

Pi vs Pi

Orange Pi 5 Plus takes on the Raspberry Pi 5



RETRO MEDIA

How to save your old disks to USB



MANJARO 23

Get the Arch distro that's easy to use

LINUX FORMAT

The **#1** open source mag

THE 5 BEST NEW DISTROS

Blast off with the next-gen in Linux tech as we test-ride smarter and safer distros



GET SECURED!

Protect your data with the best password managers

100
pages of Linux
tricks, tips
& more!

PLUS: HOW TO

- » Secure your video chats with Jami
- » Home-cook meals with Mealie recipes
- » Protect yourself from the top hacks of 2024

FASTER ETHERNET!

The low-cost options for faster Ethernet transfers

GO ANALOGUE

Wire up an old-school computer

CODE THE C64


We dive into classic 6502 processor code

LXF March 2024

FUTURE

WE'VE GOT THE OPTIONS. YOU MAKE THE CHOICE.

Car, home and travel insurance plus
broadband, breakdown cover and more.



GO.
COMPARE

Get more info or compare quotes for
a range of products at [Go.Compare](https://www.go.compare).



LINUX FORMAT



» MEET THE TEAM

This issue, we're looking at exciting new distributions, so what new distro have you tried recently that you would recommend readers try for themselves?



Nate Drake

I've had great fun taking Maestro for a spin, a Unix-like kernel and operating system written entirely in Rust. So far it's ported most Linux system calls and even has its own boot system and package manager.

I can't wait to see what they come up with next.



Les Pounder

I have a deep love for CrunchBang Linux. Since that project ended, I've been using BunsenLabs and CrunchBang++ on my old Lenovo X220. If you have old kit and want to keep using it, CrunchBang/

BunsenLabs can squeeze a little more life from it.



Michael Reed

I always keep a copy of EasyOS on a flash drive. It's an odd distribution, with a really weird take on the root user, but it can get you out of many awkward situations because it's readily expandable. Daily

driver? No way. Potential life-saver? Possibly.



David Rutland

Last month, I installed stock Debian for the first time to convert an old Chromebook into a lean, mean Linux machine. I was surprised at the ease of installation and usability. It's also super-stable, and the

ex-Chromebook hasn't been powered down since installation.



Nick Peers

I've been scouting about for a server distro for my new-build NAS-like server. I've settled on Debian Server – it's similar to Ubuntu, but supports the latest version of Podman, which I'm going to use to power all

my services. Not exciting as such, but perfect for my needs.

Cutting edges



Living in the open source world means if you want to sit on the cutting edge of development, you can. Nightly builds, compiling the kernel from source, grabbing Git repositories, if you want to try the latest, no one is going to try and stop or gatekeep you... Generally, though, that sort of thing can be exhausting, which is why it's usually just interested devs who sit on those particular cutting edges.

But there is another way. Since the dawn of Linux, helpful individuals and groups have pulled together all the right sources and binaries, and created handy distributions of everything you need in a cutting-edge OS. Back in the mists of time, Debian grew into a fundamental building block of the Linux world. More recently, Arch Linux and Gentoo Linux offered cutting-edge access; even Ubuntu turns 20 this year, and that started off wanting to offer a human spin on Linux distros.

So, this issue we're taking a look at the best next-gen distros that are offering new technologies, new ways of organising distros, new ways of packaging software, new ways of securing data and new ways of working. They're all exciting, most are brand new – though one is called Ubuntu – and we think you should give them a spin. While Mac OS stands still and Windows copies Linux, it's exciting to see the Linux desktop and ecosystem ever evolving with plenty to enjoy!

Neil

Neil Mohr Editor
neil.mohr@futurenet.com



**Subscribe
& save!**

On digital and print:
see page 16

Contents




REVIEWS




Intel Core i7 14700K.....19

Jacob Ridley thinks this could be a perfectly good stand-in for a Core i9 – well worth putting in your next build.




RixOS 2023.11.....20

Nate Drake finds himself washed up on the rock of ages that is RixOS, only to discover that it still needs work.



Emmentubtüs 5.....21

Nate Drake finds that despite being hard to pronounce, the name is the only complicated thing about this distro.



Manjaro 23.1 Vulcan.....22


Nate Drake boldly goes to take Vulcan for a test spin and finds an incredibly user-friendly and efficient operating system.

THE 5 BEST NEW DISTROS

Blast off into the future with **Matt Holder** as he explores the technology powering a new generation of Linux distributions. See page 30!




ROUNDUP



Password managers.....24

Security-conscious **Michael Reed** examines five pieces of software that can safely store and retrieve your passwords and other sensitive data.

IN DEPTH



Orange Pi vs Raspberry Pi.....44


Tam Hanna compares the highest-end single-board computer offerings from Shenzhen Xunlong and the Raspberry Pi Foundation to see which is best for you.

PI USER



Pi news 39
Introduced by **Les Pounder**, with updates on the Pi supply situation, and news of interesting Pi-based projects.

Official Pi 5 Case 40
Can the new official Raspberry Pi 5 case keep the Pi 5 cool? **Les Pounder** is on the case, stop groaning at the back!




Geniatech XPI-3566-Zero 41
We might be talking about this SBC, but all brains and no brawn is also one way to describe **Les Pounder**.

Use MQTT for sensor logs 42
Les Pounder is a sensitive soul and his new office needs sensors to record data. He'd better break out the maker kit!

CODING ACADEMY

Write and run C64-style 6502 code 90
David Bolton shows how to get started writing 6502 code and provides a useful example that sorts numbers.

LXF Shell: redirection redemption 94
Ferenc Deák continues writing shell-enhancing redirection features and hiding the odd film title for you to spot, the scamp!



REGULARS AT A GLANCE

News 6
Linux is a hero OS, a new Cosmic DE pops into existence, Unity is not so united, a Rust-covered kernel is written, plus a round-up of the latest distro releases, and opinions from industry insiders.

Kernel watch 10

Answers 11
Neil Bothwick solves Linux queries as a form of meditation. This month, he has been relaxing while working out why *LibreOffice* refuses to open spreadsheets, how to remove the need to log in, and much more.

Mailserver 14
Readers write in with requests, queries and suggestions, and Neil politely replies.

Subscriptions 16
Grab your monthly dose of Linux and save 50% off the usual price!

Back issues 62
Get hold of previous *Linux Format* editions.

Overseas subscriptions 63
Get *Linux Format* shipped around the globe.

HotPicks 83
Mayank Sharma can't believe the F1 circus is already on hiatus and he has to find a new distraction after checking out *Shotwell*, *You Tube Downloader Plus*, *Blue Recorder*, *EncryptPad*, *Buttercup*, *Resources*, *Rnote*, *heXon*, *EDuke32*, *Delta Chat* and *Proton VPN*.

Next month 98

TUTORIALS

TERMINAL: Monitor systems in style 48
Shashank Sharma doesn't look dapper in his quickly fading imitation leather jacket, but his terminal is always a sight to behold.

JAMI: Secure and private comms 50
Nick Peers takes a fresh look at an old favourite – *Jami*, a tool that aims to cover all your communication and sharing needs..

MEALIE: Recipe management 54
David Rutland dons a tall hat to prepare a portion of *Mealie* – a recipe management server with a side dish of meal planning.

OPENCBM: Rescue retro media 58
Christian Cawley has some old C64 5.25-inch disks, and wants to revisit some personal files from 40 years ago.

ADVENTURE GAME STUDIO: NPCs 64
Nate Drake invites you to relive the glory days of point-and-click adventure games by creating your very own.

EMULATION: Analogue computer 68
In our journey into analogue, **Mike Bedford** looks at another emulation approach, and how to get a real analogue computer.



ETHERNET: The wire to faster Ethernet 72
Searching for faster than Gigabit Ethernet, the ever frugal **Neil Mohr** does that compromise thing his ex told him about.

ADMINISTERIA

Administeria 74
Stuart Burns is as mad as hell and he's not going to stand for your shoddy *Docker* deployments any longer!

Starlink 76
Wherever he may roam, **Nate Drake** can enjoy fast internet, now available for a fair price.

Hatch 77
A look at an easy-to-use website builder that empowers **Mirza Bahic's** creative expression.

Protect your systems 78
Davey Winder reveals the biggest threats that hit systems in 2023 and what's likely to happen over 2024.

Newsdesk

THIS ISSUE: Linux is a hero OS » New Cosmic DE pops into existence » Unity is not so united » Rust-covered kernel is written

OPERATING SYSTEMS

Linux to save Windows 10 PCs from the apocalypse?

240 million PCs could be scrapped after October 2025, when free support ends for Windows 10. Or users could just switch to Linux.

Termination of Windows 10 support would prevent millions of devices from having a second lease of life, leaving around 240 million liable for landfill, according to research by Canalsys*. If these were all folded laptops, stacked one on top of another, they would make a pile 4,080km high!

Free support for Win 10 ends in October 2025, after which Microsoft will only support it via a paid subscription until 2028. If this mirrors past price trends, Canalsys claims it may be more cost-effective for owners to scrap their PCs.

The same research also points out that forcing users to upgrade could provide a boost to the slumping PC sales market, with an estimated growth of 8% predicted for 2024.

Canalsys does acknowledge that such PCs could be recycled, but notes, "Their incompatibility with the latest supported version of Windows massively reduces their value for refurbishing and reselling."

No doubt an obvious alternative solution to stacking junked computers on *The X Factor* has already occurred to *Linux Format* readers.

While requirements such as TPM 2.0 support may prevent a PC from being upgraded to Win 11, many Linux versions are more forgiving.


This means once Windows 10 reaches EOL, PC users can still use a current and secure OS without consigning their machine to the scrap heap. This is largely down to the support of the open source community, which actively maintains code for hardware abandoned by commercial manufacturers and developers.

The ATI R300 graphics processor series is an excellent example. First released in 2002 along with the ATI Radeon Pro, graphics cards that use them rely on an AGP (Accelerated Graphics Port) interface, rather than the newer PCIe. This means they can't run the latest games and AMD has long since discontinued active support.

Nevertheless, the open source community has released numerous updates to support the pre-R600 series hardware. Driver updates for ATI's R300 through R500 series Radeon GPUs will also be available this quarter in Mesa 24.0.

The 22-year-old ATI Radeon 9700 Pro packs a core clock speed of just 325MHz, 256MB of memory, and 19.8GB/s of bandwidth. But thanks to the open source community, it is compatible with some modern Linux distros, offering users a better alternative to being forced to send their machines to the scrap heap every few years.

Ancient hardware, such as the ATI Radeon 9700 Pro, is still supported by the open source community, which develops and maintains drivers.



DESKTOP ENVIRONMENT

Cosmic looking out of this world

Pop!_OS Cosmic DE is to be released as an alpha, with a new terminal and overhauled graphics support.

Denver-based System76 is renowned for selling powerful notebooks, desktop machines and servers with Linux preinstalled. Users can choose stock Ubuntu, but for the past six years, the company has also offered its very own Pop!_OS.

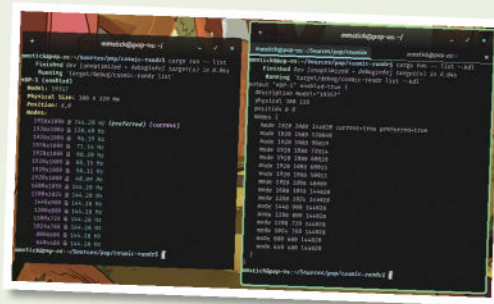
Based on Ubuntu, it uses its own Gnome-derived desktop environment, Cosmic. In a recent blog post, System 76 announced that it's on the "home stretch" in putting together an alpha for the latest Cosmic DE.

Special mention should go to Cosmic Terminal (*cosmic-term*), based on *alacritty* terminal. It features bidirectional rendering, desktop and syntax themes. It also supports GPU rendering, using glyphon and wgpu.

One use for the terminal is running *Cosmic-randr*, a command-line utility for listing and configuring displays in Wayland. It will also be used by Cosmic Settings' display settings page.

The team at System76 has also been making good use of the new screenshot applet to showcase Cosmic DE. It supports both full-screen and windows.

Cosmic DE will operate only in hybrid graphics mode. Applications will run using



Cosmic-randr, shown outputting a list of display modes (left) and listing the commands for changing display mode (right).

the dGPU (dedicated GPU) if specifically requested, but will otherwise default to using integrated graphics. The system can also turn off the dGPU to extend battery life. This feature will no doubt be put to good use to display the new 'frosted glass' effect, which can add blurred transparency to visual theming.

Work continues on other features, such as the *Tiling* applet, as well as floating window stacks and Cosmic app and applet icons.

Currently, the alpha release is scheduled for the end of March. Visit <https://blog.system76.com> for the latest news.

OPINION

NEW YEAR RELEASE!



Italo Vignoli is one of the founders of LibreOffice and the Document Foundation.

LibreOffice 24.2 Community, the new major release of the free, volunteer-supported office suite and the first to use the calendar-based numbering scheme (YY.M), is now available at <https://bit.ly/lxf312lo> for Windows (Intel, AMD and ARM), Mac OS (Apple and Intel) and Linux.

But what really sets LibreOffice apart from other office suites is that it is based on a single software engine for all environments: desktop, cloud and mobile. This architecture enables LibreOffice to offer a better user experience and produce identical and perfectly interoperable documents based on the two available ISO standards: the open ODF (ODT, ODS and ODP) for digital sovereignty-conscious users and the proprietary Microsoft OOXML (DOCX, XLSX and PPTX).

The Document Foundation does not provide technical support for users, although they can get it from volunteers on user mailing lists and the Ask LibreOffice website: <https://ask.libreoffice.org>.

LibreOffice users, free software advocates and community members can support The Document Foundation with a donation at www.libreoffice.org/donate. You can read the full update at <https://bit.ly/lxf312note>.



DISPLAY

Chromium Wayland accelerated! VA-API now works on Linux Ozone/Wayland platform.

Traditionally, VA-API (Video Acceleration API) on Linux has not been supported in Chromium for GPU-based hardware acceleration, because it targeted the libva-x11 library.

VaapiWrapper has now been updated to remove the usage of libva-x11 and the legacy VaapiVideoDecodeAccelerator. This means Linux Ozone/Wayland can share the same code path as Linux Ozone/X11. The VA-API X11 library uses the DRM library instead.

Even though this pull request has been merged, code contributors are still pragmatic:

"In so far as Linux + VA-API is not a supported platform (IOW it's a best effort), this will allow developers and users to try and make it work on their machines."

Updates like these tend to divide online opinion, with die-hard X11 adherents repeating the claim that "Wayland breaks everything." KDE developer and board member Nate Graham tackled this claim in an extensive blog post in late 2023, pointing out that Wayland was never intended as a drop-in replacement for X11.

Graham also cites X11's rather horrendous built-in UI toolkit (even when it first emerged) and the fact that apps could request the same resources to break each others' functionality.

When it comes to Linux, he also points out that it's rare for coders to develop an X11 app or Linux app per se, so much as one for Qt or GTK.

Universal cross-desktop interoperable apps may be far off, but you can read more about this merge at <https://t.ly/nvmC5>.

OPINION

DATABASE BASES



David Stokes
is a technology evangelist
at Percona.

When you look at MySQL, PostgreSQL and MongoDB, they each have a vibrant community that reflects its chosen databases. PostgreSQL has a community-driven development process that actively includes its user communities through meetings, mailing lists and online. MySQL is centrally developed by Oracle and its community seems to have diminished as HeatWave has become the main priority. And Mongo's community has peaked after its licence change flattened enthusiasm for the product.

MySQL could learn from the PostgreSQL community, in that more actively seeking community involvement pays off in more customer buy-in and goodwill. PostgreSQL is vibrant and unlikely to need to borrow anything from the other two's communities. MongoDB's community will continue to dwindle after the licence change issues. Will this affect the long-term success of the company? No. But it will reduce the number of people who are willing to commit to the community.

Another lesson for MySQL to learn from PostgreSQL is the consistent increase in performance between versions. Sometimes it is a single-digit improvement but you don't trade performance for new features of dubious utility.

GAME ENGINE

Unity bans VLC plugin

Unity management making friends and influencing people.

Unity Technologies has recently been in the news due to deciding to lay off up to 25% of its workforce. It seems that cutbacks are also taking place at the Unity Asset Store.

The store serves as a place for Unity users to sell project assets. Binaries of the open source VLC for Unity first appeared on the store in December 2019.

These serve as a bridge between the Unity game engine and the VLC multimedia engine, enabling coders to create their own media



Unity Technologies may have banned VLC plugins from its store, game devs can download the binaries from Videolabs.

player based on VLC technology in Unity-based games.

Lead developer Martin Finkel announced that in summer 2023 he received a message from Unity Technologies to say that VLC for Unity and related assets would be banned.

The supposed reason is that some package dependencies are licensed under LGPL, though Finkel and others point out that many assets with LGPL code still exist in the store, such as FFmpeg. For now, VLC for Unity has moved to Videolabs (<https://videolabs.io/store>).

HARDWARE


OpenWrt router?

OpenWrt floats the idea of its own hardware router.

Since 2004, OpenWrt has offered an open source alternative to proprietary router firmware. It can be used to revitalise old hardware or can be installed on popular devices like the Raspberry Pi.

In early January, OpenWrt developer John Crispin posted in the mailing list: "Let's celebrate this [20th] anniversary by launching our own first and fully upstream supported hardware design." The exact specs of OpenWrt OneAP-24.XY are still to be confirmed but it seems the router will be based on the Banana Pi. Pricing is aimed at below \$100.

OpenWrt will likely base the OpenWrt OneAP-24.XY on a Banana Pi board, as its specs closely match the requirements.



KERNEL

Man writes kernel in Rust

Maestro is a Unix-like kernel written from scratch in Rust.

23-year-old French programmer, Luc Lenôtre, modestly claims that computers were a "mystery" to him until recently. Still, in early January, he announced the development of Maestro, an OS written entirely in Rust.

Currently Maestro has a monolithic kernel, supporting only x86. Still, it includes 135 out of 147 Linux system calls and comprises some 48,000 lines of code.

Apart from the kernel, Maestro includes its own custom boot manager, system utility commands and package manager.

Learn more and download a test version of Maestro at <https://blog.lenot.re/a/introduction>.

Maestro includes Solfege: a boot system and daemon manager similar to Systemd.



Distro watch

What's behind the free software sofa?

LINUX MINT 21.3 VIRGINIA

Mint is one of the most popular versions of Linux, partly due to its support for multimedia codecs out of the box, as well as offering more traditional desktop interfaces via Cinamon and Mate. Version 21.3 of the OS was announced on 12th January as an LTS (Long Term Support) release. Mint now includes new features for the *Hypnotix* TV viewer, a new Actions add-on for Cinamon, and experimental Wayland support. Visit www.linuxmint.com for more information.



Linux Mint now has experimental Wayland support.

WATTOS R13

WattOS is based on Debian (in this case Debian 12 Bookworm), and incorporates the lightweight *OpenBox* window manager. This makes it ideal for installing on older machines. The latest version use the low-resource LXDE desktop and has switched to using the *Calamares* installer. It also now incorporates better hardware support out of the box, kernel v6.1 for 64-bit PC installation, and *Gdebi* for easy installation of DEB packages. You can discover more about WattOS at www.planetwatt.com/R13-details/.



WattOS is an ideal distro for installing on ageing or lower-powered hardware.

SOLUS 4.5 RESILIENCE

Solus is an independent version of Linux, built from scratch. Versions are available using a customised iteration of the Budgie desktop environment and Gnome. There's also now an experimental version running Xfce. The latest release brings updated apps and kernels, including *Firefox 121.0*, *LibreOffice 7.6.4.1* and *Thunderbird 115.6.0*. Other improvements include switching to the *Calamares* installer. The OS also now uses PipeWire by default. See <https://getsol.us>.



Solus now comes with an experimental ISO bundling Xfce.

NOBARA PROJECT 39

The Nobara Project first announced stable releases in 2022. It's a modified version of Fedora Linux, designed to be more user-friendly than its parent OS. To this end, it includes extra packages such as *Wine* dependencies and Nvidia drivers. Nobara now incorporates codec installation into system updates. The layout picker has been removed because the Nobara team claims it wants to move away from Gnome. There is still a pop-up for Nvidia drivers. Learn more at <https://nobaraproject.org>.



Nobara is a more user-friendly version of Fedora, bundling extra packages.

ZORIN OS 17

Zorin OS is based on Ubuntu and is specifically designed for people who are new to Linux. It supports running many Windows programs. The latest version has been designed to run faster on a greater range of hardware. Both the kernel and desktop have been optimised – the OS only needs 1.5GB to run now instead of 2GB. The Zorin Menu and search features have also been upgraded. You can find out more at <https://zorin.com>.



Zorin's not a Bond villain but a user-friendly Linux distro.

OPINION

WILY WESTON



Marius Vlad is a senior software engineer at Collabora.

In December, the latest version of Weston, Wayland's reference compositor, was released. Weston 13.0 comes with multiple fixes and important changes, notably the ability to load multiple back-ends simultaneously. This can be used to load VNC, RDP or PipeWire back-ends for remote access alongside the native DRM back-end.

There is also now the ability to have multiple overlapping outputs. This enables you to have a DRM output that is simultaneously streamed remotely via PipeWire, while still taking the most efficient display path. This work is not fully complete and currently requires additional code in Weston's front-end to enable; however, we are looking forward to this support being seamless in the next release. This was made possible by a large-scale rework of our internal scene-graph representation and damage tracking, something that will also enable more optimisation possibilities in the future.

Alongside all this, we have been adding multiple fixes and internal changes, such as pointer constraints/confined changes, refactoring the shells to make use of newer APIs, or Xwayland bug fixes and updates. Thank you to all the contributors who made the release possible!

OPINION

REISER'S
GUILT

Jon Masters is a kernel hacker who's been involved with Linux for over 22 years.

ReiserFS is known for two reasons, one far less savoury than the other. It was the first journaled filesystem to be added to the official Linux kernel.

Unfortunately, more people today are likely to know of ReiserFS because of the actions of its author, who murdered Nina Reiser in 2006. Following the events, the company continued its work for some time, but eventually moved on. ReiserFS has languished for years. It still has users, and, of course, legacy filesystems may be out there for mounting. It was interesting, then, to see a message from Fredrick Brennan, who reached out to Reiser from prison.

The response Brennan received may be worth broader reading, simply for the bizarre (and hopefully unique) circumstances. In the long letter, Reiser not only admits to his guilt (something he initially had denied), and the need for punishment, but he also apologises for the way he engaged with the Linux kernel community, his ReiserFS co-developers, and even for some of the fundamental design choices, which he describes in detail. The letter with its release is bound now to be a strange part of the sociological history of the Linux kernel. A copy is on the kernel mailing list:

<https://lkml.org/lkml/2024/1/18/245>

Kernel Watch

Jon Masters keeps up with all the latest happenings in the Linux kernel, so you don't have to.

Significant ongoing work is happening to support CXL (Compute eXpress Link), which supports more contemporary mechanisms for attaching IO devices to modern systems. Among its various modes, the most popular is .mem (CXL.mem), which effectively allows for serial-attached memory modules much

purposes during a system crash (crashdump). Over time, additional uses for kexec have emerged, including that of live updating machines without a reboot. While we have tools like live kernel patching, this is mostly limited to small(ish) fixes. There are cases where one might want to update – as Alex cites by way of example – a host machine running virtualisation

without tearing down the guest VM instances. His patches don't (yet) go that far, but they do intend to. He starts with a fairly innocuous (and thus uncontroversial) example of passing ftrace buffers on to the new kernel, but he plans to

“Additional uses for kexec have emerged, including that of live updating machines without a reboot.”

like PCIe plug-in cards as an alternative to DDR. While it won't replace the bandwidth and latency characteristics of DDR any time soon, CXL.mem is growing in popularity as an additional tier of memory within very large systems. The ongoing work includes updated patches to support new system calls (mempolicy, mbind and so on) intended to improve bandwidth utilisation for processes using CXL memory.

Alex Graf posted some really nice looking patches titled “kexec: Allow preservation of ftrace buffers”. Kexec is a lightweight mechanism used by Linux to chainload a secondary kernel, originally for debug

continue, citing a talk he co-presented with James Bottomley at the recent Linux Plumbers Conference (now online).

Sunny retirement

Finally this month, continued discussion of the removal of legacy architectures now touches the original 32-bit SPARC (but not the 64-bit architecture SPARC64 that is still in production), where it is intended to rip out support for sun4m and sun4d machines. As usual, this author is now feeling renewed pressure to finally ewaste some very old machines at home that haven't been powered on in the better part of several decades. **LXF**

» ON THE MAINLINE

Linus Torvalds announced the release of Linux 6.7. The new kernel includes Bcachefs, which is a filesystem inspired by the same kinds of features you will find in ZFS and Btrfs (copy-on-write, auto-repair, checksumming, data verification and so on). Other features of Linux 6.7 include an ability to disable 32-bit emulation on x86 at boot (so distros can selectively deprecate 32-bit support over time, but without removing it entirely for compatibility reasons), and a replacement for the legacy futex system call used for fast userspace locking (for example, by glibc).

It also includes many driver updates, among them support for Nvidia Ada (RTX 40 series) GPUs, and a very large update to AMD GPU support.

With the release of Linux 6.7 came the opening of the merge window (period of time during which disruptive changes are allowed) for what will become Linux 6.8 in another couple of months. The merge window itself was actually slightly delayed due to weather in the Pacific North West (PNW) region of the US, which includes Portland, Oregon, where Linus lives. Multi-day power outages tend to shift priorities more toward family matters and away from technical ones.

Among the many patches lined up for 6.8 include additional Rust support. Rust is something that Linus opined on during a keynote at the Linux Foundation's Open Source Summit Japan recently. The video is online in the usual places.

Answers

Got a burning question about open source or the kernel?
Whatever your level, email it to answers@linuxformat.com



Neil Bothwick
solves Linux queries as a form of meditation.

Q Really big numbers

I'm fairly new to Linux and run Linux Mint 21 on my desktop. Every time I try to open an XLSX sheet that I have either been sent or downloaded, I get the error message below:

```
Warning loading document abc.xlsx
The data could not be loaded completely
because the maximum number of
columns per sheet was exceeded
```

Is there a way to overcome this or know what has been lost, if anything?

Ian Madley

A Looking at the sample sheet you sent, which produced no warnings in *LibreOffice* 7.6, there are no sheets with large numbers of columns. The column limit in *LibreOffice* was 1,024 up until version 7.3. From *LibreOffice* 7.4 onwards, the column limit was increased to 16,384, so you are unlikely to hit this.

We have seen this behaviour in the past, but not recently. It appears to be a combination of the way *Excel* saves the information in the files and the way *LibreOffice* reads it, leading *LibreOffice* to think there are more columns than there really are. As long as your XLSX files do not include insane numbers of columns, you are safe to ignore this warning, as

annoying and alarming as it may be. Saving the file from *LibreOffice* should prevent the message appearing again.

Linux Mint 21 currently includes *LibreOffice* 7.3.7.2, which has the lower column limit, but still within the needs of your files. So, it is a matter of waiting for updates to install a later version without this issue. If you find it really annoying, you could uninstall *LibreOffice* from the Mint software manager and install a direct download from **LibreOffice.org**, but that is a course of action that we do not recommend. The software manager makes sure everything on your system works in harmony; installing packages outside of this system can cause problems, especially for a fairly new user.

Q Unlocking the keys

I have one problem, for which I can not find a good solution. On a remote server, I'm forced to create SSH keys that need to have a password entered at the client side before connecting. I'm aware that this protects access to my remote server. But I would like to run nightly *Cron*-automated copies using *Rsync*. Usually I use *Chiefs* to mount a remote server, but also let *Rsync* directly use SSH – it depends. I've been doing this for ages

now, but never with password-protected SSH keys. I tried something like this to use SSHFS:

```
$ SSHPASS="myVerySecretPassword"
sshfs <myRemoteServerAddress>:/mnt/ -o ssh_
command="sshpass -e ssh -l
<myRemoteLoginName>"
```

Or this to let *Rsync* run directly:

```
$ rsync -ratlv --rsh="sshpass -f $(pwd)/
pwdfile ssh "
/<myRemoteServerAddress>/ $(pwd)/
mnt/
```

Plus variants of these commands.

Klaus Kolle

A You are using *Sshpass*, which is intended to send a login password to SSH, as you would use if you were not using authentication keys. What you need is a way to send the passphrase of the SSH key, which can be done with *Ssh-agent*. This handles the transfer of private keys for the public key authentication used by SSH. You already have *Ssh-agent* installed – it is part of the **openssh** package – you just need to set it up. First start *Ssh-agent* in a terminal:

```
$ eval $(ssh-agent -s)
```

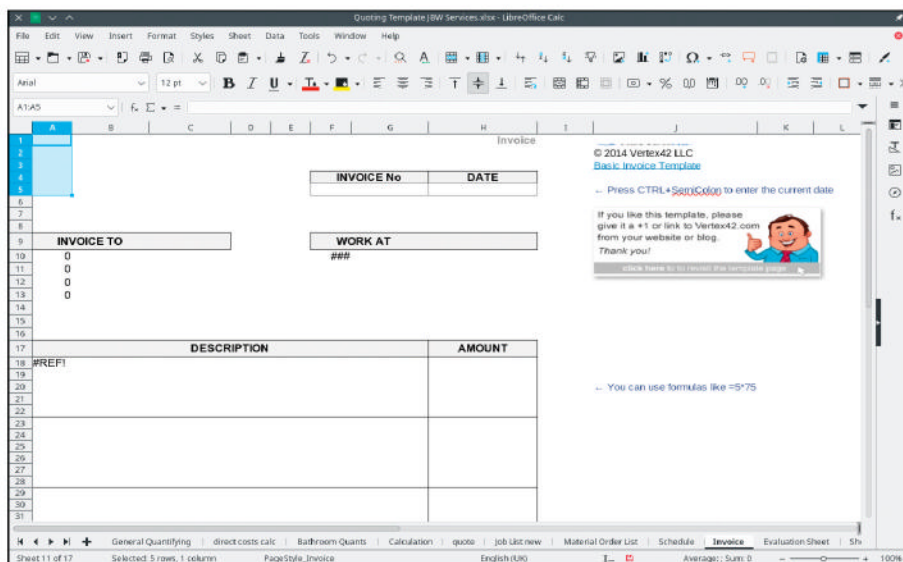
We use *eval* because *Ssh-agent* outputs some environment variables that are needed later. Then use *ssh-add* to add your keys. You can specify the individual keyfiles to add – for example:

```
$ ssh-add ~/.ssh/id_ed25519
```

Note that you are adding your private key. If you run *ssh-add* with no arguments, it adds all the keys in *~/.ssh*. You are asked for the passphrase for each key before it is added. Then you can list the keys in *Ssh-agent* with:

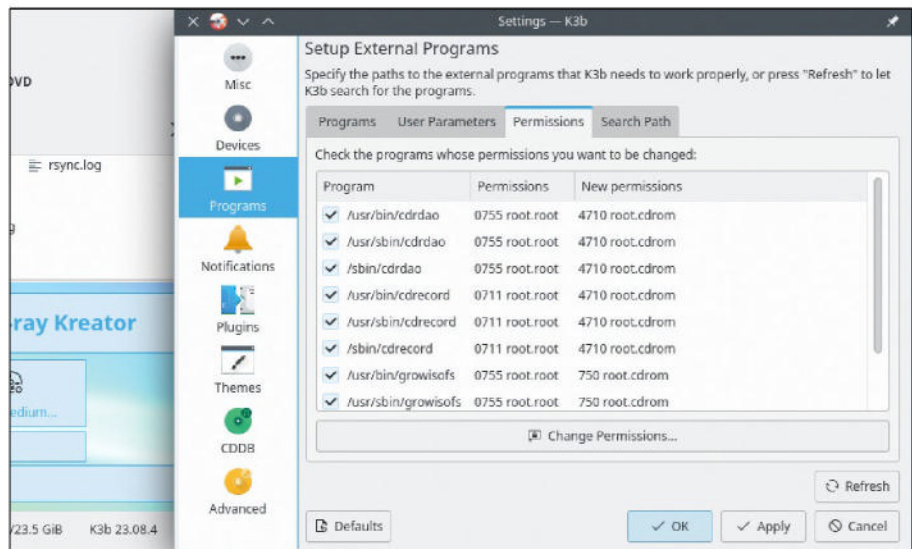
```
$ ssh-add -l
```

Because *Ssh-agent* only keeps the decrypted keys in memory and writes nothing to disk, your keys are secure. However, this means that when it stops running, the unlocked keys are lost and need to be unlocked again. So, the next time you try to use SSH, you are asked for your key, but for the last time in that session. We mentioned that *Ssh-agent* needs some environment variables,



Not an especially large spreadsheet, but some versions of LibreOffice seem to think it is.

so it only works in the environment in which those variables exist. That means that so far it only works in the terminal in which it was run. If you are using a desktop environment, you should look at the best way to start *Ssh-agent* as part of the desktop startup, so the environment variables are then inherited by all processes started from that desktop session – most desktops have a way to run commands on startup. You still need to give the passphrase the first time you use a key, but that then covers all subsequent uses of that key in that desktop session. You could even have a startup script, run after starting *Ssh-agent*, that opens a terminal and runs an SSH command – something bland like `ssh user@host` uptime will do – to a server using your keys, forcing *Ssh-agent* to unlock the key.



If K3b throws up permissions errors when accessing your hardware, you can change things in the settings.

Q Disappearing menus

I installed K3b and tried to burn a music CD, but got a message saying something like, Don't have permissions... Change in the settings menu etc...

So, I crashed around the settings but couldn't find anything about permissions. Anyway, I clicked an option on the top toolbar, which was a toggle switch that said something like View Toolbar. Or it might not have been toolbar; it might have been View Menu. I forget. Well, it turns out this toggle switch controlled the same bar/menu that the option itself was located on. You know, the File/Edit/View/etc toolbar, so clicking it made the toolbar disappear. As the option I clicked that caused this to happen was on that toolbar, I don't know how to get it back.

I tried running `sudo apt-get purge k3b` and then installing it again, but this didn't work and it reinstalled with the toolbar still missing. I haven't been able to find the folder where this program and its config file is actually located on my hard drive to see if I can edit something in the config file itself. I also haven't been able to find that pesky permissions setting either.

Connor Dixon

A

Uninstalling a package does not reset its user settings, even with the purge option; only system settings in `/etc` are affected. You can delete the user settings manually, but make sure the app is not running as it may write the settings back when it closes. Most programs store their settings in `~/.config`, so look in there for likely contenders. If in doubt, there is a trick for finding files changed by a program. First, open a terminal and run:

```
$ touch programstarted
```

Then run your program, K3b in this case, change a setting, and exit. Now go back to the terminal and run:

```
$ find ~ -type f -newer programstarted
```

This outputs a list of all files created or modified in your **home** directory since you ran `touch`, which should give a clue as to where the program is storing its files.

However, you should not need to nuke your configuration to restore the menu bar, as you can press `Ctrl+M` to toggle it. You would be able to see this on the menu entry if it were there. The permissions issue is different; you can change permissions in the Settings window, go to Programs and select the Permissions tab. Here you can change the

permissions for the programs that K3b uses behind the scenes. You have to give the administrator password, as changing permissions of system files requires root access. However, if the permissions relate to writing to your CD/DVD burner, it may just be a matter of changing the groups that your user belongs to – this also has the advantage that it works for all CD/DVD writing programs. First see who owns the device. Open a terminal and run this command and look at the output:

```
$ ls -lH /dev/cdrom
brw-rw----+ 1 root cdrom 11,0 Dec 10 19:05 /dev/cdrom
```

Here we see that the device has read and write permissions for the root user and all members of the `cdrom` group. You can see which groups you are a member of with either the `id` or `groups` command, run in the terminal. If you are not a member of `cdrom`, add yourself to it with:

```
$ sudo gpasswd --add $USER cdrom
```

You need to log out of the desktop and back in again to apply this change. If you need to change the permission settings within K3b, go to Settings > Configure K3b > Programs and select the Permissions tab. While you do not

» A QUICK REFERENCE TO... X FORWARDING

Running terminal commands on a different computer is easy – just SSH to the other computer and it is like you were there. You can also do the same with full desktops with VNC or RDP, but isn't that overkill for just one program?

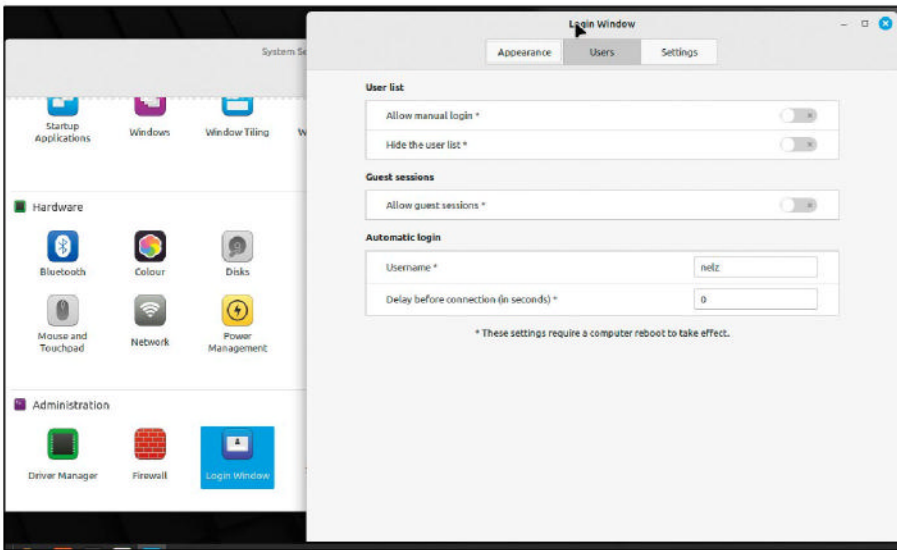
The X graphical interface is a client/server system, so it is possible to run a program on one computer and open its window on another's desktop using

the SSH connection you've set up. You need this facility enabled on the remote computer. In `/etc/ssh/sshd_config`, ensure that `X11Forwarding` is set to `Yes`. Restart `Sshd` if you changed this setting, then on the local computer connect over SSH as usual, but add the `-Y` option:

```
$ ssh -Y user@remotecomputer
```

Now run a GUI program from the command line and it opens its window

on the local computer, but the program is running on the remote system. Depending on your SSH configuration, you may need to use `-X` instead of `-Y`, but try `-Y` first as it is faster. If you are connecting over the internet, adding `-C` to compress the traffic may speed things up, although *TightVNC* or *NX* may be faster over a slow connection as they are optimised for this and X is not.



Linux Mint's Cinnamon and MATE desktops, in common with most others, let you have a user logged in on startup.

need to wipe your configuration to fix these issues, because you have been clicking around the settings while trying to fix things, it may be safer to start again.

It worked, then it didn't

You answered my question about screen brightness control on an Apple iMac in LX305. I tried the suggestion of adding `cpi_backlight=vendor` to the `GRUB_CMDLINE_LINUX_DEFAULT` in `/etc/default/grub` and running `update-grub`, and we have successfully been able to control the backlight – until now.

I discovered the folder `/sys/class/backlight/` is now empty. It happened after a Ubuntu update (which is a worry!). Hence, the question is now: how do I restore the files to the `/sys/class/backlight/` folder?

Do I use `apt install <driver name> ?` If so, I don't know what to put here – `lshw -c video` returns `AMD/ATI Whistler Radeon HD 6730M/6770M/7690M XT`. I tried `apt install <various combos of these words>`, but none work.

I also downloaded `amd-catalyst-15.9-linux-installer-15.201.1151-x86.x86_64.zip`. It unzipped to `AMD-Catalyst-15.9-Linux-installer-15.201.1151-x86.x86_64.run`. In the terminal, I typed: `$ sh AMD-Catalyst-15.9-Linux-installer-15.201.1151-x86.x86_64.run`

This started the process, but then returned an error: `Detected X Server version 'X Server _64a' is not supported.`

Chris

A The files in `/sys` (and `/proc`) are not real files; they are created by the kernel to give access to features of your hardware. There are a couple of possibilities here, indicated by the fact that this changed after an upgrade. It

would be useful to know what was upgraded – the `Apt` log at `/var/log/apt/history.log` should give a clue. It is likely this was caused by a kernel update – the kernel contains the drivers for your hardware. The first step is to see whether the brightness control files are still in `/sys` but in a different location, so try running:

```
$ find /sys -name '*brightness'
```

It may be that your previous problem, which we fixed by changing the `acpi_backlight` kernel option, has been fixed in the later kernel drivers, so this option is now undoing that fix. Try the other two options for this – `acpi_backlight=video` and `acpi_backlight=native` – to see whether either of those helps.

Finally, install the `brightnessctl` package from your package manager. This has many options for working with screen and LED brightness settings, but the following simple command shows what hardware is available for control:

```
$ brightnessctl --list
```

Trying to install software from outside of the package manager is risking problems, and the *AMD Catalyst* software is available as an Ubuntu package.

Repeated logins

I installed Linux Mint Cinnamon 21.2 yesterday and all is good so far. During the installation, I chose to encrypt the disk, and that requires me to log in at each boot time. I find it now annoying to log in so often and I wonder if there is a way to remove the login screen. If that requires removing the encryption, it's fine.

Patrick Harper

A You have encrypted your `home` directory, so a password is needed to unlock it. Otherwise, there would be little point in encrypting it. The way the

Mint installer has set this up is to use a feature of PAM (Linux's authentication system) to link logging into your account with unlocking your `home` directory, provided both use the same password. It is possible to have automatic logins to the desktop, but you'd still need to enter your password to unlock the encrypted data.

If your computer is in a secure location, such as at home, you can reduce the number of password requests by not rebooting. Suspend instead of powering down when not using the computer – this works with desktop computers as well as laptops. Depending on how the desktop environment is set up, this may not require a password when waking up from suspend. If it does, you can usually turn it off in the screensaver or power section of the configuration manager.

Naturally, you should not suspend like this if you are taking the laptop anywhere that anyone else could get hold of it – that would defeat the purpose. Sadly, if you wish to do without the encryption, you need to reinstall. Back up the contents of your `home` directory to an external drive, then reinstall, making sure the encryption box is not ticked. When setting the username and password towards the end of the installation, there is an option to have them logged in automatically. Select this and Mint boots directly to your user's desktop. You can also turn this option on (and off) on a running system by opening the System Settings, going to Administration/Login Window, and opening the Users tab. Here you can select which user is logged in automatically. **LXF**

GET HELP NOW!

We'd love to try to answer any questions you send to answers@linuxformat.com, no matter what the level. We've all been stuck before, so don't be shy. However, we're only human (although Neil's weird fingers are a bit suspect), so it's important that you include as much information as you can. If something works on one distro but not another, tell us. If you get an error message, please tell us the exact message and precisely what you did to invoke it.

If you have, or suspect, a hardware problem, let us know about the hardware. Consider installing `hardinfo` or `lshw`. These programs list the hardware on your machine, so send us their output. If you're unwilling, or unable, to install these, run the following commands in a root terminal and send us the `system.txt` file, too:

```
uname -a > system.txt
lspci >> system.txt
lspci -vv >> system.txt
```

Subscriptions: for magazine issues email help@mylinuxmagazine.co.uk

Mailserver

WRITE TO US

Do you have a burning Linux-related issue that you want to discuss? Write to us at *Linux Format*, Future Publishing, Quay House, The Ambury, Bath, BA1 1UA or email letters@linuxformat.com.

Post Office

Could you get someone to write an article explaining what the problem was/is with the *Horizon* program used by the UK Post Office? The general difficulty must be with programs working in real time for large numbers of clients. I can think of at least five reasons why your answer may be 'No'.

Dicky Clymo

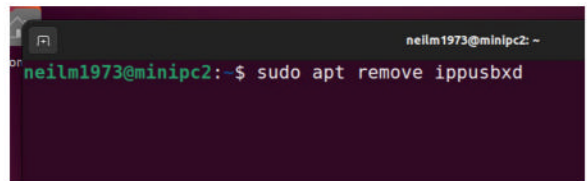
Neil says...

Is it really Linux or open source related? To be honest, I'd like to keep on the positive side of life even if it is – there are possibly Linux servers involved and open source systems, but I'm not sure we're here to analyse every IT system failure. Having said that, I did read an amazing write-up of the UK Air Traffic Control failure in 2023 (<https://jameshaydon.github.io/nats-fail>). It's fascinating to discover the huge flaws just waiting to happen in large IT systems.

I also hate to say that people have been aware of the Fujitsu issues for a long while (back to the original trial in 2018) and it's a tragedy it's taken so many years, with lost and destroyed lives, before the government finally did something about it. It's bound to happen again, too, and I'm sure the CEO involved will walk away with a golden parachute. If you want a fantastic resource, try the Post Office Trial site that followed the original trials in 2018: <https://bit.ly/lxf312post>.

Terminal learning

With reference to Ian Sheppard's letter in LXF310, where he does not think that he can learn to use the



Imagine fixing things by writing out a few words!

terminal, old age is not the limiting factor. I am in my 93rd year, but was very interested in home computers from the time that they first became available in the 1980s. It is the level of one's passion for the subject that determines the extent of one's ability to master the use of the terminal, and also to write a program.

I was happy with Windows XP, but not with its replacement, so after that I tried Linux, mainly Ubuntu. I'm now using 22.04.3 LTS. I use the terminal to install and remove programs – it's the easiest way to do it.

In recent years, the systems have got more and more complicated, sometimes resulting in problems – browsers, for example. I got frustrated with *Firefox* and *Chromium* as they wouldn't hold my settings. However, *Microsoft Edge* worked perfectly! Recently, *Linux Format* referred to the *Brave* browser. It also works, so out with *Edge* and in with *Brave*.

Another problem: a new HP inkjet printer. After a new installation of Ubuntu 22.04.3, the printer works perfectly for five or six days, then it won't work at all. A complicated renewal of the driver sometimes restores action, sometimes not. After a fresh installation of Ubuntu, it works OK for a while. I have just bought a second-hand Brother inkjet printer and it works all the time. Well, anything is better than Windows 11!

Bryan Mitchell

Neil says...


I can't blame anyone for not wanting to use the terminal, and on the whole, you don't have to. It's just that for troubleshooting and plenty of maintenance jobs, it's way more efficient. In some ways, though, it is easier than a GUI – you might have to navigate a whole list of GUI prompts, while you could copy and paste a single terminal command. But as I said, it's not a good argument, as Windows often provides Powershell fixes, which, last time I checked, is a terminal.

As for your printer trouble, apparently *ippusbxd* can cause a conflict. You can remove it with this handy

The flight plan that crashed the entire UK Air Traffic Control system; the software is, of course, proprietary.



Helpdex





Even with good software, interference can ruin your listening pleasure.

terminal line: `sudo apt remove ippusbxd`. Then reinstall your HP software.

Blue everyday

My laptop uses the Realtek RTL8723DE wireless chip. I experience audio drops on every distro I've tried with my Bluetooth headphones; this isn't the case with Windows or Android. From what I know, kernel driver btusb is responsible, right? Android uses the Linux kernel, too, so why doesn't it have the same issue?

I know Realtek and Nvidia are infamous for being poor on Linux – is there any way I can improve the situation? Audio drops get worse with more network use, so while browsing YouTube, and it ruins the whole laptop music experience.

Tyron Stillman

Neil says...

This is the usual balancing act with multiple systems and software stacks needing to all work together for there not to be issues. Then there's radio interference, too. Android devices are more tightly integrated, as it's a single Android build for a single hardware spec, plus Android uses its own Bluetooth stack, as early devices did suffer issues. The Raspberry Pi used to suffer stutters on its own RPiOS until it was updated with the Blueman front-end and Bluez stack with PulseAudio.

You don't mention the distros you've tried. Any recent mainstream distros should come with sane default software selections and configuration similar to the PiOS mentioned. So, the main culprit would be the

Realtek kernel module. We did stumble over a third-party build that's supposed to alleviate issues with newer kernels. We can't attest to its quality but it's worth a punt: <https://github.com/smlinux/rtl8723de>.

Distros everywhere

What ever happened to *Distrowatch*? I was looking at a Google Trends list of Linux distros and it was interesting to see Ubuntu was by far the most popular, Debian was second and Kali was in third place. CentOS and Arch came next, and Linux Mint was sixth. That all seems somewhat contrary to the normal *Distrowatch* list. Why not bring it back and set us straight?

Jason Bitton

Neil says...

I'm totally open to the idea of bringing back *Distrowatch*. I think we've stopped offering decent distro coverage – outside of Nate's reviews – though I'm trying to bring regular looks at new-style distros each issue. The problem for me is there's next to no way of providing a sensible ranking. **Distrowatch.com** is a fantastic resource but the ranking is problematic and sounds like it's open to manipulation, which is why you end up with oddities like MX Linux being number one for years on end. A Google Trends result is dictated by what people are actually searching for, which is better but doesn't necessarily relate to usage, downloads, installs or some other metric. **LXF**

In years gone by, we had an entire *Distrowatch* section. Should we bring it back?

Distrowatch

This month's roundup of news from the vaults of Linux distro development.

Continuing the lizard blizzard

OpenSUSE 11.4 heralds the future, partially by becoming the first major distro to come with LibreOffice.

Fresh Mint

Linux Mint 10 KDE. This release offers KDE 4.6 and Linux 2.6.35.

Do it yourself

Both Linux. Both is creating a lot of buzz for such a young distro.

On the Radar

Recently, several distros have been spotted in the wild. Here's a look at some of the most interesting ones.

Hit list

Top distros and packages for the month of February.

Category	Rank	Package
Distros	1	Ubuntu
	2	Debian
	3	Kali
	4	CentOS
	5	Arch
Packages	1	Firefox
	2	Thunderbird
	3	LibreOffice
	4	Python
	5	MySQL



shane_collinge@yahoo.com


SUBSCRIBE Save money today!

SUBSCRIBE TO LINUX FORMAT MAGAZINE

Subscribe today and enjoy great savings!

www.magazinesdirect.com/LIN/C75J

Call 0330 333 1113 and quote C75J



Save money today! **SUBSCRIBE**

» **PLUS:** Exclusive access¹ to the *Linux Format* subs area!

1,000s of DRM-free PDF back issues and articles! Get **instant access** back to issue 66 (May 2005) with tutorials, interviews, features and reviews. At linuxformat.com

1. Only available to www.magazinesdirect.com subscribers



» CHOOSE YOUR PACKAGE!

Our simple six-month best offer breaks up the annual subscription into two easily manageable payments over the course of the year.

PRINT EDITION



Only
£21.09
every 6 months

13 issues of *Linux Format* in print over 12 months, and save 50%!

DIGITAL EDITION



Only
£22.71
every 6 months

13 issues of digital *Linux Format* over 12 months, and save 50%!

OUTSIDE THE UK?
Turn to page 63 for more great subscriber deals!

Terms and conditions: Offer closes 31/03/2024. Offer open to new UK subscribers only. Pricing is guaranteed for the first 12 months and we will notify you in advance of any price changes. Please allow up to six weeks for delivery of your first subscription issue (up to eight weeks overseas). The full subscription rate is for 12 months (13 issues) and includes postage and packaging. If the magazine ordered changes frequency per annum, we will honour the number of issues paid for, not the term of the subscription. For full terms and conditions, visit www.magazinesdirect.com/terms. For enquiries please telephone +44(0)330 333 1113 (open Monday-Friday 8:30am-7pm, and Saturday 10am-3pm UK time), or email help@magazinesdirect.com. Calls to 0330 numbers will be charged at no more than a national landline call, and may be included in your phone provider's call bundle.

EXPLORE THE BEST LINUX & OPEN SOURCE SOFTWARE

From the best distros to system security, utilities and fun applications,
we've rounded up everything you could ever need



Ordering is easy. Go online at:

magazinesdirect.com

Or get it from selected supermarkets & newsagents

Intel Core i7 14700K

Jacob Ridley thinks this is a perfectly good stand-in for a Core i9.

SPECS

Socket:
V LGA1700
Process:
Intel 7 10nm
Cores (P+E):
8+12
Threads: 28
P-core: 3.4GHz
(5.6GHz turbo)
P-cache:
640KB L1,
16MB L2, 33MB
L3 (shared)
E-core: 2.5GHz
(4.3GHz turbo)
E-cache: 1.1MB
L1, 12MB L2,
33MB L3
(shared)
Unlocked: Yes
GPU:
Intel UHD 770
GPU clock:
300MHz
(1.6GHz max)
Units: 32
Display: 4, eDP
1.4b, DP 1.4a,
HDMI 2.1
Mem max:
192GB, ECC
support,
2-channel
Mem speed:
DDR5
5,600MT/s,
DDR4
3,200MT/s
PCIe: v5 and v4
20-lanes
Power:
125W PBP,
253W MTP

The Intel Core i7 14700K brings something genuinely new to the table. It offers four more E-cores than the Core i7 13700K, and it's a touch faster, for a few pounds more. With tangible benefits and near-enough price parity, it's arguably the only 14th-gen processor worth seeking out instead of the 13th-gen for your next PC build.

The 14700K comes with eight Performance-cores (P-cores) and 12 Efficient-cores (E-cores), bundled together for around £410. That's roughly on the mark for Intel's recommended customer pricing but, more importantly, the same as a Core i7 13700K, effectively killing the last-gen chip cold. While the architecture remains the same between the two – both utilise the Raptor Lake hybrid architecture – the 14700K comes with four more E-cores and 100-200MHz higher clocks on both E-cores and P-cores under boost.

The slight gain in clock speed will only amount to a minor boost in performance, as we've seen to be the case with the rest of the 14th-gen K-series processors (see reviews in **LXF310** and **LXF311**). However, those extra E-cores do amount to a tangible benefit in applications that prefer a little more horsepower, and you are clearly better off for having the extra silicon under the hood.

P versus E

It's the faster, low-latency P-cores that take the brunt of the high-end workload, and they're most important when chasing performance. Yet those extra E-cores can help out with gaming – in some instances, directly via the game engine or by offloading background threads or other tasks to free up the P-cores. There's a feature included with the 14700K/KF and 14900K/KF called Intel Application Optimization (added to kernel 5.16) that has been shown to improve utilisation of E-cores, too. However, it's limited in scope and not the killer feature you'd like it to be.

Nevertheless, the 14700K comes within touching distance of the pricier Core i9 14900K in the majority of the gaming benchmarks we've run. In *Total War: Three Kingdoms* and *Shadow of the Tomb Raider*, only a single frame separates the two. Considering that the step up in price to the Core i9 14900K is an extra £150, that's a lot for two more frames per second.

And that's a lot of money for four more E-cores (remember, these E-cores don't have Hyperthreading and only count for one thread) and some slightly higher clocks. It's maybe an odd way to think about it, but you could argue Intel's happy to hand you that same sort of bump free of charge with the 14700K. Rather, if you're planning to use your PC for more heavy-duty processing applications outside of gaming, the 14700K is still a suitable stand-in for the 14900K.

The cheaper chip is seriously impressive in tests of its multithreading muster, including scoring major points across *Cinebench R23*, *Blender* and *X264*



■ If £400 is burning a hole in your pocket, it's the processor to go for.

benchmarks. You do need to care about multithreaded performance to go for this Core i7, too, as the Core i5 14600K performs incredibly similarly in gaming.

The 14700K frequently trades blows with the Ryzen 7 7800X3D in gaming benchmarks. The benefit of AMD's chip is its high efficiency and low package power draw, neither of which are the 14th gen's strong suit. You can also pick one up for a little less than a 14700K, but otherwise there's a good case to be made for the 14700K for the multithreaded performance.

The 14700K really is not the coolest nor most efficient chip. With peak wattage of 253W and an average of over 200W, it's pretty much a match for the 14900K it so closely resembles. It also runs up to 96°C under load, which is toasty despite our test bench's chunky all-in-one 360mm cooler. The 7800X3D's 80°C max temperature feels icy by comparison.

It makes for one tough decision: do you favour multithreaded performance or efficient operation? That can come down to a few key factors, but with the two often giving as good as they get in gaming, the lower power consumption of the AMD chip does begin to sway us to the red team's offering. **LF**

VERDICT

DEVELOPER: Intel
WEB: www.intel.com
LICENCE: £410

FEATURES	9/10	EASE OF USE	9/10
PERFORMANCE	8/10	VALUE	7/10

With high frame rates in games and strong multithreaded performance, there's a strong case for picking this CPU.

» **Rating 8/10**

RlxOS 2023.11

Nate Drake finds himself washed up on the rock of ages that is RlxOS, only to discover that it still needs work.

IN BRIEF

We loved RlxOS's painless install process and simple Initial Setup, though when we peeked beneath the hood, we found some bugs, which may make it unsuitable for daily use.

SPECS

CPU: 1GHz
Mem: 700MB
 (2GB best)
HDD: 16GB
 (50GB best)
Buids: x86_64,
 AMD64

The latest release of RloxOS is code-named Silaghana, a Sanskrit word that, loosely translated, means 'solid as a rock'. This is not an obscure reference to Ashford & Simpson's 1984 song but an allusion to the nature of RloxOS, which uses an immutable filesystem. It's built from scratch and follows a rolling release model.

This means the operating system is designed to be unchangeable and read-only. Once installed, system files and directories can't be modified. This offers a number of advantages, especially for security, given that system files can't be changed during runtime. There is also a less stringent unlocked version of the OS. You can find further information about this on the project website.

The RlxOS project is a relative newcomer to the immutable party, given that its first stable release was in April 2023. Since then the project has shifted from Gnome to the Xfce 4 desktop to make it "significantly faster and lighter".

This is borne out by the 2GB ISO we downloaded to test RlxOS in a virtual machine. Although you can test the distro in a live environment, the OS is designed to be installed to a drive. Setup launched immediately and RlxOS was installed in a matter of minutes.

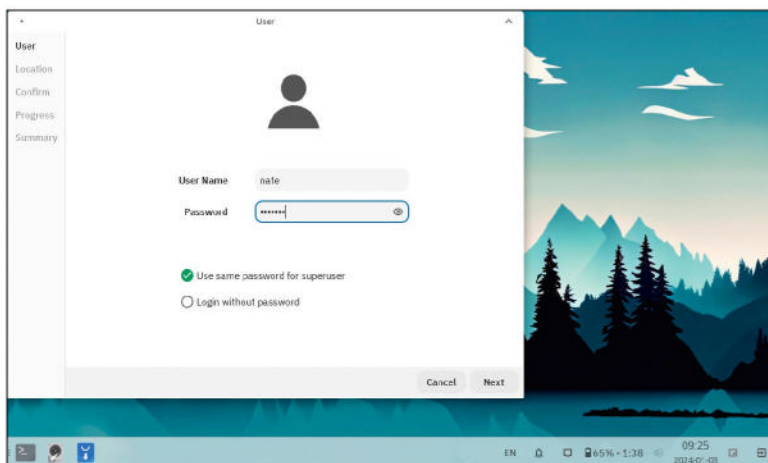
On first launch, the OS takes users to RlxOS's new Initial Setup, where you can set your time zone, define a username and password for this session, as well as decide whether to assign the same password to root.

Although this worked like a charm, when the screen locked after a few minutes of inactivity, we couldn't use this password to sign back in. This happened four times, forcing us to restart in each instance.

One puzzle of immutable filesystems is how to update the OS if system files can't be modified. Here it is handled by Swupd, a daemon that allows "temporary safe mutability" to facilitate updates via libostree.

Unfortunately, the official documentation for Swupd was inaccessible when we tried to access it on RlxOS's website, but from online research, we discovered that the Swupd daemon can store two versions of the OS: one with and one without updates. Users can also review update changelogs before applying them. We decided to check for updates using the command suggested in the release notes (`sudo swupd status`), only to be told the command wasn't found.

On the plus side, RlxOS does support *Distrobox* containers, allowing you to run tools from other distros



RlxOS's new Initial Setup enables you to choose your locale, as well as set your username and standard/root password.

via the terminal. The documentation let us down here, too, as running the listed commands resulted in a *Docker* permissions error. We were also surprised at the choice to use *Docker* over *Podman* with *Distrobox*, given that *Docker* runs containers as root, providing unrestricted access to the filesystem.

The OS also has native Flatpak support. We duly used the integrated *Firefox* browser to head over to **Flathub.org** and download the latest Flatpak for *Mozilla Thunderbird*. We next opened the terminal to install *Thunderbird* via `flatpak install` in moments.

Overall, the number of small bugs we encountered means it's hard to recommend RlxOS for daily use.

Immutability offers many advantages, such as security and more painless updates. This is why major developers are jumping on the immutable bandwagon: Fedora's Onyx, which uses the Budgie desktop, also offers an immutable filesystem, and Canonical has a Snap-based immutable version of Ubuntu coming later this year. The developers behind RlxOS have some way to go to contend with major players like these. **LXF**

VERDICT

DEVELOPER: Manjeet Saini

WEB: <https://rlxos.dev>

LICENCE: GPL-3.0

FEATURES 5/10

PERFORMANCE 7/10

EASE OF USE 8/10

DOCUMENTATION 4/10

RlxOS feels like a WIP. Although setup is simple, some features are buggy and the documentation needs work.

» **Rating 6/10**

Emmabuntüs 5

Nate Drake finds that despite being hard to pronounce, the name is the only complicated thing about this Debian-based distro.

IN BRIEF

Emmabuntüs is certainly a perfect distro for Linux beginners, offering a choice of desktop environments and easily customisable dock. There's a large number of bundled apps, justifying the 4GB ISO.

SPECS

CPU: 2GHz
Mem: 1GB
HDD: 40GB
Builds: i686, x86_64

Originally released in 2011, Emmabuntüs was part of an initiative to supply used computers to humanitarian organisations. The first of these was the charity Emmaus (see [LXF287](#)) and as the OS was then based on Ubuntu, the portmanteau Emmabuntüs was born.

Since then, the project has shifted to basing its OS on Debian's stable branch (in this case, v12.4 Bookworm), but the main site says the distro is loyal to its roots in assisting humanitarian causes, reducing waste and promoting Linux to beginners. To this end, the team has also now produced a dedicated reconditioning manual for those interested in refurbishing computers.

One perk of this OS is that it can be entirely installed without an internet connection, as all packages are contained on the disk image. The downside of this approach is a hefty ISO, the 32-bit and 64-bit versions of which weigh in at around 4GB. On the plus side, like its parent OS Debian, the project website claims the distros run with just 1GB of RAM.

On first login, you are offered the choice of the default Xfce desktop or to switch to LXQt. You're also prompted to choose your preferred keyboard layout.

Special mention should go to Emmabuntüs's Desktop Configuration, which enables you to activate *Cairo-Dock*, as well as other features such as the taskbar, workspace switcher and dark theme.

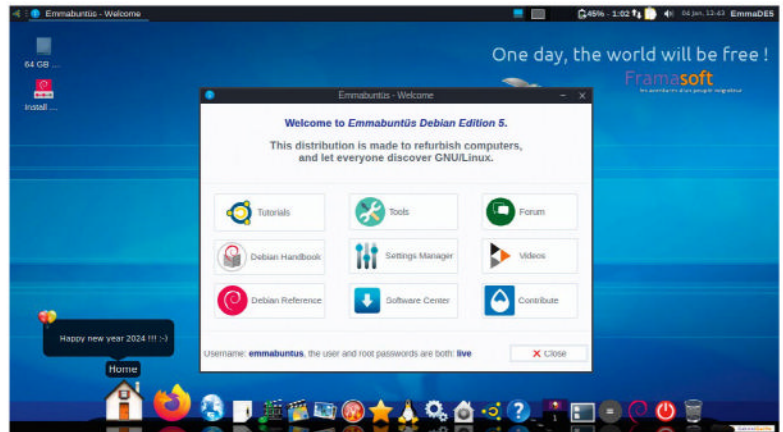
The dock comes in three different configurations: All (for those wanting to access all applications), Simple (for newbies) and Basic (for kids and beginners).

If this weren't user-friendly enough, the OS displays a handy welcome guide with links to tutorials, settings manager, software centre and dedicated forum.

Browsing through the preinstalled apps offers some justification for the hefty disk image. Web browsing is handled by *Firefox 115.6.0 ESR*. We were impressed to find *uBlock Origin* installed. The default search engine is Lilo, which, like Emmabuntüs, originated in France. Lilo donates part of its revenue to charitable causes.

Email is also now managed by *Thunderbird 115.6.0*. There are numerous productivity apps, too, including *LibreOffice*, *Gnumeric* and *AbiWord*. This theme of multiple apps for performing similar tasks is repeated in the media apps, which include *Clementine*, *Kaffeine* and *VLC Media Player*.

The most recent version of Emmabuntüs (Debian Edition 5.1.01) has switched to bundling *Zulucrypt* rather than *Veracrypt*. This apparently is because



The operating system offers a variety of docks for both experienced and new users, as well as a comprehensive welcome guide.

Veracrypt can no longer create *Truecrypt* volumes, though the advisability of using *Zulucrypt* to do so is debatable, given *Truecrypt* is no longer in development. Having slated the OS's security profile, we were pleased to see its bundled utilities include the free and open source password manager *KeePassXC*.

If you're not fond of the kitsch desktop wallpaper, Emmabuntüs's Desktop Settings offers 38 more.

When firing up the 64-bit version of Emmabuntüs in live mode, we decided to test the project's claim that the OS can run using 1GB of RAM. Via the task manager, we found the desktop at rest consumed 1.3GB, though in fairness, we had chosen the Xfce environment rather than the more lightweight LXQt.

There's little doubt that Emmabuntüs offers a very comprehensive desktop experience. The huge disk image contains every package you're likely to need, not to mention a selection of multiple programs that perform similar tasks. The customisable dock, tutorials and built-in welcome guides also make Emmabuntüs an excellent choice for Linux beginners, though they may find the sheer number of apps overwhelming. **LXF**

VERDICT

DEVELOPER: Collectif Emmabuntüs

WEB: <https://emmabuntus.org>

LICENCE: GPL

FEATURES	9/10	EASE OF USE	9/10
PERFORMANCE	7/10	DOCUMENTATION	8/10

Ideal for Linux beginners, offering a comprehensive welcome guide and documentation, as well as a customisable dock.

» **Rating 8/10**

Manjaro 23.1 Vulcan

Nate Drake boldly goes to take Vulcan for a test spin and finds an incredibly user-friendly and efficient operating system.

IN BRIEF

Manjaro is Arch done right. Not only does it have a simple installer but it can also automatically detect and configure hardware. And adding packages is simplicity itself via the package manager.

SPECS

CPU: 1GHz (dual-core)
Mem: 4GB
HDD: 30GB
Builds: x86_64, ARM64, i686

Manjaro began life as a labour of love, announced on Arch Linux BBS – the OS on which it's based. To quote its website, "Manjaro is a versatile, free and open source Linux operating system designed with a strong focus on safeguarding user privacy."

The website also claims the distro can be installed on a wide range of devices. This much is true, given that it's traditionally been a popular choice for Linux laptop vendors. Pine64's Pinebook, for instance, comes with Manjaro preinstalled. Manjaro also partnered up with Tuxedo Computers to preinstall the OS on its devices, although this seems to have given way to the Ubuntu-based Tuxedo OS.

This is one of Manjaro's main advantages over stock Arch, in that it has an extremely user-friendly graphical installer. Like Arch, it operates on a rolling release model.

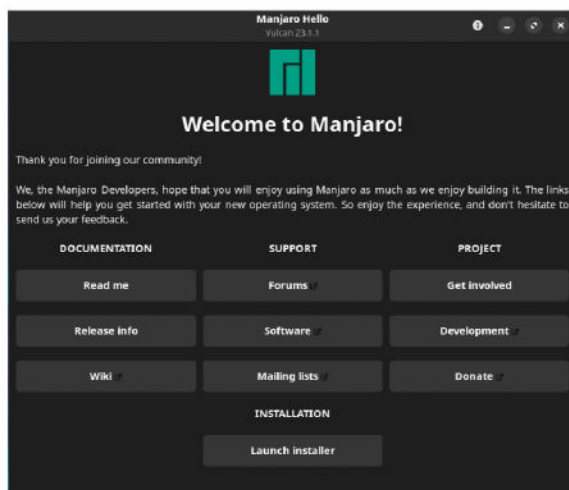
The original version of Manjaro used the Xfce desktop, but there's also an official flavour that makes use of the KDE desktop (currently Plasma 5.27 LTS) and a flavour incorporating Gnome, which is the main focus of this review. There are also community-maintained editions of Manjaro that incorporate other desktops, such as Budgie, Cinnamon and Mate.

Officially, support for 32-bit architectures has been dropped, though there's also a semi-official community project, Arch Linux32, that does support i686. Manjaro can also be installed on ARM devices, although at the time of writing, the Raspberry Pi 5 isn't fully supported.

As Manjaro is developed independently from Arch, it has its own independent repositories (although it can access the community-maintained AUR).

The packages in the stable branch are the default repositories used by Manjaro systems to provide updates and downloads to the general user base.

The unstable branch is synced several times a day from stable Arch and the packages from the Arch repo. Packages are vetted by the Arch Linux QA and the Arch



On first launch, Manjaro's welcome guide links you to the online support pages, user forums and release notes.

Linux community. This is where the latest versions of apps are stored, along with Manjaro's in-house software.

Branching out

The testing branch is snapped from unstable at irregular intervals. It functions as a testing area for packages built by Manjaro developers. This testing process means that packages will never be as bleeding-edge as Arch itself. Software usually appears in Manjaro's stable repositories weeks after Arch, but users can easily switch to the unstable or testing branches.

As with Arch, you can install packages via the CLI using *Pacman*, but you can enjoy a far more elegant experience with Manjaro's own *Pamac*. This offers an intuitive UI for searching for and installing applications, enabling you to search for packages via groups, categories or even repositories. Click the arrow next to

» HOW SECURE IS MANJARO?

Sharp-eyed readers will notice that Manjaro's website claims the OS has a strong focus on user privacy. During our research online, we found that the OS used to have a reputation for not being as secure as other Arch-based alternatives, but it seems this is now undeserved. For instance, we found that the installer does offer to encrypt partitions with LUKS during setup, although

admittedly this feature is commonplace these days.

On delving further, we found Manjaro has received specific criticism in the past for the length of time it takes packages to graduate from testing to stable branches, given they could contain unpatched security bugs.

In 2015, the team forgot to update its SSL certificate, so the main website and forum were inaccessible for a short

time. Manjaro didn't help the situation by suggesting site visitors change their system date via the command line to fix the issue.

This provoked a backlash on Reddit, with many users suggesting switching to Arch variant Antergos (then still in active development) or using install scripts to make Arch itself more user-friendly.

Manjaro doesn't look likely to change its testing model

any time soon, though in recent years there haven't been any noticeable scandals regarding unpatched packages or ancient SSL certificates.

The Manjaro wiki does provide a dedicated page on Linux security (https://wiki.manjaro.org/index.php/Linux_Security), but it doesn't list any features specific to the OS to suggest it's more secure than others.

an app description to install it, or the red trash icon to remove it. If your chosen package has optional dependencies, you can also select these from a helpful pop-up, as we found when installing *VLC Media Player*.

If this weren't user-friendly enough, the Updates tab lists all available package upgrades, enabling you to select/deselect individual entries, then click Apply. The latest version of *Pamac* also includes support for offline updates, so you can download system updates, then install them later even if not connected to the internet. The database now loads on first launch, too, improving search performance.

Pamac is also an excellent way to view Manjaro's bundled apps. Our ISO of the Gnome edition was a hefty 4.3GB, so we had high expectations that the OS would be suitable for day-to-day use out of the box.

Web browsing is now handled by *Firefox 121.0*. When opening Gnome Activities, we couldn't find anything beyond a document viewer in the Office category. Still, you can choose from either *FreeOffice* or *LibreOffice* during the install process. Manjaro also comes with the *Thunderbird* mail client (v115.6 Supernova) and our Gnome flavour came with the *Lollypop* music player.

Given the minimal number of third-party apps, we were left scratching our heads over how the ISO came to over 4GB, though this could be explained in part by the other Manjaro-specific tools incorporated into the OS.

These include *Manjaro Hardware Detection*. This usually operates undetected during the install process and serves to automatically detect and configure hardware. It includes support for hybrid graphics cards, as well as module dependencies for *VirtualBox* virtual machine installations. This is likely to be a godsend for less technically proficient users.

Manjaro also has support for multiple kernels via *Manjaro Hardware Detection Kernel*. Not only does this enable less-experienced users to juggle various kernels, but it also manages automatic updates.

One of the easiest way to add and remove kernels plus necessary modules is via Manjaro's own *Settings Manager*, simply by clicking Install. *Settings Manager* is also useful for more mundane tasks you need to do with your OS, such as managing user accounts and changing the system language.

The only category noticeably absent was wallpapers, though we found it easy enough to access this from the



Gnome 45 desktop by just right-clicking and choosing Change Background. With 22 wallpapers to choose from, you aren't short on choice.

The most recent release of Manjaro (code name Vulcan) was announced on 16th December 2023 via the Manjaro forums. The team made use of this post to discuss all that is great about Gnome 45, including a dynamic workspace indicator and better search features.

The announcement similarly praises the inclusion of Plasma 5.27 LTS in the KDE edition, with a new tiling system and better support for Wayland. The Xfce edition also receives recognition for all that's new in version 4.18, including file manager *Thunar*'s ability to highlight files as well as configure file backgrounds. This latest version of Manjaro uses version 6.6 LTS of the Linux kernel.


Naturally, it's difficult to credit an announcement like this to Manjaro specifically because all the listed changes apply to any version of Linux using the latest Gnome, KDE and Xfce desktop environments.

This is where Manjaro's forums shine, as we explored users' reactions to Vulcan (with more than a few *Star Trek* references). The forums also have a dedicated support section, categorised by software and hardware type, such as Sound, for example.

During our tests in a virtual machine, we used *System Monitor* to determine the OS was consuming 1.8GB of RAM with only the desktop running. The CPU also ran at around 10% of capacity.

Overall, Manjaro lives up to its claim of being extremely easy to use, with specialist tools to ease users into Linux installation. Admittedly, there isn't a huge number of preinstalled apps, but the integrated software manager makes this issue easy enough to fix. **LXF**

The live CD doesn't include an office suite but you can choose either LibreOffice or FreeOffice during installation.



Manjaro's own Settings Manager enables you to install and manage multiple Linux kernels, manage user accounts and set the locale.

VERDICT

DEVELOPER: Manjaro

WEB: <https://manjaro.org>

LICENCE: GPL

FEATURES 8/10
PERFORMANCE 9/10

EASE OF USE 10/10
DOCUMENTATION 9/10

Once you've downloaded the large ISO, Manjaro is quick and easy to set up with a variety of desktops and packages.

» **Rating 9/10**

Roundup

RoboForm » Keeper » NordPass
» KeePass 2 » Bitwarden



Michael Reed is very security-conscious. He only occasionally uses an OS that isn't Linux and has two locks on his bike.

Password managers

Security-conscious **Michael Reed** examines five pieces of software that can safely store and retrieve your passwords and other sensitive data.


HOW WE TESTED...

We installed everything on to Linux Mint, and all our notes about installation are based on that. However, part of the idea of a password manager is that your logins should be easily available across multiple platforms. We installed the latest, stable release of each tool, and this meant we typically had to perform a manual installation.

In each case, we imported a sizable password collection that we had exported from the *Google Chrome* web browser, giving us real-world credentials to work with.

All of the password managers come with web browser plugins, offering autocomplete functions. We tested the plugins to make sure that they worked, but we didn't notice any standout features because they all operated in the same way. For this reason, we decided not to give plugins their own separate section.

Finally, we tested each Android app to make sure they could work with that set of credentials for browsing, retrieval and editing.



Password managers can store a variety of credentials and private information, and they do so in a way that isn't tied to a particular web browser, operating system or device. Ideally, the way they work integrates with your online life for hiccup-free retrieval and storage of passwords from whatever platform you're using at the time. As well as having your passwords available to you wherever you are (including mobile), you can step up your security level by using more complex passwords, which are harder to hack, without any extra effort on your part.

Some of the systems we're looking at offer paid options, but all of them offer a usable free service. *RoboForm*, *NordPass* and *Keeper* belong in the freemium category as they are commercial software with a feature-limited free version. *KeePass 2* is a completely free, open source program that you run entirely on your own computer like any other application, but it's a bit more complicated to make its database available on other devices. *Bitwarden* sits in the middle, because it's open source software backed by a large company that happily offers commercial options for extra features.

CREDIT: Getty Images/Jon Feingersh Photography Inc

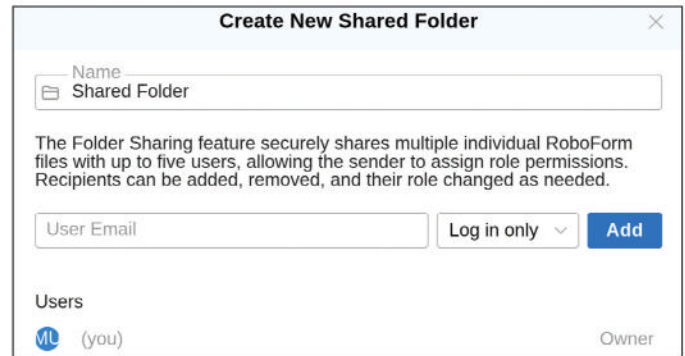
Premium (paid) plans

What do you get if you choose to pay extra for a premium package?

We'll get on to what you get for free later on, but if you want to pay for premium features, the trick is to visit each website, examine the prices for different plans, and decide which features you need or can live without. With all of the systems, it's worth checking out the latest prices before making a decision because there are sometimes special offers, which can be significant for a long-term subscription. To give an example, currently the cheapest *NordPass* plan (called Personal) is around £32 for two years, thanks to a 50% off offer.

Bitwarden is quite a lot cheaper on its single-user plan than the others. Obviously, prices could increase, but at the time of writing, it's \$10 (about £7.80) a year, which is excellent value. It also adds the ability to attach files, handy for things like scanned documents. However, its business plans are similar in pricing to the other password managers. This means *NordPass* is cheaper on its cheapest business plan right now, thanks to a sale offer.

In the case of *NordPass*, *RoboForm* and *Keeper*, the lowest-priced plan removes the limitation of access from more than one device. In the case of *NordPass*, it adds a few things such as features aimed at testing the quality of your password choices and the ability to store files as well as credentials. The *NordPass* Family plan is good value because it's just under twice the cost of single-user but offers six user accounts. *Bitwarden*'s family plan is similar but slightly cheaper, though.



I RoboForm's shared folders allow the sharing of folders of credentials, but it's a feature of the paid plans. You have to decide if you require features like this.

We wouldn't call the *Keeper* single-user plan expensive at £2.50 a month, but at £6.00 per month for the family plan, it's beginning to rack up the costs. Put in the context of overall business costs, none of the *Keeper* business plans are that expensive, starting at £22 per user per year. With services of this sort, be careful when evaluating the different business plans, because they may list the price as per user rather than for the total cost of the plan.

KeePass 2 doesn't offer any paid-for premium subscriptions at all. You get everything for free, but some high-end features and many online features aren't available.

VERDICT

ROBOFORM	7/10	KEEPASS 2	N/A
KEEPER	7/10	BITWARDEN	9/10
NORDPASS	7/10		

The *Bitwarden* personal plan is amazing value. Beyond that, shop around the sites for feature lists, prices and special offers.

Free plans

All the services are usable without shelling out, with few compromises.

B*itwarden*'s free package is the best of the bunch as it offers most of the features of the Personal plan. There is a limit of only one account, but there are no limitations on the number of supported devices. The main feature we miss in the free tier is the ability to attach files. As *Bitwarden* is open source, you can host it on your own server for free, but this doesn't save money as you must pay to unlock the extra features.

We aren't blown away with the free plan limitations of *NordPass*, *Keeper* or *RoboForm* as the feature they all take away is the ability to run the software on more than one computer at once. This can put you in a number of awkward situations, just when you're likely to need a password manager. *Keeper* is even meaner as the device has to be a mobile device. Frankly, we're reluctant to get tied down to the free packages of any of these three (beyond evaluating them), particularly as there are a couple of decent free offerings.

What you see is what you get in the case of *KeePass 2*, as it's a completely free and open source application that you run on your own computer. As there is no paid-for version, the free version is fully featured. The downside to this arrangement is that there are

Monitor for compromised passwords	?	✓	✓
2FA access to RoboForm	?	✓	✓
Use RoboForm as a TOTP authenticator		✓	✓
Cloud backup		✓	✓
Access on all devices			✓
Web access			✓
Local only mode	?		✓
Emergency Access	?	Receive	Grant and Receive

I We recommend visiting the website of each password manager and comparing the different packages to see if you can live with the compromises of the free plan.

no premium features to be bought even if you're willing to pay for them. We love open source, but having a large business offering premium features has its place, too.

VERDICT

ROBOFORM	6/10	KEEPASS 2	8/10
KEEPER	5/10	BITWARDEN	9/10
NORDPASS	6/10		

On balance, *Bitwarden*'s strong free offering, backed up by a commercial company, wins out. *KeePass 2* is completely free in its full version.

User interface and experience

The day-to-day interface.

Utilities such as password managers require a smooth user experience because it's irritating to have to fiddle around with a poor interface every time you need a password. If we're talking about the password features alone, there are two main jobs that a user would want to carry out: retrieving the passwords and applying them.

A password manager is a specialised type of database. We expect the information to be easy to retrieve and edit in a hurry, but we're happy to see extra features, if they don't get in the way of the basic functions.


In most cases, the desktop application is a containerised version of the web interface. Even if you normally use a password manager via a web browser login, we recommend keeping the desktop application installed and ready, because most of those offer at least some functionality without a working internet connection.

RoboForm

6/10

There is no standalone *RoboForm* app for the Linux desktop, but there is a web application. This uses a two-pane layout that makes good use of the screen space and seems optimised for easy access on a maximised web browser. The sidebar gives access to folders and safenotes, small text notes that can be added to your vault. The fact that the icons for a given website are retrieved in the majority of cases eases navigation and adds to the overall attractiveness of the interface. Switching over to the details view crams more information on to the screen.

It seems as though ease of use and fast access to day-to-day features have taken precedence over an interface as detailed as the one that the *Keeper* app presents us with. Overall, it's a pleasant interface to deal with on a day-to-day basis, but as it is a web app, it can't operate without an internet connection.




Keeper

9/10

The *Keeper* desktop application has the look and feel of a containerised web application, and it's almost exactly the same layout as the website. Having said that, it does add a pull-down menu bar, and both interfaces are attractive, detailed and responsive. Scrolling, selecting and searching are all snappy. The sidebar can collapse to an iconified state, saving space. We always prefer the option of right-click context menus in this sort of environment, which this app has.

Along the top of the password list, there is a highly extensive set of filters, and the list itself is switchable between icon, list and folder views, rather like a file browser.

As soon as you select an item, the information panel slides into place. This contains the expected cut and paste buttons, and it's easily switchable into edit mode. It's not the simplest interface, but it's highly functional while also being fast and nice looking.



Installation and platform support

We're interested in the Linux support as well as the multiplatform options.

Multiplatform support is an important consideration. All of these systems support Windows, Mac OS and Linux. They also have mobile apps for Android and iOS, and browser plugins for at least *Firefox* and *Chrome*.

KeePass 2 is written in Mono, so the same binary can run on Windows, Linux and other Unix-derived OSes such as Mac OS. The versions in the Ubuntu and Mint repos were out of date, so we installed it ourselves. This involves jumping through extra hoops compared to running it on Windows, such as installing Mono. There's a variety of options, as *KeePass* also has some derivative forks, such as the more native-looking *KeePassXC*.

The *Bitwarden* website begins to automatically download the AppImage archive when you visit the download page, and this worked (with the manual addition of some dependencies) when we tried it. Installing via Snap or Flatpak are other ways of getting the latest version of the desktop application.

Keeper offers separate downloads for six of the major distributions, but they actually resolve to one DEB and one

RPM archive. The DEB archive installed without problems on Linux Mint.

The *NordPass* website detected that we were using *Firefox* and redirected us to a page where we were prompted to install the browser extension. This is an acceptable way of using *NordPass*, but there is also a desktop app, which we installed using Snap.

RoboForm has desktop applications that can operate without internet access for Mac OS and Windows, but only browser plugins (*Firefox* and *Chrome*) are available for Linux.

VERDICT

ROBOFORM	6/10	KEEPASS 2	8/10
KEEPER	7/10	BITWARDEN	8/10
NORDPASS	7/10		


KeePass 2 can run on some unusual, niche platforms. RoboForm lacks a native Linux desktop app but is usable on Linux.

NordPass

7/10

The user interface of the *NordPass* app is similar to the web application. It does feature some system integration, with a status area icon on our desktop and right-click menus. The default font size is small, but the interface is zoomable. The desktop app gives read-only access to your data without an internet connection, which is better than nothing.

The interface itself features a left-hand sidebar that contains the filters for different types of stored information. There is a folder pane in this bar, if you care to create the folders yourself. The main area of the window is taken up with a scrollable list of passwords. Click on one of these and the main area changes to a summary of the most important information, with icons to reveal passwords or to copy them to the clipboard or to open the edit page. This is a desktop app that certainly covers all the basics well.



KeePass 2


8/10

KeePass 2 is a .NET app that runs under the Mono framework when on Linux. This means it's a pure desktop app, but it doesn't quite have a native look. Speed is excellent, and it adheres to conventions such as pull-down menus, right-click context menus and pop-up dialogs.

Search isn't as-you-type, as offered by some of the other apps, but it's perfectly fast once Return has been pressed.

Appearance options are extensive, and most aspects of the layout and presentation can be tweaked. When it comes to font size, *KeePass 2* seems to pick up on our system settings much more readily than the other applications, and larger font choices are available.

Some prefer the look of a modern web application for tasks such as this, but *KeePass 2* has the closest look to a standard desktop application, with the speed and the customisation facilities we'd expect.




Bitwarden

8/10

The interface of the *Bitwarden* desktop application looks like a halfway house between a desktop app and a website. It has pull-down menus and pop-up context menus, which we like to see in this type of app. It's a three-pane interface. At first the font size seemed a bit small, but we're happy to say that it's zoomable.

These things are a bit subjective, but we feel that it is faster than the similar application supplied by *NordPass*, for example, even though it is a bit slow to load. Switching between different sets of information is more or less instantaneous, and scrolling down the list of stored accounts is smooth. Using the search bar that runs along the top of the main window is fast as well.

The summary and edit dialogs are separate from each other for fast access to information combined with safety from accidental edits. A plain-looking interface that covers the day-to-day basics.



Security

Password security is vital, and breaches have been known to happen.

Bitwarden and KeePass 2 are open source. By our way of looking at things, this gives them an advantage as their internal workings can be inspected by the online community. For example, the *Bitwarden* source is regularly inspected by an independent security company with publicly published reports. Optionally, *Bitwarden* can be self-hosted, enabling you to take full control of how your data is protected.

Bitwarden, *NordPass*, *RoboForm* and *Keeper* implement a zero-knowledge architecture, meaning that the company doesn't store any details about what you have placed in your vault. They don't even store your master password. In theory, this means that if the company's server were to be breached by hackers, the information obtained would be unusable. The downside is that if you lose your password and the recovery code they send you, there is no way of retrieving your account – it's just not technically possible.

Sure enough, in the case of *RoboForm*, we lost the master password while testing it and all we were able to do was wipe

the database and start again. That said, if you do lose your *RoboForm* master password, you can get back in by installing the mobile phone app and using the biometric (fingerprint) scan to retrieve it, but you must have set this up before you lose the password, of course.

How you secure the *KeePass 2* database is entirely in your own hands, as it is self-hosted on your own machine rather than being located on a central server. The downside is that you have to do the work to keep things safe, and a professional company might have more expertise and better facilities in this regard.

VERDICT

ROBOFORM	8/10	KEEPPASS 2	7/10
KEEPER	8/10	BITWARDEN	9/10
NORDPASS	8/10		

We're fans of Bitwarden's open source model when it comes to security, but the other options are solid, too.

Import and export of credentials

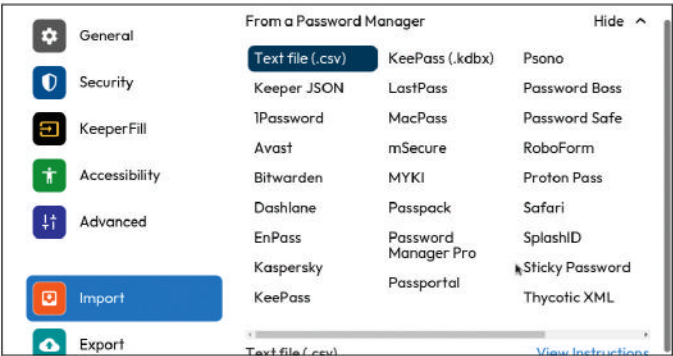
You probably have lots of login credentials that you'd like to import.

All of the password managers can import CSV (comma-separated values) files. However, that's sometimes not enough because the exact order of fields, for instance, for a CSV file can vary, so we appreciate specific CSV import filters for common web browsers and applications. All of the apps can, if nothing else, export the password list as a CSV file, which we consider the minimum acceptable functionality.

Keeper can import about 30 types of password file, but it doesn't list *Google Chrome*. In actual fact, it can import a CSV file from *Google Chrome* using the generic CSV filter. It can output your password collection as CSV, JSON and PDF files.

KeePass 2 can accept a wide variety of files in a mixture of generic and application-specific formats of file types, such as JSON, HTML and XML. The end result is a comprehensive system that can probably handle any backed-up password list you care to throw at it. The export system can support XML stylesheets for a programmable output.

Bitwarden can output in CSV, JSON and encrypted JSON. It also has a long list of import filters to rival that of *KeePass 2*.



KeePass 2 can import an extensive set of password file formats. If you don't find the software you are using listed, it might be worth trying the generic CSV import.

NordPass has a long list of supported application formats it can import, but most of these options link to instructions on how to export in CSV format from those other pieces of software. In all fairness, this should be a robust, foolproof way of working.

RoboForm can import from most of the major web browsers and password managers.

VERDICT

ROBOFORM	7/10	KEEPASS 2	8/10
KEEPER	7/10	BITWARDEN	8/10
NORDPASS	7/10		

KeePass 2 can import the greatest number of application-specific file formats, and Bitwarden is also impressive.

Mobile apps


Sign in on the move – a handy resource to have in your pocket.

Due to the open source nature of *KeePass 2*, there is more than one mobile app that can be used to interact with your password database. They all have to contend with the fact that *KeePass 2* isn't based around an online service, which means you must do some of the database synchronisation work yourself. *KeePassDroid* only supports Google Drive for storage, but this does work well. The *KeePass 2* apps don't quite match the polished experience of the commercial services, but it's possible to get a workable solution up and running.

We're not putting down the *KeePass 2* clients, but moving on to *Bitwarden*, we noticed an upgrade in the overall presentation. We like the way that each entry has an icon for the website. Within the settings, there are options for autofill, if you prefer to have *Bitwarden* carry this out rather than Android's browser. There is also a little built-in utility to generate passwords. These complex passwords might seem unwieldy, but *Bitwarden* would be handling their storage and use, of course.

The *NordPass* app has a clean, attractive layout. The built-in utilities include a password strength assessor, a data breach scanner and a password generator.

We like *Keeper's* Enable Self Destruct option, which erases all locally stored data after five failed entry attempts. We didn't appreciate being hit over the head with a pop-up reminder about the paid-for plan as soon as we logged in. It seemed to be talking down its own product (in its free version) by forcing us to press



The NordPass app has an attractive overall look, with individual icons for each login. The three-dot icon brings up a panel with most of the options you're likely to need.

the Stay Unprotected option to dismiss the pop-up. Overall, it's a basic, usable tool with a rather plain look.

As is the case with the web application, we like the aesthetics of the *RoboForm* app because it makes use of large icons representing each of the websites, and it does this without imposing a performance penalty that we can notice. If you don't like that way of working, there is also a plain-looking list format option. All of the apps feature fingerprint login, but *RoboForm* also has a PIN number option.

VERDICT

ROBOFORM	8/10	KEEPASS 2	6/10
KEEPER	7/10	BITWARDEN	8/10
NORDPASS	7/10		

Other than KeePass 2, all of the apps work in a similar way and provide basic features for when you're out and about.

The verdict

Password managers

We've chosen *Bitwarden* as the winning password manager. It stands out because it's open source software and because the free and Personal plans are both excellent. Sometimes, we like things to be completely free, but freemium has its place, and even if you stick to the free plan, you still get the benefits of a popular service with an online server. The lowest priced package is inexpensive and adds, among other things, encrypted cloud storage of files such as scanned documents. The Linux desktop app does a good job, as does the mobile app.

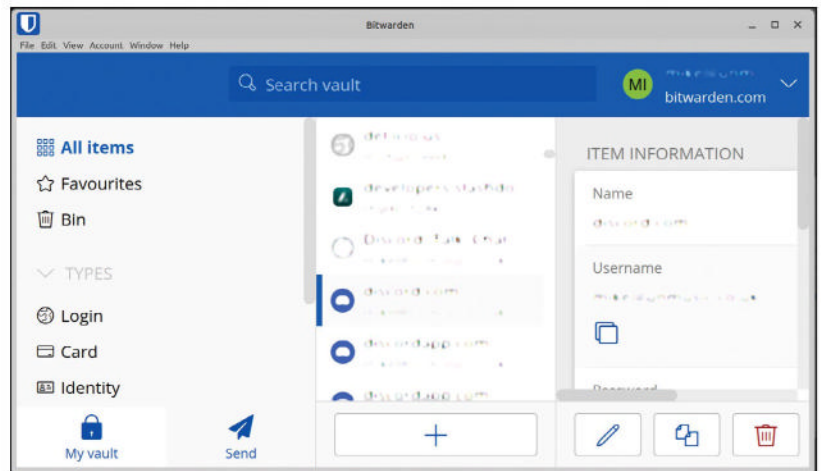
Some will like the independence offered by *KeePass 2*, because it's a standalone application that runs on most platforms. Despite being a traditional application, there are also mobile phone apps and web browser plugins available. The downside is that you may need to set up some of the facilities (such as hosting the database in the cloud) yourself, which means extra work. On the plus side, *KeePass* is a popular program and there are dozens of forks, rewrites, plugins and other add-ons on the website, giving you even more control over how you want to use it.

With the other options, the free or paid package that you choose controls the features that are available, so we're evaluating them on the merits of their lowest priced, entry-level plans and the free plan.

NordPass is one of the services that limits use to a single device at a time on its free plan, even though you can have other devices registered but logged out. It has a decent Linux desktop app, and on the paid plan, the facilities are excellent, with an ability to attach files for secure storage.

Keeper's free plan is so limited that we'd be reluctant to get tied down to it, and the single-user paid plan is a little more expensive than some of the competition. However, the system itself is very good, and we like the desktop app.

It's a shame that *RoboForm* doesn't have a Linux-native desktop app, even though the web app is perfectly usable and it has a mobile app. It lacks a facility to store files such as scanned documents, even on its paid plans. The free plan can evaluate the quality of your passwords.



1st **Bitwarden** **9/10**

Web: <https://bitwarden.com> **Licence:** AGPL-3.0 only (server); GPL-3.0 only (clients) **Version:** 23.12.0

Best of both worlds – open source with commercial hosting. Amazing value.

2nd **KeePass 2** **8/10**

Web: <https://keepass.info>

Licence: GPL-2.0 or later **Version:** 2.55

Open source desktop app. Good for technical people and highly flexible.

3rd **NordPass** **7/10**

Web: <https://nordpass.com>

Licence: Proprietary commercial **Version:** 5.10.20

Usable free plan, with good application. Extensive facilities on the paid plan.

4th **Keeper** **7/10**

Web: www.keepersecurity.com

Licence: Proprietary commercial **Version:** 16.10.10

Great app and good features. Bit pricey, with an underwhelming free plan.

5th **RoboForm** **7/10**

Web: www.roboform.com

Licence: Proprietary commercial **Version:** 9.5.6.0

User interface is good for non-experts. Lacks facilities compared to others.

» ALSO CONSIDER

1 Password has a Linux desktop app, but no free version. It has a travel mode that temporarily removes password vault data in case you're stopped and have your devices inspected.

2 LastPass has a Linux client and a free tier. The snag is that it has been subject to more than one high-profile security breach, trouncing the company's reputation. It remains to be seen whether the company can re-establish the confidence of the security community. It's also increased its prices.

3 Enpass has a Linux desktop client and uses third-party cloud storage (such as Google Drive or Dropbox).

4 Dashlane has a free version; however, it lacks a Linux desktop app.

5 PassWarden is another one that could be used on Linux, but it lacks a Linux desktop app. It has an interesting duress mode that hides some of your credentials if you are forced to hand over your password file. **LXF**

THE 5 BEST NEW DISTROS

Blast off into the future with Matt Holder as he explores the technology powering a new generation of Linux distributions.

Computing never sits still, as new hardware features arrive and software is updated to take advantage of them, so new flaws, bugs and security issues appear. OS architects devise new ways to counter filesystem errors, memory leaks and privacy issues, with Linux always leading the way.

We'll dissect the problems that next-gen distros are solving, before introducing the concepts that solve the issues and give examples of next-gen distributions you can try today.

The main players in the Linux distribution space are all committing huge amounts of developer time to creating the most reliable experience possible. Using the techniques that we've had available for at least a decade, progress has been excellent, but after all this time, advances have arguably slowed down as the experience has become better, so where are the new features being developed?

Commercial options – see Android, Mac OS and Chrome OS – use some of the techniques we'll be discussing, which have been designed to improve OS reliability. To take one example, Android uses a set of read-only partitions, so the user – who we hear causes a lot of issues – doesn't have root access to the system. This is just one way new distro architectures are improving user experiences for the better.



CREDIT: Magictorch

Designing a better distro

Unix originated in the '60s and there are better ways to work now...

Let's look at some of the changes that have become prevalent in recent years. The first is the development of init systems, which are integral to a running Linux system as they control the running of all processes. Early in the boot process, the kernel starts the init process and it continues to run as a daemon. Various options exist, all the way from the venerable SysVinit, developed for Unix decades ago. Canonical spearheaded the use of *Upstart*, before *Systemd* became the popular choice for most distros.


In the late '90s/early 2000s, the most commonly used filesystem was extended2 (ext2), soon succeeded by ext3 and ext4, which are journaled filesystems, which greatly increases the reliability when things like power cuts occur. The ZFS filesystem arrived in 2001 and is renowned for its broad feature set and how reliably it stores data. It uses checksumming, which can be used to reconstruct any data that becomes corrupt for any reason. ZFS is probably overkill for most home users, but Ubuntu makes it simple to install and use. Btrfs has been reviled for issues with some RAID functionality, but is reliable for less enterprise purposes and has extra features that we will discuss later.

Built for reliability

A cleverly designed partitioning scheme can be used to replicate what is employed by Chrome OS to provide a reliable method of software updates. Chrome uses an A-B partitioning system, and updates are applied to the partition that isn't currently used. On the next reboot, the other partition is booted from and if this fails for any reason, the previous partition is booted from again.

Immutability can be achieved in multiple ways, but one technique is to provide immutable root partitions. This means that a set of files is distributed by the project as a rock solid base for the OS. Immutability also refers to the fact that parts of the filesystem are configured as read-only, with the idea being that only the OS developers can initiate updates to files by releasing new upgrades. Using this technique, the OS's footprint can be quite small, with extra applications being added using various other techniques.

When installing packages on your favourite distro, you are trusting the maintainers to ensure that your distro stays safe. While we are not suggesting that the main distro can't be trusted, because you are giving root access to your system, every time you install or update an app, any issues in the repository could cause havoc. Adding to this concern, systems such as PPAs, which can be created by anybody, are also given root access, which is a lot of trust to give somebody you don't know. Modern packaging systems mitigate these concerns by providing sandboxes and containers to minimise access to the wider system. As is often the case in the open source world, multiple solutions exist, such as Flatpak, AppImage and Snap.



Another solution to the installation of software is a tool such as *Toolbox* or *Distrobox*. These provide tight integration with the user's session and containers that the tools are managing, so an Ubuntu container can be run, for example, on an Arch installation, and both GUI and command-line apps can be installed.

Where complete separation is required between services, it is still acceptable to run a virtual machine on a host OS. While containers are newer tech and require less overhead, there is also less separation between containers than between the host OS and a VM.

Declarative configuration has been around for years, with systems such as *Ansible* and *SaltStack* providing a way to create a set of config files that are applied to machines from a central server, allowing configuration to be changed and software to be added in a highly controlled manner, from as few as one system up to thousands. Imagine being able to configure your distro the same way. We'll discuss a distro that does just that!

BlendOS makes it incredibly simple to use the Waydroid Android emulator, by providing a simple way of installing it.

» A BETTER INIT

The *Systemd* project was originally started by a developer called Lennart Poettering. Lennart worked for Red Hat for a number of years before making the move to Microsoft. *Systemd* wasn't the first large project that Lennart worked on, with other examples being the Avahi DNS broadcasting system and PulseAudio.

The SysVinit system was designed decades ago, closer to the beginnings of Unix, and Lennart and others identified many things that could be improved, such as parallel starting of processes and a way of starting processes based on the availability of networking.

In the early days of the *Systemd* init system, the internet was awash with people complaining that the Unix philosophy of having a simple tool for a simple task was not being followed. We are not going to speak further about this philosophy or state an opinion either way, other than to say that *Systemd* has become adopted by the vast majority of distros and covers a huge number of tasks.

Using *Systemd*, the following and many more tasks can be accomplished: running jobs as the root user and logged-in user, running commands in a constrained environment with *systemd-nspawn* (think a newer and more powerful version of *chroot*), and logging is now carried out with *Systemd Journal*, rather than the older *Rsyslog*. Other tasks done by *Systemd* are the configuration of network interfaces and providing home directories that can be encrypted and moved between devices with *systemd-homed*.



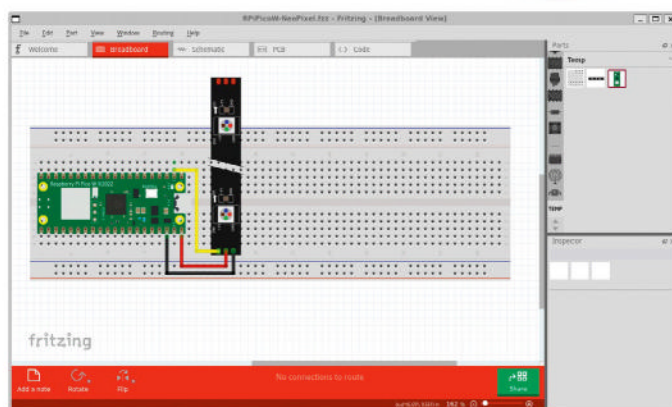
Containers and VMs

Containers will be a stalwart within next-gen distros.



Virtualisation is incredible technology that allows for a host server to run multiple operating systems, by segregating them into virtual machines. The name is very accurate as the host OS emulates all the hardware that the virtual machine requires. Over the years, there have been lots of developments, including PCI pass-through, which allows for graphics cards and other hardware to be used by the virtual machines, and CPU extensions, so the host CPU can better provide resources to the VM and the host OS therefore performs as little emulation work as possible.

Containers then became the next big thing. These are different in that they run on the host OS directly, without the need to emulate hardware. This makes the use of containers more efficient on the host system and more containers can run, compared to using VMs. When first used, containers almost appear to be magic, as configuration and data can be stored on a NAS or on the host OS itself, then these folders and files are mapped through to the container. Any network ports or other resources can also be mapped across to the container. Whenever a container needs to be updated, a new version can be downloaded and run, and the configuration and data can be accessed from the host OS in the same way as



Running from within a container, it is possible to install applications from a different operating system than the host one.

for the earlier version. The separation of configuration, data and computer are important as it makes it incredibly simple to back up configuration and data, and recover from hardware failure or move hosting. When running, containers are segregated from each other using kernel features called control groups.

Contain yourself

Numerous systems exist to manage containers, including *Docker*, *Podman* and *Kubernetes*. *Docker* and *Podman* accomplish similar tasks and are suitable for managing containers on a single machine. The *Docker-compose* tool is used to store, in a config file, the information needed to get a service up and running. This file can be as simple as needing one *Docker* container and a mapped configuration file, all the way to services that require a database, web server and many other things as well. At the other end of the scale is the incredibly powerful *Kubernetes*, which started life at Google and is now maintained and developed by the Cloud Native Computing Foundation. It is used to automate the deployment, scaling and management of huge numbers of containers, the scale of which would be needed to run enormous publicly available services.

While this introduction to containers is very interesting, how does it apply to next-generation distributions? Well, Red Hat developed *Toolbox* as a tool to allow containers to be used to set up environments for software developers. This means that any software dependencies and potentially dodgy code can run completely separate from the host OS. The other powerful thing about *Toolbox* is that it allows access to the host OS to store files, access webcams, sound servers and much more. This means that *Toolbox* can be used to run GUI applications from many different operating systems and they display on the host OS as though they were natively installed. *Distrobox* is similar to *Toolbox*, and can be installed by using the instructions in the boxout (left).

» INSTALLING DISTROBOX

Much like *Toolbox*, *Distrobox* is a way of running tightly integrated containers on your immutable OS. This means the container can access the filesystem outside of the sandbox and your hardware as well. *Distrobox* can use either *Podman* or *Docker* to coordinate the operation of the containers, and the former would be used in this instance. Further information about *Distrobox*, including a list of URLs to use to download different distros, can be found at <https://github.com/89luca89/distrobox>. Without using the *-i* flag and a URL, a Fedora container is used. On an Ubuntu host, open a terminal and enter the following commands:

```
$ sudo apt install distrobox
$ distrobox create -i docker.io/library/ubuntu:23.10
$ distrobox enter ubuntu-23-10
```

Following the first command, the packages are downloaded, the filesystem layer created and then on reboot *Distrobox* is installed. Command two runs *Distrobox*, downloads the Ubuntu 22.10 image and installs it. Finally, the third command runs the container with *Podman* and then opens a shell within it. We will now install a GUI application from the container and ensure it opens as expected:

```
$ sudo apt install fritzing
$ fritzing
```

Different steps are needed to install on immutable systems such as Fedora Silverblue – see *Distrobox*'s website for instructions.

Modern packaging

Why packaging systems of old aren't fit for purpose.

Containerised packages are a very important part of this jigsaw puzzle and allow apps to be installed with their dependencies, without the need to add to or alter the files on the root partition. For example, should an application need a particular Python library, it can be bundled with the app itself. A second reason why containerised packages are so powerful is that the same version can run on many operating systems, due to the ability to use runtimes and dependency embedding.

Thirdly, as these apps are containerised, their ability to edit the entire OS's files can be reduced. For example, an app can be allowed access to a user's **home** directory only and given access to the USB ports. Conversely, some apps are denied USB access.

Examples of containerised packaging systems are Flatpaks, Applimages and Snaps.

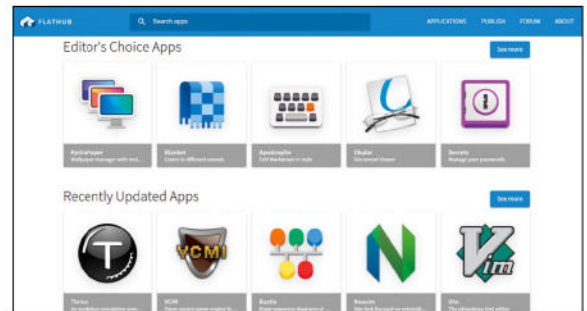


Endless OS

Released: 2014
Apps: Flatpak
PM: xe
FS: ext4, ostree
Desktop: Gnome
Features: Atomic, immutable

Flatpaks

Flatpak is the approved format for apps to be run on Fedora Silverblue and other similar systems, such as OpenSUSE's MicroOS. To run Flatpaks, one or more runtimes are required to be installed and apps are built against a runtime. When installing an app, the required runtime is installed at the same time. If libraries or other resources are required for the app to run, these are bundled into the file. When running an app from a Flatpak, they are contained within a sandbox and can only access their own resources and the runtime. Any further access to the filesystem or hardware must be explicitly granted. Flatpaks have a mechanism, known as Portals, that is used to allow the app to access files, hardware and other items from outside of the sandbox. Runtimes are distribution-agnostic, meaning that once a Flatpak has been built, it can run anywhere that the runtime can run. Flatpaks are distributed from repositories, with Flathub being the most popular. This can be added to Silverblue by using the GUI options.



Flathub is used to distribute Flatpaks. The contents can be viewed from the built-in package manager and the front-end at <https://flathub.org>.

Applimages

Applimage is a single-file format that contains all required libraries and resources for an app. It contains a filesystem that is mounted using the FUSE userspace filesystem mounting system. Applimage started in 2003 and was originally named Klik. Apps packaged as Applimages can be run on multiple Linux distros, do not require installing in the conventional sense, don't need root permissions to run and are completely portable.

Snaps

Snaps were designed by Canonical for the Ubuntu Phone OS, so there was a rival to the way Android and iOS package apps. They were extended to support IoT applications, GUI apps and server-based options. Snaps use features of Systemd and are controlled via the SnapD daemon. They are run in a sandbox with limited access to the parent OS. In recent releases, Canonical has taken the decision to package and distribute *Chrome* and *Firefox* as Snaps, due to the complexity of compiling web browsers for multiple OS versions. Much criticism has been given over this due to the time taken for some apps to open. This has been improved by changing the compression algorithm.

» PACKAGED FOR YOUR SAFETY

When using Fedora Silverblue, opening Gnome's *Software* store from the desktop displays apps that can be installed. All options that can be found from the discovery screen or from searching are Flatpaks. When installed, you have the security of these apps being containerised and they benefit from all the items discussed above. The excellent audio editor *Audacity* is available from the software centre, so search for and install this to test the system. While in the software centre,

look around to see what's available.

When using an Ubuntu distro, you can select between Snaps or the more traditional DEBs. DEBs require root access to your device while Flatpaks don't.

Applimages are generally not quite as well integrated into the OS. Once you have downloaded an Applimage for a piece of software, you need to right-click on it and change

the properties, so that the file is executable. Once executable, it is then able to open.



Filesystems

Keep your files secret, keep your files safe...



Filesystems are incredibly low level and while a choice is generally offered when installing a Linux distro, most of us select the default. On one hand, this is fine, because the options made available by the distribution makers are all suitable for everyday users. What we will discuss here is some of the features available in more modern filesystems.

Extended 4 (ext4) has a long history and is still the standard for a fair number of distros. It is incredibly stable and has journalling, which helps it automatically recover files that haven't been fully written, due to a power failure or similar issues. Other filesystems have more features and some of them can be used from a single-disk system all the way to petabyte arrays.

Now the default on a number of distributions is Btrfs, pronounced Butter-FS or Better-FS. Btrfs has a number of features that are not available with ext4 or XFS. Copy-on-write is a technique used to allow snapshots to be made of the filesystem, which are initially zero in size. When changes are made to the filesystem, the snapshot grows. Btrfs also uses checksumming to allow data to be recovered from scenarios that would

otherwise leave data corrupted. Snapshots can be sent to other Btrfs volumes, so can be received as well. In this way, block-based backups of entire filesystems can be sent over network connections for safe storage. Data restores can be carried out in a same way.

Tools such as *grub-btrfs* can be used to boot from snapshots, which is really useful should updates fail – to utilise this, you would need to ensure that you have a snapshot in place before carrying out potentially destructive operations.

A to ZFS

ZFS was originally developed for Solaris and is an incredibly well-respected filesystem with a history that goes back around 18 years. In fact, ZFS is more than a filesystem and acts as a disk management solution as well. It is highly recommended that ZFS directly controls the disks, so that it can most effectively be used to set up various zRAID options, such as a single, dual or three-disk parity array. ZFS is also a copy-on-write filesystem, has native support for checksumming, which is used to ensure data can be recovered even when issues occur on the disk, and has support for encryption and compression as well. ZFS can also be set up with other advanced features, such as being able to add SSDs/NVMe disks as a cache disk to help with speedily writing data, which is then written to the slower spinning platter disks.

ZFS can be used with Linux, but due to licensing concerns, it can be difficult to include with Linux – although Canonical has shipped Ubuntu with ZFS for years now and has clearly carried out its own legal analysis of the situation.

Tools such as the *ZFSBootMenu* bootloader can be used to boot a Linux system from a ZFS boot system. This is particularly powerful because it means that should a filesystem snapshot be taken before a piece of software or an update is installed, an earlier snapshot can still be booted from should there be any issues with the installation or update process.

While it cannot currently be recommended, Canonical's ZSys tool was designed to allow booting from ZFS snapshots and it would also ensure that snapshots were taken every time an install or upgrade operation was taken. ZSys has not seen any commits for a number of months now, unfortunately.

While looking at filesystems and functionality that our next-generation distributions can utilise, we should mention one to watch for the future. Bcachefs will be included in the 6.7 version of the Linux kernel, which means that it can undergo lots more testing. Bcachefs can also use the copy-on-write system and supports checksumming, encryption, compression and snapshots. Bcachefs describes itself as being as reliable and robust, and having the features expected of a modern operating system.



BlendOS

Released: 2023
Apps: Various
PM: bpkg
FS: OverlayFS
Desktop: Gnome, KDE, others
Features: Atomic, declarative, immutable

» WHICH MIRROR?

Anybody who has used a Chromebook has probably been impressed at how seamless it is to perform updates. Generally, the updates are downloaded and applied in the background, then instantly booted from on the next reboot. If anything fails, the boot restarts, but this time using the older version of the operating system. Chrome OS accomplishes this by having two root partitions. Updates can be applied to the partition not currently in use and then at next reboot, the partition booted from is changed.

Vanilla OS takes a similar approach and has added support for LVM, which allows for thin provisioning. Previous versions of Vanilla OS required 40GB free space to create two 20GB root partitions. This was a concern to users, so the thin provisioning system allows space to be used only when needed. This is a common technique in virtualised environments. Vanilla's in-house developed *Apex* tool allows for software to be installed on the base operating system as well as in containers, which are integrated neatly into the system.

Partition	Filesystem	Mount Point	Label	Size	Used	Unused	Flags
unallocated	unallocated			2.00 MiB	---	---	
/dev/sda1	fat32	/boot/efi		974.56 MiB	7.17 MiB	967.39 MiB	boot, esp
/dev/sda2	ext4	/boot		976.56 MiB	309.35 MiB	667.21 MiB	
/dev/sda3	btrfs		b	19.07 GiB	5.37 GiB	13.70 GiB	
/dev/sda4	btrfs	/, /opt, /system/usr, /var	a	19.07 GiB	5.37 GiB	13.70 GiB	
/dev/sda5	btrfs	/home		19.94 GiB	13.59 GiB	6.35 GiB	
unallocated	unallocated			2.00 MiB	---	---	

! This gives an example of the sort of partitioning schema used by Vanilla OS.

Immutability & atomicity

Buzz words to help stop you breaking things!

When reading about next-generation distros, the terms immutable or immutability are spoken about a lot. Immutable means something that is unchanging over time. In the Linux world, it generally refers to the fact that the base operating system is set up in a read-only fashion, and that only updates from the distribution's makers can update the operating system itself. Combining this technique with the use of sandboxed packages and containers leads to an OS that is as lean as possible and means that as few bugs as possible can affect processes running as root.

Multiple techniques are available to ensure immutability, including clever partitioning systems, which separate the basic operating system from its configuration, and Unionfs or similar systems being used to layer filesystem images on top of each other. For example, layering a writable image on top of the base OS, which contains `/etc` and other locations containing configuration, allows immutability.

Hand in hand with immutability is atomicity, the state of being composed of indivisible units. In Linux terms, atomicity refers to the fact that updates either take place or they don't. Clever techniques ensure that failed updates or package installs are detected and rolled back, so our distros don't end up broken.

Fedora

Fedora Silverblue, Kinoite and Onyx are immutable distributions providing GNOME, KDE or the Budgie desktop environments. All three use filesystem layers to install updates and new packages. The RPM-OSTree system is used to take a version-controlled filesystem image and install it to the hard drive. A reboot is required to make upgraded packages or newly installed software available. More information can be found at <https://fedoraproject.org/silverblue>.

Vanilla OS

Vanilla OS uses standard OCI images (Open Container Initiative) to provide updates to the operating system. The `ABRoot` utility is used to handle updates to the two root filesystems. The usage of OCI images ensures that

» UBUNTU CORE

The immutable version of Ubuntu will hopefully be released in April 2024. Ubuntu Core has a long history in which it has been available as a platform for IoT and other single-use purposes. For example, Ubuntu Frame uses Ubuntu Core and provides an environment to build a full-screen shell to allow a device, such as a sales terminal, to do one job reliably. Ubuntu Core uses Snaps to containerise everything, including the kernel. Constrained apps mean it is harder for a bad app to take down the system.

Ubuntu Core also keeps a copy of the last boot path, so if an update or upgrade fails (which is unlikely), it can still boot using the last known good configuration. The immutable Ubuntu distro is being based on Ubuntu Core, with the desktop environment and everything else needed constrained within Snaps, and using the built-in mechanisms to allow constrained apps to communicate with each other. This promises to be an exciting development and while Canonical is a little late to the party, it could be that using the expertise learned with Ubuntu Core results in an excellent released product. Learn more about the early stages of Ubuntu Core Desktop at <https://ubuntu.com/blog/ubuntu-core-an-immutable-linux-desktop>.

the operating system is always in a consistent state. `ABRoot` also contains its own package manager, which generates OCI images containing the user's changes, which can then be applied and enabled on a reboot. You can find out more information at <https://vanillaos.org>.

BlendOS

BlendOS is based on Arch Linux, which is a fascinating choice, given that Arch is known for being completely bleeding edge, rather than sitting on the safer side of things, like Debian. Immutability is provided using filesystem layers, which are transparently mounted on top of the root files, then on the next reboot, changes are merged into the root. Updates are provided using ISO files and overwrite the existing root files at the next reboot. `Zsync` is used to ensure that update file downloads are as small as possible. BlendOS can be downloaded from <https://blendos.co>.

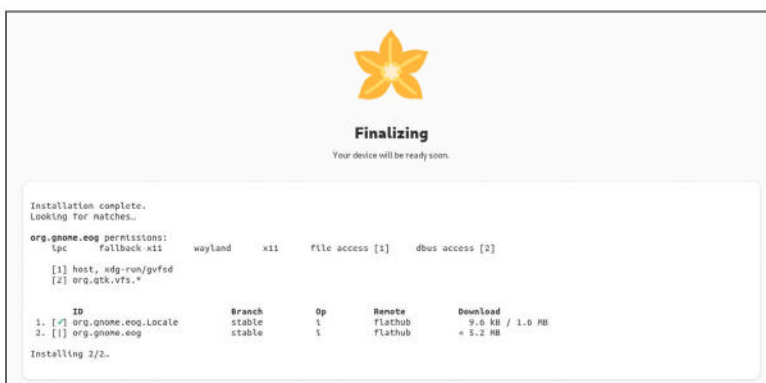
Endless OS

Based on Debian, Endless OS utilises a read-only filesystem managed by OSTree, while Flatpaks are used to install apps, and the desktop is GNOME. Endless OS can be downloaded from www.endlessos.org.



Vanilla OS

Released: 2023
Apps: OCI
PM: `apx`
FS: `ext4`, `abroot`
Desktop: GNOME
Features: Atomic, immutable



Being declarative

The future is all wrapped up in a single text file?



NixOS

Released: 2003
Apps: Nixpkgs, ICO
PM: nix
FS: ext4, btrfs, ZFS
Desktop: Any
Features: Atomic, declarative, immutable

Declarative configuration covers many areas and provides a method of defining something and then using some sort of automation to create the outcome. This sentence is deliberately vague as there's a large number of tools that take some configuration and convert it to output.

The first example is *Terraform*, which allows virtual environments to be built from a textual configuration file. *Terraform* has support for a large number of environments, so virtual hardware can be provisioned in a wide range of platforms, spanning AWS, Azure, *VirtualBox*, VMWare and more.

The second example is *Ansible*, which allows devices to be configured in a certain way from a central server. When registering devices, they can be tagged, and incredibly complicated configuration can be built using a combination of tags that link to different configuration files. A single configuration definition can also apply to many devices and groups of devices.

One powerful aspect of this sort of configuration process is that the files can be easily backed up and/or shared. Should the worst happen and you end up recovering from backups, you can set up your declarative configuration tools and run them to configure either virtual hardware or operating systems.

NixOS is another OS that uses modern techniques to make the operating system as reliable and secure as possible. It handles things very differently and also uses a lot of the same techniques as the tools that can



The NixOS installer provides all necessary options to install the system.

be used when using the Infrastructure as Code methodology. NixOS centres around the Nix package manager, which installs apps in a very different way from what we're used to. Apps can be installed on a per user basis, or for every user. When installed, each app is stored in its own directory and multiple versions of the same app can be installed. Each directory contains all of the app's dependencies and, again, multiple versions can be installed across the system. Apps installed or upgraded using this system are atomic, so the rest of the system is not broken if an upgrade or installation fails. Should a bug be found in a version of software, it can be rolled back, as the old version is not deleted straight away. Automatic updates are also easy to configure with the system. Nix can be installed on Linux distros as well as Mac OS and is also able to build or install packages from source or download and use prebuilt binaries. For small tools, source compilation wouldn't take long, but imagine building your browser or window manager for each update that is released.

So, that's Nix, but what is NixOS? The operating system uses Nix to not only build and install packages, but also to provide configurations to daemons, network configuration, package installation and more.

How does this work? Well, as we said, apps are all stored within their own directory in the Nix store. By default, this is in `/nix/store`. NixOS does not use a number of the usual directories you'd expect to see, such as `/bin`, `/sbin`, `/lib` and `/usr`, and accesses files from the Nix store instead. While there is an `/etc` directory, a lot of the files in there are just symbolic links (symlinks) to the files in the Nix store.

The benefits of a system such as NixOS are apparent, but this methodology must be maintained, as if you were to start installing software using other means, you would not be able to recover from a data issue, or set up a new device in exactly the same way. Learn more about NixOS at <https://nixos.org>.

» CONFIGURING VIRTUALBOX

Here we are installing two packages for every user of the system, then installing the *VirtualBox* guest tools on a running installation of NixOS. If you're not using *VirtualBox*, this can be omitted.

Open the Console and enter the following command: `sudo nano /etc/nixos/configuration.nix`. Now ensure that the following two sections have been added. The first already exists and needs to be edited and the second can be added at the bottom of the file, before the closing bracket. When adding the *VirtualBox* options, ensure they do not already exist in either the `configuration.nix` or `hardware-configuration.nix` files. Also, if installing on a physical device, do not add the *VirtualBox* options.

```
...
environment.systemPackages = with pkgs; [
  libreoffice
  firefox
];
...
virtualisation.virtualbox.guest.enable = true;
virtualisation.virtualbox.guest.x11 = true;
```


Cool Blue!

Matt Holder sits down with **Jorge Castro** to chat about next-gen distros.



Jorge Castro, formerly of Canonical, VMware and various other companies.

Over the last couple of years, Jorge Castro, formerly of Canonical, VMware and various other companies, has developed a keen interest in next-gen distros and is heavily involved with the Universal Blue project and Project Bluefin (<https://projectbluefin.io>). These projects build on top of Fedora and add extra functionality. You can learn more about Jorge and his projects at www.ypsidadanger.com.

LXF: What got you interested in immutable OSes?

Jorge Castro: I was growing frustrated with the lack of commitment by vendors to fix the basic reliability issues that have been plaguing the Linux desktop. Fedora was the closest to getting Chromebook-reliability but had a goal of five(!) years to get it ready. So a bunch of us decided to help accelerate that goal.

The market has rejected the current Linux client model and the desktop is stuck in a loop of failure. Linux dominates every field it's in, except this one. This bothers me. Our solution is to throw away what 'the Linux desktop' is supposed to be, adapt successful adoption patterns and leave failed ones behind.

Portal adoption, finishing the Wayland transition, new driver improvements, and Flatpak integration all can't happen if we can't deliver it to the end user. So far only Android and ChromeOS have been able to deliver reliable updates and improvements in a way normal people can consume. Systems like this level the playing field. So, I see it as fixing the delivery pipeline first.

LXF: What future concept are you most excited about?

JC: I think you'll see two main implementation patterns emerge: ostree/bootc and systemd-sysext. Both will share components and only deal with the base system. The workloads will be decoupled via OCI containers for services and Flatpak for GUI apps. There'll be vendor-specific implementations of a next-gen desktop, like Ubuntu Core, but generally it'll be about who delivers the most reliable base system with the composability users want. All of them will consume Freedesktop portals and container runtimes as first-class primitives. Universal Blue has chosen ostree/bootc.

LXF: How does Universal Blue relate to the work Red Hat is doing with Fedora Silverblue?

JC: Universal Blue consumes the Fedora Atomic OCI images then generates new base images for people to use with things like codecs and quality-of-life features. These are built using Dockerfiles and GitHub actions; it's a pattern ops people have been using on servers, so it was straightforward to find people with experience.

It started off as "let's provide base images for the community" but we discovered a new pattern on the way: it only takes a handful of ops people to maintain an awesome Linux workstation. We learned that a handful of motivated maintainers can come together

and say "Let's make a kick-ass modern Linux desktop for the next generation of open source maintainers." And that's how Project Bluefin was born.

The image is then deployed to your computer – this results in reliable upgrades and way less maintenance than a traditional setup; the maintenance happens in CI, not your local computer. And since your workload is decoupled from the OS, you can continue to use all the software you're using today. It flips the consumption model, you get everything you want, it's just stamped out in a known working state before it reaches your PC.

LXF: For all the keen gamers out there, what is Bazzite?

JC: There's a community of people who take all of the patches from SteamOS and put them into other Linux distros. Bazzite gathers all those patches and builds an image on top of Fedora. Since it's community driven, people have added support for the Lenovo Legion Go, Asus Ally, and other devices. Fedora moves quickly, so it's a great way to get new kernel updates and features that might not be in SteamOS. Our customisations are atomically layered, so you're still using Fedora; it's like setting it up for you to be a great gaming experience.

The best thing about Bazzite is it publishes builds every other day; it goes as fast as the community goes. Having a community that can publish a fix for devices immediately is a powerful pattern – users love it.

We can finally democratise getting hardware to work without janky scripts and post-install setup. We can finally get to a place where your stuff can just work. We're not there yet but this workflow makes it easier.

LXF: NixOS is so different – have you used it?

JC: I have but it doesn't solve a problem I have. I need a stable base OS and desktop that I can compose as I see fit, which Fedora already provides to everyone. We provide a Nix shortcut via the *Determinate Installer* and *Devbox* for people who want it, though. The nice thing about this new desktop model is you can just consume whatever you want from anywhere.

We prefer to reuse the same containers used in server deployments. Reusing the tech people use on servers is where the value is; the desktop is now just another container workload being sent to an edge device. I don't need to learn a new language for that. This is the kind of transparent, reliable delivery mechanism we've been missing and the clear path for us was to just copy that pattern to the desktop.


LXF: Any thoughts you'd like to leave the readers with?

JC: The consumption of open source will only ever increase, and to keep these projects healthy, we need more open source contributors. My hope is that people will use things like Project Bluefin to learn the technology, then if they want to work on open source, they have everything they need to be successful. **LXF**



Fedora Silverblue
Released: 2018
Apps: Flatpak
PM: rpm-ostree
FS: Btrfs
Desktop: Gnome
Features: Atomic, immutable


Meet the technology experts



- The world's **most** comprehensive technology website
- An **unrivalled** mix of news, opinions, reviews and features
- **All-new design**, new homepage, new features and special reports
- Backed by **over 300 years** of editorial experience

techradar
the technology experts

www.techradar.com



Let them eat lots of Pi!

Hundreds of thousands of Raspberry Pis are hitting the market, with production ramped to the max!

By the time you read this, there should be 90,000 freshly baked Raspberry Pi 5s hitting the supply lines every week. This is an uplift from the already high 70,000 weekly number Sony's production was cranking out.

Raspberry Pi CEO Eben Upton shared the promising news with us, along with pictures, which he says he took during a recent factory visit. It shows dozens of panels filled with Raspberry Pi 5 boards that are about to be tested and packed for shipping. Each panel in the picture contains nine Pi boards.

"Uplift to 90ku is mostly due to bringing more test heads online in the auto test pods," Upton explained. "Sony really are quite remarkable."

Sony runs the factory in Pencoed, South Wales (see **LXF182**), where Raspberry Pi boards are manufactured. Upton noted that the process is highly automated, with everything from the testing to the packing being done by machine.

After a rough couple of years during which all models of Raspberry Pi (except

the Pico microcontroller) were in short supply, it is now relatively easy to find most SKUs in stock, although the latest Raspberry Pi 5 models are still coming in and out of stock. Hence the recent bump in production supply.



If there's nine Pis per shelf, we're looking at £32,000 of Pi.



Les Pounder works with groups such as the Raspberry Pi Foundation to help boost people's maker skills.

» PEERLESS PIN NUMBERS

I've been part of the Raspberry Pi community since its inception. Back in 2011, we heard whispers of a new single-board computer (SBC) and gasped in awe as we heard the spec – and the price. Before the Raspberry Pi, SBCs were typically £100+ and used mainly in embedded applications. A £35 Linux computer that could also dabble with electronics? What wasn't there to like?

There have been attempts to dethrone the Raspberry Pi and, sure, some of them have more RAM, a better CPU and extra features. So, what keeps us coming back to the Raspberry Pi? The GPIO. Those 40 pins (26 in the original model) are the gateway to learning new skills, and for a decade they have been used to build an amazing and eccentric mix of Pi projects.

The Raspberry Pi was updated in 2014 with the B+ model and a new GPIO standard. Hardware Attached on Top (HAT) was also introduced, and it has since been the foundation for many boards and projects.

In late 2023, the Raspberry Pi team introduced HAT+ on the Raspberry Pi 5 (<https://datasheets.raspberrypi.com/hat/hat-plus-specification.pdf>). This new specification is still very much in the works, and the change means that many of our favourite HATs are currently incompatible with the Raspberry Pi 5, but this is only temporary. Many developers are working to make their products work with the latest flagship Raspberry Pi.


We are looking forward to seeing the first HAT+ based boards in 2024 and learning more about this new standard.

CREDIT: Raspberry Pi Foundation

ProPico

Better by design.

Maker and developer Dmytro has created a seriously upgraded design for the Raspberry Pi Pico. Sticking to the original form factor, new features include larger flash size, a reset button and a DC-DC converter PSU, plus losing the micro-USB. Find out more: <https://github.com/diminDDL/ProPico>




CREDIT: Dmytro, Raspduino Uno

Cyberdeck

Survive anything!

Survive the apocalypse with the PiDex. With a 32GB flash drive that stores a copy of Wikipedia, it's fitted with an OLED screen and a toggle power switch, and it's packed inside a waterproof box. The main display is a 7-inch touchscreen. Plus Ethernet and dependable 3.5mm headphone jack. Find out more: www.youtube.com/watch?v=bsT4CR9Jw-c



Official Pi 5 Case

Can the new official Raspberry Pi 5 case keep the Pi 5 cool?
Les Pounder is – brace yourselves readers – on the case...

IN BRIEF

The official Raspberry Pi 5 case continues the design lineage of its predecessors and can keep the flagship Pi 5 cool even under pressure, thanks to an integrated fan that uses the provided fan connector. The injection-moulded case is strong and light, and a removable lid provides access to the GPIO and HAT support.

The official Raspberry Pi case was introduced during the tenure of the Pi 3B, and the white and raspberry pink case has protected many Pis since. The Raspberry Pi 5 has now received its first official case and on the outside it continues its lineage. The case clips together without the need for tools, and the build quality is excellent. There is even a button with a light-pipe for the Pi 5 power button and LED.

Inside we have a plastic insert that holds a small 5V fan just above the new 2.4GHz quad-core Arm Cortex A76 CPU. We need active cooling because the Raspberry Pi 5 can get a bit hot without it. So, is a case a smart move? Let's find out by benchmarking it!

We attached the included heatsink to the BCM2712 SoC. There is only one heatsink in the box, yet the PMIC chip can get just as hot as the CPU. The fan connects to the new fan header, freeing up GPIO access. We also have to say that the fan is quiet. The official case fan for the Pi 4 sounded like angry wasps were trying to escape their plastic prison, but this one is super-quiet.

Keep a lid on it!

We left the Pi 5 in the case with the lid closed for five minutes, and then ran a *Stressberry* test. At idle, the Raspberry Pi 5 sat at 47.7°C and the fan was off; under stress, running at 2.4GHz, the maximum temperature was 70.3°C and the fan kicked in at 50°C. This is well under the 82°C thermal throttle point. This means we can overclock the CPU and get some free performance. Removing the lid is possible, and this brings in a little more fresh air. We repeated the test with no lid, just the fan insert and the fan connected. The idle temperature was spookily 47.7°C. Under stress, we saw a max of 67.5°C – 2.8°C under the previous test. The conclusion: at idle, it makes no difference to the temperature, and under stress it makes little difference, so keep the lid on unless you need to access the GPIO.

The fan insert can be removed and that means we can use the official active cooler. This cooler has a beefier heatsink that cools the SoC, PMIC and the RP1





This two-tone white and raspberry pink case is a delight. The injection-moulded plastic is strong and will protect our Pi 5 from harm.

chip. It also has a fan. With this in place, we saw an idle temperature of 46.6°C (1.1°C lower than the stock case fan) and under stress we hit 68.6°C (1.7°C lower than the stock fan). So, the active cooler does provide some benefit, but only by a fraction. That said, the active cooler is only around £5 and well worth the money.

Can we connect a HAT? With the fan insert in place we can – just use a header extension to ensure that your HAT isn't resting on the fan. And obviously take the lid off. HAT compatibility is important because Raspberry Pi has announced a range of HATs designed for the Pi 5. One of these is the M.2 HAT, which brings support for 2230 and 2242 SSDs.

Is the official case worth the money? Short answer is yes and you should buy one. It protects your investment while looking good. It is well built and provides easy access to the ports and the GPIO. The included fan does a good job of cooling the Raspberry Pi 5; we'd just like to see a larger heatsink to cool the PMIC and RP1.

Regarding cases, our preference is for Pimoroni's Pibow, but we have yet to get hands-on with one for the Pi 5, so for now the official case is the one to beat. 



If you want to give the fan a little more air, or need to use the GPIO, the lid easily comes off without the need to remove the fan.

VERDICT

DEVELOPER: Raspberry Pi Ltd
WEB: <http://raspberrypi.com>
PRICE: £9.90

FEATURES	8/10	EASE OF USE	8/10
PERFORMANCE	8/10	VALUE	9/10

A great case that looks cool while also keeping the flagship Pi 5 cool under stress. This is the case to buy – for now.

» **Rating 8/10**

Geniatech XPI 3566 Zero

All brains and no brawn applies equally to this SBC and to **Les Pounder**.

SPECS

SoC: Rockchip RK3566
CPU: Quad-core ARM Cortex A55 up to 1.8GHz
GPU: ARM G52 2EE
NPU: RKNN NPU AI accelerator, 1Tops@INT8, Caffe, TensorFlow, PyTorch and more
RAM: 512MB to 8GB
Storage: 8GB to 128GB eMMC
Ports: Mini HDMI, 1x USB-C
 USB 2.0 OTG, 1x USB-C Host, MIPI-CSI
Comms: Wi-Fi Ampak AP6256 2.4/5GHz, Bluetooth 5.0
GPIO: 40-pin, UART, SPI, I2C, PWM, digital IO
Power: 5V 2A via USB C
Size: 65x30mm
OS: Debian 11, Raspberry Pi OS emulated

Geniatech's XPI-3566-Zero aims to be a Pi Zero 2 W killer. Based on the RK3566, it's a quad-core ARM Cortex A55 running at 1.8GHz that offers multiple RAM and eMMC options, up to 8GB of RAM and 128GB of eMMC storage. Here we're testing the modest 2GB RAM, 8GB eMMC option.


If you confused this for a Raspberry Pi Zero, you'd be forgiven; the key difference is that there is no microSD card slot, and the USB ports are USB-C. The underside sports a couple of buttons for reset and setting the boot mode, with an antenna connection for the onboard Wi-Fi and Bluetooth module. It also won't fit inside a Pi Zero case – well, not without modification to accommodate the USB-C port.

Because there is no microSD card slot, the only means to flash the onboard eMMC is via USB, and for that we need a companion app on our desktop PC. We need an updated driver from Geniatech support, then using the *RKDevTool* to flash the firmware, the default language choice was Chinese. The fix here was to edit the config file and under Language set Selected to 2, for English. After that, we could carry on the install process and finally boot the board.

Speed bumps

Performance is weak, as you would expect. The onboard eMMC is fast enough, with a read speed of 113MB per second; however, 720p video playback dropped over a third of the frames and offered flickering while taking screenshots. We used a spare antenna (one is not supplied) and tested the Wi-Fi performance to our PC, connected to the network using Ethernet. We managed a paltry 36.5Mbps per second.

We ran a benchmark to check the CPU temperature when idle, and under an all-core stress test. Suffice to say, the Geniatech XPI-3566-Zero can get a little warm when under stress. Idling, it ran at 42°C and 400MHz; under stress, the CPU jumps to 1.8GHz and hits 81.7°C. This is with no cooler – we'd be looking to add at least a heatsink. As for power draw, at idle it sipped away at



The Geniatech is similar in size but sports a better specification to a Pi Zero – just don't expect to be able to use its GPIO as easily.


0.97W (5.13V * 0.19A). When under stress, this jumped to 2.45W (5.11V * 0.46A).

We'd love to say that accessing the GPIO was easy. Essentially we can set the status of a GPIO pin and work with it with only a few lines of *Bash*. The problem is that we need to do the maths to convert the GPIO pinout reference into something that the OS understands.

We managed it, and performed a simple test of the GPIO with a push button and an LED. It worked, but we'd love to see alternative means to access the GPIO. Our first thought would be Python. There are installation candidates for RPi.GPIO and GPIO Zero, but we'll cut to the chase and say that they don't work. There is probably some low-level config that can be used to make it work, but that beyond the scope of this review.

The GPIO may be 40 pins, but don't expect to drop the best Raspberry Pi HATs directly on to it. There is largely no support for Raspberry Pi HATs, unless you want to get your hands dirty. If so, make sure to document the work on GitHub so that others can benefit.

If you need the Raspberry Pi Zero form factor with a little more horsepower, the Geniatech XPI-3566-Zero is a viable alternative. But if you need the GPIO, get a Raspberry Pi Zero 2 W, **LXP**



An external antenna can be connected but you need to supply your own.

VERDICT

DEVELOPER: Geniatech
WEB: www.geniatech.com
PRICE: £48

FEATURES	9/10	EASE OF USE	6/10
PERFORMANCE	6/10	VALUE	4/10

A decent attempt that focuses on RAM, storage and CPU speed over anything else. It's better than the Raspberry Pi Zero 2 W on paper, but we'd take a Pi Zero 2 W every time.

» **Rating 6/10**

NODE-RED

Credit: <http://nodered.org>

Start using MQTT for sensor logs

Les Pounder is a sensitive soul and his new office needs sensors to record data. He'd better break out the maker kit!



OUR EXPERT

Les Pounder is associate editor at Tom's Hardware and a freelance maker for hire. He blogs about his adventures and projects at <http://bigles.com>.

The Internet of Things (IoT) is hard, right? Wrong. It is easy thanks to MQTT and *Node-RED*. For this tutorial, we are building a simple Raspberry Pi Pico W-powered sensor that sends data to a Pi-powered dashboard running on *Node-RED*. The protocol that we're using is called MQTT and there are more details in the boxout (*opposite page*).

There is only one sensor connected to the Pico W. This is a common DHT11 temperature and humidity sensor, which can be picked up cheaply and is often bundled in electronics kits. The DHT11 has only three connections to the Pico W: 3V, Data and GND (red, yellow and black wires in the diagram). Please refer to the high-res image in the download for this issue.

While holding the BOOTSEL button, connect your Pi Pico to your computer. Go to <https://bit.ly/lxfupdate> and download the version of MicroPython for your Pico or Pico W. Open your file manager, go to the downloaded file and copy it to the root of the **RPI-RP2** drive. This flashes the new firmware to the Pico. Using your distro's package manager, install *Thonny*. For the latest Ubuntu release we have to use a Snap package:

```
sudo snap install thonny
```


Open *Thonny* and connect the Pico to your machine. Go to Tools > Options and select the Interpreter tab. Set the interpreter to MicroPython (Raspberry Pi Pico) and set the Port to match the location of your Pico. Click OK. *Thonny* now connects to the board and we can start writing code.

Writing code for the Pico W

First we need to import a library of prewritten MicroPython, specifically for MQTT on MicroPython: *Umqtt*. Open a new blank file in *Thonny* and visit <https://bit.ly/lxf312python>. Copy and paste the contents of the file browser page for this link and save it to the root of the Pico as **simple.py**.

Now create a new file in *Thonny* and import more modules of code. First is *machine*, used to interface with the Pico's GPIO. Then *dht*, a module for use with the DHT temperature sensor. Next is *simple*, our MQTT module. Lastly we import *network* and *time* to provide network connections, and to pause the code:

```
import machine
import dht
```



The circuit for this project requires just three wires to connect the DHT11 to the Raspberry Pi Pico W.

```
from simple import MQTTClient
import network
import time
```

Now create two variables to hold your Wi-Fi SSID and password:

```
WIFI_SSID = 'YOUR_WIFI_HERE'
WIFI_PASSWORD = 'WIFI_PASSWORD'
```

Add a section for the MQTT broker details – we're using the public HiveMQ broker, so remember that your data is publicly visible. We need the broker address, port, client ID (our Pico) and the topic:

```
MQTT_BROKER = 'broker.hivemq.com'
MQTT_PORT = 1883
MQTT_CLIENT_ID = 'pico_sensor'
MQTT_TOPIC = b'sensor_data'
```

Create an object that tells our code where to find the DHT11 sensor. This is on pin 16, which we set as an input. Create an object to work with the DHT11 sensor using the *DHT* module:

```
dht_pin = machine.Pin(16, machine.Pin.IN)
dht_sensor = dht.DHT11(dht_pin)
```

To connect to the Wi-Fi, make a function that uses our SSID and password. The function is called **connect_wifi** and we first use the *network* module to start the Wi-Fi chip:

```
def connect_wifi():
    sta_if = network.WLAN(network.STA_IF)
```

YOU NEED

- > Pico W
- > Breadboard
- > DHT11 sensor
- > 3x male-to-male jumper wires
- > Network
- > Code: <https://github.com/lesp/LXF312-PicoW-MQTT/archive/refs/heads/main.zip>

```
print('Connecting to Wi-Fi...')
sta_if.active(True)
sta_if.connect(WIFI_SSID, WIFI_
PASSWORD)
```

Now we need to create a new function that handles connecting to the MQTT broker using the details we created earlier:

```
def connect_mqtt():
    client = MQTTClient(MQTT_CLIENT_ID,
MQTT_BROKER,port=MQTT_PORT)
    client.connect()
    return client
```

This reads the temperature via the DHT11 and stores it in an object called **temperature**:

```
def read_sensor_data():
    dht_sensor.measure()
    temperature = dht_sensor.temperature()
    return temperature
```

Connect to the Wi-Fi and MQTT broker:

```
connect_wifi()
mqtt_client = connect_mqtt()
```

Inside a **while True** loop, we read the sensor data and store the value in a variable called **temperature**:

```
while True:
    temperature = read_sensor_data()
```

Publish the sensor data to MQTT using the MQTT_TOPIC. We then print a message to the Python Shell, along with the current temperature. The code sleeps for one minute before the loop repeats:

```
mqtt_client.publish(MQTT_TOPIC, str(temperature))
print('Data sent!', temperature)
time.sleep(60)
```

Save the code as **temperature_logger.py** to the Raspberry Pi Pico W. Click on Run > Run Current Script (or the green button) to run.

Now we are sending the code, we need a means to read it. For that we go to our Pi and install Node-RED.

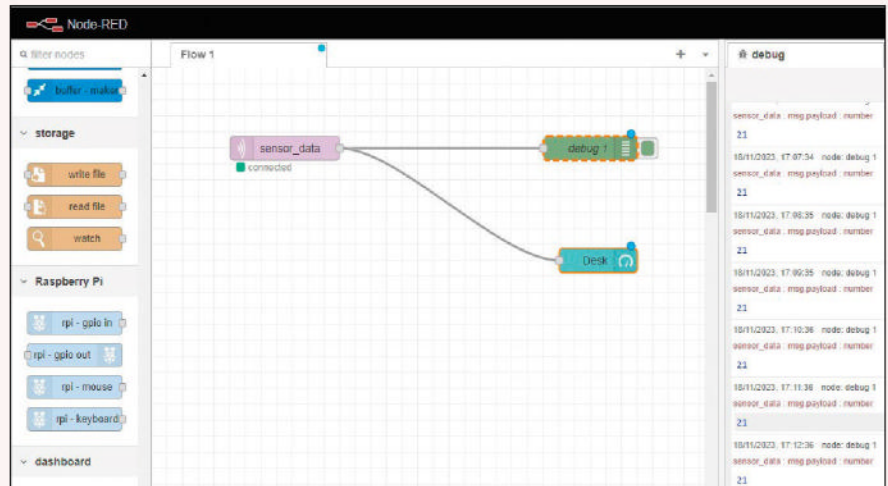
Setting up Node-RED

Node-RED is a graphical programming language that uses nodes joined by wires to create a flow. The nodes are the lines of code, the wires join the output of one node to the next (or multiple). The flow is the entire body of code.

To install Node-RED on a Pi, follow the guidance at <https://nodered.org/docs/getting-started/raspberrypi>. There is a one-line *Bash* script that automates the process. In a terminal, start the Node-RED service.

```
$ node-red-start
```

On the Pi, open the browser to **127.0.0.1:1880** and run through the startup tips. Look for an mqtt in node from the list on the left side. Drag this into the flow. Double-click on the node and click on the pencil next to Add New Mqtt-Broker. In the new screen, name the broker HiveMQ and then enter the server (**broker.hivemq.com**) and click Add. In the next screen, set the Topic to **sensor_data** (this matches what the Pico is sending out) and click Done to save and use. Drag a debug node from the left side and join the right side of the mqtt in node to the left of the debug node by connecting the grey dots together. Click on Deploy and then click on the bug icon on the right side of the



screen. This is the debug console and it shows that MQTT data is being received.

We'll use a dashboard to make the data look better. Click on the three horizontal lines in the top-right of the Node-RED window. Select Manage Palette. Click on the Install tab and search for node-red-dashboard. Click Install and the new nodes are installed.

Scroll down to the dashboard nodes and drag a gauge node. Connect the output of the mqtt in node to the input of this node. Double-click on the gauge node and click on the pencil icon next to Add New Dashboard Group. In Properties, set the name to Office Temperature and then add a new dashboard tab called Office. Click Add to save. In the gauge node properties, set the Label to Desk, and the units to °C. Then set the range to between 0 and 40.

Click on Deploy to run the code. Open a new browser tab to **127.0.0.1:1880/ui** and you can see the live temperature data in a dashboard. **LXF**

Node-RED uses nodes (lines of code) joined by wires to create flows (projects) that take the output of one node to be the input of the next.

» MQTT AND UMQTT


MQ Telemetry Transport is a lightweight publish-subscribe protocol created by Andy Stanford-Clark and Arlen Nipper as a means to send data over low-bandwidth communication devices. Its use exploded with the Internet of Things and now it powers many projects, including Facebook Messenger.

MQTT is simple to use, lightweight and it runs on pretty much everything, even our £8 Raspberry Pi Pico W. This is thanks to the micropython-lib maintainer who wrote Umqtt, which we used. We didn't install it as they envisaged. Ideally we would use *Mip* to install. *Mip* is basically *Pip* (Python packaging) for MicroPython, but the install method we used saved a paragraph in our tight word count.

Our Raspberry Pi Pico W is publishing data to MQTT using a topic as a means to filter data. Think of it like YouTube. Our Pico is publishing content, HiveMQ is YouTube (broker) and we subscribe to the content (topics) that we want to see. Multiple devices can subscribe to a topic and there are many clients to read and send data over MQTT, even Android and Apple apps.


MQTT is an interesting and versatile protocol for simple messaging projects, and if you would like to know more, head over to the official MQTT site at <https://mqtt.org>.

» **GET YOUR Pi FILLING HERE** Subscribe now at <http://bit.ly/LinuxFormat>



Orange Pi vs Raspberry Pi

Tam Hanna compares the highest-end single-board computer offerings from Shenzhen Xunlong and the Raspberry Pi Foundation to find out which one is best suited for you.



An old French proverb claims that appetite comes with eating. Similarly, Eben Upton's deployment of the first-generation Raspberry Pi opened the floodgates on the single-board computer (SBC) market. As single-board computers became more affordable, more people purchased them. This increase in volume led to more market participants, which brought prices down further and led to even more single-board computer usage.

As the market is now more than 10 years old, we can see the effects of maturity. Notably, new trends have emerged; while the classic, credit-card format process computer is still going strong, smaller and larger boards are also becoming available. This, obviously, is an effect of market breadth – in a broader market, more offerings can survive, thereby permitting designers to select components that are ideally suited to their needs.

We're comparing the highest-end offering from Shenzhen Xunlong, the Orange Pi 5 Plus, against the new Raspberry Pi 5. Both of these SBCs are incredibly performant and offer more computing power than workstations could provide a decade or so ago. Nevertheless, each vendor specialises in different applications, so read on carefully to determine which SBC is better suited to run your world.

Increasing demand by process computer applications led to a CPU power arms race – both Shenzhen Xunlong and the Raspberry Pi Foundation try to provide as much computing power as possible. In the case of the Raspberry Pi 5, a four-core system-on-a-chip (SoC) is used; this time, the instruction set architecture (ISA) is ARM Cortex-A76. Combined with a clock rate of up to 2.4GHz, performance has doubled over its predecessor.

Shenzhen Xunlong takes a different approach by sourcing the SoC from RockChip. It's a diverse octo-core unit implementing the big.little architecture. The CPU provides four compute and four economy cores; the operating system schedules between them.

As the proof is in the pudding, the best approach is to use the SysBench benchmark. Our tables show the results for single-core, quad-core and octo-core runs.

SysBench CPU (events per second): single thread

Raspberry Pi 5	RPi 5, cooled	Orange Pi 5 Plus
1,476.15	2,724.44	2,741.43


SysBench CPU: four threads

Raspberry Pi 5	RPi 5, cooled	Orange Pi 5 Plus
5,887.46	10,912.94	10,927.82

SysBench CPU: eight threads

Raspberry Pi 5	RPi 5, cooled	Orange Pi 5 Plus
N/A	N/A	14,482.21

The availability of eight cores translates to higher throughput for parallelised applications. SysBench shows the different compute performance; bringing in the four economy cores brings just 4,000 additional events. The Raspberry Pi Foundation provides an (optional) active heatsink that mounts to the rest of the board via two plastic pins shown in the photo (above).



The attaching pins for the official Raspberry Pi 5 active cooler.

While many reviewers reported slowdowns without the fan, we could not reproduce these in our tests. Running SysBench CPU 10 times in sequence using a shell script did not yield a significant slowdown – the tables below show the first and tenth results on the Orange Pi 5 Plus and a Raspberry Pi without heatsink.

SysBench CPU (events per second):

	RPi 5, uncooled	RPi 5, cooled	Orange Pi 5 Plus
1st	10,921.81	10,911.97	14,489.69
10th	10,925.32	10,942.87	14,393.84

Performing such benchmarks is complicated as the heatsink comes with glue – once installed, removing it requires the replacement of the conducting adhesive.


Furthermore, remember that these tests ran on an open desk in a chilly laboratory. Put the Raspberry Pi into a tight enclosure (or block heat flow) to get potentially poorer results.

Another exciting aspect is that the trend to bundle high-performance AI accelerators has ended. The GyrFalcon AI accelerator on Orange Pi 4 variants is no longer available. Instead, the SoC has ISA optimisations that permit more efficient execution of ML tasks.

In terms of remanent memory, the Raspberry Pi 5, in theory, is available in 1GB, 2GB, 4GB and 8GB variants. At the time of writing, however, only the 4GB and 8GB versions of the boards are produced.

For the Orange Pi 5 Plus, 4GB, 8GB and 16GB variants are available; a 32GB RAM version has been announced but is yet to be available. The table shows the achieved memory bandwidths on our samples.

A dedicated built-in NMVe offers speedy storage out of the box.



QUICK TIP

Recent research has shown that some process computer operating systems check the way power is supplied to the SBC. Should you supply your circuit via the GPIO header, be sure to perform tests with the target operating system early on during prototype bring-up to prevent costly mistakes at a later point.

SysBench Memory (MiB transferred):

Raspberry Pi 5	Orange Pi 5 Plus	Orange Pi 5 Plus
Four threads	Four threads	Eight threads
89,219.32 MiB (8,920.89 MiB/sec)	86,157.11 MiB (8,613.66 MiB/sec)	102,400.00 MiB (11,914.12 MiB/sec)

Factor in expansion

Single-board computers are often not used stationary. When deployed in a drone, size and weight affect range and total cost of ownership. In the case of our comparison, the Raspberry Pi is significantly smaller, as can be seen below.

SBC	Weight	Size
Orange Pi 5 Plus	86.5g	100x75mm
Raspberry Pi 5	46g (69g with heatsink)	85x56mm

Shenzhen Xunlong makes good use of the extra space. The port complement of the Orange Pi 5 Plus is richer. The Ethernet cluster consists of two Gigabit interfaces and three HDMI ports. Two are intended for displays, while the third ingests HDMI signals.

Five USB ports are spread across the board: the two at the back are USB 2.0, while the three at the front provide USB 3 data rates. Furthermore, the USB-C port in the front can be used to

control a third display. For convenience, there's also an analogue headphone jack, a microphone and even an infrared remote control receiver.

Sadly, the spacing of the ports

can be suboptimal for process computer applications. On the other hand, having front-facing USB ports greatly simplifies desktop replacement use – a task for which the processor is more than strong enough.

Shenzhen Xunlong offers advanced peripheral interface functionalities: first of all, the Orange Pi 5 Plus has multiple memory interfaces, such as the eMMC module, which mounts to the top of the board. It provides a vibration-robust interface and can be handy in automotive applications.

Memory and operating system installation is usually done via a microSD card. It then executes a script that flashes the actual operating system to the eMMC

module – when the process is done, the microSD card can be discarded and reused on the next machine.

An additional interface used in the workstation space is PCIe. Shenzhen Xunlong's process computers have had PCIe interfaces for some time; the Raspberry Pi 5 marks the entry of the Raspberry Pi Foundation into this club. A difference in philosophies is visible in this area: while the Orange Pi has slots on the top and the bottom, PCIe hardware will connect to the Raspberry Pi 5 via an FPC cable and an – as-of-yet-unfinished – external expansion board.

The slot on the top of the Orange Pi implements the E Key form factor, while the slot at the bottom uses the M Key form factor commonly used for SSDs.

The FPC cable adds flexibility in terms of physical layout. However, it comes at a cost: the official data rate is limited to 5GT/s. A website hosted at

<https://bit.ly/LXF312Pi> claims that a configuration change enables faster transfers – in practice, care must be taken. Higher-speed transfers might work in the lab, only to fail when confronted with the real EMI environment found in an automotive system.

Battery time

Accurate timekeeping is helpful. Power management solutions, such as the one outlined in the application note AN4121 by Microchip (www.microchip.com/en-us/application-notes/an4121), bring down standby power consumption but lead to time loss due to the disconnection of the process computer's power supply.

Workstations avoid this problem by the use of a CMOS battery, a feature that is available on both the Raspberry Pi 5 and the Orange Pi 5 Plus.

Interestingly, the different philosophies in design are visible. In the case of the Orange Pi 5, the RTC schematics – available at <https://bit.ly/lxf312rtc> – closely follow the approaches known from the industry.

The circuit's most obvious consequence is that the RTC battery cannot be recharged. This permits using any battery – whether salvaged from a workstation or purchased at the computer store, the process computer is off to the races.

In the case of the Raspberry Pi Foundation, an ML2020 manganese battery is used. It has a lower nominal capacity (45mAH versus around 200mAH) but gets charged via the process computer.


This has advantages and disadvantages: while the frequent recharging ensures that the RTC never gets empty, connecting a non-rechargeable battery to a Raspberry Pi can lead to nasty results. Furthermore, the lower nominal capacity limits the amount of time that the process computer can stay offline.

Network philosophy

While achieving Gigabit Ethernet rates was a problem at the beginning of process computer-dom, nowadays all vendors have a good grip on the technology. Due to this, *lperf* runs do not yield any surprises.


Nevertheless, a difference in psychology and approach can be seen. While the Raspberry Pi Foundation provides a Bluetooth and Wi-Fi transmitter, Shenzhen Xunlong takes a more open approach.

The above-mentioned top PCIe slot takes any wireless module – in addition to the Wi-Fi and Bluetooth combination card sold by Shenzhen Xunlong,



The eMMC module provides robust, fast and reliable onboard storage.

A wider range of outputs is attractive for specific roles.



users are free to use modules from other vendors. This extends to 4G, 5G and LPWAN modules, as long as drivers for Linux are available.

IO bandwidth can be an issue. Remember that the SoC's use of process computers usually comes from the mobile phone space, where peripheral bandwidth is unimportant. Due to this, the total amount of bandwidth available is generally somewhat limited.

A representative test for bandwidth concerns using an external SSD – for these steps, we are using the rather pedestrian Kingston XS1000. For reasons of reproducibility, 10 runs were averaged, and the following shows the raw results of the SysBench IO benchmark: the Raspberry Pi 5 achieved around 500Mb/s, while the Orange Pi 5 Plus saw 360Mb/s. This, however, is not completely representative – microbenchmarks of storage are always a challenge.

Sniffing for bandwidth starvation can be handled via the Gigabit port. Invoking *Iperf* with a parameter set of `iperf -c 192.168.1.74 -t 300 -i 2` instructs the client to emit information to the command line periodically. That way, the current amount of bandwidth can be judged – a bandwidth starvation test can be performed by simply running *SysBench* after starting *Iperf*. This test was performed on both contenders; no significant bandwidth degradation was visible in either case.

GPIO interaction


Rising process computer prices and a fall in the price of x86 hardware have led to the two computer types ending up in the same price bracket. Advocates of process computers then point to the benefits of the GPIO interfaces.

Both systems come with a 40-pin GPIO header. In the case of the Orange Pi 5 Plus, colour coding of the individual header pins means that supply and IO pins can be easily kept apart.

An exciting surprise awaits when interacting with GPIO pins using C. In the case of the Raspberry Pi 5, the PiGPIO library found here <https://abyz.me.uk/rpi/pigpio/library> is considered the quasi-standard. Sadly, it is not able to interact with the Raspberry Pi 5 GPIO transceiver if the image **2023-10-10-raspbian-bookworm-arm64-full.img** is used. Instead, developers have to fall back to the LibGPIOD API documented under www.ics.com/blog/gpio-programming-exploring-libgpiod-library – this leads to significantly higher performance, but programming for it is considerably more difficult.

In the case of the Orange Pi 5 Plus, on the other hand, the situation could not be more straightforward. A program such as the following can be run, leading to the output of a 500kHz square wave:

```
#include <wiringPi.h>
int main (void) {
    wiringPiSetup();
    pinMode (27, OUTPUT);
    for (;;) {
        digitalWrite (27, HIGH); // On
        digitalWrite (27, LOW); // Off
        digitalWrite (27, HIGH); // On
        digitalWrite (27, LOW); // Off
    }
    return 0;
}
```



Concluding this, a remark is necessary: the Raspberry Pi Foundation seems to have shifted the focus of its developer activities to Python, and Python programs would work out of the box.

He who wrangles with LibGPIOD can get up to 20MHz worth of performance.

Questionable distribution

Purchasing either process computer is a bit complicated. For the Raspberry Pi 5, the 'small eats first' approach leads to the usual problems. At the time of writing, an OEMsecrets search of all reputable distributors shows no listings (check www.oemsecrets.com/compare/SC1111 to find a distributor near you).

Instead, small retailers sometimes get limited amounts of stock, which sell out quickly. Nevertheless, the manufacturer-suggested retail prices are for the 4GB and 8GB versions are \$60 and \$80 respectively.

In the case of Shenzhen Xunlong, the main sales channels are Amazon and AliExpress. At the time of writing, the prices start at around \$100, excluding shipping, for the 4GB model.

Your Pi guy

With the Raspberry Pi 5, the Raspberry Pi Foundation refreshes an existing (and well-loved) concept, adding features that the competition has had for some time.

Nevertheless, the Raspberry Pi 5 remains a powerful process computer – even though the GPIO issues are an unexpected problem, the community support as a whole is excellent.

With the Orange Pi 5 Plus, Shenzhen Xunlong clearly states its intent to play in the desktop/thin-client market. Significantly increased compute performance and the availability of a 32GB version put desktop-related tasks (think secretary, not military engineer) within reach. Furthermore, the port arrangement is ideally suited to desktop use.

Choosing between the two is, by and large, a question of taste. The richer communication capabilities and the higher performance of the Orange Pi 5 Plus make for a more comfortable desktop-replacement scenario; furthermore, the PCIe interface is more stable.

On the other hand, the Raspberry Pi 5 has a larger community, which can be crucial in saving work time when small one-off projects are to be realised. **LXF**

QUICK TIP

While PCIe is also used for interfacing graphics cards, do not dream of connecting a GeForce Ti to your SBC. The reason for that is that the workstation graphics port usually has 16 parallel lanes, while the interfaces used here expose only one or two lanes.

Monitor systems in style

Shashank Sharma doesn't look dapper in his quickly fading imitation leather jacket, but his terminal is always a sight to behold.



OUR EXPERT

Shashank Sharma is a trial lawyer in Delhi, India. He's been writing about open source software for 20 years, and lawyering for 10.

No matter what your system specification is, even if you are running a 12-core machine with hundreds of gigabytes of RAM or more, you will experience bottlenecks. And you will want to use a system monitor to see just what nefarious process is hogging your resources. This is the law. It's also why there are so many resource monitors on offer for Linux distros. For most Linux users, *Top* is the de facto system monitor, and you can think of *Btop* as a super-charged variant.

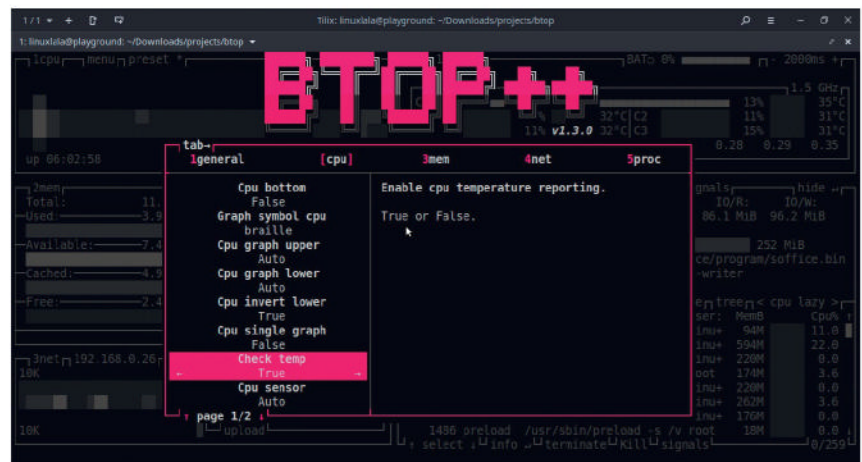
Written in C++ and released under the Apache Licence 2.0, *Btop* itself boasts a rich pedigree. The project can trace its roots back to earlier iterations such as *Bashtop* and *Bpytop*, written respectively in Bash and Python.

At first glance, *Btop* is just another resource monitor that can be used to view usage and stats for processor, memory, disks, network and processes. However, the interface and features set it apart from almost all other similar projects. For one, you can control every aspect of the utility using the defined keybindings, or even your mouse. You can also filter processes, switch to a tree view, and even terminate or kill a process from within the *Btop* interface itself.

The colourful graphs and the large number of themes make it a delight to work with, and you can configure every aspect from within the tool itself, using the game-inspired menu system.

Easy installation

Btop is offered in the software repositories of many desktop distros. Refer to the project's GitHub page for a list of the supported distros and the version of *Btop* on offer. Most modern terminal emulators should be able to handle *Btop*'s feature set. Linux users can use *Btop* to also monitor AMD and Nvidia GPUs out of the



Complete mouse support, where all lists are scrollable and buttons – identified with a highlighted key – are clickable is rare for CLI utilities.

box, assuming you have the correct drivers for your model installed on your system.

If your distro doesn't feature *Btop*, or offers an older version, there's a myriad ways to install it. If you're comfortable with the idea and have configured your system for Snap packages, you can install *Btop* using the Snap package with `sudo snap install btop`.

If you also want *Btop* to be added to the menu on your graphical desktop, you have to install the larger **btop-desktop** package with `sudo snap install btop-desktop`. You also have to run `sudo snap connect btop:removable-media` or `sudo snap connect btop-desktop:removable-media`, depending on the Snap package you choose to install.

Btop also offers an installation script to help ease the installation. Head over to the Releases page and download the latest compressed tarball.

```
$ cd ~/Downloads/projects
$ tar xjf btop-x86_64-linux-musl.tbz
$ cd btop
$ ls
bin      CHANGES.md  install.sh  Makefile    setuid.sh
uninstall.sh
btop.desktop  Img  LICENSE  README.md  themes
$ ./install.sh
```

You're prompted to enter the sudo password when you run the install script, and then *Btop* is installed,

complete with relevant entry in the System Tools menu on your chosen GUI. Should you ever tire of the utility, you can similarly run the **uninstall.sh** script to remove *Btop*.

From the terminal emulator of your choice, you can now run *Btop* with the **btop** command, or by clicking Btop++ under Menu > System Tools on your desktop.

Start monitoring

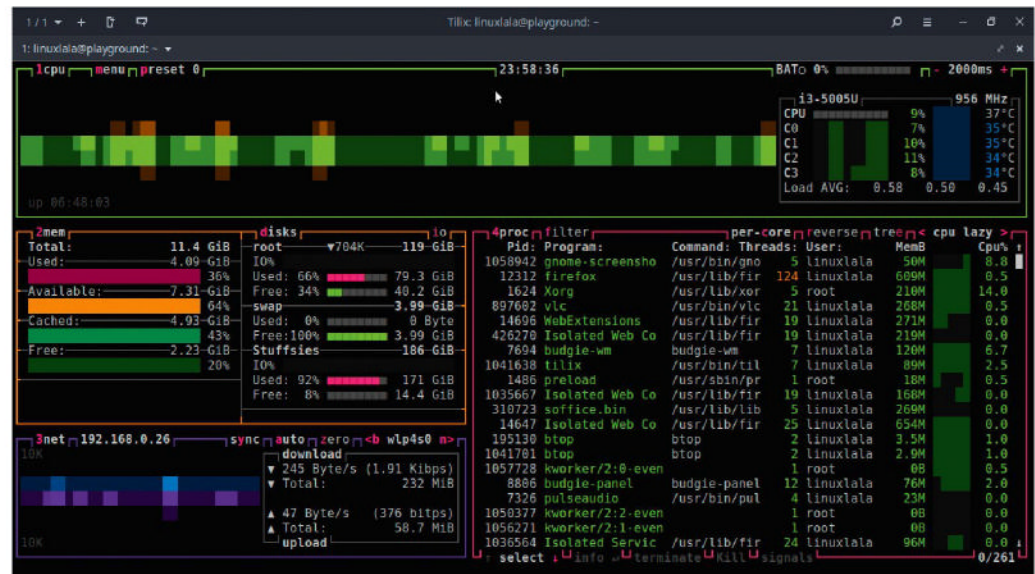
The default interface shows the CPU usage box across the top. Memory, disk usage and network traffic boxes are under the CPU box on the left, and you'll find a list of active processes on the right. *Btop* refers to these individual panes as boxes, and you can toggle each of them on and off using keybindings.

If you dislike the default layout of the different boxes on *Btop*, you can cycle through some of the different presets by pressing **p**. You can press **?** for a complete list of defined keybindings, but we'll list some of the most useful ones to help you get started.

Keybinding	Purpose
1	Toggle CPU box on/off
2	Toggle Memory box on/off
3	Toggle Network box on/off
4	Toggle Process box on/off
5	Toggle GPU box on/off
d	Toggle Disks view in Memory box
f or /	Filter the process
F1, Shift+h or ?	View the Help screen
Esc or m	Access the main menu
F2 or o	Open the Options screen
Left-click	Select process in the Process box
Mouse scroll	Scroll any lists

You can double-click a process in the Process box, or select it using the arrow keys and then hit Enter to view detailed statistics. After you select a process in the Proc box, you can also use the Info, Terminate, Kill or Signals buttons at the bottom of box. If you don't see these buttons on your emulator, try expanding the size of the emulator window, or switch to full-screen to get the most out of *Btop*'s interface.

If scrolling through the list seems too cumbersome, you can also press **f** or **/** to filter for the process you are looking for. After selecting a process in the Proc list, you can click Info or press Enter for detailed stats, which are displayed in an independent box on top of the Proc box. You must click the Hide button on the



extreme right of this box or press Enter again to exit the detailed stats box. To kill a process, you can select it from the list and press **K** or click the Kill button at the bottom of the Proc box. *Btop* asks you for confirmation before killing the process.

Some of the defined keybindings might not work with *Btop* on your system, depending on the distro, because they may already be in use. You can tweak these by editing the config file. Refer to the boxout (below) to learn how to tweak *Btop* to your liking.

We've had our eye on *Btop* and its previous Bash- and Python-based incarnations for some time, and while the previous iterations have been criticised by some as being slower than traditional utilities such as *Top* and alternatives such as *Htop*, our experience was far more pleasurable. The tweakable nature of the tool, and the use of colours, coupled with its ease of use and installation make *Btop* a wonderful CLI utility, and one we're confident most users would enjoy. **LXF**

Btop relies on colours to help you quickly identify the keybindings for the various boxes and buttons spread across the interface.

» CONFIGURING BTOP

You can change all tweakable options of *Btop* from its interface. Press Esc, and select Options or press **o**. The config is split across sections such as General, CPU, Mem, Net and Proc. Press Tab to go to the next tab, or click the button to access the relevant settings.

If you want to change the theme or colour scheme, or enable Vim keys to move around *Btop*, choose which boxes to display by default, and so on, head to the General tab of the Options screen. Some entries have a True/False setting. You'll notice right and left arrows when you select such an entry in the Options menu. You can use the right/left arrow key or click the little arrows to toggle the settings.

For other entries in the options screen, you'll notice an E symbol on the right when you select them. This means that you can edit the entry to write a custom value. To do so, select the entry in the list, and then either press **E** or click the E button. You can now type the new value for the selected entry. Hit Enter when you're done, and Esc to close the Options menu. *Btop* automatically saves the new settings in the `~/.config/btop/btop.conf` file.

» **ENHANCE YOUR TERMINAL-FU** Subscribe now at <http://bit.ly/LinuxFormat>

JAMI

Credit: <https://jami.net>

Secure and private comms made easy

Nick Peers takes a fresh look at an old favourite – Jami, a tool that aims to cover all your communication and sharing needs.



OUR EXPERT

Nick Peers would love to say he only uses open source, private messaging systems, but convincing others of the merits of *Signal* (never mind *Jami*) is a tough sell at times.

Looking for an alternative messaging and communications platform that isn't leakier than a colander and doesn't demand personally identifiable information when you sign up? You need *Jami* (<https://jami.net>). It may not attract the same levels of attention as the big boys, but reduced visibility usually means less unwelcome attention – helped by the fact that *Jami* runs on a distributed network, which means no central servers exist as tempting targets. Connections are peer-to-peer and protected by end-to-end encryption. You don't even need to provide any personal info to set up an account.

Traditionally, *Jami* has suffered a little in terms of performance and stability, but those issues are being ironed out, and with each new major release, it's becoming better and slicker. Back in November, *Jami* rolled out its latest major new release – Eleutheria – that resolves long-standing issues with connectivity and performance.

Since we last featured *Jami* (**LXF271**), it's undergone radical changes – the most notable of which is swarms, a new technology that ensures fully synced history across all devices you're logged into, plus guaranteed message delivery, even when there's a connection dropout. Since its introduction in September 2021, when swarms were restricted to one-to-one chats, it's expanded to be usable with groups of any size.

The new release unveils a new built-in extension store, too, along with numerous UI improvements, including the ability to edit messages after they've been sent. With all this added goodness, what better time to revisit *Jami* and see if it's the perfect fit for all your chat needs?

Getting in a Jami

Jami works on all major platforms, and that includes mobile (iOS, Android and e/OS via F-Droid). Visit <https://jami.net> and click Download to reveal links for Linux, Mac OS and Windows.

The Linux page reveals direct download links for major distros, including Ubuntu, Mint, Fedora and Red Hat, but we recommend clicking the >_ button instead

```
nick@nick-mint: ~
File Edit View Search Terminal Help
ng.gpg > /dev/null
nick@nick-mint:~$ sudo sh -c "echo 'deb [signed-by=/usr/share/keyrings/jami-archive-mi.net/nightly/ubuntu_22.04/ jami main' > /etc/apt/sources.list.d/jami.list"
nick@nick-mint:~$ sudo apt-get update && sudo apt-get install jami
Get:1 http://security.ubuntu.com/ubuntu jammy-security InRelease [110 kB]
Hit:2 http://archive.ubuntu.com/ubuntu jammy InRelease
Get:3 http://archive.ubuntu.com/ubuntu jammy-updates InRelease [119 kB]
Hit:4 http://archive.ubuntu.com/ubuntu jammy-backports InRelease
Ign:5 https://mirrors.ukfast.co.uk/sites/linuxmint.com/packages victoria InRelease
Hit:6 https://mirrors.ukfast.co.uk/sites/linuxmint.com/packages victoria Release
Get:7 https://dl.jami.net/nightly/ubuntu_22.04 jami InRelease [3,596 B]
Get:9 https://dl.jami.net/nightly/ubuntu_22.04 jami/main amd64 Packages [3,086 B]
Fetched 236 kB in 4s (54.2 kB/s)
Reading package lists... Done
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
The following additional packages will be installed:
  firebird3.0-common firebird3.0-common-doc jami-daemon libfbclient2 libjsnccp25
  libmysqldclient2 libodbc2 libpq5 libqt-jami libre2-9 libtommath1 libyaml-cpp0.7
Suggested packages:
```

We recommend installing Jami through its own repo – that way, you get the latest updates when they're rolled out.

to reveal the commands required to install *Jami* via Snap (`sudo snap install jami`) or its own repo, ensuring it remains up to date going forward. The following works on Ubuntu LTS-based distros, including Mint:

```
$ sudo apt install gnupg dirmngr ca-certificates curl
--no-install-recommends
$ curl -s https://dl.jami.net/public-key.gpg | sudo tee /
usr/share/keyrings/jami-archive-keyring.gpg > /dev/
null
$ sudo sh -c "echo 'deb [signed-by=/usr/share/
keyrings/jami-archive-keyring.gpg] https://dl.jami.net/
nightly/ubuntu_22.04/ jami main' > /etc/apt/sources.
list.d/jami.list"
$ sudo apt-get update && sudo apt-get install jami
```

Open the Jami

Once installed, open the app via your launcher (Mint users can find it under Internet). You're prompted to sign up for a new account or sign in with an existing account if you've already installed *Jami* on another device, such as your mobile.

Click Join Jami to start from scratch. You're told you can choose a username – up to 32 characters in length, and a mixture of alphanumeric characters and the dash symbol are supported. Although it's not explicitly clear, this is a purely optional step; your account is assigned a profile ID – a randomly generated

QUICK TIP

For a technical guide to *Jami*, visit the site's documentation at https://docs.jami.net/en_US/ – there's little practical explanation for how to use *Jami*, but lots of technological detail, including how *Jami*'s distributed network works to connect you securely to your contacts.

long string of hexadecimal code – that can be used to connect to others.

The profile ID is impossible to remember but can be copied to the clipboard or shared with friends in close proximity via a QR code. Unless you're desperate to remain anonymous, it's easier to generate a username. This can be obvious – John-Doe, for example – or you could make use of your password manager's username or passphrase generation capabilities to create something random such as Stable8064 or undergrad-guru-gloss (*Bitwarden* generated these examples for us). Your username or profile ID is the only piece of information you need to communicate to contacts.

If you decide to go with a more memorable username, *Jami* checks whether it's been used by anyone else as you enter it. This is the only piece of information *Jami* stores centrally; the rest of your account details only ever reside on any devices running *Jami* to which you connect your account.

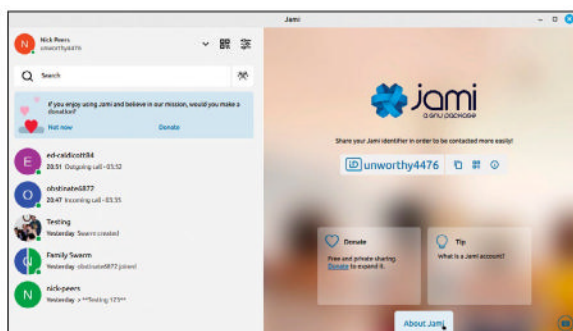
You'll see an Encrypt option beneath the Join *Jami* button. Click this and you're given the opportunity to encrypt your account with a password for additional security. The usual rules apply here, so your password manager's generator tool is a good place to set one up. Make sure you save that password somewhere secure – like said password manager – because if you forget the password, your account is toast. It's easy to create another account, but still a hassle if you then need to communicate your new details to all your contacts.

The Customize button to its right enables you to add an optional Display Name and photo. These details are only ever shared with people you add to your contacts list, so behind the obscure username you can provide a bit of identifiable information if you wish. Once you've set up your account as you like it, click the Join *Jami* button.

Spread Jami around

If you've already got an account set up, click I Already Have An Account. Two options appear, inviting you to import from another device or an archive backup. The step-by-step guide (page 52) reveals how to import your account to a new device using one that's logged in, while the archive option (see box, above-right) is a reminder of another way in which *Jami* distinguishes itself from the competition – you can't simply spread out to new devices using your username and password. As with a forgotten password, if you end up with *Jami* installed on no devices, the account is lost.

Once logged in, you find yourself at the main *Jami* screen, which has undergone quite the facelift over the



Jami's UI improvements make sharing your Jami username or profile ID with others much easier from the Linux client.


» BACK UP YOUR PROFILE

It's easy to lose your *Jami* username because it's not enough to know your username and password to restore your account to any device; you need physical access to that account from elsewhere, too. Lose this, and you lose your account (it's easy to set up another, but you have to start again with all your contacts, which probably entails contacting everyone to explain what happened and beg them to share their username or profile ID with you).

With this in mind, *Jami* offers an option that allows you to create a backup archive of your account, which doesn't just serve as a fail-safe way of getting your account resurrected, but can also be used in place of a direct connection to extend your account to a new device.

To create this archive in the Linux client, open the settings and scroll down the Manage Account screen to locate the Backup Account button you need to click. Choose a suitable backup location and give it a descriptive name, before clicking Save. If necessary, enter your account password (which has the bonus of encrypting said backup) and click Export. Then, when the file has appeared, click the X button to close the dialog.

Now, when needed, transfer the file to your new device and click I Already Have An Account, then choose the Import From An Archive Backup to add your account to the new device.



An archive backup can be used to add your account to a new device as well as restore a lost account.

past few years. If friends, family or colleagues have signed up, ask them for their username or profile ID, then type it into the Search box, where it should pop up under Search Results. Look for a tiny green icon appearing over their avatar, indicating they're online, which means they're hopefully ready and waiting for your message. Click their avatar to launch a conversation, which takes you to *Jami*'s chat screen.

Here you'll find all the tools you need to communicate and share with the person in question – the annotation (page 53) runs down what's available.

When you instigate a conversation like this, the person you've messaged is automatically added to your contacts list. However, that doesn't mean you're now connected – they receive a notification that invites them to accept, refuse or block your request. Refuse and block seem to do the same thing – you disappear from their contacts list, and you have no inkling that they're ignoring you, but you're free to try to contact them again (this prompts a new invitation).

If they accept, they're whisked to the conversation page and are free to engage you using the same tools – a mixture of text-based chat, file attachments, voice and video messages and live video and audio calls.

Jami attracts swarms

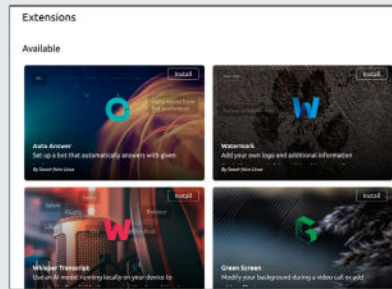
The technology underpinning *Jami* has evolved over the past few years, with conversations now based on swarms. These can encompass private one-to-one chats or be expanded to include as many people as you like. Swarms work in a similar way to one-to-one messaging; instead of searching for an individual, click Start Swarm next to the Search box and you're

QUICK TIP

If you encounter issues with video, navigate to Settings > Audio And Video > Video to enable hardware acceleration, which offloads the processing to your GPU if it's capable enough. This is also where you can choose which webcam to use if more than one is connected.

» ADD NEW TOOLS

Jami has long supported extensions that provide extra features. With the new Eleuthera release, plugging in new extensions got a lot easier thanks to a built-in extension store. Access is via the Settings button, next to your account name in the top-left corner of the desktop client. Simply select Extensions to reveal a small selection of popular plugins.



There are currently only five extensions available for Jami, but they're all worth investigating.

Auto Answer enables you to create custom 'I'm away' messages, while Watermark makes it possible to attach a logo, text and the current date/time to video messages you send. Whisper Transcript provides real-time subtitles in video calls, and Green Screen allows you to replace or blur out your messy office's background. The final offer is an audio reverb filter for anyone jamming online.

We expect the number of extensions you can install directly from here to grow over time; in the meantime, though, you'll find a handy Install button for quickly adding other extensions (supplied as JPL files) to Jami.

Once installed, look out for an Extension button in the call layout – clicking this should activate the plugin and enable you to use it.

prompted to give your swarm a name, description and – if you wish – a thumbnail photo. Click Create The Swarm and a new messaging window appears. Use the Add Participants button next to the Search button to select people from your contacts to add to the swarm.

These people aren't automatically joined to the swarm – as when you first made contact, they're given a separate invitation to join, so they're under no obligation to be a part of it. You can keep an eye on who's joined up and not by clicking the Details button in the top-right to reveal a slide-out panel letting you

view members, see which files have been shared in the swarm and access settings. Here you can change the icon colour, temporarily mute the conversation or choose Remove Conversation to leave the swarm.

You'll also see a Default Host (Calls) entry, currently set to None. This is important because it enables you to select which dedicated device will host all future calls in the swarm – leaving it as is puts the onus on the device launching the call to be able to handle it. As there are no longer any restrictions on the size of your swarm, you're limited basically by technical capabilities, so it's obviously a good idea to choose the most powerful device with the fastest connection – typically your desktop PC connected directly to the router via an Ethernet cable.


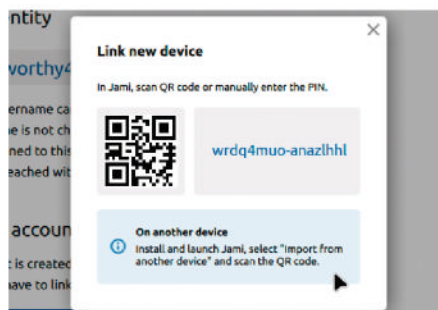
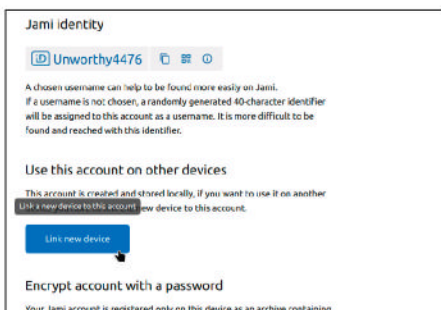
Video jamming

Jami works beautifully as a simple messaging platform, but it's equally capable at hosting everything from intimate one-to-one video and audio chats to full-blown conferences. You can place one-to-one calls by selecting a contact in the left-hand pane to bring up the messages window, then click either the phone button or camera button in the top-right to initiate an audio or video call. Your contact is notified you're trying to make contact, and if they accept the call, the call window pops up, showing either your contact's video feed (with your own video in a small window) or a blank screen if you're communicating via audio.

Roll your mouse over the screen and a row of buttons appears along the bottom of the main video feed. Most are self-explanatory – the three in the middle enable you to temporarily mute your mic (or choose a different input device if required), hang up the call, and disable or enable your webcam, which reveals that you don't need to hang up to switch between an audio-only call and a video call, just switch your webcam on or off as required.

There are seven further buttons on the right, providing additional tools and options – if the screen isn't wide enough, one or more of them may be hidden

LINK TO ANOTHER DEVICE



1 Instigate

You can't simply connect a device to an existing account using your username and password; you must introduce it via a device you're logged into. Instructions vary depending on your original device, but look under Account Settings for a Link New Device button or a Link Another Device link, then click or tap it.

2 Obtain PIN

If your account is password-protected, you're prompted to enter this before you can continue. Once done, Jami prepares to export your account before displaying an alphanumeric PIN and optional QR code, which you need to keep visible on screen while you set up Jami on your second device.

3 Make connection

Switch to your new device and choose the option to link an existing account or connect from another device. You should see options for entering the PIN manually or scanning in the QR code using your device's camera. Make your choice, then enter your account password again (if applicable) before connecting.

behind a 'more options' button (a vertical ellipsis) that, when clicked, displays them in a vertical list. From left to right, these controls enable you to switch audio output device, plus invite other people to your conversation (thus turning it into a conference call or group chat). Next to this is an option to show the chat window, which appears beneath the video feed if the *Jami* window is wide enough; otherwise, it overlays it (click the back window next to your chat partner or swarm's name to hide it).

Sharing Jami


Next to the chat window button is the Share Screen button – click the \wedge above it to reveal four sharing options: the entire desktop, a specific application window, a screen area (basically click and drag a rectangle to select part of the desktop), and files. The latter option opens and displays the file in question – if it's a video, it streams it.

The next button along enables you to change the screen layout. This only works when there's two or more people involved. It offers you options for hiding your own video feed, then displaying one person as the main view, with others shown in small thumbnails horizontally along the top or vertically down one side. There's also a Hide Spectators option for removing those who haven't got their cameras on from the feed to make room for others.

The final option enables you to record your chats as audio or video files. Click once to start recording – you should see the button flash – and again to stop. By default, *Jami* saves video streams (in WebM format) in a dedicated folder in your personal **Videos** directory, but you can change this, plus configure the quality of recordings by using the slider to push the bit rate all the way up to 5Mb/s (subject to connection speed and your PC's own hardware capabilities), and make *Jami* automatically record all calls by clicking the settings button and navigating to General > Call Recording.

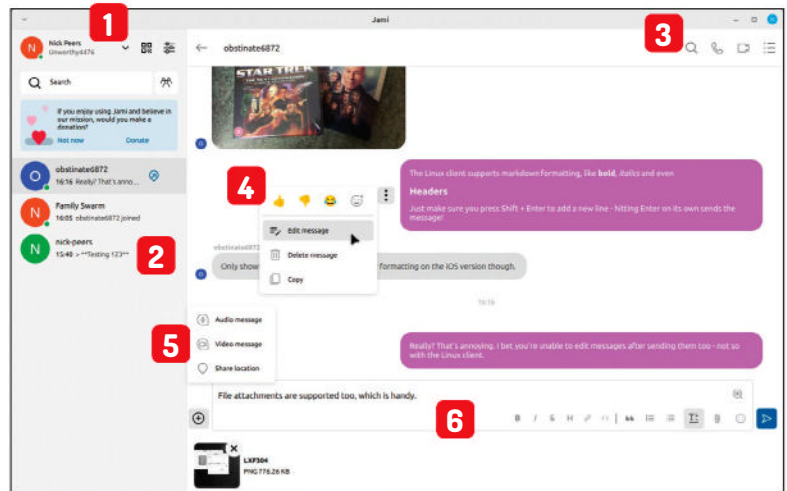
Group chat tips

We've already seen how you can change the layout when chatting with two or more people. When video chatting, you take on the role of host, and it's up to you



When participating in group chats, take the time to select what kind of layout you would like your audience to see.

CHAT WITH FRIENDS IN JAMI



1 Account details

From here, you can run multiple accounts from the same client or access your account settings.

2 Contacts list

All connected contacts and swarms can be viewed here in chronological order. Click one to switch to it.

3 Voice and video chats

Like all good messaging platforms, Jami enables you to take part in voice and video calls using your webcam.

4 Message window

Any shared text messages and files are accessible through the main chat window. Click on a message to edit it.

5 Add more content

You can leave voice and video messages if your contacts aren't online, plus share your location via a map.

6 Text entry box

Click inside this box to send text messages – you can format them, add file attachments, and more.

to mix and combine video feeds for everyone with the hope of minimising the bandwidth usage and preventing screen tearing and other glitches.

If you roll your mouse over someone's feed, you'll see a group of tiny buttons pop up, giving you various options. The crown button enables you to set (and unset) any participant as a moderator – they gain similar powers to yours, bringing up the same panel of buttons to give them control over part of the chat.

These buttons include being able to make a user 'full screen' – this simply enlarges their video to make them the focus of the chat, with other video feeds relegated to horizontal or vertical thumbnails across the top or right-hand edge of the video feed. You can also remove people from the chat by forcibly using the Hang Up button.

One final tip: as things stand, all group chats require one person to instigate the chat and invite people into it. If you'd rather set up a public conference room where contacts can join without specific invitations, set up a rendezvous point. You can temporarily convert your own account to a rendezvous point via Settings > Account > Call Settings, but a better idea is to set up a dedicated account for the job – see the Quick Tip (right) for details. **LXF**

QUICK TIP

You can create multiple identities if you wish and store them all on the same device, switching between them as necessary. After setting up your first account, click your username in the main *Jami* window and select Add Account+. Once done, use this menu to switch accounts as you require.

» **STAY IN CONTACT WITH US AND...** Subscribe now at <http://bit.ly/LinuxFormat>

Make a Mealie of it

David Rutland dons a tall hat to prepare you a portion of Mealie – a sophisticated recipe management server with a side of meal planning.



OUR EXPERT

David Rutland has finally moved on from his student days of living on jacket potatoes and baked beans. His friends and family are relieved.

Everybody eats. The chemical reactions upon which life depends break down complex molecules and release the energy needed to build other complex molecules, build bodies and build brains (*careful now!—ED*).

Contrary to certain historical beliefs, man can live on bread alone – but it's nutritionally better if it's a keema naan with extra pea flour. Even so, the finest curry house cuisine fails to inspire if you have the same thing every day.

So, it should come as no surprise to discover that recipes are some of the oldest examples of writing that still exist today. Visit the Yale Peabody Museum of Natural History, and you'll be able to read 25 recipes for stews and broth etched into stone tablets that have survived for four millennia – at least you would be able to if you could read ancient Sumerian cuneiform.

Recipes serve as both inspiration and a set of instructions – simultaneously solving the ever-present conundrum of what you can have for lunch, along with

```

david@vubuntu:~$ cd mealie/
david@vubuntu:~/mealie$
david@vubuntu:~/mealie$ docker-compose up -d
[+] Running 24/1
  mealie 23 layers [=====] 0B/
0B Pulled 179.0s
[+] Running 3/3
  Network mealie_default Created 0.1s
  Volume "mealie_mealie-data" Created 0.0s
  Container mealie Started 0.4s
david@vubuntu:~/mealie$ docker-compose ps
  
```

Docker Compose makes it easy to install Mealie. If only it made it easy to acquire the skills to cook properly.

» OUT AND ABOUT WITH MEALIE

Perusing recipes on your laptop is all very well when you're relaxing on the sofa, flicking through your own virtual *Mrs Beeton*. But you'd look silly cruising down the aisles of Lidl with a 15-inch clamshell perched on the child seat of your trolley. Thankfully, *Mealie*'s web interface is responsive, and adapts to suit any size of device.

Visit your recipe instance on a standard mobile phone, and the layout shifts so your ingredients list is near the top, and you'll notice that next to each item is a checkbox. This makes it super-simple to ensure you have everything you need before you leave the shop – tap the checkbox to mark that the required item is in your basket.

You can pin a *Mealie* browser shortcut to your mobile desktop, but sadly, there are no official apps available.

Various third-party developers have filled the gap on Android with apps that send a recipe from your browser directly to *Mealie* – which is a decent quality of life improvement. And there's the excellent *Mealient*. While it only has around 1,000 downloads at the time of writing, *Mealient* is worth a look, and gives you a mobile-optimised front-end that's perfect for using with one hand while you balance three boxes of wine and half a chicken in the other.

Despite a cursory inspection of Apple's App Store, we were unable to find anything for iOS.

the necessary follow-on queries of what do you need to make it, and how do you cook it?

The recipe book industry is huge, but for every celebrity or Michelin-starred chef's hardback, there are thousands of websites, giving away their step-by-step guides for free, along with a heart-wrenching origin story, and a bunch of sponsored links.

When you find a dish you like, you can bookmark it for later, save it as a PDF, or even print it off and add it to that dusty folder at the back of the bread bin.

If you lack the foresight to do that, you can also find dozens of recipes on the pages of your favourite supermarket's in-store magazine.

But data on the web is transient. The recipes dished up by a supermarket magazine – and hosted on a corporate website – aren't there for your convenience, or for long-term traffic, or even for SEO. They exist to drive foot traffic to a particular range of store products over a short period.

Recipes that featured on Asda's Good Living website at Easter will have left the site by Christmas, and there are other reasons why online recipes may become inaccessible. Maybe the site moved to a new URL structure, perhaps the owner lost interest and let the domain lapse. Sometimes – as in the case of <http://bbcgoodfood.com> – some members-only or so-called premium recipes are now hidden behind a paywall.

Sure, you can save every recipe as a PDF, but organising them is horrendous, and they contain reams of superfluous cruft. Printing recipes out takes up space – and without a master's degree in advanced librarianship, it can be even harder to locate an

ingredients list for the Persian pomegranate chicken you had on Valentine's Day.

Keep your recipes close...

Self-hosted recipe managers are more convenient than any of the options above. They're online, so accessible from anywhere, but because they're running on your hardware and in your home, you have control of both the content and the availability.

If you don't want the original chef's life story, you don't have to include it. Fancy adding an image of the lasagne you made with your niece instead of a stock photo? You can do that, too. And if you've made your own adjustments or amendments, you can adapt and rate the recipe, so you can refer to your own version later.

Best of all, it's not going to disappear unless you want it to, and your recipe collection will always be exactly where you left it.

There are several open source and self-hosted recipe managers available, and we think we've tried out all of them. While Nextcloud's *Cookbook* app is particularly praiseworthy – and certainly worth a try if you already have a Nextcloud instance but don't fancy spinning up another service right now – the best of these is *Mealie*, which offers a host of features and advantages beyond those we've listed so far, including multiple users, meal planning and a rating system, plus it's instantly easy to use.

Getting the ingredients

You can install *Mealie* on any computer that runs Linux, but we have a soft spot for using the Raspberry Pi range of SBCs as servers. A low-cost VPS is another good option.

You'll be connecting to your server (Pi or otherwise) over the internet, so the first thing to do is buy a domain name. It's not a huge investment – a .uk costs around a fiver per year. Cheaper domains are available, but they may renew at a higher price, so be careful.

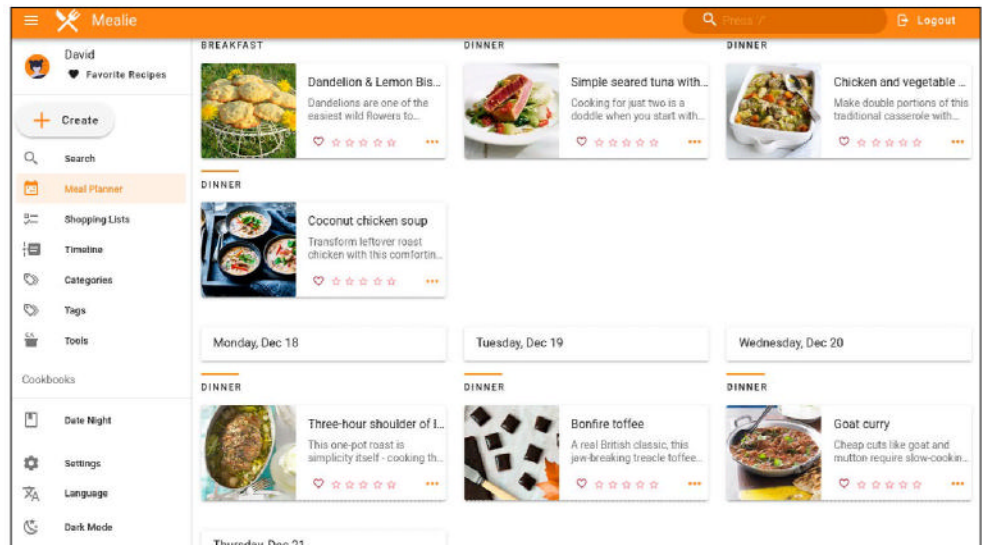
In your registrar's Advanced DNS section, create a new A record. Set the host as @, and the value as your domestic IP address (or VPS address). The TTL (time to lease) should be as low as possible.

Connect to your router using your web browser, and locate Port Forwarding or Port Mapping. Add one entry for port 80, and another for port 443 – these should forward to the local IP address of your server (see the *Pi-hole* feature in **LXF311** to learn how to give your server a static local IP address). Save the changes, then on your desktop computer, open a terminal and connect to the server using SSH (Secure Shell), and as usual, update and upgrade the system:

```
$ ssh username@server.local.ip.address
$ sudo apt update && sudo apt upgrade
```

The easiest way to install, update, and manage *Mealie* is with *Docker* and *Docker Compose*.

Docker is available in all of the standard distro repositories, and in the AUR. Install it using your



preferred package manager, then add your user to the *Docker* group:

```
$ sudo usermod -aG docker your_username
```

Log out and back in again for the changes to take effect. Visit the *Docker Compose* GitHub releases page (<https://github.com/docker/compose/releases>), and copy the link address for your server architecture type. In the terminal, use *Wget* to download it:

```
$ wget https://github.com/docker/compose/releases/download/v2.23.3/docker-compose-linux-aarch64
```

Now move it into your path and make it executable:

```
$ mv docker-compose* /usr/local/bin/docker-compose
$ sudo chmod +x /usr/local/bin/docker-compose
```

Create a directory for *Mealie*, move into it, and use *Nano* to create a new *Docker Compose* config:

```
$ mkdir mealie && cd mealie
$ nano docker-compose.yml
```

In the new file, paste the following:

```
version: "3.7"
services:
  mealie:
    image: ghcr.io/mealie-recipes/mealie:v1.0.0-RC1.1
    container_name: mealie
    ports:
      - "9927:9000"
    deploy:
      resources:
        limits:
          memory: 1000M
    volumes:
      - mealie-data:/app/data/
    environment:
      - ALLOW_SIGNUP=true
      - PUID=1000
      - PGID=1000
      - TZ=Europe/London
      - MAX_WORKERS=1
      - WEB_CONCURRENCY=1
      - BASE_URL=https://yourdomain.com
    restart: always
    volumes:
      mealie-data:
        driver: local
```

You should change the `BASE_URL` to your domain or subdomain. Save and exit *Nano* with Ctrl+O then

Using the random function in Mealie's Meal Planner can be a dice roll, but provided you have properly tagged your recipes, there shouldn't be too many surprises.

QUICK TIP

Mealie uses a calendar-like view to help you plan your meals; toggle through the calendar by clicking the arrows on the top of the page. In editor mode, you can use the random recipe buttons or manually add an entry.



QUICK TIP

Maintain a shopping list – you can add items directly to the list or link a recipe and all of its ingredients to track meals during the week.

Ctrl+X. Next, bring up *Docker Compose* with:

```
$ docker-compose up -d
```

You should be able to visit your *Mealie* instance in your browser at <http://your.server.local.ip:9927>.

And you can log in with the default credentials: **changeme@example.com** and MyPassword.

Click on the hamburger menu in the top-left, then Change Me, to set a proper username and password.

If you're only accessing *Mealie* over your local network, you can stop here; otherwise, install Apache along with some useful mods, and fire it up:

```
$ sudo apt install apache2
```

```
$ sudo a2enmod rewrite http2 proxy proxy_http proxy_https proxy_wstunnel
```

```
$ sudo systemctl start apache2
```

```
$ sudo systemctl enable apache2
```

You also need *Certbot* to easily manage your TLS:

```
$ sudo add-apt-repository ppa:certbot/certbot
```

```
$ sudo apt update
```

```
$ sudo apt-get install python3-certbot-apache
```

Create a new conf file for Apache:

```
$ cd /etc/apache2/sites-available/
```

```
$ nano mealie.conf
```

In the new file, enter:

```
<VirtualHost *:80>
```

```
ServerName you_domain.tld
```

```
ProxyPass "/socket" "ws://localhost:9927/socket"
```

```
ProxyPassReverse "/socket" "ws://localhost:9927/socket"
```

```
ProxyPass "/" "http://localhost:9927/"
```

```
ProxyPassReverse "/" "http://localhost:9927/"
```

```
</VirtualHost>
```

Save and exit, then enable the new site with:

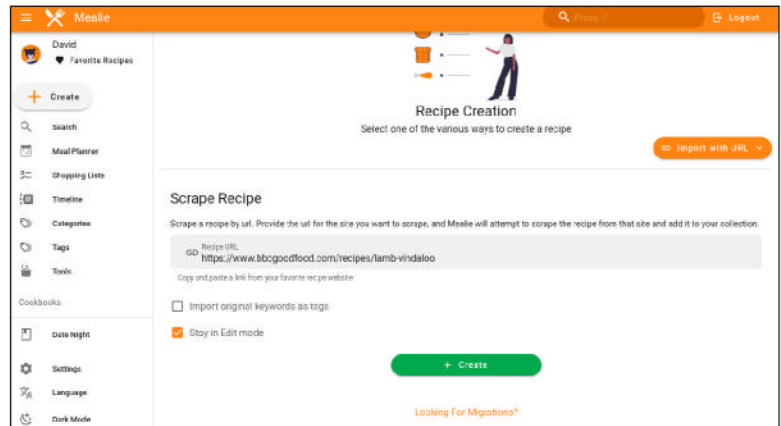
```
$ sudo a2ensite mealie.conf
```

Run *Certbot*:

```
$ sudo certbot
```

Fill in the quiz and select your site from a list.

Certbot automatically fetches and deploys security



You wouldn't download a curry, would you? Recipe scraping isn't piracy because recipes can't actually be copyrighted. The images and accompanying text can, though.

certificates and keys, before amending (or creating new) conf files.

You can now securely visit your *Mealie* instance using your very own domain name.

First, catch your rabbit

Mealie doesn't come with any recipes already included – which is a shame, because it's always nice to have new recommendations. Instead, you're greeted by an orange menu bar; buttons for categories, tags, tools and foods; and a search box.

We imagine that most of your go-to recipes come from a wide range of online sources, and *Mealie* can helpfully scrape the relevant data from any page that uses the standard recipe schema markup. This covers most recipe websites, because it makes recipes SEO friendly and more likely to rank higher on Google search results.

Some of our favourite fully compliant recipe hunting grounds include BBC Good Food, Slimming World and Hello Fresh. Other sites comply partially, but you may need to take additional steps.

Once you've found the recipe you want to save to *Mealie*, copy the URL to your clipboard, and head to your *Mealie* instance.

Click the hamburger menu, then Create > Import. Paste the URL into the box, and hit the big green Create button. Until you get used to the quirks of various recipe sites, it's worth checking Stay In Edit Mode. You can also import the original keywords as

» PLAN FOR THE FUTURE AND REVISIT PAST MEALS

There are two types of people in the world: those who plan their weekly food intake in advance, and make a single wallet-busting trip to the supermarket, and those who wing it, deciding what to feed a family of four on the fly, and finalising the day's menu just before leaving the house.

If you fall into the first category, *Mealie*'s there for

you with its excellent Meal Planner tool that helps you create menus for your day, week, month, or even year.

Click on Meal Planner in the side bar and select a date range. Press Edit, and for each day, you can select recipes for breakfast, lunch, dinner or a side dish. If you'd rather leave your culinary fate to the gods of chance, roll the

dice and have *Mealie* select random recipes for you.

A test of this feature means we're having dandelion and lemon biscuits for tomorrow's lunch, and coconut chicken soup for tea. Next Tuesday, our main meal is a delicious serving of bonfire toffee – which serves as an excellent demonstration of the importance of proper tag use.

Mealie's timeline offers another way to navigate your recipes, giving you a reminder of every recipe you've added, altered and rated. If you really need to remember what you had for dinner in order to recreate your tenth wedding anniversary, this feature can save your bacon – as well as your sausages, and your extra special liver and onions, too.

tags, although in our opinion, this can result in a cluttered interface.

In mere seconds, *Mealie* scrapes the desired page, and presents the complete recipe for a Black Forest fool in its own interface.

You should see the original featured image at the top of the page, followed by useful details including the recipe name, number of servings, prep time, cooking time, and a thankfully brief description taken from the metadata.

After this, you'll see an itemised list of ingredients and recipe steps.

It's worth paying attention here, because occasionally parts are missing. If this is the case, revisit the original website, and copy the relevant section. *Mealie*'s ingredients and steps sections come with a convenient Bulk Add tool, meaning that as long as each ingredient or step is on a new line, you can dump in the list and leave *Mealie* to sort it out. You may also want to separate or elaborate on steps, add tags, list specialist equipment or assign categories. Check everything again and click Save.

The other way to add a recipe is to create it manually. This may involve copying out from an old recipe card, using OCR software, or listing the steps for your own original creation.

To get started, click the hamburger menu, then Create > Create. Give your recipe a name, and fill in the fields. It's difficult to go wrong.

The last creation activity we're going to touch on is the cookbook. The search bar is fine if you're looking for something specific, and tags are a good way of organising your recipes and finding inspiration quickly. You type **salmon** to find all of the recipes with salmon in the title, for instance, and hit the Seafood tag to quickly view all your recipes with an aquatic ingredient.

Cookbooks enable you to gather certain tags, recipes and categories in one place. Creating a cookbook adds an entry to the sidebar, and all the recipes with the tags and categories chosen are displayed in the cookbook.

This is a cool feature that means you don't need a profusion of messy tags, and you can quickly find inspiration for a romantic meal from your Date Night cookbook without undue agonising or excessive searching. Click Create > Cookbook, and give your cookbook a name and description. Select the tags and categories, and you're good to go. It would be nice to be able to add individual recipes to the cookbook, but that's a feature for another day.

Get the family involved

Cooking, like eating, should be a social activity, and can be a miserable affair if you're in the kitchen on your own, tied to

the sink and slaving over a hot stove. Similarly, hoarding your culinary knowledge deprives others of potentially wonderful experiences.

A good way to get others involved is to give them an account on your *Mealie* server – that way, they can add and edit their own recipes, and it gives them a sense of ownership over what emerges from the steam behind the oven door.

Mealie offers user and group management, and you can afford different privilege levels to each.

To create a new group, go to Settings > Groups > Create. Give your group a name, then choose what rights the group has. You can restrict recipes to users in the same group, show or hide nutritional information, set the group to private, and prevent group members from commenting on recipes.

You can add a user through Settings > Users > Create. Choose which group they belong to, and dish out the rights they'll have on the server. You can choose to allow them to add new users, make them a group manager, or even a server admin. There's also Enable Advanced Features checkbox, although we haven't worked out what that does.

In recent releases, your *Mealie* instance requires you to log in to see any recipes, but you can, and perhaps should, open it up to the world.

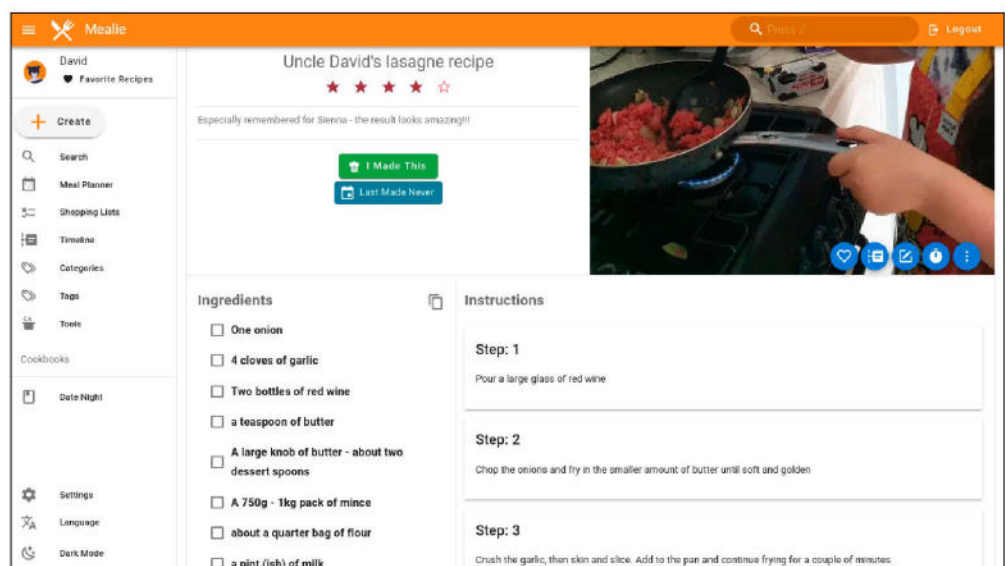
The proof of the pudding

After a while, you'll have thousands of recipes – some of which you actually made! How did it go? At the bottom of each *Mealie* recipe page is a comment field. Need more garlic next time? Add a note here. If your sister can't stand the smell of truffles, make a note of it for next time. If your Burns Night flaming Irn Bru trifle set fire to the curtains, you should probably make a note to have an extinguisher on hand next time.

For a quick reference to how good a recipe (or your execution) is, you can click on the five-star rating scale from any page on which the recipe is displayed to instantly rank it. Bone Apple Tea! **LXF**

QUICK TIP

Mealie's cookbooks are like a saved search results page – use them to group similar recipes together so you can quickly look them up later on.



Mealie is even better if you access the front-end through a kitchen tablet, and you can even snag passing children to help you out with the hard work!

» **DON'T MAKE A MEAL OF BUYING US** Subscribe now at <http://bit.ly/LinuxFormat>

Rescue and back up your retro media

Christian Cawley has some old Commodore 64 5.25-inch disks, and wants to revisit some personal files from 40 years ago.



**OUR
EXPERT**

Christian Cawley owns an original Commodore 64, an Amiga 1200, and a BBC Micro, but incredibly he doesn't have enough time to use them all. Who would have guessed?

Digging through some old boxes, this author found a shoebox full of 5.25-inch floppy disks. Back in the '80s and early '90s, having owned a Commodore 64, with a 1541-II disk drive. A lot of time was spent on this computer, with activities ranging from playing games (the usual) to creating very basic digital art and using software such as *Shoot-'Em-Up Construction Kit (SEUCK)* and *Graphic Adventure Creator*. The aim was to create something at least playable, although that wasn't always the case (if ever).

We were intrigued by the possibilities of these disks. What old gem or embarrassing archive would we find?

Why archive old disks?

It won't take long for you to find a collection of downloadable C64 disk images in the D64 format online. Various sites offer disk images and tape images for use in emulators. Of course, the legal status of these ROM files is dubious (see *boxout*, *opposite*).

But your own data is yours to archive as you see fit. Save game files, spreadsheet data, art, text documents – you can convert it into whatever format you like. Similarly, you can run public domain software on the original media and convert it to a new format.

Our own collection of disks is heavy on the personal backups of games with a fastloader. But among all these backups and save states are files we were curious to revisit. You might have a similar reason to convert your 5.25-inch physical media into disk images.

Do old disks work?

Perhaps the most significant roadblock in archiving old media is whether you can get the disk to work. Various challenges can prevent the volume from being opened successfully, and the further you go back through the generations of media, the greater the risk of failure.

When it comes to 5.25-inch magnetic disks, storage is paramount. Keeping the disks in a sealed box is preferable, and various issues can arise if this was not the case. Proximity to magnets or electromagnetic sources can damage disks, but be aware that a loss of magnetic orientation is something that can happen naturally over time.

Additionally, bent disks can be a problem, so avoid trying to load one. Flattening for a prolonged period of

```
atomickarma@atomic-pangolin: ~
atomickarma@atomic-pangolin:~$ sudo apt install libusb-1.0-0-dev
[sudo] password for atomickarma:
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
The following additional packages will be installed:
  libusb-1.0-doc
The following NEW packages will be installed:
  libusb-1.0-dev libusb-1.0-doc
0 to upgrade, 2 to newly install, 0 to remove and 0 not to upgrade.
Need to get 259 kB of archives.
After this operation, 1,940 kB of additional disk space will be used.
Do you want to continue? [Y/n]
```

A USB drivers library should be installed ahead of OpenCBM. This ensures the XUM1541 adaptor communicates over USB.

time (between heavy books, for example) can work, but it's a gamble.

Another risk is the presence of dust and dirt between the magnetic disk and the plastic outer sleeve. This can invade via the main opening or the notch, so keep that covered up. Inserting a dirty disk in a drive can lead to both the disk being scratched by dirt and the drive itself becoming damaged. And, of course, you have the risk of fingerprints on the opening.

Repairing a 40-year-old disk drive is not impossible, but it probably isn't a wise use of your time. While you can find replacements online (typically on eBay), they often prove to be expensive. Some are in excellent condition, whereas others are significantly cheaper because they need fixing.

Make some preliminary checks

You need a specialist piece of kit to connect a 1541-series disk drive to your Linux PC. (This includes the original Commodore 1541, the 1541-II and the 1570/71 models.) Before you do that, however, you should check whether the drive works.

This is a simple test. All you need is the Commodore drive, its power supply and a disk; we recommend using a disk you're not particularly fond of. Connect the PSU to the disk drive and power it on. Wait for the green LED, then insert the disk; flip down the locking mechanism to secure the disk in place. Whether a computer is connected or not, the red or amber LED should activate as the disk is detected and read.

Once the amber LED blinks off, eject the disk. It should pop out without any problems. If you

QUICK TIP

Power supplies for old equipment can be unreliable. New, modern PSUs and reconditioned supplies can be bought online, with an eBay-bought power block costing up to £40.

experienced any clicking, there may be an issue with the disk drive's read head, or the disk itself; various YouTube videos and websites can help with troubleshooting.

Find a suitable interface

Now it's time to get your own serial-to-USB connector. Commodore used a proprietary serial connection standard for its data cables, but a conversion method has been available since the 2000s. With the XUM1541 interface standard and *OpenCBM*, you can capture a disk image of the inserted C64 disk on a Linux system.

Better still, because the disks have a very low capacity (170Kb per side), the imaging process is almost instantaneous. In fact, the slowest part of the process is arguably loading and swapping the disks!

So, where do you find a XUM1541, or ZOOM1541, device? Your best bet is eBay but take care to buy from sellers with a good reputation and don't pay too much – we'd suggest £40 is about normal.


What is OpenCBM?

Suitable for use with the XU1541 adaptor, XUM1541 (you may find it listed as ZoomFloppy) and XA1541/XM1541 adaptors, *OpenCBM* is designed to access serial devices used on 8-bit Commodore computers. CBM is an acronym of Commodore Business Machines, the main subsidiary of Commodore International until the company went bust and its inventory and IP was picked up by Escom in 1995.

OpenCBM features a kernel module/driver that can be used with various other integrated tools to access devices such as the Commodore 15x1 disk drives. Additionally, it is a collection of tools that interact with Commodore drives, ranging from status checks to formatting tools and, of course, reading the disks.

You can install *OpenCBM* by adding a new repo or installing manually. Options for the latter option are available for Debian, OpenSUSE, Raspbian and Ubuntu. For example, the Ubuntu 23.04 steps are as follows:

```
$ echo 'deb http://download.opensuse.org/
repositories/home:strik/xUbuntu_23.04/' | sudo tee
/etc/apt/sources.list.d/home:strik.list
$ curl -fsSL https://download.opensuse.org/
repositories/home:strik/xUbuntu_23.04/Release.key |
gpg --dearmor | sudo tee /etc/apt/trusted.gpg.d/home_
strik.gpg > /dev/null
```



Long lost digital art (or whatever this is) might be found while creating disk images of old C64 disks.

```
$ sudo apt update
```

```
$ sudo apt install opencbm
```

Steps for the repo option can be found on the OpenSUSE-hosted *OpenCBM* project page: <https://bit.ly/LXFOpenSUSE> **OpenCBM**. (At some point it was possible to build from source, but the *OpenCBM* page now links to a dead tarball.)

With *OpenCBM* installed, you also need to install the driver for your chosen cable. These pages focus on the XUM1541 interface, which is installed with the libusb 1.0 packages:

```
$ sudo apt install libusb-1.0-0-dev
```

Now you're ready to connect everything together: your PC, the XUM1541-based cable/adaptor, the Commodore 1541 disk drive, and its power supply.

Detect the 1541 drive

The first thing to do with the disk drive powered on and connected is to open a terminal and detect the device

– `cbmctrl` can be used to first reset the device:

```
$ cbmctrl reset
```

You can then check the status of the disk drive with:

```
$ cbmctrl status 8
```

The **8** represents the assigned drive number. The 1541 and 1541-II disk drives will return `73,cbm dos v2.6 1541,00,00` whereas a 1571 will return `73,cbm dos v3.0 1571,00,00`. If the response matches what you see here, the disk drive can be considered to be working.

Create a disk image

A disk image is a single file that essentially clones the contents of physical media. If you own a C64 Mini, you may already be familiar with this concept; the Mini



Public domain disks featured a host of utilities, from disk copying tools to, um, search creation software.

QUICK TIP

Various Commodore 1541-compatible serial-to-USB interfaces are available to buy, also on eBay.

» IS ARCHIVING DISKS LEGAL?

If you're interested in getting data off old media, you're probably also aware that it can be a problem when it comes to copyright. While you're essentially making a backup for your own use, the laws concerning this differ around the world. So, what are your options?

Personal data, save game files, art files, word processing documents, spreadsheets, your own in-development programs and similar are all yours to do with as you please. As such, if you own such disks, you can comfortably make copies of them in D64 format on your Linux PC using the *OpenCBM* software and libusb driver.

You also have the option of public domain (PD) software. Typically distributed via bulletin boards on the nascent internet, PD software could also be picked up for the price of a disk in computer shops, or on the cover of publications such as *Commodore Disk User*.

PD software was a forerunner to open source, and its creators claimed zero ownership or intellectual property rights over the code.

As well as appearing in completed, published format, public domain software might also have been found in magazines and books as type-ins. So, if you own titles with pages of code to type up and run on your Commodore 64, and you saved these to disk, you can copy them to D64 format without any concern.

Should you want to type up some games, Usborne's classic 1980s computer books are available in PDF format – see <https://usborne.com/gb/books/computer-and-coding-books>.

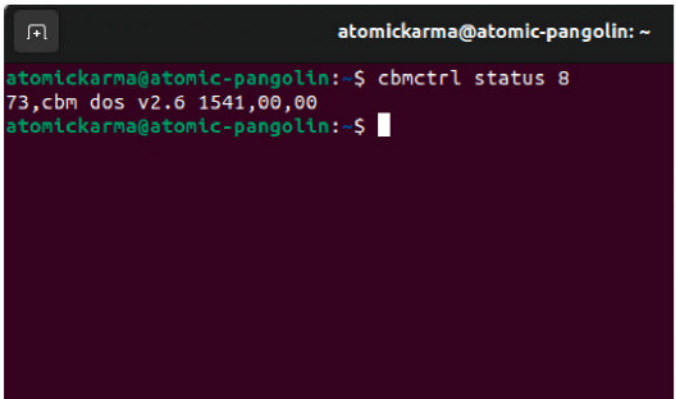


TURN A COMMODORE 5.25-INCH FLOPPY INTO A READABLE DISK IMAGE



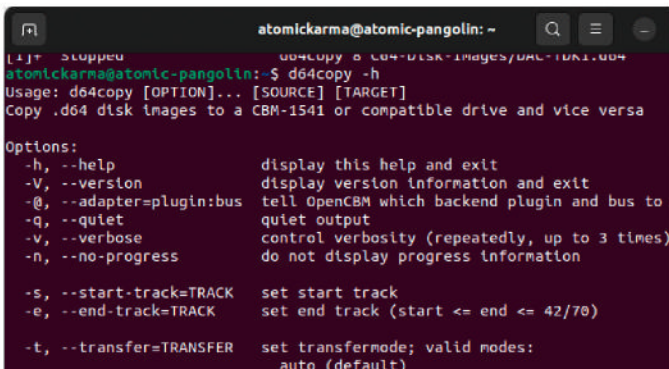
1 Hook up the disk drive

Your Commodore 1541-series disk drive should be connected to its PSU, which should be receiving power. The serial cable should be connected to the serial port on the disk drive, and to the XUM1541 (or variant) interface/adaptor. This in turn should be connected to your PC's USB port.



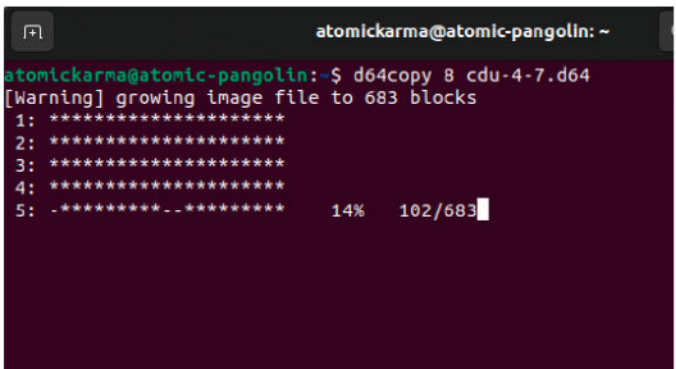
2 Power on and detect

After installing the *OpenCBM* software and XUM1541 driver on your Linux PC (see main article), power on the disk drive. Detect the device is receiving data by opening a terminal and resetting with `cbmctrl reset`. Next, check the status of the drive with `cbmctrl status 8` (correct output in image).




3 Prepare the disk image

You are now ready to create your first disk image. The simplest way to do this is by using the `d64copy` tool, although others (`cbmcopy` and `imgcopy`) are also included within the *OpenCBM* software. While they are used in a similar way, you should use `-h` or `--help` if you want to see the full list of command-specific options.



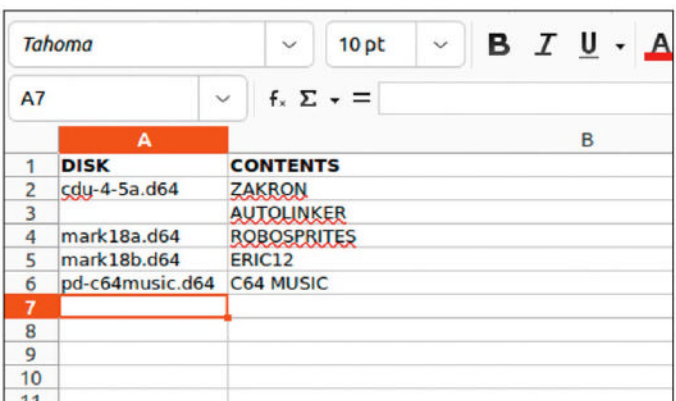
4 Create a disk image

Standard Commodore 64 floppy disks can be read and a disk image created on Linux. Simply input a command in the format `d64copy 8 image.d64`. Although the default command works in most cases, various options can be checked with `d64copy -h`. Remember to change `image.d64` to your preferred filename each time.



5 Launch the D64 image in VICE

Confirm the disk image is working by launching it in the Commodore 64 emulator *VICE*. The easiest way to do this is to attach the disk image. Go to **File > Attach Disk Image > Drive #8** (or #9 if preferred) to browse the disk contents, then use **Attach/Load** to run it.



6 Create an index

Keep a record of what is in each disk file. This could be done with a pen and paper, text editor, a database file or a spreadsheet. Giving each disk image a unique, relevant name will help in this process, and the contents of each image file should be listed alongside the name.

remake of the C64 features licensed titles saved as PRG files. Later firmware revisions include support for D64, as well as tape (TAP) and cartridge (CRT) images.

Creating a disk image of your 5.25-inch floppies is straightforward. The disk drive should be powered on, the XUM1541 or other interface correctly connected, and the 'ready' LED illuminated. A basic command uses the *d64copy* tool in *OpenCBM*:

```
$ d64copy 8 image.d64
```

This first calls the *d64copy* utility, specifies device 8 (your disk drive), and creates a file called **image.d64**.

For example, we want to make a disk image of a *Commodore Disk User* coverdisk from 1991, specifically Volume 4 Issue 7. We've given the disk image a name that corresponds with this for easy indexing (see below).

```
$ d64copy 8 cdu-4-7.d64
```

By default, the image is created in your **home** directory. A specific filepath can be set, for example:

```
$ d64copy 8 C64-Disk-Image/cdu-4-7.d64
```

Image creation should take a couple of minutes.

OpenCBM offers a wealth of options for the *d64copy* feature alone. With this, you can specify a start track (-s) and an end track (-e), set a transfer mode (-t) and if necessary disable the default warp mode (--no-warp), which slows down transfers but possibly makes them more accurate (which could be useful for an unreliable disk). You can also specify retries (--retry-count=n).

The *d64copy* tool also features support for double-sided disk reading from a 1571 drive. The option for this is -2, as follows:

```
$ d64copy -2 8 image.d64
```

Note that there is no local file protection; you can overwrite a disk image if you don't change the filename each time.

One thing of importance to note: if you created the disks using a fastload tool (for example, software included with The Final Cartridge and the Action Replay cart), these are more difficult to image. As such, try the *cbmcopy* and *imgcopy* tools also included in *OpenCBM*.

View disks in VICE

VICE (see [LXF267](#)) is an emulator for Commodore 64 and 128 computers (as well as PET, CBM-II, VIC-20 and Plus-4 systems). Not only does it emulate the C64, but you can use the disk attach tool to browse the contents of each disk image you make.

To do this, go to File > Attach Disk Image > Drive #8 and browse to the directory in which the disk images are stored. The image contents are listed in the right-hand pane.


When you find the disk image you want, click Attach/Load. You can also use the Autostart button if you don't want to bother with the disk loading command in the C64's main screen:

```
LOAD ""*,8,1
```

After that, it's a case of finding the right software to load the files. For example, in 1990 or so, we created an image in *OCP Art Studio*. In this scenario, you would need both the original application and your disk image

file, and frankly you're on your own with that.

One thing to note is that while the software may be running on your top-end Linux laptop, or even a 15-year-old Pentium 4 system, it will still be slow to load. Do not be impatient and rage-quit the emulator. In most cases, the contents of the disk image will load eventually. When they don't, you can tweak the CPU speed by clicking at the bottom of the emulator window and setting a new maximum speed. (It doesn't hurt to click the Warp button to speed things up, though...)



The very functional contents menu of a Commodore Disk User free coverdisk of public domain software.

Keep a record

When you pick up a 5.25-inch disk, you can generally tell what's on it. The label should detail the main programs or files stored on the medium; but what happens when you save the contents as a D64 file?

While you can view the contents from within *VICE*, it is a good idea to maintain a list of contents. For example, you might keep a library of photos, taking a snap of each disk on your phone. Our solution is to add disks to a spreadsheet; each entry has a label name that matches the disk image name, a list of contents, and a note as to whether the disk uses compression or archiving. You might follow suit, or organise disks by type: games, coverdisks, data disks, and so on. As long as you can easily reference each disk's contents, they can remain at your fingertips for years to come.

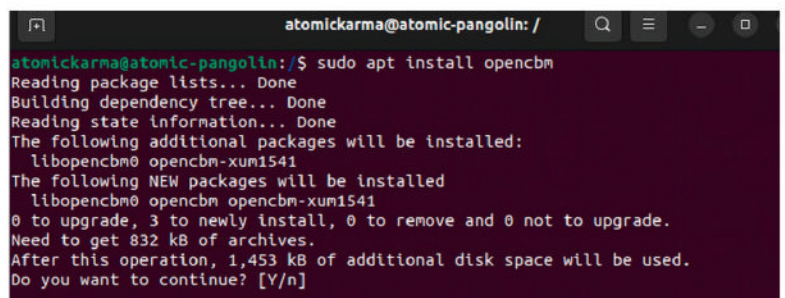
Using C64 disk images

What's the next step? Well, it really depends on what your intentions were when you started, and what sort of data you've found.

In our own case, we found graphic files, save game files, and even archives of old *SEUCK* data – that we're tempted to finish off. The brushes are certainly shareable, as are the sprites. This is really just the tip of the iceberg, though. You could share your creations online, copy assets to a USB stick and load them up in suitable software on a C64 Mini, or simply sit back, satisfied that whatever you were doing 40 years ago can be revisited whenever you want. Or, you might think about tackling those old Commodore 64 cassettes, although that's a whole different project. **LXF**

QUICK TIP

Disk imaging tools exist for most legacy computer systems. For example, the *Greasewearze* interface can be used with Amigas, BBC Micros, Amstrads and even MS-DOS floppy disks. Steps differ for each system, as does the firmware used, but they follow the same basic method.



Installing OpenCBM includes the opencbm-xum1541 driver, which ensures the disk drive communicates over your chosen XUM1541-based USB adaptor.

» **WE'RE ALREADY ARCHIVED, SO...** Subscribe now at <http://bit.ly/LinuxFormat>

BACK ISSUES » MISSED ONE?

ISSUE 311 February 2024

Product code:
LXFDB0311



In the magazine

Stay safe online by sending net nasties to *Pi-hole*, and read about the rise and fall (and rise again) of Basic, and the roll-out of the rolling-release Rhino Linux distro. You can also learn how to strengthen your shell history, set up a home CCTV system, add an inventory to your point-and-click adventure, and much more. Plus, check out our packed reviews, news and Raspberry Pi sections.

ISSUE 310 January 2024

Product code:
LXFDB0310



In the magazine

Learn Linux with our in-depth guide to getting started with Elementary OS, or explore a more security-centric distro as we dive into Tails. Plus, find out how to perfect the art of OCR, finish your *WordPress* website, recreate the BBC Micro, and make your own point-and-click adventure. And no issue of *Linux Format* would be complete without news, reviews, Pi projects and much more besides.

ISSUE 309 December 2023

Product code:
LXFDB0309



In the magazine

Manage your apps by harnessing the power of *Podman*, and discover how random numbers are used by the kernel to keep your data safe. Plus, find out how to master printing in Linux, customise your *WordPress* website, tweak your kernel scheduler, use Python to scrape the web, process your smart-home data, and much more. And make sure you read our *Roundup* of Raspberry Pi killers...

ISSUE 308 November 2023

Product code:
LXFDB0308




In the magazine

Bring together all your movie, TV, music, audiobook and photo streaming needs under one roof. Plus, we take an in-depth look at the Pop!_OS distro, explain how networking works, and guide you through tutorials on *WordPress*, fast file management, reviving the BBC Domesday Project and more. And don't forget to check out our packed reviews section, plus we examine the brand-new Raspberry Pi 5.

ISSUE 307 October 2023

Product code:
LXFDB0307



In the magazine

Power up your Pi with awesome AI projects, and take control of your phone with the privacy-focused /e/OS. Plus, find out which open source personal cloud server is best for you, read all about the forthcoming Wi-Fi 7, and discover how to transform images into GIFs via the terminal, manage a VPS, boost your Steam Deck, and much more. And there's our usual news, reviews and sysadmin tips.

ISSUE 306 September 2023

Product code:
LXFDB0306



In the magazine

Discover how AI can (and can't) help with your coding, plus add Pi-style GPIO to your PC, access your home server from anywhere, and write your own text adventure and *Breakout*-style games. Also, *Adminsteria* returns with news and reviews for sysadmin types, we take an in-depth look at how Azure stacks up against AWS, there's a *Roundup* of media-creation distros, and plenty more...

To order, visit **www.magazinesdirect.com**

Select **Single Issues** from the tab menu, then select **Linux Format**.

Or call the back issues hotline on **0330 333 1113**
or **+44 (0)330 333 1113** for overseas orders.

Quote the product code shown above and have your credit or debit card details ready.

USA? EU? THE MOON?

UK readers
turn to
p16

SUBSCRIBE!

Don't wait for the latest issue to reach your local store – subscribe today and let *Linux Format* fly straight to you. Faster, cheaper and with DRM-free archive access!



3 GREAT
WAYS TO
SUBSCRIBE

Print, digital-only,
and print+digital
bundles!

» USA
From \$125.99
For 13 issues

» REST OF THE WORLD
From \$125.99
For 13 issues

» EUROPE
From €110.99
For 13 issues

IT'S EASY TO SUBSCRIBE!

Visit www.magazinesdirect.com/linux-format

Call +44 0330 333 1113

Lines open Monday-Friday, 9am-5pm, UK time.

*We don't actually deliver to the Moon. Yet.

Create NPCs and their conversations

Nate Drake invites you to relive the glory days of point-and-click adventure games by creating your very own.



**OUR
EXPERT**

Nate Drake is a tech journalist specialising in cybersecurity and retro games. The first thing he did when discovering Linux in 2004 was play *Beneath a Steel Sky*. Be vigilant.

In the first part of this series, we covered setting up the first room of your point-and-click adventure, including the engine, palette and walkable areas. We also interacted with objects, by looking at them and adding them to the inventory. The player can now use the wrench inventory item on some machine pieces to assemble a broken robot.

With the robot up and running, it's time to give it a voice. Conversations are a natural part of any point-and-click game, providing useful hints. NPCs can also sometimes be persuaded to give players useful items.

Crafting characters

In the series so far, we've covered two of the main types of things with which your player can interact: items and hotspots. However, AGS also supports a third type: characters.

A character contains elements of both hotspots and objects. They can't be picked up but they can be interacted with to have conversations. They can also move between rooms and carry their own inventory.

Look to the project tree at the top-right and expand Characters. You'll see there the name of the main player character. By default, this is cRoger, but if you've been following this series, you may have changed his name to one of your own – cNate, for example.

Double-click on your character name. Under Selected Character Settings, you'll see a notification saying this is the player's character. This means when the game starts, you control this character and action begins in the character's StartingRoom, as defined in the Properties grid at the bottom-right.

This is where we can add our lovable robot as a character, but at this stage, there's a problem: currently, the fixed robot is an object. This is because in the previous tutorial, all we needed oFixedRobot to do was appear when the player used the wrench on the broken robot parts.

Luckily, this is easy to fix. First, return to Room1 in the project tree via double-clicking. Next, choose Edit This Room's > Room > Objects > oFixedRobot. Right-click and choose Delete. Click Yes to confirm.

Next, return to the project tree and right-click Characters. Choose New Character. Go to the properties grid and select RealName. This is the name

that the editor will use to refer to your new character, so set this to something simple like Robot.

Next, you need to set the ScriptName. This is the name by which the character will be referred to in the actual game script, so set this to cRobot (or cRobot1 if you want to have more than one robot in your game).

We need to change the sprite for this new character to resemble the robot. Click Sprites in the project tree to make sure you still have the Fixed Robot image. If not, download it from <https://bit.ly/lxf312robot>.

Next, in the project tree, right-click on Views > New View. In the new window that opens, check the box marked Show Preview, then on the button marked Create New Loop. From here, we could create an animated sprite with all the bells and whistles, but in this case, the robot is standing steady, so under Loop 0 just click Create New Frame.

This displays the default image of the blue cup. Double-click this to open the Select A Sprite window, then double-click again on the fixed robot image.

In the project tree, under Characters, click on cRobot. In the properties grid, scroll down and click the ... next to NormalView. From here, you can now select View2 in the left-hand pane, then click Use This View.

Roaming robots

While scrolling through the cRobot properties grid, you may have noticed there's a value for StartingRoom. By default, this is set to (None). You'll also see StartX and StartY properties to indicate the coordinates where the character should appear. (Remember, AGS makes this easy for game creators by displaying the current Mouse Coords at the top-left of the room editor.)

Usually, you would define a starting room and coordinates for your new character, but in this case, the fixed robot is only meant to appear when the player assembles the broken robot parts using the wrench.

Unlike objects, characters don't have a Visible property, so we have to be more inventive to keep our robot out of sight until it's needed. Go to the Room1 tab and click Edit This Room's > Objects > oBrokenRobot. Click on the Events button (lightning symbol) in the properties grid, then on the ... next to Interact Object.

As things stand, if you try to run the code as it currently is, it throws up an error because oFixedRobot

QUICK TIP

You can also change the StartingRoom value for your player character. As you add more rooms to your game, this can be a good way to test out new rooms quickly. We'll add another room in the final part of this series.

no longer exists. You can fix this firstly by removing the following line:

```
oFixedRobot.Visible = true;
```

Now add the line:

```
cRobot.ChangeRoom(1, 71, 137);
```

The **ChangeRoom** command is simple to follow. First, it specifies that cRobot's starting room is Room1. It then specifies the X and Y coordinates respectively of where to place cRobot. Feel free to tweak these.

Naturally, using this command isn't necessary if you want the character to be there as soon as the player enters a room – simply tweak the StartingRoom property accordingly and drag the character to where you want it to be using the mouse. We've included this here as characters sometimes only appear in rooms when certain conditions are fulfilled.

Chinwagging characters

Take a moment to press F5 and run the code so far. Click on the wrench to have your player pick it up, then select it from your inventory to use it on the robot parts. The robot character now appears.

Currently, if you click on the robot to interact with it, your player just exclaims, "Got nothing to say."

Different point-and-click adventures had varying approaches to conversations. LucasArts titles, such as *The Curse of Monkey Island*, employed a dialogue tree. This enabled you to choose subjects of conversation, as well as pick from a selection of responses.

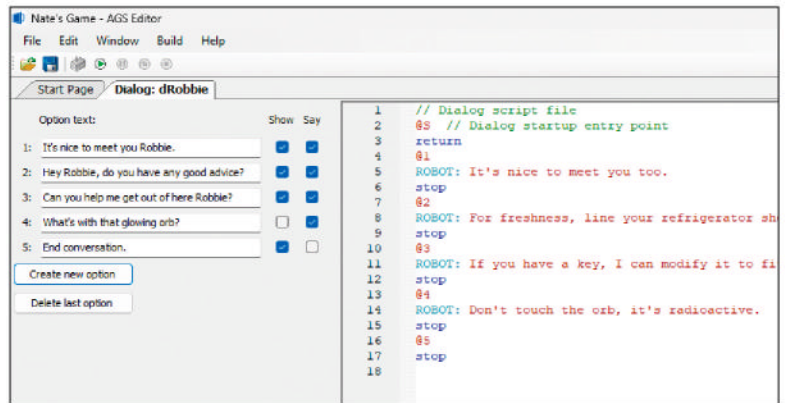
This makes for a much more interesting game, as it means you can have multiple, different conversations with the same character. Saying the right thing at the right time can even trigger in-game events, such as the character giving the player vital information or an item.

To get your robot talking, first return to the project tree and right-click Dialogs. Choose New Dialog.

Before we begin, it's important to establish some terminology. It helps to think of a dialogue as a topic of conversation. This is a set of options the player can talk about, though not all dialogues/topics are available at the very start of the game.

Each topic can contain some initial text, which is displayed before the options are presented to the player. When the player selects an option from a topic, the topic's dialogue script is run.

Let's start by giving this conversation a memorable name. Go to Name in the properties grid. We've used



dRobot but feel free to use dRobot1 if you plan to have multiple conversations with different robots.

We'll now give the player a conversation opener to use on the robot. Click the button marked Create New Option. For now, let's keep it simple, so add:

-What's your name?

Take a look at the **Dialog Script** file on the right-hand side. From here, you can specify what the characters say using very simple scripting language.

The **@S** indicates the start of the dialogue, so you can insert a line after this to have your player greet the other character, for example:

```
NATE: Hello!
```

Note here it's essential to use the ScriptName value of the character but without the **c**. For instance, **cRobot** would be **ROBOT**. Start a new line after **@1** and add a response for the robot:

```
ROBOT: Greetings! I am Robbie the Robot.  
stop
```

Note the **stop** command tells AGS to end the conversation and continue the game, whereas **return** simply goes back to the list of conversation options.

Now we just need to launch this conversation when the player tries to speak to the character. In the project tree, go to Characters > cRobot. Choose the Events button in the property grid, then click the ... next to Interact Character. (Note there is a Talk To Character option but this isn't available in the BASS interface we selected earlier.)

Insert the code **dRobot.Start();** to have AGS launch this conversation. Next press F5 to save and run the game so far. Once the robot is assembled, click on it to

Untick 'Show' to hide a conversation option. You can enable it later as the player performs certain actions.

QUICK TIP

Download a copy of all code used so far in this series from <https://github.com/azuregate/lxfpythontextadventure/tree/main/part3>. This includes the new Robbie Robot character and complete dialogue scripts. As before, we encourage you to create your own sprites rather than rely on Nate's terrible artwork.

SHOW AND TELL

When creating dialogue, you may have noticed that each time you click Create New Option there are two tick boxes marked Show and Say. By default, both are checked.

You can untick the Say box for options that, for whatever reason, you don't want the player or character to say. For instance, for protracted dialogue, you may want to add an End Conversation option they can click to stop the dialogue script.

If selected, the Show tick box immediately presents that dialogue option to the player when a conversation starts. However, most point-and-click games add topics when the player completes certain in-game tasks. To learn how to do this, open the dRobbie dialogue and add a new option asking "What's with that glowing orb?" Make sure Show is unticked.

In this case, we'll make this option appear only once the player has already inspected the glowing orb. Return to the Room1 tab and select the **hGlowingOrb** hotspot. Click Events, then on the ... by Look At Hotspot. Add the following line to function **hGlowingOrb_Look()**:

```
dRobbie.SetOptionState(4,eOptionOn);
```

The **SetOptionState** variable determines if a conversation option is available. In this case, **4** represents the fourth option in the **dRobbie** dialogue script, where the player asks about the glowing orb. The code uses **eOptionOn** to tell AGS that this conversation option should now be available after the player looks at the orb. Read more about dialogue at: <https://adventuregamestudio.github.io/ags-manual/Dialog.html>.



QUICK TIP

Adding a gifted item to an NPC's inventory isn't strictly necessary, as you can simply remove it from the player's inventory. Still, this can come in useful if you want the character to return the item at a later stage in the game.

say hello and ask its name.

Now you're familiar with questions and responses, exit the running game and return to the dRobot dialogue tab. Click Create New Option to add a couple more questions and responses, for example:

–Where am I?

–How do I get out of here?

Change the **stop** value to **return** in the dialogue script and add some answers of your own, such as:

```
// Dialog script file
@S // Dialog startup entry point
NATE: Hello!
return
@1
ROBOT: Greetings! I am Robbie the Robot.
return
@2
ROBOT: Somewhere in space.
return
@3
ROBOT: Just pick an exit and walk.
return
```

Nested conversations

In point-and-click adventures, you usually have to navigate conversations in order to obtain the right information or item. As things stand, the robot offers some very sassy answers to the final two questions but responds positively when you ask its name.

For the purposes of this tutorial, we'll assume this is the key to the player developing the conversation. Return to the project tree and right-click Dialogs once again to create a New Dialog. In the properties grid, give this a memorable name, such as dRobbie. Click Create New Option three times to give your player some conversation choices, such as:

–It's nice to meet you Robbie.

–Hey Robbie, do you have any good advice?

–Can you help me get out of here Robbie?

Edit the **Dialog Script** file to add some answers:

```
// Dialog script file
@S // Dialog startup entry point
return
@1
ROBOT: It's nice to meet you too.
```



In this version of the script, the player only needs to choose the correct conversation option to receive a gift from the NPC.

```
stop
@2
ROBOT: For freshness, line your refrigerator shelves
with press-and-seal wrap.
stop
@3
ROBOT: Take this master key. It should open any lock.
stop
```

In the project tree, double-click once again on dRobot under Dialogs. Modify the first response:

```
@1
ROBOT: Greetings! I am Robbie the Robot.
goto-dialog dRobbie
```

Now instead of simply returning the player to the original dRobot dialogue choices, the game invokes the dRobbie dialogue. Save and run the game by pressing F5 to test out your new dialogue script.

Dialogue functions

Now it's time to work some regular AGS script into dialogue to allow in-game events.

If you typed in the above code for the dRobbie dialogue script, you'll see that if the player asks for help, the robot tells them to take a master key. Currently, though, this does nothing in the script.

To get started with this new item, we first need to download a new image for the item the character is giving our player. In this case, use the sample image available from <https://github.com/azuregate/pointandclicklxf/blob/main/masterkey.png>. Once this is downloaded, double-click on Sprites in the project

FUN WITH FONTS

During your player's interaction with the character, you'll have seen that the speech for both the player and the NPC is written in the same red font.

To change the colour of a character's speech, head over to the project tree and under Characters, click on the one in question – cRobot, for example. Look at the properties box at the bottom-right and click the drop-down menu next to SpeechColor to pull up a palette.

To view AGS's default fonts, simply expand the Fonts option in the project tree. AGS has three built in:

Font 0 is the standard font used in message boxes.

Font 1 is the speech font, used with LucasArts-style speech.

Font 2 is an outline font for the speech font.

As speech in LucasArts-style games is drawn directly on screen, outline fonts can make the text easier to read, no matter what colour the screen is behind it.

AGS supports TrueType (TTF) and SCI fonts (Sierra's font format). SCI fonts are faster to render than TTF fonts but only support 128 characters, which means that they are only useful for English. There's no font value for character's speech, but you can change the speech font used by the game from within a dialogue script using the global **SetFont** variable, such as:

```
Game.SpeechFont = eFontFont2;
```

Add an indent to let AGS know this is regular script.

Visit www.adventuregamestudio.co.uk/wiki/Fonts for more information on working with fonts in AGS.

tree, then on Import New Sprite(s) From Files. From here, navigate to master key image and click Open.

Return to the project tree and right-click on Inventory Items to choose New Inventory Item. Look to the properties box. Insert a meaningful Description, such as Master Key, and a name, such as iMasterKey.

Click on CursorImage and Image, then on the ... to select the sprite for the new master key. Note that the value **PlayerStartsWithItem** is set to **False**, so the master key isn't in the inventory when the game starts.

While you're here, click on Events and then on Look At Inventory Item to add a useful description, such as:

```
Display("I should be unable to unlock any door with this.");
```

Return to the project tree to access the dRobbie dialogue or click on the relevant tab to see the dialogue script. We now need to modify the third dialogue item to have the master key appear in the player's inventory:

```
@3
ROBOT: Take this master key. It should open any lock.
player.AddInventory(iMasterKey);
NATE: Hey, thanks!
stop
```

Note the indent. This tells AGS that this is script, not dialogue. Remove the indent to return to conversation, as shown when the character thanks the robot.

Give and take

So far, we've devised a very trite dialogue between the player and an NPC whereupon they almost immediately give up a very valuable item. However, most point-and-click games require some give and take, whereby characters only give up items in exchange for items or some vital information.

To make this interaction more realistic, let's assume that instead of simply giving the player a new key the robot has the ability to modify the key already in their inventory to become a master key. Firstly, let's change the dialogue in dRobbie to inform the player of this:

```
@3
ROBOT: If you have a key, I can modify it to fit any lock.
stop
```

Return to the project tree and double-click on cRobot under Characters. In the properties grid, click Events, then on ... next to Use Inventory On Character.

```
function cRobot_UseInv()
{
if (player.ActiveInventory == iKey) {
player.LoseInventory(iKey);
cRobot.AddInventory(iKey);
player.AddInventory(iMasterKey);
cRobot.Say("Working...Check your inventory for your master key!");
}
}
```

Most of the code here is familiar. Firstly, the script checks that the inventory item in question is iKey, and if so, it's taken away from the player's own inventory. As cRobot is a character, the key is added to theirs.

Next, the iMasterKey inventory item is added to your player's inventory. The Say command simply has the robot character spontaneously speak without entering a dialogue tree.

Press F5 to test the code and you'll probably notice a flaw, as the player can present their key to the robot and swap it for a master key without talking to Robbie.

Fortunately, AGS has a number of functions and properties related to dialogue scripts. One of these is **HasOptionBeenChosen(X)**. This checks that the player has chosen a certain numbered option (X) in a script.

We need to be sure the player has chosen the third option in the dRobbie script. We do this by changing the line **function cRobot_UseInv()** above from:

```
if (player.ActiveInventory == iKey) {
to:
if (player.ActiveInventory == iKey && dRobbie.
HasOptionBeenChosen(3)) {
```

The **&&** operator represents **and** in the **if** statement. You'll also see the statement is wrapped in brackets.

Press F5 to run the game and assemble the robot. If you now try to use your key on the robot, nothing happens. Talk to Robbie and ask if he can help. You can now swap out your regular key for a master key.

You can also use this to modify the dialogue, so you don't have to greet the robot and ask its name every time you talk to it. Return to cRobot under Character. Click Events and click to modify Interact Character:

```
function cRobot_Interact()
{
if (dRobot.HasOptionBeenChosen(1)) {
dRobbie.Start();
}
else
{
dRobot.Start();
}}
```

This checks if the player has asked the robot its name in dRobot. If so, it jumps to the dRobbie script.

Stay tuned

You now have a working room containing objects and hotspots with which you can interact. The game also supports conversation with a character, from whom you can get info and items, depending on your actions.

In the next and final part of the series, we'll explore advanced room functions, such as loading new rooms, as well as how to export and save your game. **LXF**

The player can now give their key to Robbie in exchange for a master key, but only if they've said the right things.



» **ADD LXF TO YOUR INVENTORY!** Subscribe now at <http://bit.ly/LinuxFormat>

Emulate an analogue computer digitally

In our continuing journey into analogue computing, **Mike Bedford** looks at another emulation approach, and how to get a real analogue computer.



OUR EXPERT

Mike Bedford is fascinated to see how analogue computers of the '60s and '70s can demonstrate chaotic systems, even though chaos theory wasn't popularised until much later.

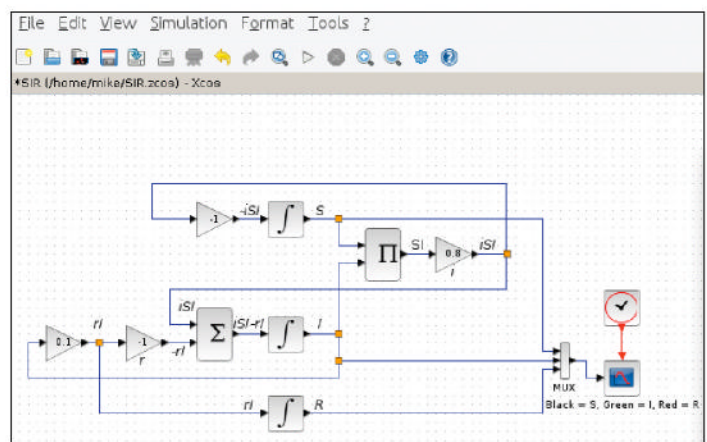
As we saw last month, although analogue computers might seem ancient, they gave digital computers a run for their money until the early '80s. Their particular strength was solving differential equations, so they were used for simulation, primarily in scientific and technical applications.

In the first part of our exposé of electronic analogue computers, we turned to emulation to see how they worked and how they were used. Specifically, we emulated a simple analogue machine – not too dissimilar from the ones that were used in education – but that approach limited the types of problem that could be solved. So, now we're going to see how to emulate a more sophisticated machine, to allow us to use analogue techniques to solve some more complicated and more interesting problems.

And finally, for those who want to try out some real analogue hardware, as opposed to emulating analogue computers digitally, we'll introduce some exercises to help you better understand the underlying electronics, and even point you towards a real open hardware machine that you can buy.

A different approach

The emulator that we used last month was built using the *Falstad* circuit simulator, and this served to provide some insight into analogue computing hardware. The same approach could have been extended to create a larger machine with more potentiometers, summers and integrators. But while this would have allowed problems involving more differential equations to be solved, it would still have been limited. In particular, while it allowed a variable to be multiplied by a constant – this being the function of a potentiometer – it couldn't multiply two variables. And while lots of problems, like the examples in part one, don't need a multiplier, many others do. We could have used *Falstad* to add a multiplier, but while the circuits for potentiometers, summers and integrators are pretty simple, the same isn't true of an analogue multiplier. If you're really keen, you could research the subject and



The Susceptible – Infected – Recovered model of an infectious disease, implemented in Xcos, requires multiplication, which wasn't available in last month's Falstad emulator.

try expanding our *Falstad*-based emulator, but we're going to pursue a different and simpler way forward.

Our approach maintains the concept of programming by connecting various functional units together – after all, nothing less can be considered as emulating the analogue computing paradigm – but we are making some changes. Most fundamentally, having emulated an analogue computer's patch panel last month, we're abandoning that here. Instead, we're wiring functional blocks together on screen to produce something that looks closer to the wiring diagrams that define how a problem is solved. Essentially, therefore, it eliminates the need to figure out the patching from the wiring diagram, but that's not really fundamental to the analogue approach.

Emulation in Scilab

The software we're using is *Scilab* (www.scilab.org), which is a numerical computation framework and associated programming language. However, we are not programming in that language but, instead, we're using the associated Xcos package, which allows functional blocks to be wired up on screen. This is clearly a good analogy to analogue computer programming. And because Xcos includes high-level blocks such as integrators, it doesn't require these functions to be created by wiring together electronic

components, as we did in creating the *Falstad*-based emulator last month.


This isn't to say that the blocks provided in *Xcos* are the same as the ones in an analogue computer. For example, the integrator in *Xcos* is exactly that – it integrates a single input variable. An analogue computer's integrator, on the other hand, integrates the sum of its gain 1 and gain 10 inputs, and inverts the output. *Xcos* allows a group of blocks and their interconnections to be converted to a so-called superblock for future use. This way, it would be possible to create new blocks that better emulate an analogue computer's summers and integrators.

The main *Xcos* blocks we're using are the integrator, summer, gain, constant and multiply. If you need an analogue computer integrator, use an *Xcos* integrator, although if you're connecting up a wiring diagram for a real analogue computer program, you might have to prefix it with a summer, which might itself have to be prefixed by one or more gain blocks configured for a gain of 10. You could also route its output via a gain block with its gain set to -1 to invert the output, but in many cases, it's easier to diverge somewhat from a true analogue computer wiring diagram by using inverters only where you actually need them.

Unlike an analogue computer's integrator, an *Xcos* integrator doesn't have an initial condition input; instead, it's defined by one of its parameters (double-click on a block to see and change its parameters). Parameters also dictate the upper and lower limits for the output – the equivalent of the +15V and -15V of our *Falstad* emulator, although you can often use the default values of +1 and -1. For an analogue computer summer, use an *Xcos* summer, although you might need to feed some of its inputs via gain blocks set to 10. Again, we recommend using a following gain block as an inverter only if you genuinely need that.

Use an *Xcos* gain block set to a value between 0 and 1 for a potentiometer if it's connected between the output of one summer or integrator and the input of another. However, when you need to emulate a potentiometer connected to a supply voltage to provide a constant, use an *Xcos* constant block. And use an *Xcos* multiplier block as an analogue computer multiplier – what could be easier?

So, with that introduction, we're ready to program our first problem in what we could refer to as our



analogue computer-ish emulator. The model we're going to patch up simulates the progress of an epidemic. It's a simple model that assumes perfect mixing of people in the population, and that when people have recovered they're no longer infectious and are immune from re-infection. It defines the number of people who are susceptible (S), the number of people who are infected and hence infectious (I), and the number of people who have recovered (R), according to the following differential equations:

$$\frac{dS}{dt} = -iSI$$

$$\frac{dI}{dt} = iSI - rI$$

$$\frac{dR}{dt} = rI$$

You'll see that the model includes the term SI , which requires analogue multiplication, so we couldn't have solved it using the *Falstad* emulator. There are two constants: i , which is inversely proportional to the time taken to become ill once exposed to the disease, and r , which is inversely proportional to the recovery time. You can see exactly how to wire it up in the screenshot (*opposite*). What you do need to know, however, is that the simulation starts off with almost the complete population susceptible and just a few infected. So, you should set the initial condition of the S integrator to 0.9 (if you accepted the default limits for all the integrators of +1 and -1), set the initial condition of the I integrator to 0.1, and use the default initial condition of 0 for the R integrator. Of course, the actual population isn't actually 1 – we're normalising the value to the range 0–1, as mathematicians would say. This means the

Analogue meets quantum in D-Wave's Advantage quantum annealing computer.

QUICK TIP

Functional blocks that were available in some analogue computers but that we don't use in our sample exercises include dividers and diode-based function generators. However, these and more can easily be used in *Xcos*, and they'll allow you to solve otherwise unsolvable problems.

» 21ST CENTURY ANALOGUE COMPUTING

Surprisingly, you might think, given that it largely died out four decades ago, research into analogue computing is alive and well. It's one of several technologies being researched to give digital computing an extra lease of life when year-on-year performance benefits come to the end of the line.

Modern analogue computing wouldn't aim

to replace digital computers, though, but to work hand-in-hand alongside them. As such, they'd be more akin to the hybrid computers of the '70s and '80s that we looked at last month. However, the digital computer would go much further in streamlining the analogue approach. As well as controlling the analogue computer and reading its outputs, digital

technology would work out what analogue building blocks were needed to solve a problem and would even wire them up automatically.

Using analogue hardware in this way can be likened to using a coprocessor, like the floating-point coprocessors of old, or today's use of GPUs for heavy number crunching. But while an analogue coprocessor could provide a

significant performance boost for some types of exercise, the inherent accuracy limitations of analogue circuitry would remain. But here again, the collaborative approach could reap benefits. The analogue hardware could get close to the solution before handing over to the digital hardware, which would home in on a more accurate result.



QUICK TIP

An embedded computer is generally thought of as a digital innovation. However, the first electronic analogue computers were designed for exactly this sort of task. Possibly the first such machine controlled the 1944 V-2 rocket, which was the first ever long-range guided ballistic missile.

actual population, and the values of S, I and R, are the normalised values multiplied by the actual population.

Our next example is one that you might like to try yourself, and although it's not significantly more complicated than the epidemic exercise, it gives much more interesting results. Just in case you run into difficulties, we're providing a screenshot (*far-right*), but we recommend that you try it out first without cheating. The model is defined by the following equations:


$$\frac{dx}{dt} = -y - z$$

$$\frac{dy}{dt} = x + ay$$

$$\frac{dz}{dt} = b + z(x - c)$$

You don't need to know much about these equations except that they define the Rössler attractor, that x, y and z are the coordinates in three-dimensional space, and that a, b and c are three constants that should be set to 0.1, 0.1 and 14 respectively. And although we're leaving you to figure out the wiring diagram, we will say that you should use limits of +25 and -25 for the x and y integrators, and limits of +50 and 0 for the z integrator.

The initial conditions are fairly non-critical, although we can say that using 10 for all three integrators is OK. You can plot any or all of the variables against time, and they certainly don't vary in a perfectly repeatable pattern, but to see the classic view of the Rössler attractor, you should also use an Xcos 3D graph (CSCOPXY3D), which enables you to view and manipulate the results in three-dimensional space. Finally, we recommend that you extend the simulation time (Simulation > Setup > Final Integration Time) from its default to 500. If you're not familiar with chaotic attractors, or chaotic systems more generally, you might like to read up on the subject, but even if you don't feel like delving into the maths, we trust you'll recognise that the Rössler attractor is an excellent example of how a simple analogue computer setup can



The Analog Thing (THAT) offers access to real analogue computing for little outlay.

result in some truly complicated behaviour.

Real hardware

If using an analogue computer emulator is good, surely using real analogue hardware is even better. And the good news is that it's easy to get a feel for analogue computer circuits. First, let's build a summer and put it through its paces, then do the same with an integrator. Although they're not parts of the summer or integrator circuits, to do anything useful with them you need to wire them to a couple of potentiometers to provide variable voltage inputs.

You need a small breadboard, a handful of patch leads, and a test meter – but these can be bought for as little as £10 in total – plus the few cheap components that are shown on the schematic (*bottom-right*). The components are all fairly non-critical but do bear in mind the following comments. Op-amp chips can be either through-hole or surface mounting – you need a through-hole one. Any eight-pin dual op-amp chip other than the LM358 we used will almost certainly work so long as its maximum supply voltages are high enough, but most support the necessary +9V and -9V. Also check the pin numbers – nearly all dual op-amp chips are the same, but if not you'll have to rearrange the breadboard wiring.

Make sure you place the chips the right way round – note the dot adjacent to pin 1 and the semicircular indentation between pins 1 and 8. For a proper analogue computer, the choice of capacitor type would be important, but for this experiment, any capacitor type will work so long as it's not polarised – that is, not electrolytic or tantalum. The only snag is that with some capacitors, the integrator output might change slowly, even with a 0V input.

Regarding the potentiometers, types suitable for breadboard mounting are available but you could use ordinary ones by soldering them to short wire leads that can then be plugged into the breadboard. If you use an ordinary panel-mounting potentiometer, follow our diagram to ensure you use the correct terminals; if

» ANALOGUE GOES QUANTUM

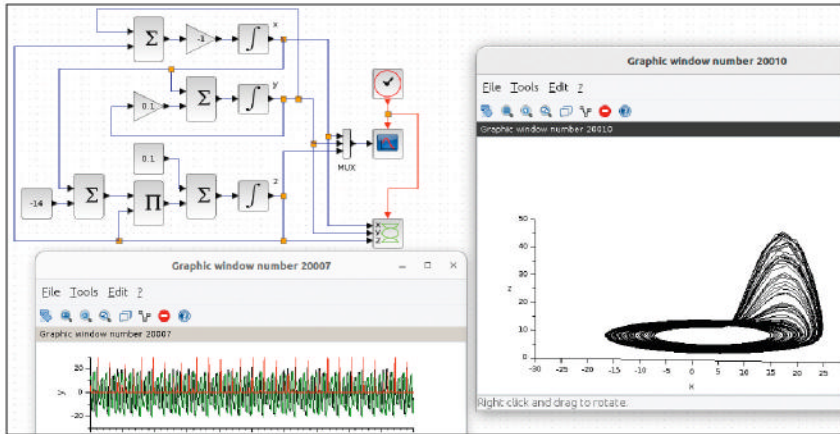
D-Wave Systems was the first company to offer a quantum computer commercially. If you've never delved into the subject, be warned that it's strange in the extreme, as quantum effects just don't correspond to our common sense view of the world. For example, quantum computers have quantum bits, or qubits, which can store a zero and a one simultaneously. A qubit

remains in that state until you try to read its value, at which point it adopts a value of zero or one at random. And if that's not strange enough, quantum computers also rely on an effect called entanglement, which is so odd that even Einstein referred to it as "spooky action at a distance".

The holy grail of most quantum computing research

is to build a practical general-purpose machine, one that can carry out any computable problem, like today's digital computers can, but almost unbelievably faster. This isn't what D-Wave Systems has developed, though, as it clarifies: "D-Wave Systems uses a process called quantum annealing to search for solutions to a problem. In nature, physical systems tend

to evolve toward their lowest energy state: objects slide down hills, hot things cool down, and so on. This behaviour also applies to quantum systems. To imagine this, think of a traveller looking for the best solution by finding the lowest valley in the energy landscape that represents the problem." This has been described as an analogue system.



This model, comprising just three integrators, three summers and a multiplier, plus a few gain blocks and a constant block, illustrates something of the power of analogue computing.

you use a breadboard or PCB-mounting type, consult the datasheet. Ideally use linear potentiometers rather than logarithmic ones, so that their rotary positions more closely correspond to their outputs.

Not everyone has used breadboards before so, especially since this tutorial is primarily concerned with software – you can get the diagram in the code pack at www.linuxformat.com/archives?issue=312. The same diagram can also be used for the integrator, so long as you remove the 100kΩ resistor marked R1, and put a 100nF capacitor in its place, and also replace the two 100kΩ resistors, R2 and R3, with 1MΩ resistors.

Unlike the summers and integrators that we emulated last month, our circuits here only have two inputs, both with a gain of 1, but that's sufficient to try out some hardware because we don't recommend you build on this to create a real analogue computer. Also, we're simplifying the integrator by not allowing an initial condition to be set, which means the output is always 0V when you first apply power.

Wire up the circuit as shown in the diagram, but don't connect the batteries while you check your wiring. Now connect the batteries and check that nothing's getting hot, especially the op-amp. If anything is overheating, this suggests a wiring error, so immediately remove the battery and correct the fault before trying again.

Now turn on the test meter and set it to measure voltage. Adjust both potentiometers to fully anticlockwise. You should see 0V because the circuit is adding together the two 0V outputs from the potentiometers. Adjust one or both of the potentiometers and you should see an output that is a negated sum of the two potentiometer outputs. Bear in mind, though, that the op-amp output will never be higher or lower than the +9V and -9V supplies, so you'll get the same result with both potentiometers fully clockwise that you'd get with only one fully clockwise.

When you've put the summer through its paces, disconnect the batteries and convert the circuit to an integrator, as we've already described. Go through the same procedure that you used with the summer, but this time, when you apply a voltage via one or more of the potentiometers, the op-amp's output should

ramp down to the minimum value, at a speed that is proportional to the sum of the potentiometer outputs. You could also modify the circuit by rewiring one of the potentiometers to -9V instead of +9V. Now, an output from one of the potentiometers causes the op-amp output to ramp down as before, but the other causes it to ramp up.

And finally, if this introduction has whetted your appetite, here's how you could try out a real analogue computer, as

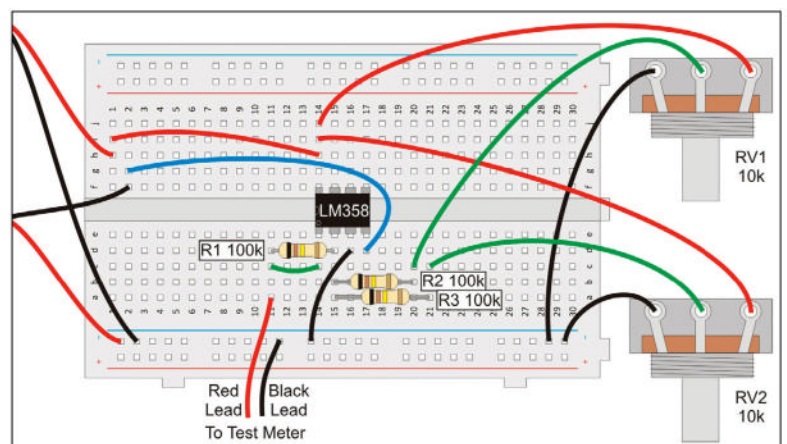
opposed to emulating one or tinkering with simple summer and integrator circuits. Our recommended solution is to buy a THAT (The Analog Thing), which is a newly designed open source, not-for profit analogue computer (<https://the-analog-thing.org>). THAT has four summers, five integrators, eight potentiometers, two multipliers and two comparators. It also has a panel meter that you can use to monitor the outputs of any of the functional units, although you'll probably want to use it with an oscilloscope to view results graphically. Alternatively, Raspberry Pi oscilloscope projects abound, and would provide a low-cost alternative – see, for example, <https://github.com/fhdm-dev/scopy>.

THAT costs €503.20, or €453.78 for educational institutions. At first sight, that might sound a lot, especially for something that's just a curiosity, but don't be too hasty to write it off. Back in the early '80s, home computers costing around £100 first appeared, and so the home computing boom was born. Those machines, starting with the Sinclair ZX80, were really just curiosities, too, and when inflation is taken into account, they cost only slightly less than THAT. Alternatively, thanks to the open hardware philosophy, if you feel confident to source your own components and to wield a soldering iron, you could even build your own. What more excuse do you need to jump on board an analogue revival? **LXF**

QUICK TIP

Who said computer music is a digital phenomenon? *Acid in the Style of David Tudor*, the 2009 album by Florian Hecker, was created using the combination of a Buchla modular analogue synthesiser and a Comdyna GP-6 analogue computer. To be fair, though, we're tempted to adapt the famously misquoted *Star Trek* line as, "It's music Jim, but not as we know it."

Copying this simple breadboard arrangement is the easiest way of building our summer and integrator circuits.



» TURN YOUR LINUX SKILLS UP TO 11! Subscribe now at <http://bit.ly/LinuxFormat>

ETHERNET

The wire to faster Ethernet

Searching for faster than Gigabit Ethernet, the ever frugal **Neil Mohr** does that compromise thing his ex told him about.



**OUR
EXPERT**

Neil Mohr won't spend a penny on anything new and fangled – his children rarely visit.

Ethernet is 50 years old. We ran a feature wondering why it hadn't died off back in **LXF310**, and the simple answer is that it just gets faster. Most home users have been pootling around with Gigabit since the early 2000s, things for consumers are not really any faster 24 years on. In data centres, 40GbE and 100GbE are common, with even 200/400GbE being experimented with.

Here in the home office, uploads may top 20Mbps/s and downloading at 108Mbps/s over fibre, moving PDF files around, remote desktoping, SSHing to servers, and accessing the odd networked VM locally. It tends to be things that Gigabit is fine for, but faster is always better, so what are the options?

For a long while, 10GbE was the only game in town if you wanted to move from Gigabit, and it was hellishly expensive. So, around 2016/17, slower 2.5GbE and 5GbE were devised as intermediary options, but even so, these have remained expensive until recently. The good news is that 2.5GbE switches and network interface cards (NIC) have dropped to commodity prices. We're playing with a five-port 2.5GbE switch that cost less than £60 and a PCIe add-in card that was £20. Meaning you can upgrade your local network for about £100 and likely keep the same cabling, too.

So, we're going to drop a 2.5GbE NIC in a basic server PC running Ubuntu, and get that up and running, outlining any issues we might encounter. Then we'll hook up the switch and connect another standard desktop PC with 2.5GbE. We'll run some standard *iperf* speed tests, do some real-world transfers, and see what performance we get. We'll also play around with cable options and see if some of the more dubious-looking cables we found in the server dungeon's cable chest reduce performance or not.

For testing, we sourced the imaginatively named TP-Link 2.5 Gigabit PCIe Network Adapter for £18 from a large rainforest-themed retailer (other retailers are available). The

```
lxf@lxf-desktop:~$ iperf -c 192.168.1.37
Client connecting to 192.168.1.37, TCP port 5001
TCP window size: 85.0 KByte (default)
[ 1] local 192.168.1.38 port 35998 connected with 192.168.1.37 port 5001
[ ID] Interval      Transfer    Bandwidth
[ 1] 0.0000-10.0199 sec  2.74 GBytes  2.35 Gbits/sec
lxf@lxf-desktop:~$
```

The network bits are simply batting along the cables.

specs state Linux is supported, but we followed the link symbol to the small print and there's a most general of waivers stating you might need to update your OS driver, which doesn't fill us with optimism. Heading off to the TP-Link product page, we end up on the Support page, and the FAQ "How can I install the driver of TX201 on Linux?" piques our interest.

This links to what we have to say is a comprehensive driver guide, which explains that this is a Realtek RTL8125 controller that has kernel source support from 2.4.20 onwards. The short version being if you drop this in, it'll just work. Sure enough, firing up Ubuntu and using the following lists the controller:

```
$ lspci | grep Ethernet
```

If you want the full capabilities of a network port, you need to know its device name. The easy option is to open your network settings and read from there (see screenshot, opposite page, top-right). This used to be ethX but naming conventions have changed. Otherwise, we recommend using:

```
$ nmcli device status
```

For some reason, our Ubuntu Ethernet ports are called enp6s0 and enp7s0, the new enp6s0 being shown as disconnected at the moment because no network cable is attached. If you want to check its full capabilities, try:

```
$ ethtool enp6s0
```

It should show **2500baseT/Full**, meaning that Ubuntu knows it's capable of 2.5GbE. So, we know it's installed and working.

If you've ended up buying a network card that's not supported by the Linux kernel – to be honest, in this day and age, that's a hard thing to do – you're on your

The box that makes it all work together.




own. The first thing is to find out the IC controller the network card uses; DuckDuckGo that with Linux driver, and ideally the manufacturer supports a driver you can `sudo modprobe <modulename>` on to your system, but be aware that any 'out of tree' drivers are removed when the kernel is updated. So, you may lose network connectivity after an update, which is annoying.

If the manufacturer doesn't offer a driver, shame on them. It could be that some kind soul has written one and there's a GitHub repository or a maintained PPA, but we're in the area of general handwaving, and you'll need to follow any instructions on compiling source or installing third-party repositories to get the driver installed.

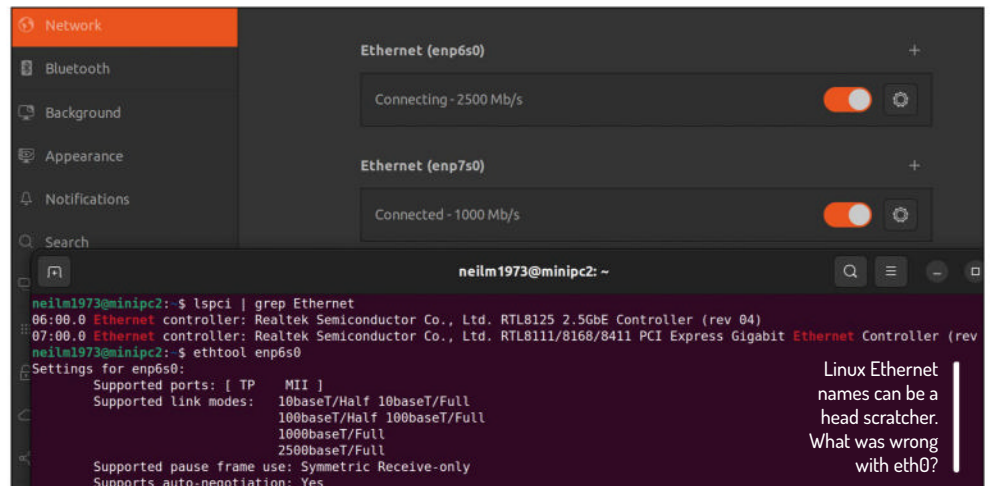
We also need a 2.5GbE switch to hook up our new speedy NICs together. We procured a Zyxel MG-105, which is a five-port unmanaged affair. The usual RRP is £80 but we managed to snag one for £60, as they occasionally seem to pop up on offer. If you have heavier requirements, eight- and 12-port models are available. Our main concern, however, is simply connecting our home server and main work desktop for high-speed transfers. An outside third option would be a Wi-Fi 7 router that could outstrip Gigabit speeds (that's around 115MB/s, with a theoretical maximum 125MB/s) but our feeling is that's unlikely for the foreseeable in a home or SMB environment.

To begin, we tested a basic large file transfer from the NAS server to the desktop and back:

Action	Speed
Read	255.5MB/s
Write	270.5MB/s



TP-Link's 2.5 Gigabit PCIe adaptor will set you back less than £20.



We then ran an *iPerf* performance test and recorded these results: 2.74GB in 10 seconds at 280MB/s

We're using new CAT5e wiring over short runs here, but we are also interested in how longer runs over perhaps older or maybe less reputable cables might hamper performance. So, we dug out a couple of dodgy 10m lengths that have been mouldering in the cable box and didn't see any change in speed in either direction.

Certainly, 2.5GbE switches still remain somewhat pricey but updating your existing network shouldn't break the bank. A shift to 2.5GbE isn't a case of everything has to go for the majority of your network devices; Gigabit speeds should remain more than enough – not to mention Wi-Fi. It's only the devices that have to do any heavy network lifting that are likely needing to be considered at all. These will be file storage NAS boxes and any desktop workstations that need to access them – don't forget laptops either. If you have a laptop, suitably fast 5+Gb/s USB-C port adaptors are available from around £30, but as always, check for Linux compatibility beforehand! **LXF**

QUICK TIP

If an add-in card doesn't suggest Linux compatibility, search for a driver using its controller's IC number from the specs.

» CABLES MATTER

The good news is that Ethernet cables are generally pretty straightforward; the CAT (literally Category) standard defines a physical and electrical system for connecting Ethernet ports. Effectively, CAT5e – ratified in 2001 – is the most common. It supports 2.5GbE and a run length of 100m, and is relatively cheap. CAT5 was ratified in 1995 and is effectively the same; it's just that CAT5e tightened up requirements, although a CAT5 cable could very well qualify as CAT5e. The main limitations are that CAT5e is unshielded and only tested to carry 100MHz signals.

When moving to 10GbE for reliable transmission, CAT6(A) cables are the order of the day. The key improvement is support for 250/500MHz transmission rates. CAT6 only supports 10GbE run lengths up to 55m, which is why CAT6A was introduced. CAT6A uses tighter twisted pairs – this does result in 50% thicker cables – that boost run lengths to 100m for 10GbE speeds. CAT6 also restricts bends to 4x the cable diameter – so no kinking – and the outer jacket can't be stripped back more than 13mm.

» **GET US FASTER THAN THE SHOPS** Subscribe now at <http://bit.ly/LinuxFormat>

Do Docker like an adult!

Stuart Burns is as mad as hell and he's not going to stand for your shoddy Docker deployments any longer!

In the world of ephemeral services, *Docker* images provide a great way to have disposable services on a 'quick in, quick out' scenario. With that ease of use, bad practice can creep in. Here we're discussing some of the ways to optimise production usage of *Docker* in a single-host scenario.

While it is simple enough to just run *Docker* from the command line using various switches, you are missing a big time-saving trick. Replacing those command lines with *Docker Compose* and a YAML *Docker* config file is a much better way to achieve a consistent outcome. Using *Docker Compose* also cleans up after itself, removing resources when it is brought down, so it makes complete sense to use it.

Using a single *Docker Compose* file, you can include an entire software stack defined in one file that can be deployed, version-controlled and consistent. You could use a *Bash* script but it's not the '90s any more. You do, however, need to install the *Docker Compose* tool to use it. For example, in Ubuntu, it's as simple as:

```
$ sudo apt install -y docker-compose
```

People get a bit unnerved with the whole YAML creation process, but in reality, it provides the same settings as you'd find in a similar *Docker* command line usage. The content essentially breaks down into a few key groups: version, services to be included, networks and ports, and mount points as an absolute minimum.

A very simple Apache setup is shown below (it's one we use when doing quick and dirty HTTP services):

```
version: '3.9'
services:
  apache:
    image: httpd:latest
    container_name: apache
    volumes:
      - /encrypted/apache/data:/usr/local/apache2/htdocs
    ports:
      - 9118:80
    restart: unless-stopped
```

To use this example in your environment, you are only required to change the left-hand side of the **volumes** statement, changing it to whatever the desired local path is. Once done, save the file (using

whatever name you desire, as long as it ends in .yaml), and use the command:

```
$ docker-compose up
```

Once the container is up and running, browsing to **localhost:9118** should show a basic Apache installation. It's really useful to use *Docker Compose* with the **-d** switch to make sure the container runs in the background, or you have lost the use of your terminal for the duration.

At the end of the day, it's just a case of where the bits from the command line go in terms of the YAML file. Building more complex examples is a case of adding in more of the same.

So, with everything up and running, cancelling a running container is as simple as running the command below (in the same folder):

```
$ docker-compose down
```

This shuts down the container and removes the network it created (if you didn't specify one in the YAML file). The great thing is that if you keep your configuration files stored on the persistent volume (the folder on the computer's local disk), it can be brought back to life by just repeating **docker-compose up -d**.

A quick pro tip about building *Docker Compose* files is not to try to create them manually in *Nano* or *Vi*, but to use a YAML-aware IDE. The syntax highlighting makes it simple to see errors and highlight them before you try to deploy and get an error. YAML can be gnarly!

Now let's double up on our *Docker Compose* goodness and add a second *Docker*-ised application. All the applications are declared in the services point, so a simple double application stack would look like this (the official *WordPress* deployment YAML file):

```
version: '3.1'
services:
  wordpress:
    image: wordpress
    restart: always
    ports:
      - 8080:80
    environment:
      WORDPRESS_DB_HOST: db
      WORDPRESS_DB_USER: exampleuser
```

Here we're using the example WordPress Docker Compose to create the entire stack in one easy step.

```
Stopping wordpress_db_1 ... done
sysadmin@dockhost:~/dockerfile/wordpress$ docker-compose up -d
Starting wordpress_db_1 ... done
Starting wordpress_wordpress_1 ... done
sysadmin@dockhost:~/dockerfile/wordpress$
```


NAME	CPU %	MEM USAGE / LIMIT	MEM %	NET I/O	BLOCK I/O	PIDS
	0.05%	191.2MiB / 15.5GiB	1.20%	693kB / 74kB	227MB / 57.4MB	28
	0.09%	401.8MiB / 15.5GiB	2.53%	403kB / 60.1kB	56.5GB / 28.4GB	27
	0.03%	128.9MiB / 15.5GiB	0.81%	129kB / 285kB	28.5GB / 28.4GB	30
	0.07%	145.4MiB / 15.5GiB	0.92%	3.3MB / 132kB	254MB / 71.7MB	28
	0.09%	165.3MiB / 15.5GiB	1.04%	475kB / 2.61MB	252MB / 39.5MB	29
	0.01%	9.18MiB / 15.5GiB	0.06%	11.1kB / 3.79kB	13.9MB / 4.1kB	82
	0.02%	89.9MiB / 15.5GiB	0.57%	211kB / 8.23MB	134MB / 8.19kB	22
	0.00%	319.1MiB / 15.5GiB	2.01%	4.68MB / 159kB	396MB / 26.1MB	29

```
WORDPRESS_DB_PASSWORD: examplepass
WORDPRESS_DB_NAME: exampledb
volumes:
- wordpress:/var/www/html
db:
image: mysql:8.0
restart: always
environment:
MYSQL_DATABASE: exampledb
MYSQL_USER: exampleuser
MYSQL_PASSWORD: examplepass
MYSQL_RANDOM_ROOT_PASSWORD: '1'
volumes:
- db:/var/lib/mysql
volumes:
wordpress:
db:
```

Using this would allow the admin to deploy WordPress from scratch. The great thing is that you can create several different stacks and they are separate from the other stacks (creating a new file in a separate folder to do this, using `docker-compose up`). The network ports need to be tweaked so as not to clash.

Taking it one step further, if the administrator were to create and post a Docker container to a repository, it becomes trivial to provide a Docker Compose file to another party, and the environment just downloads and runs as it did for those who built it. Docker's 'build once, run anywhere' in action.

We have just touched the surface here, but we hope to have illustrated that while Docker itself is great, Docker Compose adds a whole new level of capability. There are some options we haven't discussed that are shown above, but you can research and tweak those to your own needs. **LXF**

Running everything Docker-ised uses a tiny amount of RAM compared to full VMs.

```
> stu > ! example.yml
version: "2.1"
services:
  sonarr:
    image: lscr.io/linuxserver/sonarr:latest
    container_name: sonarr
    networks:
      - example1
    environment:
      - PUID=1000
      - PGID=1000
      - TZ=Etc/UTC
    volumes:
      - /encrypted/docker/sonarr/config:/config
      - /encrypted/tvseries:/tv #optional
      - /downloads:/downloads #optional
    ports:
      - 8989:8989
    restart: unless-stopped
sabnzbd:
```

A simple example of how an appropriate IDE can make things a lot easier to avoid mistakes.



Stuart Burns is a Linux administrator for a Fortune 500 company specialising in Linux.

» GEEKS OF THE WORLD UNITE!

For unimportant reasons, I spent a lot of the Christmas period refactoring a production Docker environment.

Sometimes you only get to understand why best practices work when you hit a problem that was created because best practices weren't followed.

In other news, the Broadcom purchase of VMware was completed. For most, it means a lot more expenditure. So what, you may ask? Well, a lot of smaller companies are looking at Linux-based Proxmox to get out of the inbound mega-bills.

One item of note is that the UK government is still trying to make the world of peer-to-peer communications easier to eavesdrop on. It's a bit of a sleight of hand by those idiots in power, but the long and the short of it is that rather than trying to break cryptography (you can't half break it, as I have mentioned in the past), the Tory government is now trying to pass measures that will prohibit the big providers from introducing better encryption and privacy without getting the OK from the government. This has tinges of the NSA elliptic curve scenario all over again – the US tried to introduce a fatally flawed cryptographic algorithm back in 2006.

Why am I putting this in here? Because it's time for us geeks to write to our MPs to educate them that cryptography security must be absolute. Secondly, it's laziness on the part of the government. It has other ways to collect this information. To check out the government's words of wisdom, read the paper at <https://bit.ly/3SaONxS>.

» DOCKY DASHY

All this talk of Docker containers running can make keeping them all together a pain. We were in this scenario and stumbled across the Dashy container (<https://bit.ly/48Fx2Ms>).

It is designed to be a simple bookmark/monitoring service. Yes, bookmarks are common fare in browsers but having

a site with them in means rather than having to add the same URL to several machines, it's possible to have a good group of links. On the plus side, it's also a Docker image. There are many such tools but this one stands out because of its simplicity and ease of use. One quick tip, though: adding icons to the links

makes it better looking. However, you need to host the images somewhere, so do like we did and set up a Docker-ised Apache server (it only takes up a few megabytes of RAM – the joy of containers!) and have it serve up all those cute icons. A good selection of icons can be found at <https://bit.ly/3RKJvrw>.

Starlink

Wherever he may roam, **Nate Drake** can enjoy fast internet at a fair price.

IN BRIEF

A base install includes: Starlink antenna, base stand, router, 15m data cable and 1.8m power cable.

Dish: Electronic phased array

Orientation: Motorised self-orientating

Rated: IP54

Snow melt: Up to 40mm/hour

Temp range: -30°C to 50°C

FoV: 100°

Power: 50-75W

Size: 513 x303x544mm

Router

Wi-Fi: IEEE

802.11a/b/g/n/ac, Wi-Fi 5

Radio: Dual-band 3x3 MIMO

Security: WPA2

Rated: IP54

indoor use

Mesh: Up to 3

Starlink Mesh nodes

Ethernet:

Optional extra

After moving to the countryside and experiencing repeated issues with traditional broadband, we decided to give Starlink's residential service a try. Elon Musk's SpaceX has been in the news frequently of late – and the service is improving all the time. The portability feature makes it possible to deploy satellite internet elsewhere for a small monthly fee, plus Starlink hopes to provide an off-grid cellular service in the next few years.

One of the big caveats of Starlink is that it's not yet available in every area of every country. Fortunately, the main website immediately prompts you to enter your address, to check if coverage is in place or planned for your area, then lists various subscription plans.

Given the cost of launching satellites, it's hardly surprising that the price of a basic subscription has risen in recent years; currently it's £75 (US\$120) per month for a Standard plan, which offers "unlimited" data. Plus £10 monthly hardware rental or a one-off £449 fee, refurbished units are available for £199.

There is a caveat, however, in that Standard data usage is subject to a fair use policy. Full details of expected speeds are available on the website but in brief these are a download speed of 25-100Mbps, an upload speed of 5-10Mbps, and a latency of 25-60ms.

If you find this too restrictive, you can also pay for a certain amount of Priority data. For instance, paying an extra £21 per month nets you 40GB of Priority speeds. The expected performance of Priority data is much greater, with download speeds of 40-220Mbps and uploads of 8-25Mbps, though latency is the same. If you are looking for the best small business router and ISP, Starlink for Businesses offers separate plans.

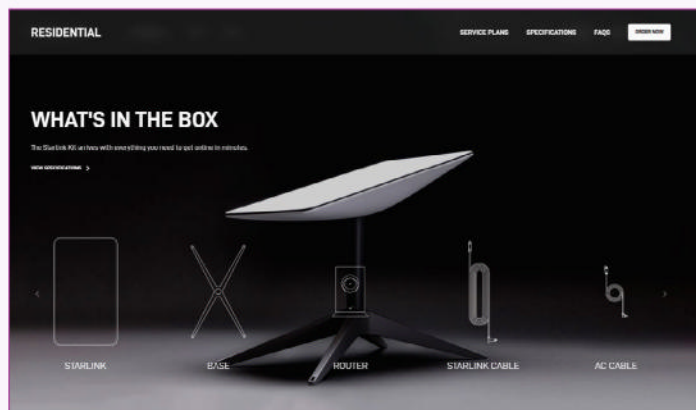
All the fun of the fair

Starlink warns that it tries to allocate data for its Standard plan in a "fair and equitable manner", and if you exceed this, it may seek to restrict your speed. The examples given in the fair use policy include streaming, video calls or online gaming, though puzzlingly, the same policy also promises to treat all traffic neutrally.

After one month and nearly 400GB of data, we experienced no issues with the policy, but this may be because we use VPNs on all devices, encrypting traffic so Starlink can't distinguish specific content like streaming video.

We are luckier than most as we have an outside power socket and an open lawn. The dish slots into the stand effortlessly, then, like any of the best Wi-Fi routers, it's simple to connect the cable to the device.

When it comes to placing the dish, we recommend using the *Starlink* app's built in obstructions tool, which



You need a dedicated Starlink receiver propped up in your back garden.

can scan the sky to check whether there's anything in the way of the dish receiving signals from SpaceX's low-Earth-orbit satellites.

Placing Starlink's standard rectangular dish on a front lawn as we did, with cables snaking from it on to your porch, simply isn't sustainable. In the first case, it can lead to obstructions, as in our case, given that the house was directly between the dish and the satellites.

The Starlink website can come to the rescue here. Apart from selling more advanced satellite dishes, you can also buy a standalone pipe mount or various wall/roof mounting kits. The other issue, of course, is that of the router. It's very difficult to store outside safely, so unless you're very comfortable with DIY, you need to do as we did and shell out to have a man come and drill holes in the wall to stow the router safely inside while still connected to the dish.

In order to avoid the issue of network congestion, we decided to run our tests at around 11.00pm using the *Starlink* app. The Starlink Speed for downloads was 71Mbps, while we achieved uploads of 15Mbps. This was entirely consistent with the promised speeds for a Standard plan of 25-100Mbps for downloads and 5-10Mbps for uploads. **LXF**

VERDICT

DEVELOPER: SpaceX

WEB: www.starlink.com

PRICE: From £75 per month (plus £10pm rental)

FEATURES	9/10	EASE OF USE	9/10
PERFORMANCE	7/10	VALUE	7/10

Starlink's satellite service is easy for subscribers and you can set up the hardware in minutes. Be prepared to spend more time and money for optimal performance, though.

» **Rating 8/10**

Hatch

An easy-to-use builder that empowers **Mirza Bahic's** creative expression.

IN BRIEF

A fantastic choice for aspiring creators eager to explore web development without breaking the bank. With a solid collection of customisable templates, a charming Hatchbot AI, and a forever-free plan, it's an ideal starting point for those venturing into site-building. However, its limitations for high-traffic sites and the absence of real-time support might not suit larger-scale businesses seeking robust performance and advanced features.

In the dynamic realm of web development, where coding often stands as a barrier to entry, Hatch emerges as a game-changer. Although it positions itself alongside some of the best website builders, it differentiates itself by focusing on the empowerment of tech-curious creators, helping create their online projects without much effort or technical know-how.

The Hatch Free plan comes at no cost and includes hosting with a subdomain (**yourname.hatch.one**), a Hatch watermark, limits of three published projects and three pages per project, unlimited community projects, 2GB of media storage, and standard customer support.

For those ready to take their creative journey further, the Hatch Pro plan, at \$9 per month, offers an array of advanced features. This includes Hatch hosting, a free custom domain name (with the option to bring your own domain), a removable watermark, an expanded limit of 20 projects with unlimited pages, unlimited community projects, 10GB of media storage, and priority customer support.

Within Hatch's creative haven, designers, whether casual or seasoned, can explore interactivity, animation, physics and generative AI using intuitive drag-and-drop kits. As part of the creator community, you can make your projects remixable, which means fellow members can duplicate and modify content in any way they want.

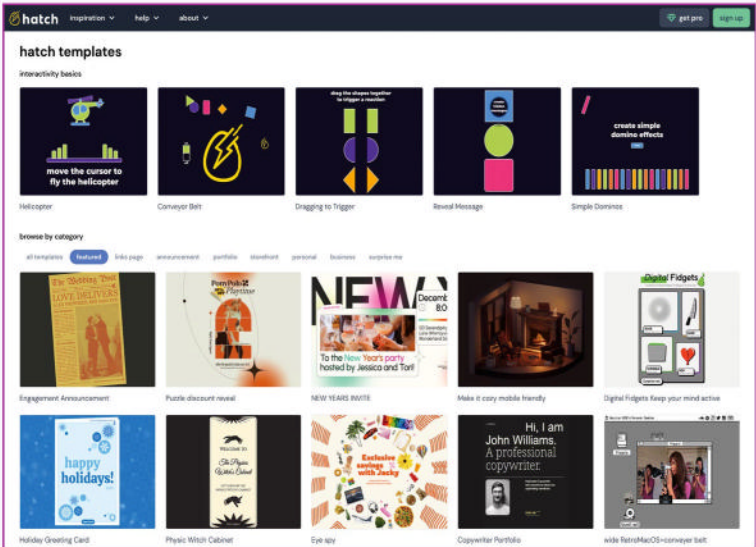
We must also mention Hatchbot AI, the star feature here. This generative web builder can transform your descriptions into web page elements without any coding know-how. It harnesses the power of GPT, enabling you to describe the design and functionality you want, then it brings it to life. The creative possibilities are impressive.

Additionally, Hatch enables the integration of diverse media, including thousands of GIFs, stickers, photos, videos, animated text, emojis and content from popular platforms such as YouTube, SoundCloud and Spotify.

Notably absent from Hatch plans is information on resources or the type of hosting provided. Advanced ecommerce tools, performance enhancements such as CDNs, SEO features, web analytics, robust security with SSL certificates, scalability options, server reliability details and uptime guarantees are all missing with Hatch.

Also, besides blank statements about caring about your security and privacy, Hatch's website doesn't share any in-depth details about the security features it uses to keep your site and its data safe. This is concerning.

Once signed up, you arrive at Hatch's template gallery, where you can choose from over 50 eye-catching templates. However, if you prefer to start from scratch,



Lowering the entry to promoting your life online, with easy template, drag-and-drop design.

that's an option, too. Whichever you pick, you land on Hatch's easy-to-use dashboard, where you can tweak the template to your heart's content.

We tested our Hatch site's performance with our trusted tool, *GTmetrix*, which helps us assess loading speed, page optimisation and other crucial performance metrics. The performance was rated with a rock-solid B (82%), which isn't perfect but still presentable. Although the site took some time to fully load (11 seconds), it didn't significantly impact our experience.

As far as support goes, there's an array of handy YouTube video tutorials, while if you have a question or need basic guidance, there's a short FAQ section on the website. For more in-depth insight, you can seek advice on Hatch's community forum.

The only way you can get in touch directly with Hatch's customer support staff is via contact form. If you've opted for a paid subscription, you'll benefit from priority support, ensuring quicker assistance. **LXF**

VERDICT


DEVELOPER: Hatch
WEB: <http://hatch.one>
PRICE: Free (\$9pm Pro)

FEATURES	5/10	EASE OF USE	8/10
PERFORMANCE	7/10	VALUE	9/10

Tech-minded creators seeking a budget-friendly solution with extensive customisation will find Hatch an excellent choice, but it may not suit businesses or high-traffic stores.

» **Rating 7/10**

CREDIT: Hatch



PROTECT YOUR SYSTEMS

Davey Winder reveals the biggest threats that hit systems in 2023 and what's likely to happen over 2024.

You might assume that cybercriminals would focus their attentions on the biggest organisations – after all, those are the ones best able to afford huge ransoms. In fact, smaller targets can and do get attacked all the time.

“Regardless of the size of the business,” says Matt Cooke, strategist at cybersecurity specialist Proofpoint, “if they have a bank account, or sensitive information to steal, they are very much at risk.” Proofpoint has found that 72% of organisations with fewer than 500 employees have dealt with a material loss of sensitive information within the past 12 months (see <http://proofpoint.com> for more information).

Indeed, small and medium-sized businesses are prime targets for malicious actors. Recent research from cybersecurity specialist Trellix (<http://trellix.com>) reveals that companies with 51-200 employees are the most common victims of ransomware, representing a third of all attacks in Q1 2023. “Smaller, less prominent organisations often lack the resources to implement robust security measures,” explains Trellix VP Fabien Rech, “either from a financial or skill-based standpoint.”

To protect your company and data, therefore, it's vital to be aware of the threats that you're most likely

CREDIT: Magictorch

A real-world targeted phishing attempt; the 'boss' is often posed asking a subordinate to buy gift vouchers.



to encounter, and to direct your (likely limited) resources accordingly, so as to minimise exposure to these security risks.

Phishing for cash

One important thing to understand is that your small business will mostly face the same range of threats as larger ones. This is because, by and large, cybercriminals are creatures of habit. They use tried and trusted methods to attack your networks and data, sometimes with an indiscriminate 'spray and pray' approach, other times with carefully targeted attacks. In both cases, by far the most common threat facing businesses of all sizes is phishing.

The challenge with phishing is that it comes in many forms. Michael Skelton, VP of security operations at Bugcrowd (<http://bugcrowd.com>), notes that, while most associated with email, phishing can equally take place via SMS, online chat services or even live telephone calls. The most common theme is using social engineering to trick employees into revealing sensitive information, especially login credentials, which can then be used for further mischief.

"Once armed with this information," Skelton explains, "cybercriminals can bypass security measures and gain access to systems where they can initiate attacks like ransomware or data breaches."

There is some positive news: while phishing is a widespread threat, most campaigns are far from sophisticated. Many attempted attacks can be defeated simply by recognising the red flags, such as unexpected invitations to open an untrusted attachment or log into an unfamiliar untrusted website.

"Staff cyber-awareness training is particularly important in combatting phishing," says Jamie Akhtar, CEO and co-founder at CyberSmart (<http://cybersmart.co.uk>). "It doesn't have to cost a lot or take hours of employees' time; a basic grounding should help them to spot most threats."

Sometimes, however, cybercriminals target specific employees of a business – an attack type sometimes known as spear phishing or whaling. Such tailor-made attacks can be much harder to spot, and the goal might not be merely to steal information, but to trick trusted staff into literally giving away large sums of money. "Attackers impersonate executives within businesses, urging employees to quickly transfer funds from one account to another," warns David Emm, principal security researcher at Kaspersky.

"These emails are often expertly crafted, and due to the urgency conveyed by the scammers, targets frequently fail to question their authenticity, and hand over information or transfer the funds as requested." The risk can be mitigated by implementing processes to cover financial transfers, such as requiring sign-off from more than one named individual in the business, and getting direct confirmation through a trusted communications channel – such as a face-to-face meeting or an internal telephone extension.

Ransomware the small stuff

When ransomware hits the news, it's normally because high-profile organisations have been hit with demands for enormous sums. However, a report by Management Today (<http://managementtoday.co.uk>) found that

» TOP THREE THREATS FOR 2024

5G and IoT devices

"As smart tech and 5G networks expand, hackers will have more opportunities to target vulnerable IoT devices. Small businesses can counter this by installing firewalls, using VPNs to encrypt traffic and regularly updating router firmware. You should also segment your network to remove single-point-of-failure vulnerabilities, and ensure IoT devices run the latest versions of OSes and security tools." Jamie Akhtar, CyberSmart

Advanced persistent threats

"Our researchers have found that threat actors appear to be scaling existing tactics seen in enterprise-targeted phishing campaigns for less robust SMB environments. These include the use of compromised infrastructure in phishing campaigns; regional targeting by state-aligned actors for financial theft; and managed-services providers being targeted via phishing, introducing the threat of supply-chain attacks." Matt Cooke, Proofpoint

AI-powered attacks

"The next evolutionary step in phishing attacks is likely to be AI-powered campaigns making full use of large language models. You can expect fraudulent communications to become more sophisticated, and more targeted at scale. You can mitigate this with regular awareness training, multi-factor authentication, and ensuring employees use direct sources for logging into corporate services." Michael Skelton, Bugcrowd

SIZE ISN'T IMPORTANT


"One important thing to understand is that your small business will mostly face the same threats as larger ones."

more than 60% of global ransomware attacks between January 2020 and July 2022 were on small businesses.

"Small businesses are especially at risk to ransomware attacks," notes Lewis West, head of cybersecurity at tech recruiter Hamilton Barnes (<http://hamilton-barnes.com>). "Cybercriminals know that smaller businesses are much more likely to pay the ransom, as they are not as equipped to absorb the attack as larger businesses might be."

"Where purse strings are tight, time offline is money lost," agrees Azeem Aleem, director of Northern Europe and UK for Sygnia (<http://sygnia.co>). "Many small businesses opt to pay the ransom quickly in a knee-jerk reaction – unaware that the threat attacker may then choose to go in for a second attack, or only release some of their data."

If you do get hit, rather than paying up right away, consider engaging a security expert trained in ransomware negotiation. "These experts investigate how the threat actor gained access, often straight from the horse's mouth," Aleem says. "They use various data-extraction and negotiation techniques, similar to a real-world hostage scenario. This can buy more time to



Enforcing two-factor authentication policies can radically reduce successful attacks.

contain the threat, help the business negotiate a lower ransom and help with remediation.”

However, even if you do reach an agreement with the attacker, there's no guarantee that your data won't still be sold online – nor any real assurance that working decryption keys will be provided. It's far better to mitigate the ransomware risk in the first place.

“There are a couple of things SMBs can do to protect themselves,” Jamie Akhtar says. “First, invest in upskilling staff. Most ransomware attacks start with human error, and training makes this less likely. Second, use data backups for anything important – and have a plan for what to do if the worst does happen.”

Supply-chain attacks

While you may be doing all you can to protect your own business from online attacks, you have far less control over the defences of your suppliers. Supply-chain attacks are growing in frequency and seriousness; indeed, what was referred to as the biggest data theft of 2023 was achieved using a supply-chain attack. During the first half of the year, the CIOp ransomware group exploited a zero-day vulnerability in the *MOVEit* file transfer tool, used by organisations of all sizes, from small businesses to government agencies, to steal data on a huge scale. It has been reported that more than 2,000 organisations were caught up in the ongoing security breaches, and more than 62 million people were impacted by data thefts.

“Cybercrooks have learned that finding and exploiting a vulnerability in a widely used software solution will give them an exponential push,” observes Dirk Schrader, CISO and VP of security research with Netwrix (<http://netwrix.com>). “One single vulnerability will give them an inroad into many IT estates.”

There may be limits to how far you can audit and secure your tools, but being aware of the risk can help you reduce your exposure. “First of all, ensure your own cybersecurity is in order. Second, write cybersecurity

requirements into RFPs and contracts with smaller suppliers,” Akhtar recommends.

Fileless attacks

Fileless attacks commonly involve an adversary hijacking existing software on a device for malicious purposes. Such attacks are called fileless because they take place in memory; for example, a fileless attack might be launched via a script, running inside a genuine, trusted app – or even on a cloud server. Often, no trace of the malicious software is written to the hard disk, either on your desktop or in the data centre.

“This lack of footprint or signature makes it hard for most antivirus software to detect a threat, and hackers are aware of this,” said Robert Smith, cybersecurity product manager at cloud services provider M247 (<http://m247.com>). “This is why 70% of the malware attacks we see are fileless and operating in the cloud.”

So, how can you best mitigate against fileless attacks? “Endpoint detection and response (EDR) can provide much-needed protection to businesses,” Smith says. “EDR uses behavioural analysis to monitor your data and what tasks you normally carry out on your device, and can thereby detect when there are changes to your daily activities caused by malware.” Advanced EDR protection can help small businesses block attacks that would otherwise be almost impossible to detect.

Crypto bros

“The popularity of cryptocurrencies and their typically high transaction volumes have made cryptocurrency exchanges very attractive for small businesses,” says Sygnia's Azeem Aleem. This is because they allow companies to transfer sums of money of any size very quickly, with minimal or no fees.

“Unfortunately,” Aleem continues, “companies handling crypto may be targeted by tailored attacks, performed through quick, anonymous transactions, which are not easily detected until it's too late.” To reduce the attack threat, Aleem says, small businesses need to understand how these crypto attacks unfold. Kaspersky has put together a useful resource illustrating some of the different cryptocurrency attacks (see <https://bit.ly/lxf312crypto>) that have taken place over the past year. These include the use of Trojanised hardware wallets, server hacks, fake Chrome extensions, and clipboard-injected malware.

Overall, Kaspersky advises organisations to “treat all crypto-related offers, emails, letters and innocent

SEE THE ERROR OF YOUR WAYS

“Invest in upskilling staff. Most ransomware attacks start with human error, and training makes this less likely.”

questions with maximum suspicion, and always use security software tailored for crypto investments on all relevant devices”.

Underinvestment

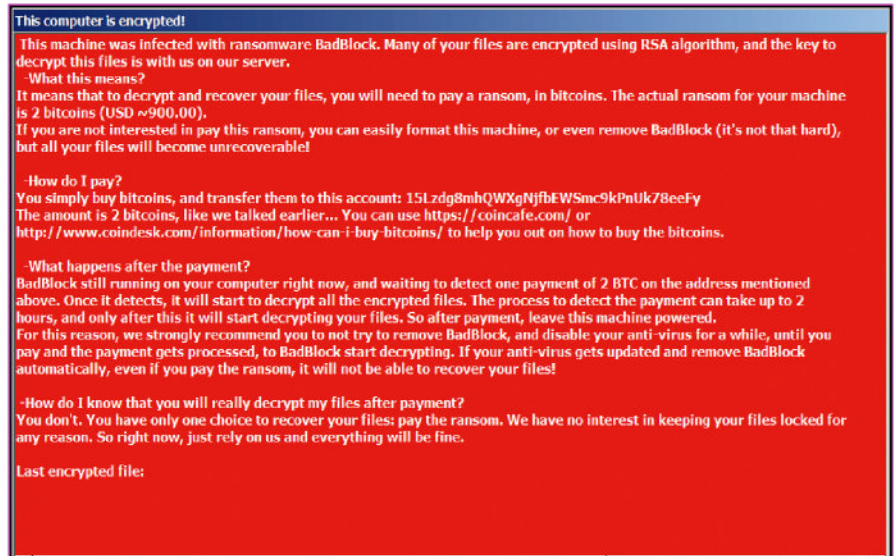
The final threat category is, perhaps, the hardest to overcome – especially in times of high inflation, when businesses are facing bottom-line challenges at every turn. Simply put, the issue is a shortage of investment in cybersecurity. CyberSmart’s Jamie Akhtar shares an alarming statistic: fully one third of the small businesses surveyed for an upcoming report had either decreased their cybersecurity investment since the cost of living crisis in the UK, or admitted that they had never really invested in it at all.

“Economic uncertainty has made investment in cybersecurity tricky for many small businesses,” Akhtar says, “but it’s not an optional cost, and it doesn’t have to be expensive. Many SMBs could benefit from consolidating and refining their security estate to focus on the must-haves.”

And so we come back to the question: how should a small business best allocate its resources to mitigate cyberthreats? Believe it or not, the answer might not involve any new spending on technology at all. Part of the solution is simply to prioritise stopping threats before they ever reach their intended victims: “Start with a proactive, not reactive mindset,” suggests Michael Skelton. “Assume that the question is when, not if, an employee’s account is going to be compromised. Then explore the security features of your cloud-based offerings, ensuring that two-factor authentication, appropriate lockout policies, and any other relevant mitigations are enabled for your organisation.”

To support your defensive efforts, experts recommend a human-first approach to cybersecurity. “Today’s attacks target people, not just technology,” points out Proofpoint’s Matt Cooke. “Cybercriminals have found new ways to exploit the instincts of curiosity and trust that lead well-intentioned people to play into the hands of the attacker.”

“Culture eats technology for breakfast,” agrees Rupert Lee-Browne, founder and chief executive of payments fintech Caxton (<http://caxton.io>). “Fundamentally, our vulnerability at the hands of the scammers is a cultural issue. We need to build a strong



culture within the business of understanding where the risks are, in order to beat the criminals.”

How can that be achieved? Over to Lewis West, head of cybersecurity at Hamilton Barnes: “Training employees to recognise attackers’ methods is an effective way to reduce the number of successful cyber-attacks, especially phishing attacks,” says West. “Companies should be embarking on in-house training for their teams, to improve their online safety knowledge and flag the types of scenarios to which

You’re unlikely to see a screen like this on Linux, or such a high ransom demand – but you never know.

TRAIN OF THOUGHT

“Training employees to recognise attackers’ methods is an effective way to reduce the number of successful attacks.”

staff could fall victim: phishing, malware, ransomware, denial of service attacks and more.”

Ben Aung, chief risk officer at Sage, sums up the message in three bullet points:

“**Make cybersecurity visible and part of your business goals** – this includes being clear on how it relates to them, and ensuring cybersecurity is being talked about and promoted by leaders.

“**Focus on the basics** – setting long and strong passwords, enabling two-factor authentication and reporting suspicious emails might sound like obvious activities, but they can all have a huge impact for smaller businesses.”

“**Implement an easy and quick way for people to report cybersecurity issues** – a clear, consistent process reduces confusion and ensures that everyone feels safe to raise concerns and act on them. There is no such thing as over-reporting in cybersecurity.”

Even though your SMB probably doesn’t have the same security resources as large enterprises, you can fend off most attacks by taking an intelligent approach to cybersecurity that focuses on the most common threats and the most-exploited vulnerabilities. **LXF**



■ Cybercriminals don’t need access to your premises...

tom's **HARDWARE**

GET ALL THE ESSENTIAL BREAKING NEWS FOR THE TECH ENTHUSIAST!

No matter if you're building a PC, buying a laptop or learning about robots, Tom's Hardware has all the comprehensive knowledge you need.



**Scan & Subscribe
for free!**



GET READY TO BECOME A TRUE PYTHON EXPERT

Make Python work for you with tutorials on coding with Django, Flask, Pygame and even more useful third-party frameworks.



Ordering is easy. Go online at:

magazinesdirect.com

Or get it from selected supermarkets & newsagents

HotPicks



Mayank Sharma

can't believe it's already that time of year when he has to find a distraction as the F1 circus goes on hiatus. After he's done picking HotPicks, that is.

Shotwell » YouTube Downloader Plus » Blue Recorder
» EncryptPad » Buttercup » Resources » Rnote
» heXon » EDuke32 » Delta Chat » Proton VPN

PHOTO MANAGER

Shotwell

Version: 0.32.3 Web: <https://wiki.gnome.org/Apps/Shotwell>

There's a reason why *Shotwell* is the default photo manager on virtually every distribution that uses the GNOME desktop. It has all the features you would expect from an everyday photo manager, with the added benefit of some very handy photo-editing functions.

The app is now available as a Flatpak, which means you can use it on any desktop environment. Use [flatpak install flathub org.gnome.Shotwell](#) to install the app on any distro that's set up to install Flatpaks.

On first launch, the app automatically imports photos from the ~/Pictures folder, although you can import photos from other places as well. It automatically groups photos by date and tags while importing them.

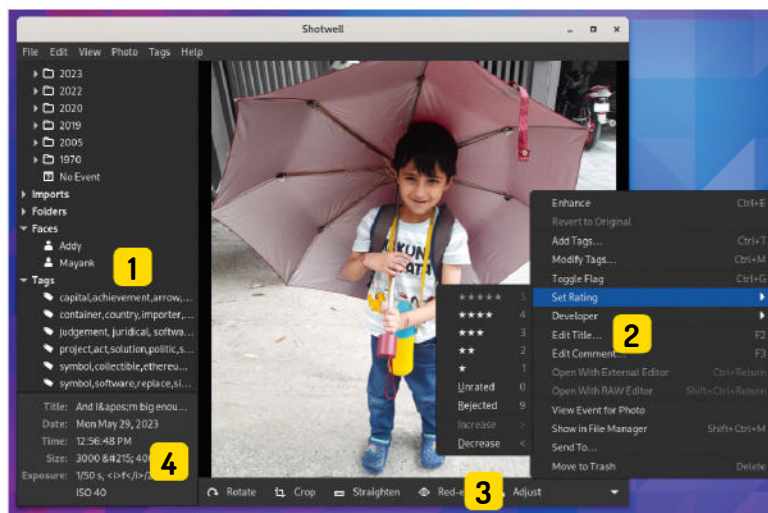
The app has a straightforward two-pane interface. The left pane helps you find images by breaking them down by details such as date, tags and such. The thumbnails for the images that meet the criteria selected in the left-panel are displayed on the right. Double-click on an image to work on it.

In addition to images, *Shotwell* also works with all the popular video formats. However, when you double-click on a video, the app passes it on to the external video player installed on your distro.

When you are viewing a particular photo, the bottom panel lists all the tools that you can use to edit the image. You can crop, straighten and rotate images, as well as remove red-eye from portraits. There's a very helpful Enhance button that automatically improves the appearance of the image. However, the app also offers advanced users the option to manually adjust the colour and tone of the image.

You can also rate the images and assign tags to the photos to help sort your library. The app also enables you to mark faces

LET'S EXPLORE SHOTWELL



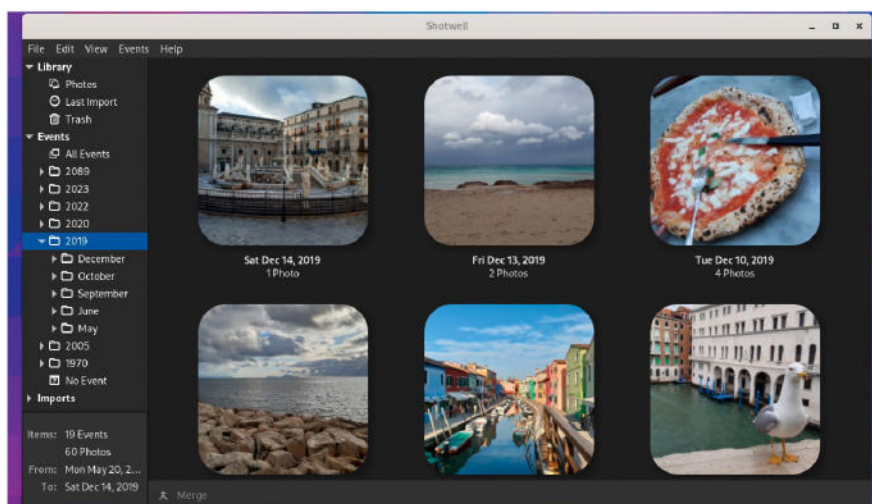
1 Image sorting toolbar
You can use the different sections listed here to sort your images by date, faces and tags.

2 Name and rate images
The right-click context menu offers the option to add a title and comment to the selected image, along with a rating.

3 Edit image
The bottom panel lists all the non-destructive image-editing tools, including a zoom slider.

4 Image details
Here you get all the relevant metadata about the image. Head to View > Extended Information to view even more details.

in the images, which again helps with the sorting. Once you're done editing, you can use *Shotwell* to publish the images in your Flickr or Google Photos account, or in a *Piwigo* installation.

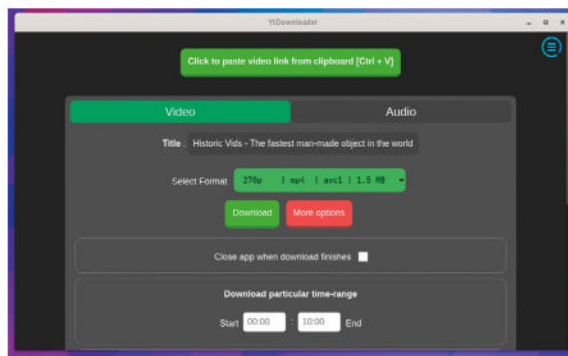


Shotwell works well for sorting through small personal image libraries, and casual image editing.

VIDEO DOWNLOADER

YouTube Downloader Plus

Version: 3.15.1

Web: <https://ytdn.netlify.app>

Do you need to keep a video from a streaming website for occasions when you don't have access to the internet? If you do, *YouTube Downloader Plus* will do the trick. Despite its name, the app can download videos of different qualities from hundreds of sites, including the popular ones such as YouTube, Facebook, Instagram, TikTok, Twitter (*who?—ED*), Twitch and so on.

The app is available as a Flatpak and can be installed with the following command: `flatpak install flathub io.github.aandrew_me.ytdn`.

The app's user interface is pretty straightforward. On launch, it offers a single button that fetches the link to the video you want to download from the clipboard. There's no option to paste a URL or drag and drop one into the app. Just copy the URL so that it's in the clipboard and hit the only button in the app to download the video. Optionally, you can also download

an entire playlist. Just copy the link to the playlist instead of the video, and select Download Playlist from under the app's hamburger menu.

After parsing the URL, the app gives you several options to download the video via a pull-down menu. In addition to the different resolutions, the list includes options to download the video with different audio and video containers. Helpfully, it also displays the size of the download along with the different combinations. However, the actual size of the downloaded video in our tests was always a little bigger than what had been specified in the app.

For even more flexibility, you can also choose to download a segment from the video instead of the whole thing by specifying start and end timestamps. Similarly, you can also use the app to download just the audio from the video. Again you have several options to grab the audio in multiple formats.

Don't be fooled by its name – YouTube Downloader Plus can grab video from dozens of streaming services.

SCREEN RECORDER

Blue Recorder

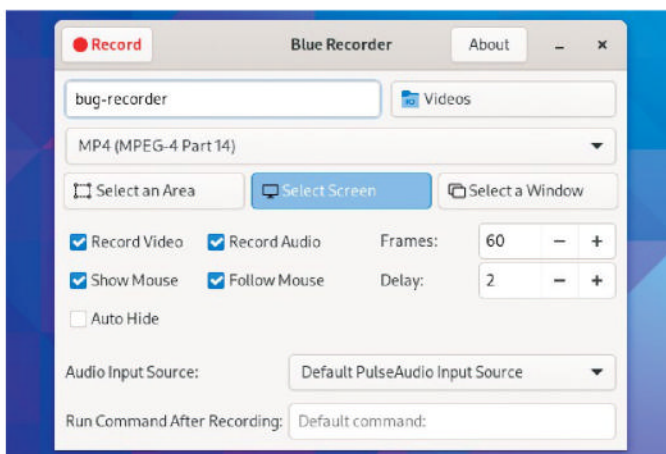
Version: 0.2.0 Web: <https://github.com/xlmnxp/blue-recorder>

Gnome Shell's built-in screen recording feature is good for capturing the odd clip, but it can't stack up against a dedicated screen recorder if you need to record longer sessions. *Blue Recorder* is one such app that gives you greater control over all aspects of the recording, including video quality, output format, sound capture, frame rate and more.

Blue Recorder is available on Flathub and can be installed with the following command: `flatpak install flathub sa.sy.blurecorder`.

The app exposes all its tricks in a single pane. You begin by specifying a name for the recording and where you want to place it. By default, *Blue Recorder* uses the timestamp when the recording was initiated as the name.

Next, use the drop-down menu to select a format for the recording. By default, the app uses MP4, but it also supports MKV, WebM, AVI, WMV, NUT and GIF formats. The app defaults to recording the entire screen, but you can also ask it to record a particular window or a specific area of the screen.



Blue Recorder captures both audio and video at 60fps straight out of the box. But you have the option of adjusting the frame rate and even recording the video without audio.

A useful option is the ability to show and follow the mouse in the recording. There's also a pull-down to help you select the audio input source. You can also define a delay to help you prepare the screen before the recording commences. If all you need to do is hide the app's interface, use the Auto Hide option to save yourself a click.

Once it's done recording, the app takes a couple of seconds to process and save the screencast. Use the Play button in the app to replay the recording, or navigate to it using your distro's file manager.

The good thing about Blue Recorder is that it can record screencasts in Wayland on both Gnome and KDE.

ENCRYPTED TEXT EDITOR

EncryptPad

Version: 0.5.0.3

Web: <https://evpo.net/encryptpad>

Although it's not a good idea, many people still continue to store passwords in plain text files. While we won't force you to trust and offload the credentials to a third-party password manager, we'll suggest you at least encrypt those text files for an additional level of security. Better still, use an encrypted text editor such as *EncryptPad* to safeguard the sensitive info.

The cross-platform *EncryptPad* has binaries for Ubuntu and Linux Mint that you can install from a PPA:

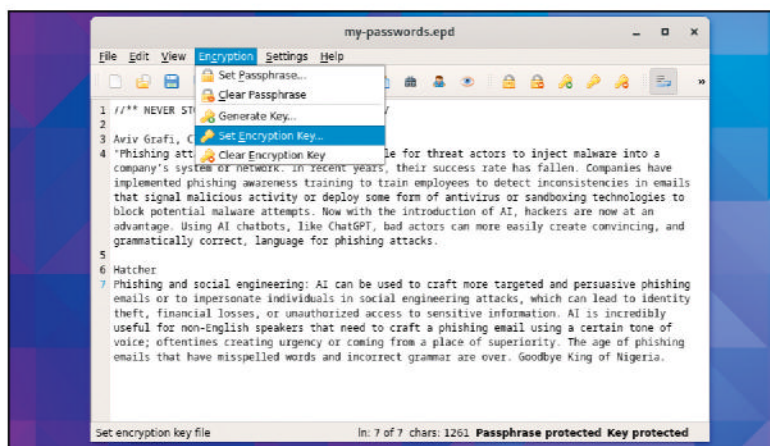
```
$ sudo add-apt-repository ppa:evpo/main
```

```
$ sudo apt update
```

```
$ sudo apt install encryptpad encryptcli
```

Users of other distros can use the official Appliance. Just grab it from the project's Download section, make it an executable from the file manager or with the `chmod +x` command, and then double-click to launch the app.

On launch, the app looks like a standard vanilla text editor. By default, *EncryptPad* doesn't have a password or a key, as highlighted in red in the status bar.



Use the editor to open a regular text file, or enter some text in the app. When you are done, you can lock it by heading to Encryption > Generate Key. Browse and select the path where you want to store the key, then specify the filename of the key. You're now prompted to enter a passphrase, after which the app generates a key for the text file.

When you now save the file, it prompts you to enter the name for the text file, which has the .epd file extension. After you've specified the filename, the app prompts you to define a passphrase to add another layer of protection.

To open the encrypted EPD text file, you're prompted for both the passphrase as well as the key with which you've locked the file.

Besides plain text files, you can also use *EncryptPad* to encrypt any file, including pictures and archives.

PASSWORD MANAGER

Buttercup

Version: 2.24.3

Web: <https://buttercup.pw>

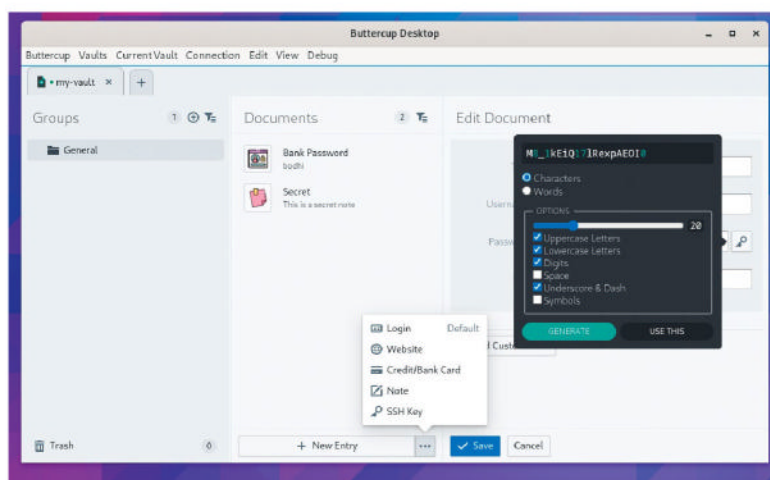
Remote cloud-based password managers are convenient, especially if you regularly switch computers. But trusting your passwords to third parties isn't everyone's cup of tea. If you are one of those, you can use *Buttercup*, a cross-platform app that gives you the option to store your passwords locally.

The app is available as a distro-agnostic Appliance. Grab the latest release from its website, and make it an executable from the file manager or with `chmod +x`.

On first launch, you're asked to create a new vault. This can be a local file, or can be placed in your Dropbox, Google Drive or a WebDAV location. When creating a local vault, you have to first specify its name and location in your filesystem, and give it a password.

You are then shown a three-pane interface. The first pane lists the broad categories or groups for the credentials. By default, there's a General group, but you can add your own, such as Personal or Work, to segregate your passwords.

After selecting a group, you're given the option to store your credentials in the second pane. In addition to



regular usernames and passwords, you can also use *Buttercup* to keep website login details safe, as well as details of your credit cards.

The app also has a password generator that you can use to create strong passwords. By default, it generates passwords that are 20 characters long and have a mix of uppercase and lowercase letters, along with digits, underscores and dashes. You can customise these options to generate longer or shorter passwords.

Buttercup protects all your credentials with a primary password. Make sure this is strong enough to resist brute force attacks, and isn't used elsewhere. You need to enter this password every time you want to access your vault.

Besides passwords, you can also use *Buttercup* to protect other sensitive information, such as notes and SSH keys.

SYSTEM MONITOR

Resources

Version: 1.2.1 Web: <https://github.com/nokyan/resources>

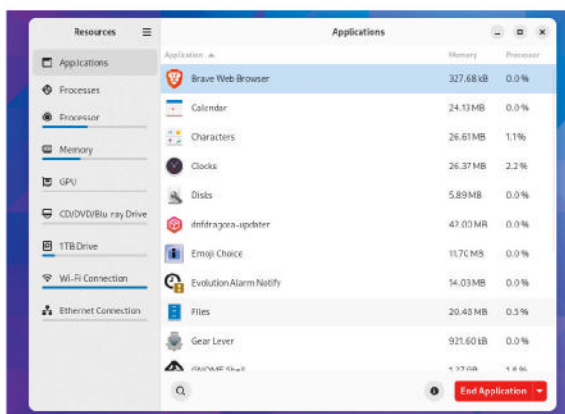
Virtually all distributions ship with a system monitor that does the job it's meant to do. However, they aren't very intuitive to a new user, not used to associating apps with processes. If you're looking for a system monitor that's more user-friendly, take a look at *Resources*.

The app is available on Flathub and can be installed on top of any distro with `flatpak install flathub net.nokyan.Resources`.

When launched, *Resources* shows an alphabetical list of all the running apps by their name, rather than by their process name. You can also reorder the list by memory and processor usage.

To view details about a particular app's resource usage, select it and hit the Info button at bottom of the interface. In addition to showing the app's processor and memory usage, you can also view its disk interactions and other details, such as whether it is running as a containerised Flatpak, in a pop-up.

Surprisingly, *Resources* doesn't show AppImages in its list of apps, but these do show up in the list of



Resources is a beautiful-looking system and resource manager, which does a good job of balancing form and function.

processes. The End Application button next to the Info button kills the selected app.

You can also use *Resources* to monitor CPU, GPU, memory, network interface and storage device usage.

Besides usage details, all components display lots of additional useful information. For instance, the Processor tab displays its temperature, along with details such as its architecture, maximum frequency and more. The Memory tab tells you how many free memory slots are available (if any) and the type, form factor and speed of the memory chips installed.

Resources might not make much sense to advanced users, given that it has a heavier resource footprint than the default system monitors. However, despite this handicap, we think it would prove quite useful for newcomers to Linux.

NOTE TAKING

Rnote

Version: 0.9.4

Web: <https://rnote.flxzt.net>

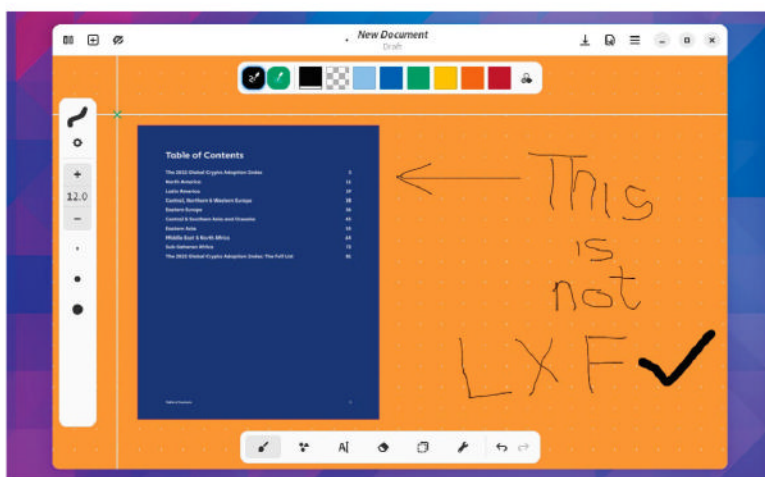
There are plenty of open source note-taking apps, but ones that allow handwritten notes are a rarity. *Rnote* is one such app that you can use to write freehand notes, and even annotate images or documents.

Rnote is available as a Flatpak and can be installed with `flatpak install flathub com.github.flxzt.rnote`.

By default, the app launches with an infinite canvas, though you do have the option to change the layout to a semi-infinite or fixed size. Even though it's designed to work on drawing tablets, it works flawlessly with the mouse as the primary input device.

The canvas has tools along the left-hand side, top and bottom. From the tool palette at the bottom, you can switch between the primary functions, such as the freehand brush, shape drawer, selector and more. You can use the top palette to customise the colour of the strokes.

Each of these functions has several options that are displayed in the palette on the left of the canvas. For instance, if you have selected the brush, you can



choose three styles, each of which has its own configurable parameters, such as size. Similarly, if using the selector, you have options to select elements in the canvas, which you can resize, rotate, move and copy.

Besides the on-canvas parameters, *Rnote* has several other customisable options and settings. Using these you can alter the background colour, and choose from a variety of patterns such as lines and grids.

In addition to freehand notes, you can also use *Rnote* to annotate all types of documents and pictures. Use the Import menu to pull these into the canvas, or paste them from the clipboard. When you're done, you can save the canvas into the native *.rnote* format or export it as a PDF or SVG.

In addition to freehand scribbling, you can use *Rnote* to enter typed text in various fonts, alignments and typefaces.

ARCADE ACTION

heXon

Version: 20231111**Web:** <https://luckey.games>

HeXon is a simple space shooter that's difficult to master. The gameplay involves flying your spaceship and destroying enemies, while collecting upgrades.

To play, download the ZIP archive from the project's website at <https://luckeyproductions.itch.io/hexon#download>. Once downloaded, extract the ZIP archive with `unzip hexon.zip -d hexon`, then switch to the extracted `hexon` directory and launch the game with `./hexon`.

The game, which supports up to four players, begins with your character walking into a hangar to select your spacecraft. As soon as you do, you're taken directly to the battlefield. Use the WASD keys to steer your craft. The number pad fires lasers in the direction of the keys. For instance, the 7 key fires in the top-left direction, while the 3 key fires lasers towards the bottom-right.

Mastering flight does take some getting used to, because all edges of the hexagonal space arena are connected to their opposite sides. So, instead of



bouncing off the edge, you reappear in the arena on the opposite side, making for puzzling gameplay.

You must avoid the Notyous, which appear in the form of razors and spires. Each has their own unique characteristics. Destroying razors earns you five points, while destroying spires rewards you with 10.

The game also has pickups in the form of apples, hearts and ChaoBalls. Apples give you 23 points, and collecting five in a row upgrades your weapons. Hearts help heal half the damage to your vessel, and collecting five in a row both restores you to full health and upgrades your shields. When you fly into a ChaoBall, it turns all enemies in its proximity into ChaoMines that you can push into spots buzzing with enemies, destroying them all.

Besides the keyboard, you can also play heXon with a PS3 controller, which works over Bluetooth but needs to be initially connected via USB.

FIRST-PERSON SHOOTER

EDuke32

Version: 20231208**Web:** www.eduke32.com

Duke Nukem 3D is one of the marquee first-person shooters of the '90s that helped popularise the genre. The game was hailed for its gameplay and the not-so-subtle satirical takes on the extravagant and unreal action heroes of the '90s. EDuke32 lets you experience this all over again.

It is a source port of the '90s hit that you can use to play the original and its expansion packs on modern hardware. It's available as Flatpak and can be installed with `flatpak install flathub com.eduke32.EDuke32`.

While EDuke32 is free, you need a copy of *Duke Nukem 3D* or its mods to play the game. The app also works with the freely available one-episode shareware version of the game. When you have these, copy the **DUKE3D.GRP** file in the `~/Documents/EDuke32` directory and you're good to go.

When you fire it up, EDuke32 first brings up the launcher that lets you set up some of the essential elements. For instance, you can choose the resolution, which can go as high as 3,072x2,304. You can also use the launcher to choose between software rendering



and two different hardware-accelerated renderers. Besides OpenGL, there's the Polymer renderer that boasts of graphical details like improved textures, coloured fog, dynamic coloured lighting, and more.

When you've finished setting things up in the launcher, hit Start to fire up the game. The developers have also reworked the controls system, offering improved compatibility with the modern combination of keyboard and mouse play style.

They have also given the configuration menus a 21st-century upgrade, stuffing them full of new options to help you configure the game. Also, while the gameplay remains true to the original, there's an improved in-game console with added features.

EDuke32 remains faithful to the '90s classic, but gives it a much-needed upgrade to match the sensibilities of modern first-person shooters.

INSTANT MESSENGER

Delta Chat

Version: 1.42.2

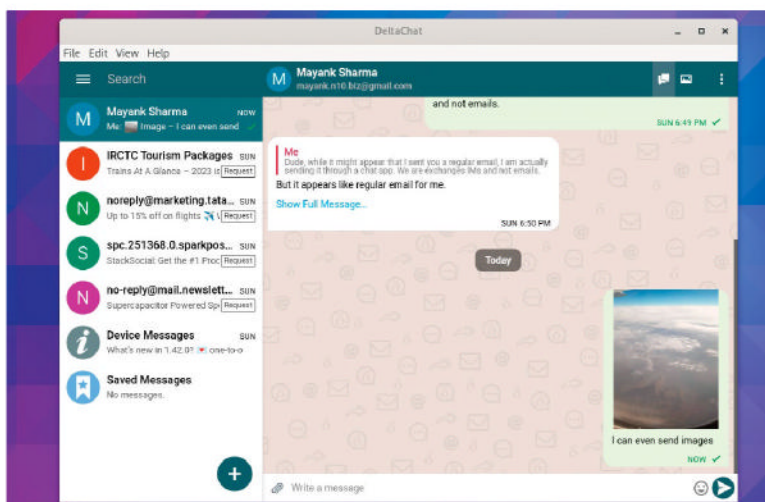
Web: <https://delta.chat>

There's no dearth of instant messengers; some focus on convenience and interoperability, while others focus on privacy. *Delta Chat*'s USP is that it doesn't use a centralised server to exchange messages. Instead, it uses email to send and receive encrypted messages.

The app has RPM and DEB packages, but the most convenient installation is via distro-agnostic Flatpaks and AppImages. Either grab the AppImage from its website or install the app from Flathub with **flatpak install flathub chat.delta.desktop**.

When you first launch *Delta Chat*, you must log in to your email account, like hooking up your email with an email client. For known email providers such as Gmail, Fastmail and such, the app can automatically fill in settings. If you're using an exotic service or running your own server, you must enter its settings manually.

You also need to create an app-specific password to let the email provider know that *Delta Chat* isn't a bot. The process for doing this differs based on the email provider, and the app points you to the relevant



documentation. Once you have the app-specific password, you have to use it along with your email address to log into *Delta Chat*.

Once you've got past the hurdle of logging in, the rest is easy. The app looks and behaves like any chat app. Since it just uses email, you can add friends by their email address. Click on the big + icon to either enter an address manually or browse through your address book to add an existing contact.

You can then converse with them as you do in any IM. To your contacts, the conversation appears like regular emails, which they can respond to. And for you, they appear like conversations in any IM client.

Although *Delta Chat* isn't your regular IM client, it supports all the features you'd expect from one, such as group chats and audio/video calls.

VPN

Proton VPN

Version: 4.1.0

Web: <https://protonvpn.com>

Proton has a good track record when it comes to open source security products. Its Proton Mail service is popular with privacy-conscious users, particularly for its open source web client and the *Mail Bridge* app. The company has extended its same open source promise to its Proton VPN service, which has a very useful and unlimited free tier as well.

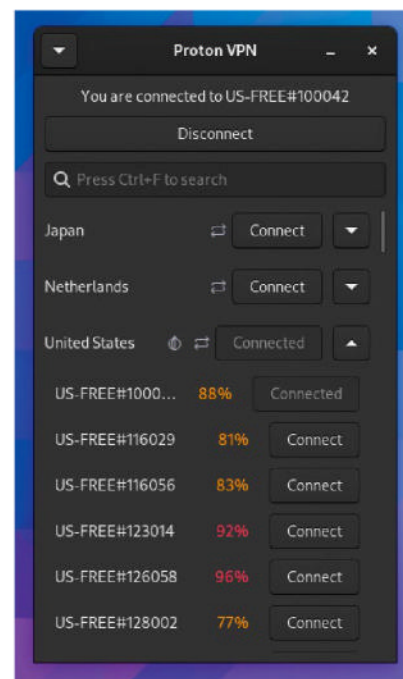
To get started, you first have to create an account with the service. While its paid accounts are competitively priced, it does let you create a free account as well. When creating an account, which involves nothing more than giving it your email address, the service creates a strong password for you, although you can replace it with your own.

After you have created an account with the service, you need to install its open source client. Proton has official binary clients for Deb- and RPM-based distributions. Experienced campaigners can also configure OpenVPN or WireGuard to route traffic through Proton VPN. However, the most convenient mechanism is to use the unofficial distro-agnostic



Flatpak package with the **flatpak install flathub com.protonvpn.www** command.

When you fire up the app, you have to log in with the credentials that you used when creating your account with the service. You're then shown a list of servers. You can just hit the Quick Connect button to delegate the server selection to the app, or you can select a server manually from the list. Users of the free service can only select one of three countries, namely Japan, Netherlands and the US.

You can use the Connect button next to a country, or use the pull-down menu to select a specific server within the country. Servers are listed along with their load, although the ones with the lowest load aren't available to free users. **LXF**



The open source Proton VPN client offers several useful features to free-tier users, such as uncapped usage and a kill switch.



**Can you help inspire the
next generation of coders?**



Code Club is a nationwide network of volunteer-led after school clubs for children aged 9-11.

We're always looking for people with coding skills to volunteer to run a club at their local primary school, library or community centre for an hour a week.

You can team up with colleagues, a teacher will be there to support you and we provide all the materials you'll need to help get children excited about digital making.

There are loads of ways to get involved!
So to find out more, join us at **www.codeclub.org.uk**

EASY 6502

Credit: <https://skilldrick.github.io/easy6502>

Part One!
Don't miss
next issue,
subscribe on
page 16!

Write and run C64-style 6502 code

David Bolton shows how to get started writing 6502 code and provides a useful example that sorts numbers.



OUR EXPERT

David Bolton used to write Commodore 64 games for a living back in the mists of time, and taught himself Z80 by writing a 6502 assembler in Z80.

How long do you think it took to modify a 2,200-line C program that runs on Windows to have it run on Linux? Just one hour. The main changes were switching some string routines such as **strcpy_s** to **strcpy**, altering paths for includes and changing the high-speed timing code. It might take between a week and a month or longer to do that in assembly. And moving that C code from, say, Ubuntu on x86 to Raspberry Pi on ARM would take almost no time at all. In assembler, it would take a complete rewrite; more months of work.

Having tried to put you off modern assembly, let's try some old-fashioned assembler from 40 years ago. With roughly 60 instructions, the 6502 CPU is simpler than modern CPUs, which can have several hundred.

What is a 6502?

Pretend it's 1983 and you have bought a brand new Commodore 64 and assembler cartridge that will


'compile' 6502 assembler into machine code. In this article, we are learning some 6502 assembly language, writing a short program in it and then running it on a web emulator.

The 6502 has just three registers: A (known as the accumulator), X and Y. Each register is 8-bit, so can hold a number between 0 and 255. It's easier to deal with values and addresses in base 16 – hexadecimal – so a byte holds a value between \$00 and \$FF.

You can find a list of all the 6502 instructions and details about addressing modes and lots more at www.masswerk.at/6502/6502_instruction_set.html, but here are a few useful things you should know about the Commodore 64 and the 6502 CPU.

On the Commodore 64, the processor is actually a 6510 (a 6502 with a bit of hardware mapped in). It has 6,553,610 bytes of RAM, but some RAM addresses are used for the hardware sprites. In hexadecimal, 6,553,610 is \$10000 and the memory addresses run

from \$0000-\$FFFF. However, the Commodore 64 charset graphics are mapped at \$1000-\$2000, sprite hardware is mapped at addresses D000-D02E, the sound hardware at D400-D7FF, parallel port at DC00-DCFF, and so on. So, your program should use RAM in the region between \$2000 and \$CFFF, which is 45,055 bytes – enough for most games. And, of course, page zero, which is the first 256 bytes, and runs from \$00-\$FF. By putting variables and data in page zero, the instructions to access them are only two bytes long – the opcode is the first byte and the address is the second byte. Accessing



An Easy 6502 screen with a short listing, showing the pixels on the right.

```
Preprocessing ...
Indexing labels ...
Found 3 labels.
Assembling code ...
Code assembled successfully, 79 bytes.
```

An Easy 6502 results window after a successful assemble.

memory in addresses after page zero is slower because it requires two byte addresses and those instructions are three bytes long.

Machine code, assemble!

To do this we need a 6502 assembler. There are many of these on the web. If you search, you'll probably end on <http://6502.org>, which is a terrific resource, with links to assemblers, source code examples, emulators and more. It's an http site, though, so your browser might complain.

The Assemblers page has links to many assemblers, but for this, it's easiest getting started by using a JavaScript web assembler, editor and debugger called Easy 6502 (<https://skilldrick.github.io/easy6502>).

Sort out your code

Our ultimate aim is to create a game that'll run on an emulated Commodore 64, but for now we'll run it on Easy 6502. The 6502 program is around 70 lines long and will sort an array of bytes using an inefficient bubble-sort. While Quicksort would be better, the code for that is longer and harder to understand. The original bubble-sort code came from 6502.org but has been changed a lot to run on Easy 6502.

Be aware to get it to assemble with Easy 6502, the labels must have a colon on the end.

The original had lines like this:

```
SORT8      LDY #$00
```

But for Easy 6502, it must be:

```
SORT8:     LDY #$00
```

You'll see this with many 6502 assemblers. Some need the colon on labels, some don't.

On Easy 6502, the addresses at \$200 map to the pixels in the black square, with \$200 the address of the first line, \$220 the line below, and so on. Our very simple program (see *screenshot, left*) puts bytes 1-6 into addresses \$200-\$205. There are 16 colours with values \$00-\$0F.

The **LDA #** value puts the value into the accumulator. **STA** then stores the accumulator value in the specified address. Note the difference between **LDA #\$01** and **LDA \$01**. The first puts the value 1 into the accumulator, the second loads the contents of memory location 1 into it.

```
define Count 6
```

This defines a constant called **Count** with the value 6. Then you can write code like:

```
LDA #Count
```

This loads the value 6 into the accumulator. Unfortunately, Easy 6502 doesn't have any way of specifying data bytes. These are bytes that are not machine code. In this case, they're the data that we want to sort.

Normal assemblers let you specify bytes or even strings. This program needs seven bytes of data. The first byte is the count and the remaining six bytes are

the unsorted values. If you used an assembler that specified those data bytes, then there would be no problem, but for Easy 6502 it takes an extra 20 instructions to put the values in successive memory locations \$40, \$41 and so on. The first value is the count (6) followed by the six values.

Assembling the program

Assembling is just a matter of clicking the Assemble button and checking that the panel below shows it assembled OK (see *screenshot, above*).

Before you run or step through the program, click Disassemble. This opens a window showing the generated machine code. Sometimes the assembler doesn't output an error but flags it in a different way. The opcode **LDA Data,Y** was incorrectly put in and it output **LDA \$FFFF,Y** instead of **LDA \$40,Y**. This is because there is no zero page **LDA ,Y** instruction. It was a typo and should have been **LDA Data,X**.

You can do a **LDX \$40,Y**, though, but only X can be loaded using Y in page zero. On the 6502 Opcodes page, you can see there are 16 instructions that use zero page,X but only two for zero page,Y. There is a 16-bit address,Y but we're using zero page to make the program shorter and faster.

Addressing modes

The 6502 has several addressing modes, including indirect, but those aren't used here. The main addressing mode used in the sort is called X-Indexed Zero Page. For instance, you'll see **LDA Data,X**. This means load the accumulator from the address formed

» WHAT IS 6502 ASSEMBLY CODE?

The 6510 in a Commodore 64 is a variant of a 6502 CPU and it runs 6502 programs. This is generated by compiling C code or writing 6502 assembly code and using an assembler.

The 6502 has three 8-bit registers called A, X and Y, and can address 64Kb of RAM. There are 60 or so unique assembly instructions but half a dozen or so different ways of accessing the value. Memory is accessed with a two-byte address – for example, \$4567. A \$ in front of a number means it's in hexadecimal (base 16).

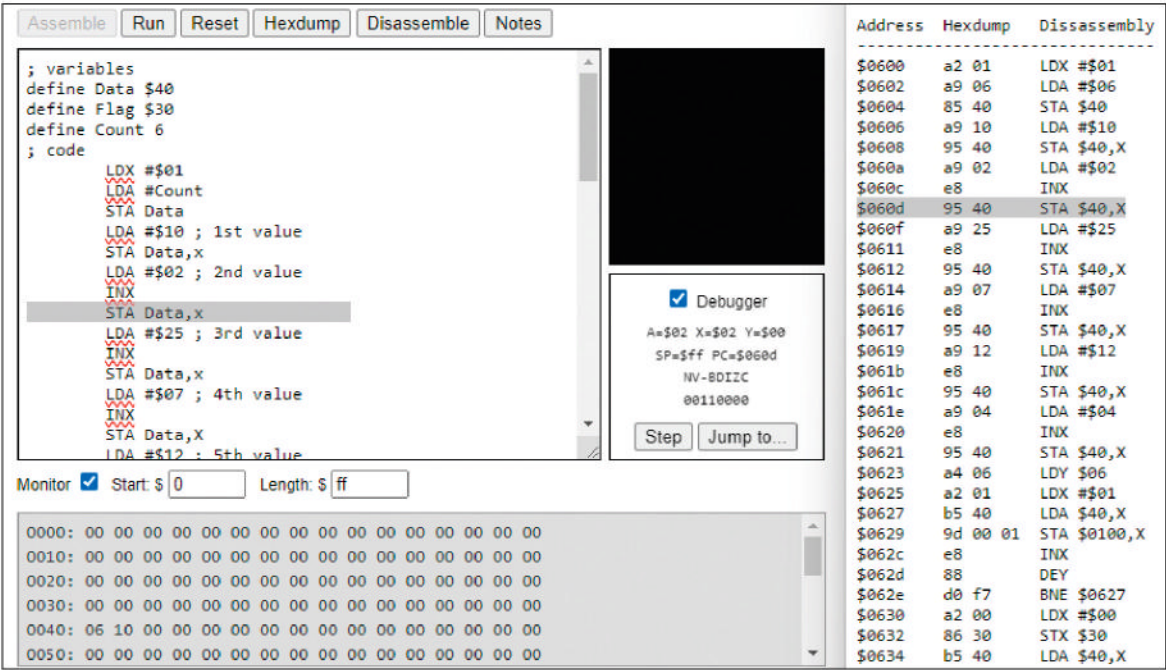
For instance, you put the value 12710 value into the accumulator with **LDA #\$7f** or load it from an address \$20 in page 0 (the first 256 bytes of RAM) with **LDA \$20**. Some of the other addressing modes use indexing with the X or Y registers. So, if X has the value \$10 in it, then **LDA \$30,X** is doing the same as **LDA \$40**. Moving a value from a register to memory is done with **STA**, **STX** or **STY**.

You can transfer the value in A to X or to Y with **TAX** and **TAY**, and the opposite direction with **TXA** and **TYA**.

The CPU carries out each instruction one by one, only deviating if a branch instruction is true, or at a **JMP** or **JSR** (jump to subroutine). With a clock running at 1MHz, it executes somewhere between 500,000 and 250,000 instructions per second.



Easy 6502 in the middle of a debugging session with manually highlighted source plus executing code and disassembly window on the right.



QUICK TIP

6502 programs are all about getting values into the accumulator, manipulating it, then storing the result in RAM. X and Y are mainly used for indexing but can be used to temporarily hold the accumulator.

by adding the contents of X to Data. Data was defined as \$40, so is a page-zero address. If Data is \$40 and X=2, this loads the accumulator from address \$42. If you see an instruction like `LDA ($40),Y`, it is called a Zero Page Indirect Y-Indexed. Fetch the two bytes at \$40 and \$41. Those form a 16-bit address. Now add the contents of Y to this and that is the address from which the byte is loaded.

What a debugger

Although you can just click the Run button, if it doesn't work, you need to get your hands dirty debugging. To debug it, after a successful assemble, tick the Debugger checkbox. Also make sure you tick the Monitor checkbox. You should also click the Disassemble button and keep the disassembly window handy. When you click the Step button, there's no way of seeing which line of code is being executed other than the PC= address and cross-referencing it against the disassembly window.

In the image (above), you can see the disassembly window on the right. The source line and equivalent in the disassembly window have been manually highlighted. By looking at the monitor window at the bottom, you can see the 0040 line starts with 06 10. After clicking Step, the 02 is written into \$0042. Under the pixel window is the debugger pane.

SP (stack pointer), PC (program counter) and NV_BDIZC are the CPU flags.

This block of code shows the values by copying them to the screen. As long as Y is non-zero, it loops to SHOW1. INX adds 1 to X and DEY subtracts 1 from Y.

```
LDY Data
LDX #$01
SHOW1: LDA Data,X
      STA $200,X
      INX
      DEY
      BNE SHOW1
```

When Y reaches 0, the code drops through and you should see some coloured pixels (see magnified screen dump, opposite page). Colours go from \$00-\$0F and the six colours here are F, 2, 5, A, 7 and 4 in that order.

A similar block to display the pixels is at the end of the program, but this time the pixels are shown two lines below (address \$240) and in sorted order.

Source listing

You can view this on GitHub with the name sort.6502 (<https://github.com/David-H-Bolton/Projects>).

```
; variables
      define Data $40
      define Flag $30
      define Count 6

; code
      LDX #$01 ; Initialise data
      LDA #Count
      STA Data
      LDA #$0F ; 1st value
      STA Data,x
      LDA #$02 ; 2nd value
      INX
      STA Data,x
      LDA #$05 ; 3rd value
      INX
      STA Data,x
      LDA #$0A ; 4th value
      INX
      STA Data,X
      LDA #$07 ; 5th value
      INX
      STA Data,X
      LDA #$04 ; 6th value
      INX
      STA Data,X
      ; Now show the values
      LDY Data
      LDX #$01
SHOW1: LDA Data,X
      STA $200,X
```

An Easy 6502 register window magnified during a debugging session.



» WHY IS ASSEMBLER NOT USED SO MUCH NOW?

One popular programmer myth is that to make your program run faster, you should rewrite it in assembler. This was really only true 40 years ago, when home computers first appeared.

Running at 1MHz, programs written in Basic were slow, so games were written in assembler.

Why not write in assembler now? Because programs written in optimised C/C++ or even C# run faster than anything you can write in assembly. Games programmers

switched to C in the early '90s and later on to C++.

Other reasons not to write in assembler are:

1 It's slow and tedious to write. Most code just moves values between memory and registers, or vice versa. There are no structures, strings or objects, no while or for loops; it's all bytes. Writing 10 lines of assembler takes the same time as 10 lines of C. Those 10 lines of C do way more than 10 lines of assembler.

2 You need to know the operating system that the program will run on and how to call into it to read keys, handle mouse clicks and read or write to files, request memory, and so on.

3 C/C++ and so on are portable. Assembly is only good for one CPU. To move an assembly language program from, say, an x64 CPU to ARM needs a complete rewrite. With C/C++, it takes just a recompile and minor changes.

```

INX
DEY
BNE SHOW1
; Now sort the values
SORT8: LDX #$00
      STX Flag ;Set Flag to 0
      LDA Data ; Fetch Count
      TAY      ; Transfer to Y
      INX      ; X is now 1- 1st element in list
      DEY      ; Dec Count in y
NXTTEL: LDA Data,X ; Fetch Element
      INX
      CMP Data,X ;Larger than next element?
      BCC CHKEND
      BEQ CHKEND
      ; Yes. Swap Elements
      PHA      ; Save Value on stack.
      LDA Data,X ; Get Next value
      DEX      ; AND STORE IT in previous byte
      STA Data,X
      INX
      PLA
      STA Data,X
      LDA #$FF ;Set flag
      STA Flag
CHKEND: DEY ;Reached the end of the list?
      BNE NXTTEL ;No. Fetch next element
      BIT Flag ;YES. Flag still off?
      BMI SORT8 ;No. Go through the list again
; Now show the values again two lines below
      LDY Data
      LDX #$01
SHOW2: LDA Data,X
      STA $240,X
      INX
      DEY
      BNE SHOW2
      RTS      ; Finished all sorted
; end of program

```

The tasty source

The above code performs a bubble sort. It indexes through the list, comparing each element with the one

after it. If the value is bigger than the one following it, then they need to be swapped and the **Flag** variable is set. When they are all ordered, the flag is no longer set and it finishes.

After a CMP instruction, the carry flag is set when the value in the accumulator is greater than the value it's being compared against. It's clear if the value is greater than the accumulator.

NXTTEL: LDA Data,X ; Fetch Element

INX

CMP Data,X ;Larger than the next element

BCC CHKEND

BEQ CHKEND

If the code gets to here, then the accumulator value was greater.

Swapping is done by pushing the current element in the accumulator on the stack. Then the next element is read. X is decremented and it is stored in the current element's position. Then X is incremented, the value popped off the stack and saved in the next element.

Before sorting, the values in memory are:

0040: 06 0f 02 05 0a 07 04

And after the sort they are:

0040: 06 02 04 05 07 0a 0f

Remember, the first byte is the count (6); only the bytes at \$41-\$46 are sorted.

After a swap, the value in the accumulator is 255 and it loops back and runs the sort again. Only when it did no swaps will it drop through and finish.

Going C64

Easy 6502 is a handy way to test 6502 code, even if you don't have a proper 6502 computer or an emulator. However, it is a bit limited. Next month, in the second article in this series, we'll use a 6502 assembler to generate code for a Commodore 64 and then run it on an emulator. **LXF**

QUICK TIP

Put all variables in page zero. It speeds up your program and reduces code size. You only have 254 bytes in page zero to play with on a Commodore 64, as addresses 0 and 1 have special uses.

An Easy 6502 output window after a successful run.



» FOR MORE OUT OF DATE STUFF... Subscribe now at <http://bit.ly/LinuxFormat>

LXF SHELL

The LXF Shell in... the redirection redemption

Ferenc Deák continues writing shell-enhancing redirection features and hiding the odd film title for you to spot, the little scamp!



OUR EXPERT

Ferenc Deák is sure that the usual suspects of programming languages fit for writing a shell have been exhausted, so he sticks to C++.

QUICK TIP

The code for the LXF Shell can still be found at <https://github.com/fritzone/lxf-shell>.

Previously in our LXF Shell series, we reached a stage where we could execute applications (LXF310), and looked at how to properly redirect the output of an application (LXF310). For this latest instalment, we are planning how to implement the input redirection for the shell, and as a final touch, how to bring all these together and create a shell that can sequentially execute applications by piping input and output between them.


Input redirection

Linux input redirection in a shell enables you to change the source of input for a command or program. It is accomplished with the `<` operator, by adding `<` to redirect the standard input (stdin) of a command or program from a file instead of the keyboard after the command, and specifying a filename whose content is read and used as input for the command to be executed. For example: `wc -l < input.txt`.

In this example, `wc -l` reads its input from the `input.txt` file instead of waiting for keyboard input, and counts the lines in that file. Input redirection can be particularly useful when working with scripts and batch processing, because it enables you to automate tasks and process large volumes of data without manual input.

From a programming point of view, input redirection happens in a similar way to output redirection, using pipes, as the following program exemplifies. To save precious space, we have omitted the error checking; however, the example code found at <https://github.com/fritzone/lxf-shell> has all the necessary checks and comments, so be sure to check it out.

```
int main() {
    int pipe_fd[2] = {-1, -1};
```



This diagram represents how pipe interactions can happen in Linux.

```
pipe(pipe_fd);
pid_t child_pid = fork();
if (child_pid == 0) {
    close(pipe_fd[1]);
    dup2(pipe_fd[0], STDIN_FILENO);
    close(pipe_fd[0]);
    char program_name[] = "wc", *args[] = {program_name, NULL};
    execvp(program_name, args);
} else {
    close(pipe_fd[0]);
    FILE *input_file = fopen("input.txt", "r");
    char buffer[4096] = {0};
    size_t bytesRead = 0;
    while ((bytesRead = fread(buffer, 1, sizeof(buffer), input_file)) > 0)
        write(pipe_fd[1], buffer, bytesRead);
    close(pipe_fd[1]);
    wait(NULL);
```

```
}
}
```

This very short program uses pipes for inter-process communication, spawning a child process to execute the **wc** command. The child's standard input is redirected to read from a pipe, and the parent process reads the content of **input.txt** and writes it to the pipe. The child process, executing **wc**, prints the number of words in the input file to stdout (being the result of **wc**). As explained in our previous articles, we use **fork** and **execvp**, and wait to manage the processes.

Double indirection

Our next goal is to bring all this together into our shell, and extend it with the capability to redirect both the input and output of the executed commands. The bases have been well laid by now, so implementing it is a piece of cake.

We need to extend the **execute** function with a new parameter, let's call it **const std::string& stdinInputFile**, representing the file from which the input is read.

We need to check whether there is a need to redirect the input, and create the corresponding pipes, at the same time checking for errors, and this is done as follows:

```
const bool needsStdinRedirect = ! stdinInputFile.empty();
int pipeStdinFd[2] = {-1, -1};
if(needsStdinRedirect){
    if (pipe(pipeStdinFd) == -1) return 1;
}
```

In the child process (the true branch of the **if (pid == 0)** check), we need to close the corresponding stdin pipes:

```
if(needsStdinRedirect) {
    close(pipeStdinFd[1]);
    dup2(pipeStdinFd[0], STDIN_FILENO);
    close(pipeStdinFd[0]);
}
```

And in the parent process, we need the following code to add the required functionality:

```
if(needsStdinRedirect) {
    close(pipeStdinFd[0]);
    FILE *input_file = fopen(stdinInputFile.c_str(), "r");
    if (input_file == NULL) return 1;
    char buffer[4096] = {0};
    size_t bytesRead = 0;
    while ((bytesRead = fread(buffer, 1, sizeof(buffer), input_file)) > 0)
        write(pipeStdinFd[1], buffer, bytesRead);
    fclose(input_file);
    close(pipeStdinFd[1]);
}
```

This piece of code makes it properly read the input from the input file, and redirect it where it is supposed to be. With all this in mind, our redirecting shell is done, and it functions exactly as the diagram (see *opposite page*) describes:

Popeye, the shellorman

Now that we have all the input and output redirections in place, it is time to move on to our next target in our ever-evolving shell: piping the commands. As the boxout (*above-right*) explains, the piping and chaining of commands is done with the **|** symbol, so our next

» BLOWING YOUR OWN PIPE

Piping is a powerful feature of the Linux shell that promotes the composition of simple, specialised commands to perform complexly chained tasks efficiently. It enables you to build intricate workflows by combining the strengths of individual commands.

It is a mechanism that allows the output of one command to be used as the input for another. It enables you to chain multiple commands together, creating a pipeline for data to flow from one command to the next. The pipe symbol (**|**) is used to denote the connection between commands in a pipeline. Here's a basic explanation of how piping works:

1 Command execution

When you enter a command in the Linux shell, the command is executed, and it produces output (stdout) that is typically displayed on the terminal.

2 Pipe symbol (|)

The pipe symbol (**|**) is used to connect the output of one command to the input of another. It tells the shell to take the output from the command on the left and use it as the input for the command on the right.

3 Data flow

The data produced by the first command is sent through the pipe to the second command. This enables you to perform a series of operations on the data without explicitly creating temporary files or storing intermediate results.

4 Chaining commands

You can chain multiple commands together in a pipeline, creating a sequence of operations. For example:

```
bash
command1 | command2 | command3
```

Here, the output of **command1** is passed to **command2**, and the output of **command2** is passed to **command3**.

5 Use cases

Piping is commonly used for tasks such as filtering, sorting, searching and transforming data. For instance, you might use **grep** to filter lines containing a specific pattern and then use **sort** to sort those lines.

step is to extend our application to handle that. There is just one very important thing to note: for the moment, we are not planning for input/output redirection and piping to be used together in the same command, because our shell is still toddling around in its baby shoes. Please be aware that this article aims to provide educational value, and our intention is not to intimidate readers by concocting code with excessive complexity that would leave everyone dazed and confused. The fact is that for the moment, there is no consensus among shell developers concerning this, so the command **ls -l > 123.txt | wc -l** gives different results when run on various shells (for example, **Zsh** prints the actual number of lines, **10**, for the build directory of our projects, **Bash** and **Fish** print **0**, while **Csh** just whines about an ambiguous output redirect), so we took the decision not to support this kind of scenario, and in practice this means that we need to split the execution path



THE BEST OF BOTH WORLDS

If you have a background in C and/or C++, you might have noticed that we tend to mix both C and C++ APIs and concepts in this article. Before programming language zealots of either camp jump upon us, accusing us of interbreeding their favourite with the enemy, we should inform you that this approach is intentional. By using a careful combination of selected C and C++ features and notions, we bring forth dozens of advantages, and foster a harmonious synergy that enhances both efficiency and flexibility for our program.

While C excels in low-level system programming and is renowned for its

efficiency, raw power and minimalistic design, C++ provides a higher level of abstraction, facilitating the creation of complex applications, and since modern problems require modern solutions, the latest C++ standards introduced constructs that greatly enhance the safety and security of modern applications. By seamlessly blending these languages, we can harness the power of C for performance-critical components and employ C++ for more intricate, object-orientated aspects of our shellish project, while benefiting from a more concise and shorter syntax for certain scenarios.

Moreover, the Linux kernel, predominantly written in C, exhibits a natural affinity towards C programs, and almost all lower-level libraries and APIs we use have a C interface. Incorporating C++ into this environment can be effortless, provided it adheres to certain guidelines. This amalgamation permits us to enjoy the benefits of C++ features such as containers, RAII-style resource management, or everyone's favourite, `std::string`, which can be a blessing when it comes to handling character data. The strategic interweaving of C and C++ in a project is not merely a pragmatic choice; it is a well-advised one.

depending on whether the input identifies itself as a chain of commands waiting to be executed by means of piping, or just a normal command, which might have input/output redirection and parameters, and then handle both of them accordingly. Apologies for the long-winded sentence.

The following code is a short C++ function that identifies whether the input belongs to the first or to the second category.

```
bool containsPipeSymbol(const std::string& input) {
    bool insideQuotes = false;
    for (char c : input) {
        if (c == '"' || c == "'") insideQuotes = !insideQuotes;
        else if (c == '|' && !insideQuotes) return true;
    }
    return false;
}
```

It might seem complicated, but its purpose is to identify whether a string contains the `|` symbol or not, if and only if that `|` symbol is not to be found in single or double quotes, in order to avoid a situation like `echo "|"` being identified as the piped commands `echo "` and `"`, and this function is as good as it gets.

Because the path that executes commands with redirection and parameters must be familiar to our

readers by this point, we are now presenting you with the alternative path: the one that creates a multitude of pipes and runs various applications by passing the information between them using the aforementioned pipes.

```
bool runPipedCommands(const
std::vector<std::string>& commands) {
    const std::size_t numPipes = commands.size();
    int pipefds[2 * numPipes] = {-1}, status = 0;
    for(int i = 0; i < numPipes; i++)
        if(pipe(pipefds + i*2) < 0) return false;
    int j = 0;
    auto cmd = commands.begin();
    while(cmd != commands.end()) {
        if(const pid_t pid = fork(); pid == 0) {
            auto nextCommand = cmd;
            ++ nextCommand;
            if(nextCommand != commands.end())
                if((dup2(pipefds[j+1], STDOUT_FILENO)) == -1)
                    return false;
            if(j != 0)
                if((dup2(pipefds[j-2], STDIN_FILENO)) == -1)
                    return false;
            for(int i = 0; i < 2 * numPipes; i++)
                closeFd(pipefds[i]);
            auto split = splitStringByWhitespace(*cmd);
            const char** arguments = new const char*[split.
size() + 1];
            for(size_t i = 0; i < split.size(); i++) arguments[i] =
split[i].c_str();
            arguments[split.size()] = nullptr;
            execvp(arguments[0], const_cast<char *
const*>(arguments));
        }
        else
            if(pid < 0) return false;
        ++ cmd;
        j+=2;
    }
    for(int i = 0; i < 2 * numPipes; i++) closeFd(pipefds[i]);
    for(int i = 0; i < numPipes + 1; i++) waitpid(-1, &status,
0);
    return true;
}
```

When not compressed to destroy readability, the shell's code is pretty lengthy and easy to follow.

```
474 //+
475 //+ Brief Will run each of the commands in the vector in a way that the output of the current command
476 //+ will be used as input for the next command, like the piping in Linux shell works
477 //+
478 //+ @param commands the commands
479 //+
480 //+ \return true in case of success, false otherwise
481 //+
482 bool runPipedCommands(const std::vector<std::string>& commands)
483 {
484     const std::size_t numPipes = commands.size();
485     int pipefds[2 * numPipes] = {-1};
486     for(int i = 0; i < numPipes; i++)
487     {
488         if(pipe(pipefds + i*2) < 0)
489         {
490             log_err << explain_pipe(pipefds + i*2) << std::endl;
491             return false;
492         }
493     }
494     auto cmd = commands.begin();
495     int j = 0;
496     while(cmd != commands.end())
497     {
498         if(const pid_t pid = fork(); pid == 0)
499         {
500             //check if we are not running the last command
501             auto nextCommand = cmd;
502             ++ nextCommand;
503             if(nextCommand != commands.end())
504                 if((dup2(pipefds[j+1], STDOUT_FILENO)) == -1)
505                     return false;
506             if(j != 0)
507                 if((dup2(pipefds[j-2], STDIN_FILENO)) == -1)
508                     return false;
509             for(int i = 0; i < 2 * numPipes; i++)
510                 closeFd(pipefds[i]);
511             auto split = splitStringByWhitespace(*cmd);
512             const char** arguments = new const char*[split.
size() + 1];
513             for(size_t i = 0; i < split.size(); i++) arguments[i] =
split[i].c_str();
514             arguments[split.size()] = nullptr;
515             execvp(arguments[0], const_cast<char *
const*>(arguments));
516         }
517         else
518             if(pid < 0) return false;
519         ++ cmd;
520         j+=2;
521     }
522     for(int i = 0; i < 2 * numPipes; i++) closeFd(pipefds[i]);
523     for(int i = 0; i < numPipes + 1; i++) waitpid(-1, &status,
0);
524     return true;
525 }
```

We apologise that at this stage the code sequences are already starting to become very complex but, well, writing a shell is not a straightforward assignment. To shed a little light on the situation, here is a breakdown of what is happening in the previous piece of code. The function takes a vector of strings, where each string is a command to be executed (with or without parameters). This vector was created in the main function by means of a newly introduced function: **splitStringWithDelimiter**, which as its name suggests, splits a string where the tokens are bordered by a specific delimiter.

When the function **runPipedCommands** enters, first it creates an array of file descriptors for the pipes, with the number of pipes equal to the number of commands in the following manner: as we remember from the previous tutorial, the pipe call fills in a two-element array with two new file descriptors – the first one (**index 0**) can be read from, and the second one (**index 1**) can be written to. So, the system call `pipe(pipefds + i*2)` fills up the pipe array `int pipefds[2*numPipes]` with a pair of read/write pipe pairs, starting from the first position where the destination index is the address of `pipefds` for the `i==0` in the loop, then moving on to the next one (the second element), with the destination index being the address of `pipefds+2*sizeof(int)`, this intricate value being obtained by doing pointer arithmetics on an integer pointer (to which the integer array is translated automatically by the compiler) in the code sequence `pipefds+i*2`, and so on, until all the `pipefds` array contains the correct number of pipes for each command.

Then it starts iterating over the commands, and for each command, it creates a new process using the **fork** system call, as we have seen in the previous parts, where we needed this forking to execute commands.

The real magic of piping happens in the child process, with an intricate set of **dup2** calls. If the command we are executing is not the last command (verified by the check `if(nextCommand != commands.end())`), it duplicates the write end of the current pipe to standard output using the following system call: `dup2(pipefds[j+1], STDOUT_FILENO)`. Similarly, if it's not the first command (verified by the check `if(j!=0)`), it duplicates the read end of the previous pipe to standard input with `dup2(pipefds[j-2], STDIN_FILENO)`. As a slightly irrelevant side note, the check `if(j!=0)` is equivalent to `if(j)` but we preferred to use the more verbose and explicit form of it to make the code a little less confusing.

After the function closes all the unnecessary pipe file descriptors as required, and splits up the command into separate tokens, creates an argument array, and executes the command using **execvp**, this sequence of code is identical to the one found in the previous parts of this tutorial.

After the loop, the parent process waits for all child processes to complete using **waitpid** by waiting for `numPipes+1` processes to ensure that all child processes have finished.

With this in place, we can conclude that the piping and redirection features of our shell are done; certainly,

```
lxfsh$ wc <Makefile >Mkf.counts
lxfsh$ cat Mkf.counts
      181      585      5279
lxfsh$ wc Makefile
      181      585      5279 Makefile
lxfsh$ ls -l Makefile
-rw-rw-r-- 1 fld fld 5279 nov.   8 13:13 Makefile
lxfsh$
```

reaching this stage in our series of tutorials should enable you to tinker more with the code, and come up with your own implementation of ideas, suggestions and features, or simply contact us to let us know what feature requests you would like to be incorporated in the shell's code.

Back to the future

For our next instalment, we are going to be delving into something even more intriguing, namely how to implement a plugin architecture that will enable us to have a set of built-in commands for our shell, each performing the operations it was meant to do. We will exemplify this with the `cd` command, to give our shell something of a playground in which it can leave the confines of its build directory.

However, before we do that, we need to consider another very important aspect of promoting proper programming practice. As our tiny shell grows and grows, and more and more code is incorporated into it, and it becomes more and more complex, the need arises for us to be able to ensure that it always operates correctly.

There is more than one way of assuring the correctness and validity of our shell, and one of these methods is the inclusion of unit testing early in the development process. So, we are going to dedicate half of next month's instalment to the sole purpose of properly preparing the programming playground to perpetrate the prime proper programming practice, namely unit testing.

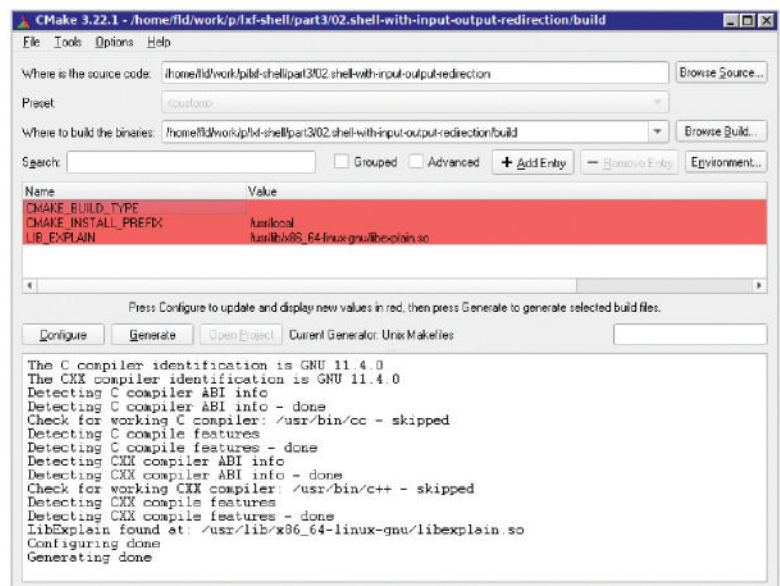
Happy coding! **LXF**

The identical counts validate the code and verify that the redirection works as expected.

QUICK TIP

If you're not up to speed on C coding, take a trip to our friends at Codecademy: www.codecademy.com/learn/paths/c.

Properly set up, `cmake-gui` helps a lot with the building process of the shell.



» **DON'T MISS ANY OF THE SERIES!** Subscribe now at <http://bit.ly/LinuxFormat>

NEXT MONTH

LXF313
will be on sale
Tuesday 5th
March

OLD-SCHOOL HACKING 101

Jonni Bidwell drags us kicking and screaming back to the '80s to rediscover the true meaning of hack.

That's a nice filesystem...

Be a shame if something nasty happened to it! We explore what makes a filesystem, what can go wrong, and why you'd change.

Coding your CAD

We've let Tam Hanna loose with *OpenSCAD* (as he wrote the book on it) and we might regret this, but he's using it to print cigars...

Easy retro gaming

Dedicated distros help you get retro gaming with ease – we put five of the best to the test, so we can relive our gaming past.

RAID your SSD

Have a pile of SSDs lying around? Perhaps it's time that you combined them into one speedy storage pool...

Content of future issues subject to change. Especially if Jonni's canal boat batteries run dry...

LINUX
The #1 open source mag **FORMAT**

Future Publishing Limited,
Quay House, The Ambury, Bath, BA1 1UA
Email contact@linuxformat.com

EDITORIAL

Editor-in-chief Neil Mohr
Art editor Fraser McDermott
Production editor Katharine Davies
Group editor-in-chief Graham Barlow
Group art director Warren Brown
Editorial contributors
Mirza Bahic, Mike Bedford, David Bolton,
Stuart Burns, Christian Cawley, Ferenc Deák,
Nate Drake, Nick Peers, Les Pounder,
Michael Reed, Jacob Ridley, David Rutland,
Mayank Sharma, Shashank Sharma, Davey Winders
Cover illustration Magictorch.com
Raspberry Pi is a trademark of the Raspberry Pi Foundation.
Tux credit: Larry Ewing (lewing@isc.tamu.edu).

Content production Adequate Media Limited

ADVERTISING

Commercial sales director Clare Dove
clare.dove@futurenet.com
Advertising director Lara Jaggon
lara.jaggon@futurenet.com
Account director Andrew Tilbury
andrew.tilbury@futurenet.com

INTERNATIONAL LICENSING

Head of print licensing Rachel Shaw
Linux Format is available for licensing and syndication.
To find our more contact us at licensing@futurenet.com
or view our content at www.futurecontenthub.com

NEW SUBSCRIPTIONS & BACK ISSUES

Web www.magazinesdirect.com
UK 0330 333 1113 World +44 (0) 330 333 1113

EXISTING SUBSCRIPTIONS

Web www.mymagazine.co.uk
UK 0330 333 4333 World +44 (0) 330 333 4333
Subscription delays: Disruption remains within UK
and international delivery networks. Please allow up
to seven days before contacting us about a late
delivery to help@magazinesdirect.com

CIRCULATION

Newstrade & retail category director Ben Oakden

PRODUCTION AND DISTRIBUTION

Group head of production Mark Constance
Production manager Nola Cokely
Senior ad production manager Jo Crosby
Digital editions manager Jason Hudson

THE MANAGEMENT

Managing director technology group Paul Newman
Global head of design Rodney Dive
Commercial finance director Tania Brunning
Printed by Wyndeham Peterborough, Storey's Bar
Road, Peterborough, Cambridgeshire, PE1 5YS
Distributed by Marketforce, 5 Churchill Place,
Canary Wharf, London, E14 5HU
www.marketforce.co.uk For enquiries email:
mfcommunications@futurenet.com

Linux® is the registered trademark of Linus Torvalds in the US and other countries.
GNU/Linux is abbreviated to Linux throughout for brevity. Where applicable,
code printed in this magazine is licensed under the GNU GPL v2 or later. See
www.gnu.org/copyleft/gpl.html. All copyrights and trademarks are recognised
and respected.

Disclaimer All contents © 2024 Future Publishing Limited or published under
licence. All rights reserved. No part of this magazine may be used, stored,
transmitted or reproduced in any way without the prior written permission of the
publisher. Future Publishing Limited (company number 2008885) is registered in
England and Wales. Registered office: Quay House, The Ambury, Bath BA1 1UA. All
information contained in this publication is for information only and is, as far as we
are aware, correct at the time of going to press. Future Publishing Limited cannot
accept any responsibility for errors or inaccuracies in such information. You are
advised to contact manufacturers and retailers directly with regard to the price of
products/services referred to in this publication. Apps and websites mentioned in
this publication are not under our control. We are not responsible for their contents
or any other changes or updates to them. This magazine is fully independent and
not affiliated in any way with the companies mentioned herein.

If you submit material to us, you warrant that you own the material and/or have
the necessary rights/permissions to supply the material and you automatically
grant Future Publishing Limited and its licensees a licence to publish your
submission in whole or in part in any/all issues and/or editions of publications, in any
format published worldwide and on associated websites, social media channels
and associated products. Any material you submit is sent at your own risk and,
although every care is taken, neither Future nor its employees, agents,
subcontractors or licensees shall be liable for loss or damage. We assume all
unsolicited material is for publication unless otherwise stated, and reserve the right
to edit, amend and adapt all submissions. All contents in this magazine are used at
your own risk. We accept no liability for any loss of data or damage to your systems,
peripherals or software through the use of any guide. Notes: the giant gleaming Tux
sat as a beacon of hope in the midst of a dystopian landscape was poo-pooed, so
retro-style rockets it is!

We are committed to only using magazine paper derived from
responsibly managed, certified forestry and chlorine-free
manufacture. The paper in this magazine was sourced and
produced from sustainable managed forests, conforming to
strict environmental and socioeconomic standards.



Future is an award-winning international media group and
leading digital business. We reach more than 57 million
international consumers a month and create world-class
content and advertising solutions for passionate
consumers online, on tablet & smartphone and in print.

FUTURE Connectors. Creators. Experience Makers.


Future plc is a public company quoted on the London Stock Exchange
(symbol: FUTR)
www.futureplc.com

Chief Executive Officer Jon Steinberg
Non-Executive Chairman Richard Huntingford
Chief Financial and Strategy Officer Penny Ladkin-Brand

Tel +44 (0)225 442 2444

CREDIT: Magictorch

ipso. For press freedom
with responsibility
www.linuxformat.com



The Electronic Frontier Foundation is the leading nonprofit organization defending civil liberties in the digital world. Founded in 1990, EFF champions user privacy, free expression, and innovation through impact litigation, policy analysis, grassroots activism, and technology development. We work to ensure that rights and freedoms are enhanced and protected as our use of technology grows.

EFF.ORG

ELECTRONIC FRONTIER FOUNDATION

Protecting Rights and Promoting Freedom on the Electronic Frontier



**THE
BRAIN
TUMOUR
CHARITY**

A CURE CAN'T WAIT

**BRAIN TUMOURS
MOVE FAST.
WITH YOUR HELP,
WE CAN TOO!**

We're working to create a future where brain tumours are curable.
We urgently need your help to accelerate research.

Text DEFEAT5 to 70507 to donate £5, please help us to find a cure.

thebraintumourcharity.org

© The Brain Tumour Charity 2020. Registered Charity in England and Wales
(1150054) and Scotland (SC045081)



Registered with
**FUNDRAISING
REGULATOR**