22.08.1 calculator  
KDE partition manager 22.08.1 partition manager  
LibreOffice 7.4.2 office suite, Qt interface version  
Lubuntu Update Notifier 0.4 software update notifier\*  
LXImage-Qt 1.1.0 image viewer and screenshot tool  
LXQt Archiver 0.6.0 archive manager  
Muon 5.8.0 package manager\*  
Noblenote 1.2.0 note taker\*

PCManFM-Qt 1.1.0 file manager  
PulseAudio 16.1 audio controller  
Qlipper 5.1.2 clipboard manager\*  
qPDFview 0.4.18 PDF viewer\*  
QTerminal 1.1.0 terminal emulator  
Qtransmission 3.0.0 BitTorrent client, Qt interface version\*  
Quassel 0.14.0 IRC client\*  
ScreenGrab 2.2.0 screenshot tool\*  
Skanslite 22.04.3 scanning utility  
Startup Disk Creator 0.3.15 (usb-creator-kde) USB boot disk maker\*  
VLC 3.0.17.4 media player  
Wget 1.21.3 command line webpage downloader  
XScreenSaver 6.02 screensaver and screen locker\*

\* Indicates the same version used in Lubuntu 22.04 LTS.

\*\* supplied as a snap, so version

depends on the upstream package manager and will update automatically.

Other than updated versions, there have been no changes to the mix of applications provided in Lubuntu 22.10.

Like the earlier LXQt releases, Lubuntu 22.10 does not come with a webcam application, email client, CD/DVD burning, photo editing or video editing software, although these can be easily added from the repositories, if desired.



It does seem from the release announcement that there was some consideration of moving away from Firefox as the default

browser. There have been user complaints of slow start-up times, but the development team has decided to stick with the Firefox snap package. In my own testing, the first start of Firefox can take a few extra seconds, but after that it opens quite quickly.

This Lubuntu release continues to use PulseAudio as its audio controller, while Ubuntu 22.10 has moved to PipeWire. It will be interesting to see over time if Lubuntu sticks with PulseAudio or not.

## CONCLUSIONS

If Lubuntu 22.10 is any indication



# REVIEW



of how the rest of the development cycle will go, then it looks like we can expect only a few minor tweaks between now and the LTS release in April, 2024. Overall, that is not a bad thing as the LXQt desktop is working really well these days and most users don't seem to see a need for big changes.

## EXTERNAL LINKS

Official website:  
<https://lubuntu.me/>



**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 Erik

# Zorin OS pro 16.1

Website: <https://zorin.com/os/>

Price: \$39 USD

Blurb: *"Zorin OS is the alternative to Windows and macOS designed to make your computer faster, more powerful, secure, and privacy-respecting."*

After having a really bad experience with the free version of Zorin OS 16, I thought I should give the pro version a go.

There are no "informationals" whilst installing, so the install proceeded quite quickly.

A bit worrying was that the installer did not respect my options, choosing not to install updates whilst installing, or taking part in the "census", it still reached out to the internet and did some stuff. :( The second time, I disconnected the cable and watched it get stuck in loops as it tried to call home. This kind of behaviour is NOT on. I don't care where it was connecting to, I made

it clear not to. (Privacy respecting, my Tuchus!).

Looking around the install, I noticed that there was a system folder (/var/lib/flatpak) taking up a lot of space. This was, to my mind, the secret sauce folder. It housed Flatpaks of Gnome and KDE, also some applications for them (the largest folder in my installation).


So let us talk about the main feature of the OS, the layout chooser. All the layouts, except whatever this is... worked out of the box.

The left hand menu, be it plank or docky or whatever, is missing. Though my screen is a crappy 1366x768, I could not adjust some of the window sizes, like the appearance window seen here.

As you can see (right), choosing orange in light mode gave me orange, but choosing orange in dark mode netted me a muted salmon... and also a weird overlay like a night light, on both light and dark.

All of the "layouts" are re-skins of Gnome (not XFCE like the free version), so I'm not sure where the KDE flatpak fits in. The only reason I can see for this is there are a few KDE applications installed, like KdenLive and Krita, by default, though they are flatpaks too, meaning I'm left scratching my

beard again. For all the applications installed, it did not have a decent top installed, atop, btop, htop – only plain top, nor inxi or neofetch, but things like FreeCAD and LibreCAD, both. This makes it feel less curated and more like "let's throw things against the wall and see what sticks"; even default Firefox just comes with a modified



# REVIEW

home page, but all the crap you don't want as a Linux user. (Clearly, the targeted demographic here is Windows 'converts'... XD). There were some nice touches, like being able to right click-the taskbar to find the layouts application or open the system monitor.

If you have a modern computer, Zorin actually runs reasonably well,

unless you plan on opening the fonts application – then... you wait. Even opening it sequentially... you wait. It may be because they have lots of spaghetti fonts installed. There were also nice fonts installed as well. I didn't check to see if the fonts application was a Flatpak too, but the system relies heavily on Flatpaks, I saw seventy (70) out of the box.

The "Kooha" application also did not launch for me, but it was an application I was not familiar with, so it was neither here nor there.


At rest, Zorin actually used a lot less memory than expected, and when the dust settled, it used less hard drive space too. Though it settled at 19GB, I could not install it

on a 30GB SSD, it would run out of space and give warnings about 107MB of space remaining before crashing and rebooting the system. Checking the website, I saw that 40GB was needed for Pro. You need to be aware of this, or your installation experience will put you off Zorin.

## SNAP, SNAP OR SNAP?

The software application was filled to the brim with software, though with some software sources being Zorin only. (I am assuming these are deb files compiled on and by Zorin themselves.) Some were not labelled, and I assumed these to be plain .deb-files, but they were Flatpaks, so I'd say there was still a bit of work to be done. The software centre does give you the option of installing different versions of software, but sometimes there are no real choices at all. I did appreciate the choice of office during the installation process, but if you are going to do that, you need to give them all a fair chance, no?

Being spoiled by Voyager Linux, I found it mildly annoying that I could not turn off my trackpad at will, so I



# REVIEW

do not touch it whilst typing, but I saw no such option in the mouse settings. Nor did I easily see where to change the default mouse pointer. On a white (light) theme you have a black pointer, and on a black (dark) theme you have a black pointer, not ideal. This could just be a Gnome thing, but one would think that this is one of the basic things you need, especially since corner tapping the trackpad is not a thing in most Linux distributions. (Tuxedo

has it).

There is a selection of backgrounds that should appeal to most people included, and not mimicking Apple is a breath of fresh air. I realise most Linux users will theme immediately, but the distro is aimed at Windows converts, who are still sheep at first.

Overall the distro behaves quite well, though there are a few

annoyances that need attention. If you are a Gnome fan, you should feel right at home. It may not feel as polished as other Gnome distributions that have put a lot of effort in, but it is standard fare meant to ease Windows users, not so much on the Mac front. The Zorin developers are putting effort in, so if you are able to support them, do so please.




**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.



# BOOK REVIEW

Written by Greg Walters

# Python For Kids



## Python for Kids 2nd Edition A Playful Introduction to Programming

Jason R. Briggs

Release Date: November 15, 2022  
No Starch Press

368 Pages  
ISBN-13 978-1718503021  
\$33.49 Paperback Amazon.com (US)

I love teaching people things. Especially Python. No Starch Press has, with author Jason R. Briggs, updated his book on teaching kids Python called 'Python for Kids Second Edition'. I'll give you a hint, I LOVE IT!

Jason Briggs brings the concept of Python to kids (and adults who want to learn) in a way that won't bore or frustrate new Python programmers. Chapter 1 takes the reader through the process of installing Python on Windows, Mac, Linux and even Raspberry Pi machines. Then he shows how to use IDLE, the builtin IDE for Python. He then moves the reader from basic calculations, Operators, and variables, into strings, lists, tuples, and even dictionaries. Then he jumps the reader into Turtle graphics, all within the first four chapters.

Jason then goes through more Python basics like if | else and loops, then tackles Classes and Objects. Ending up part 1 is a chapter on using Tkinter (and you know how much I like Tkinter!)

Parts 2 and 3 start teaching the reader how to write some basic games (again using Tkinter) in a structured and logical way. I love the fact that Part 3 is named "Mr. Stickman Races for the Exit". Jason is my kind of programmer!

All in all, this is a great book to start for first-time programmers on how to use a programming language, and especially Python. Kid or Adult, this is a wonderful book for beginners.



## Table of Contents:

### Part 1: Learning to Program

- Chapter 1: Not All Snakes Slither
- Chapter 2: Calculations and Variables
- Chapter 3: Strings, Lists, Tuples and Dictionaries
- Chapter 4: Drawing with Turtles
- Chapter 5: Asking Questions with if and else
- Chapter 6: Going Loopy
- Chapter 7: Recycling your code with Functions and Modules
- Chapter 8: How to use Classes and Objects
- Chapter 9: More Turtle Graphics
- Chapter 10: Using tkinter for Better Graphics

### Part 2: Bounce!

- Chapter 11: Beginning your First Game: Bounce!
- Chapter 12: Finishing your First Game: Bounce!

### Part 3: Mr Stick Man Races for The Exit

- Chapter 13: Creating Graphics for the Mr. Stick Man Game
- Chapter 14: Developing the Mr. Stick Man Game
- Chapter 15: Creating Mr. Stick Man
- Chapter 16: Completing the Mr. Stick Man



# LETTERS

If you would like to submit a letter for publication, compliment or complaint, please email it to: [letters@fullcirclemagazine.org](mailto: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](https://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)


## VENTOY - THE REAL SOLUTION?

I have a simpler solution for the problem which was described in the Ventoy letters in FCM 187. The issue happens because Ventoy creates two partitions - one bigger exFat, and another small FAT. Older TVs are usually compatible only with FAT (maximum)/NTFS (some of them) partitions. So, to revert back the situation, we have to delete the partitions and re-create the FAT partition. This can be easily done by the "Disks" application which comes with Ubuntu installations.

Step 1) Select the partition and delete it using "-" (red minus) sign.



Step 2) It will ask you to confirm. Select "Delete" to proceed with it.



Step 3) Repeat the same with the other partition. The Result will be free space occupying the complete USB disk.




Step 4) Next click on "+" ( plus sign ) to get the "Create partition" screen as below.



Step 5) Click Next in the above screen to proceed to the "Format

Volume" screen. Here you can select the name of the volume. Additionally, you must change the selection to "FAT" for maximum compatibility.



Step 6) The restoration is complete as shown below.



Manish Raj

## 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](mailto: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 standing in for one of our technicians at a large tyre factory. They were a nice fat target for viruses, malware and whatnots, as is the case with very large companies. To help combat the issue, they had, I think, Mail Marshall running as well, at the head office. Obviously this parsed all the content in an email and blocked emails containing certain words. You then got a notification, telling you it had been blocked. You then had to contact your local IT (in

this particular case, me), to contact HQ IT to release the email after they had inspected it. You would also get a notification if someone sent you a virus, presumably so you could inform your contact. This was a major issue, as people would send you messages to release the email with a virus in it. No matter how many times you explained that you cannot release a virus, the more they would insist. Having basically "grown up" in IT, I cannot fathom what goes through a user's head when they do stuff like that. I had one guy go to his HOD and they both tried to gang up on me to get the virus released, because they "needed it". They would not listen to reason, it was just IT being spiteful. I have even had people wanting Nigerian prince emails released, because it was "important". I don't even open emails from people at vendors I don't know. I get sent an email from cathy.jones@vmware; I don't know Cathy Jones; it goes straight to junk. I don't open it. This also cuts down on email fatigue. This BS, where you "can't miss an email", is nonsense, you don't do anything

urgent via email as emails can get lost, etc. If it is urgent, you pick up the phone, or send a message on Teams/Slack or whatever, to say I just sent you an urgent email. I agree we need trust in our lives, but not blind trust.

Q: I'm not too sure about this, I have Ubuntu 20, and I have apparently hung the Software Updater. After a few minutes, I killed the process. Restarted the Software Updater and it stalled on "checking for updates". I'm almost positive the last one was quick. How can I fix this?

A: You could try updating via the command-line to see what errors you get (it may be a PPA?). Alternatively, you could switch update mirrors. May I suggest that you start updating, and getting a cuppa, and have a chat with the neighbour, and coming back. If it is installing kernels or the like it may take a bit longer, same if the mirror is busy.

Q: At the top of the screen, there is date and time window in there is a bell icon, if you click on the clock you see: "Pending Update Snap Store". If I click on updates, I get Snap Store 41.3-60 48.2MB Update. If I then click on 'Update All', I get Unable to update "Snap Store": (null): cannot refresh "snap-store": snap "snap-store" has running apps (ubuntu-software) This seems circular to me?

A: Close everything, open a terminal and type:

```
sudo snap refresh
```

then type in your password and let it complete and reboot and you should be good.

Q: I am a beginner to ubuntu. SDA is full and SDB is empty. SDB is not mounted yet. I need detailed instructions since I am a beginner of ubuntu.

A: You did not say for what? If you want to format and mount sdb, simply use the disks application; if you want terminal help, you should say so. If your drive

## Q&A

is too full to do any of it, boot from the live image and clear some space. You need to say exactly what you want; I can't see your screen? I will do a piece on disks in the magazine, just for you.

**Q**: I have a problem with some nohup output. When I run a job, I have a lot of output data. The output nohup.out becomes too large and everything slows down. Someone suggested I redirect the output. I'm not sure what to do.

**A**: You could use: nohup "your command" > /dev/null 2>&1& which just sends the output to a black hole. There is a nice write up on: <https://unix.stackexchange.com/questions/3886/difference-between-nohup-disown-and-148698#148698> - where they explain nohup better than I can in a Q & A.

**Q**: I am looking to get tab-completion on my command-line aliases. I heavily use aliases, as I am new to Ubuntu and Linux in general, and I hate typing long commands. Like: alias sau='sudo apt update'. BUT aliases do not work with tab completion with the names of apps that you don't know, so sudo apt install ink [tab], [tab]

will bring up inkscape, but not if I use an alias.

**A**: Aaah, never tried it myself, I type fast enough, but it seems there is someone out there who likes you. See: <https://github.com/cykerway/complete-alias>. This seems to be what you need. Just pay attention to the dependencies.

**Q**: I have a no-name bluetooth USB dongle in a USB port on my black-box-special PC, which lights up when I insert it, but that is about it. How can I get Ubuntu to see it? Oh yes, I have 22.10 on a four-year old machine, not sure what's inside.

**A**: You can try this: [https://help.ubuntu.com/community/BluetoothSetup#Manual\\_Discovery](https://help.ubuntu.com/community/BluetoothSetup#Manual_Discovery), I have no bluetooth devices, so I have never used it. If you want to know "what's inside", install inxi and run: inxi -a

**Q**: I have mongodb running on Ubuntu 20.04 and I have a few Python scripts running that use it. However, lately most of them have been failing, and I don't know what has changed. I have not upgraded Ubuntu (yet!). Is it Ubuntu or Python?

**A**: Mongodb has changed; the "mongo" command has been superseded by the mongosh command, which breaks previous compatibility with the mongo command. <https://www.mongodb.com/blog/post/introducing-the-new-shell>. It has nothing to do with Ubuntu.

**Q**: Hey guy. I am trying to install Lubuntu on an older Core2Duo machine. I created a boot USB, and can get it running from there. However when I try to install it, I get almost all the way to the end and it fails installing the bootloader. I get the error: "The bootloader could not be installed. Erm, what now?"

**A**: I had a similar issue with Linux Lite 6.0. When 6.1 launched, the issue was fixed and it installed flawlessly, so I cannot tell you what the issue was. Try getting a later version of the install image (you may have to wait, unfortunately). Alternatively, may I suggest blanking the disk with a gparted bootable USB and trying again? (unless you know how to use fdisk?).

P.S. I just tried in a VM, and it installed first time. Maybe try that,

it may be an error with the installation medium.

**Q**: I take my Ubuntu Laptop to work, just to keep work and private/semi-private stuff apart, so I don't do anything not work related on my work laptop. Thing is, when I go for a quick coffee break, I close the lid, but I only want it to lock the laptop, not go to sleep. I have caffeine and caffeine-indicator installed, but it does nothing. In windows it is in the power options, but I don't see it on Xubuntu?

**A**: It is not an option out-of-the-box as I recall, but you can\* change it. See: <https://ubuntuhandbook.org/index.php/2020/05/lid-close-behavior-ubuntu-20-04/>

**Q**: I have a new headset and I cannot use it on Ubuntu 22.04 with Skype or anything. The other party does not hear me at all. I installed audacity and I record, but I get nothing. The mic is not disabled in the settings, I checked. This is the headphones <link removed>, and this is the laptop <link removed>. My old broken headset worked here before. The only difference between the 2 is the USB cable.

**A**: Guys, if you send links, try not to make them Amazon shopping links, link to the provider. The USB cable on the newer model drives only the disco-disco LEDs. If you are sure that you can record with the on-board microphone, look for a hardware switch on the cable for the headset. That model definitely has one. Make sure it is not in the off position, effectively muting the unit. There is no other reason, except for a defective unit, that it would not work.



**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.



# TABLETOP UBUNTU

Written by Josh Hertel



BACK NEXT MONTH



Josh Hertel is a husband, father, mathematics educator, tabletop gamer, techie, and geek.

<https://twitter.com/hertelj1>



# UBUNTU GAMES

Written by Erik

Website: <https://www.kasedogames.com/ixion>

Price: 25-35 USD (source: <https://isthereanydeal.com/game/ixion/info/>)

Blurb: "Authorised by DOLOS AEC, you will guide the Tiqqun space-station onwards in a perilous journey to find a new home for humanity. Station infrastructure, population management, exploration, survival, these responsibilities fall to you. Will you find hope out there in the dark? Good luck, Administrator."

Before I get into this game, I would like you to experience what got me – hook line and sinker. The soundtrack! <https://www.youtube.com/watch?v=RhTXyUERugQ> Should the link be removed due to copyright, search for; "IXION Vanir's legacy". Do crank that volume to 11 and I suggest headphones, something that can handle bass, as the first bass drops hard!

Ixion is not a native Linux game though, which is a crying shame. If you want to play it, you need proton (<https://www.protondb.com/app/1113120>)

So as you already know, the soundtrack is amazing, coupled with the beautiful views of the earth and other planets, you get whisked away – right into the fantasy. Immersion levels during travel, and cut scenes, will let you forget about reality completely. I am not even going to say any more on this topic, as you need to experience it first hand.

The main game play loop reminds me of Frostpunk, but there are enough differences for this to be its own game. At heart, this is a colony management sim, but there is no micro. You start in one of the six sectors of the ship, and you open more as you progress. You have to think about placing structures as you have limited space and/or resources. One of the bonuses is that the game does not penalise you for destroying a building. You get credited 100% of the materials. Though the game does not involve micro-

management, the things like water and food and power are still present. The game also features a skill tree and quests. Yes, you get to build miners, cargo ships, science ships, and probes that let you play inside the craft and outside, in the solar system.

This is all fine and dandy, but the game rounds itself out with a kick-ass (can we say that?) story line. Your mission is to take humanity to their future on an exoplanet. To be honest, whenever I hear the word exoplanet, I think rogue planet. **<Spoiler alert>** Skip to the next paragraph, if you don't want any spoilers.<Spoiler starts> However, your jump goes wrong and instead of jumping through space, you jump through time. You go from being the heroes of humanity, straight to being the villain.<Spoiler ends>

There is the usual resource management, quests with outcomes that fit the old adage of: you can have something well done, cheap, or fast, but never all three. You also need to manage your



“trust rating” as the administrator of the vessel as well as the people on the vessel’s happiness and work optimization. If you overwork your people, there will be accidents, which have an influence on things like morale. Nobody wants strikes and riots. Then again, you can employ the BBC, I mean propaganda centre (we are splitting hairs, aren’t we?) and nightclubs to improve morale.

The game was the highest rated and most played demo in the steam next festival, which alone should tell you something. Though the reviews on Steam are “mixed”, the complaints are that people find it difficult, most are “stuck” on level two. This was the same with

Frostpunk. You cannot play this game like you do Starcraft; it is not that type of game, you need to consider every move, like a game of chess. You need to weigh up your mission, with the needs of your population and with the resources at hand. For instance, you can keep building storage spaces, but each one requires power and workers, on top of the building cost. Sometimes you have to turn buildings off, to get power or people back to do other things, and this is where many players miss the plot completely.

The game play is actually fun and you are always chasing the next story point in the great mystery. Resource collection and

management does not feel like a major grind. The game loads pretty quickly on your SSD, unlike another release this month that feels like it’s loading from a floppy – yes, I’m looking at YOU – Knights of Honor II!

If you need to buy yourself a Christmas present this year, consider this awesome game. Local pricing puts the game at a respectable price too, considering Knights of honor II is double the price for half the fun. However, the game is a bit linear, with the quests providing multiple paths allowing for some replay-ability.

Though I have not completed the game (it came out this week

and I really have no time to play), I still recommend it based on what I have experienced so far. Yup, sounds nuts, but I am a sucker for story-rich games. The game does have a few paper cuts that need fixing: for instance, you can move anything between sectors, except water. Why not? Anyway, it does not take away from the game, so I will stick to calling these issues paper cuts.



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





# PATRONS

## MONTHLY PATRONS

Alex Crabtree  
 Alex Popescu  
 Andy Garay  
 Bill Berninghausen  
 Bob C  
 Brian Bogdan  
 CBinMV  
 Darren  
 Dennis Mack  
 Devin McPherson  
 Doug Bruce  
 Elizabeth K. Joseph  
 Eric Meddleton  
 Gary Campbell  
 George Smith  
 Henry D Mills  
 Hugo Sutherland  
 Jack  
 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

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

**2023:**  
 Floyd Smith  
 Richard Almeida  
 Gavin Thompson  
 Raymond Mccarthy  
 Michael Grugel

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

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

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

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



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



<https://paypal.me/ronnietucker>



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



# HOW TO CONTRIBUTE

## FULL CIRCLE NEEDS YOU!

A magazine isn't a magazine without articles and Full Circle is no exception. We need your opinions, desktops, stories, how-to's, reviews, and anything else you want to tell your fellow \*buntu users. Send your articles to: [articles@fullcirclemagazine.org](mailto: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](mailto:letters@fullcirclemagazine.org)  
Hardware/software **reviews** should be sent to: [reviews@fullcirclemagazine.org](mailto:reviews@fullcirclemagazine.org)  
**Questions** for Q&A should go to: [questions@fullcirclemagazine.org](mailto:questions@fullcirclemagazine.org)  
**Desktop** screens should be emailed to: [misc@fullcirclemagazine.org](mailto:misc@fullcirclemagazine.org)  
... or you can visit our **site** via: [fullcirclemagazine.org](http://fullcirclemagazine.org)



## FCM#190

### Deadline:

Sunday 05th Feb. 2023.

### Release:

Friday 24th Feb. 2023.



## Full Circle Team



**Editor** - Ronnie Tucker

[ronnie@fullcirclemagazine.org](mailto:ronnie@fullcirclemagazine.org)

**Webmaster** -

[admin@fullcirclemagazine.org](mailto:admin@fullcirclemagazine.org)

## Editing & Proofreading

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

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

## Getting Full Circle Magazine:

### For the Full Circle Weekly News:



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



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



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



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



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

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