

FORMA

TUX-MNT20

o 2.10 int 2.1 PLUS! LOW-COST SERVERS REPLACE GOOGLE MAPS FASTER STEAM GAMING

The #1 open source mag

SAFER CODE

Using dynamic

analysis pipelines

CUSTOMISE MINT 20!

Get the system of your dreams with the freshest distro and a tricked-out custom desktop

67 pages of tutorials & features

Emulate the Atari 800 for fun and old-school games

Create a classic pin-hole camera using a Raspberry Pi

Build more secure websites with full-strength certificates



CODING OUTRUN

Recreate classic racing games with clever 3D tricks

WHO IS MOST AT RISK? WHAT CAN PREVENT IT?



PLEASE TAKE THE SURVEY gmul.ac.uk/covidence















MEET THE TEAM

This issue we're customising our Linux desktops. What's your go-to desktop of choice and do you have a top customising tip for our dear readers?



Jonni Bidwell

If you like i3 or Sway (or just minimal terminal emulators), spend some time making a semidecent .Xresources file, then use git to track changes to it and other dotfiles. Now not even you will be able to break your configuration!



Les Pounder

My love is Openbox which stems from using Crunchbang Linux. The dark aesthetic is soothing and I can tweak Conky to show pertinent information. Tweaking the Conky settings is made easy, they all live in a text file

which I edit, save and then reload Conky to see the effect.



David Rutland

I use the i3 window manager. It's infinitely customisable via its config file, keyboard driven, and has insanely low overheads. No need to use the mouse to close a window or start up a program. A bonus is that users of other OSes will find your machine almost impossible to use.



Shashank Sharma

My current desktop, which I've been a fan of for many years, is Budgie. I've configured it with a disappearing panel at the bottom featuring commonly used apps. I also can't live without workspaces, and use four, with

keyboard shortcuts to switch between them.



Alexander Tolstov

I personally prefer Plasma, although as the author of Hotpicks I also run a second machine with Gnome where I can better test the best apps from the GTK and GNOME ecosystem. In Plasma I love to

tinker with the Configurable Button widget. I've already built a mini-dashboard of various switchers!

Customised heart



One of the strengths of the open source world is its inherent ability to be customised. I think that's a real attraction for users that grew up using computers pre-2000s, while the locked-down nature of modern devices. Windows and macOS forces anyone that wants to tinker in depth - or just learn how systems works - naturally towards open source operating systems.

Even on the most superficial level having total choice about how your

desktop looks and works is key for many Linux users. Just look at the Gnome Wars and how impassioned people get about fundamental changes to desktop design! As we have a fresh release of Linux Mint with its custom Cinnamon desktop, we thought it was high time we took a deep look at desktop customisation.

We'll dive into getting the most from Linux Mint, tweak its Cinnamon desktop to death, then look at how you can swing other cool desktop options into action. Of course this will work with almost all distros, so you can customise to your heart's content.

Elsewhere we look at low-cost virtual private servers, see if GIMP is still the best paint package out there, get more from Steam gaming, and build a pinhole Pi camera. We also have some cool coding guides and loads more fun stuff, so enjoy!





Contents



REVIEWS

AMD Ryzen 7 5700G 21 The AMD APU makes a more-thanwelcome return for Chris Szewczyk.



Pop!_OS 21.04 2	3
A desktop distro that offers a different way	y
of working Neil Bothwick spends a night	
on the tiles with Pop!_OS.	

NixOs 21.05	24
A complete working distro created from a	а
single configuration file. It sounds different	nt,
and Neil Bothwick discovers that it is.	

Linux Mint 20.2	26
Jonni Bidwell is running low on mint	
metaphors, but the promise of some	
budget mojito mix expedited this review.	

Griftlands	28
Management loves a tough negotiation	but
not with the staff, so Sam Greer will be	
getting a visit from the LXF shock troops	5



CUSTON CONTRACTOR OF CONTRACTO

PLUS!

ROUNDUP



IN-DEPTH



It's a virtual server world 54 From a caravan in Wales, David Rutland takes you on a touring holiday through the world of virtual private servers so you can run your own (a server, not a caravan).

CONTENTS

Pi USER

Raspberry Pi news 44 Introduced by Josh Lowe. The new Pi recycling scheme is launched, we look at a new Argon NAS and Pi goes to the movies!

Les Pounder wants another slice of Raspberry Pi and he needs it real 'qwiic'!

Drawing with Scratch 46 Les Pounder shows how we can use simple maths to draw patterns and shapes.

Build a pinhole camera 48 Interchangeable lenses help us to experiment with that historical curiosity, the pinhole camera. Mike Bedford shows you how.

Build a better Nextcloud 52 Christian Cawley demonstrates how to get the most out of your Raspberry Pi-powered Nextcloud server.





CODING ACADEMY

Build a dynamic security pipeline 88 Dynamic Analysis Security Testing takes centre stage in the third and final instalment of our Web Application Security series with Tim Armstrong.

Classic pseudo-3D racing effects 92 With a bag of 10p pieces in hand, Andrew Smith whisks us back to the classic arcade days to recreate pseudo-3D racing games.



REGULARS AT A GLANCE

Google takes shots at the handling of kernel security issues, StackOverflow returns the results of its dev survey, all the details on LibreOffice 7.2, Mozilla Common Voice improves and Russia makes RISC-V CPUs.

Kernel watch	
--------------	--

Answers 12 Sorting out Deskreen for the Pi, read-only USB sticks, Windows 98 emulation, kernel update GRUB madness, DVD booting UEFI tricks and Onap NAS fixes.

Subscriptions	
Back issues	
Get hold of previous Linux Format e	ditions,

but act fast because they soon sell out!

Overseas subscriptions 63

HotPicks 81 Alexander Tolstoy hasn't got time to go and fix dodgy Russian International Space Station modules, he's too busy testing modules like PDF Mix Tool, DroidCam, Fractal, Polyfoto, ImprovedTube, LSD, Pam usb, Billiards, Blues Brothers, EasyEffects and Latte Dock.

DVD pages	96
Next month	98



DVD pages Jonni Bidwell has everything that you need to know on Mint 20.2 and Elementary 6.0.

TUTORIALS

TERMINAL: Replacing man 58 A fondness of reading is a good habit to inculcate, but Shashank Sharma believes it requires more than spending quality time with man pages.

COPYQ: Clipboard 2.0 60 Nick Peers dons his best hiking gear and goes hunting for a simple single clipboard tool to rule them all. Lo and behold, CopvO is happy to oblige.

EMULATION: Atari 800 64 Les Pounder goes back to a time before E.T. cartridges were buried in the desert and 48KB was considered to be plenty of RAM.

CERTIFICATES: Web security 68 Stuart Burns walks you through what certificates are, how to set up a certificate server running on your local network, and how to make use of it.

QGIS: Open source mapping 72 If you like Google Maps you're going to love QGIS. Mike Bedford shows you how to get started and introduces some of its base mapping functionality.

STEAM: Better Linux gaming 76 Michael Reed looks at the tips and tricks you can use to run a larger variety of games under Steam, including those designed to run on Windows.



Newsdesk

KERNEL

Google slams Linux kernel

The tech behemoth believes the Linux kernel is in need of major investment to improve security.

ees Cook's recent blog post, Linux Kernel Security Done Right (https://bit.ly/lxf281googleblog) makes for interesting reading. Cook, who's part of Google's Open Source Security Team, spells out what he see as the major security issues facing the Linux kernel, while also offering advice on how to fix it.

While noting that "Linux still remains the largest collaborative development project in the history of computing," and complimenting its large community. Cook says, "What's still missing, though, is sufficient focus to make sure that Linux fails well too," and that "When flaws do manifest, it's important to handle them effectively."

He highlights areas that need addressing, and starts by calling for substantial investment to make sure the code is as robust as possible. This should stop bugs from appearing as often as they currently do.

As Cook says, "Rather than only taking a onebug-at-a-time perspective, preemptive actions can stop bugs from having bad effects," and that Linux will need to adapt to do this, especially with Linux being written in the C language. This, the author warns, will mean that it'll continue to "have a long tail of associated problems."

In the blog, Cook says that the stable bug-fix only release of the kernel comes out with about 100 new fixes every week. This high rate means that downstream vendors are faced with three choices: either to ignore all fixes, prioritise the "important" ones, or apply them all.

Applying them all, Cook argues, is the only realistic option – but taking this approach has some serious implications. While it means all important bug fixes are applied, it can also bring

problems – and many vendors are unable to test the updates before applying them to check for any problems.

So, what can be done? Cook suggests this is a "simple resource allocation problem, and is more easily accomplished than might be imagined: downstream redundancy can be moved into greater upstream collaboration." Essentially, he's calling for more engineers to review the code and fix bugs earlier, as well as test the kernels during development. According to the blog, the Linux kernel and its toolchains are "underinvested by at least 100 engineers, so it's up to everyone to



Google has some big ideas for boosting Linux kernel security, but how realistic are they?

GOOGLE'S GREAT EXPECTATIONS Employing more engineers may be something a huge company like Google can do, but for smaller projects and businesses, this "simple resource allocation problem" may not be as simple..."

bring their developer talent together upstream. This is the only solution that will ensure a balance of security at reasonable long-term cost."

This is, of course, easier said than done. Employing more engineers may be something a huge company like Google can do, but for smaller projects and businesses, this "simple resource allocation problem" may not be as simple as Google makes out.

SOFTWARE

LibreOffice 7.2 released

The latest version of the venerable office suite is now available. Will more businesses choose to use it?

ibreOffice has reached version 7.2, and it comes at an interesting time for the office suite. Over the course of the pandemic, many people have found themselves working from home, and the prospect of shelling out for a *Microsoft Office* licence (or *Office* 365 subscription) was an unappealing one.

This should have been an open goal for LibreOffice, but many organisations are sticking with Microsoft's offering, while Google Docs appears to be the free alternative of choice, despite its limitations (and taste for telemetry).

LibreOffice 7.2 is the ideal opportunity for the Document Foundation to get its office suite in front of new people. Perhaps the biggest feature that can help with this is improved compatibility with *Microsoft Office*, to ensure that *Office* XML files can be opened in *LibreOffice* as accurately as possible. As such, there is now improved interoperability with .docx, .xlsx, .pptx and older. doc files. As the Document Foundation points out in its release announcement (see https://bit.ly/ lxf28Ilibreoffice). "Microsoft files are still based on the proprietary format deprecated by the ISO in April 2008, and not on the ISO approved standard, so they embed a large amount of hidden artificial complexity." Much work has gone in to making it work with *LibreOffice*'s Open Document (ODF) standard, but being able to open *Office* files, this could help to move more people over to *LibreOffice*.

Version 7.2 also comes with a build for Macs running ARM-based M1 chips. However, the Foundation warns that "because of the early stage of development on this specific platform, binaries are provided but should not be used for any critical purpose." Still, it's another positive step that could win over new converts. To download *LibreOffice 7.2*, visit https://bit.ly/ lxf281libreofficedownload.



LibreOffice still has some way to go to be the de facto office suite.

DEVELOPMENT

Devs love Linux (but prefer Windows) Latest Stackoverflow results offer interesting insights.

he latest Stackoverflow survey results are in (you can see them all at https://bit.ly/lxf281stackoverflow). Over 80,000 developers from 181 countries participated in the survey earlier this year. The findings reveal how they work, including the apps and operating systems they rely on.

While Linux, for example, is a more popular operating system for developers than macOS, with 25.32 per cent of respondents using a Linux distro compared to 25.19 per cent , Windows continues to be the most popular by quite a margin, with 45.33 per cent of respondents primarily using that.

Interestingly, this is the first year that people have been able to select Windows Subsystem for Linux (WSL), which enables devs to run Linux subsystems within Windows – and a not-insubstantial 3.29 per cent responded that they primarily use that.

The results also show that Rust is the most loved language, while COBOL is the most unpopular. AWS is the most loved cloud platform, ahead of Google Cloud Platform and Microsoft Azure. Meanwhile, JavaScript is the most commonly used programming language (for the ninth year running), followed by HTML.

The results also highlight the lack of diversity amongst developers. They are overwhelmingly male (91.67 per cent compared to just 5.31 per cent who identify as women), straight (84.52 per cent) and white (58.43 per cent). There's still a long way to go for people and organisations who want to see more diversity among developers.

OPINION

SCORES ON THE DOORS



Keith Edmunds is MD of Tiger Computing Ltd, which provides support for businesses using Linux.

When you're responsible for managing Linux systems, it can be helpful to step back and objectively assess how healthy those systems are. Helpful, but not easy. When we're caught up in the dayto-day running of those systems, it's sometimes hard to see the wood for the trees.

Here at Tiger Computing we've drawn on our experience of managing commercial Linux systems for the past 20 years and designed the Linux Reliability Scorecard. It consists of 14 questions, all with yes/no answers. On average, they take just over 60 seconds to answer.

The result is a downloadable PDF that will tell you exactly which areas need focus. Maybe it'll confirm what you already know, or perhaps it'll highlight some areas you'd not considered before.

Either way, it's a productive way to fill a coffee break and it might serve as evidence to your management that more work is needed. And if you complete the scorecard in the next month, you can get a free printed copy of my best-selling book The Linux Solution.

You can find the scorecard here: https://scorecard. tiger-computing.

computing.

co.uk.



NEWSDESK

OPINION

HERD IN The cloud



Matt Yonkovit

is Percona's Head of Open Source Strategy and a member of SHA (Silly Hats Anonymous).

Going 'cloud native' involves building applications in new ways. Traditional applications are generally designed with a two- or three-tier architecture, but cloud native applications utilise newer innovations like microservices and containers.

Each microservice component uses its own database instance to store data over time. Depending on your design, you can use ephemeral instances. These are created, used and then deleted when no longer required. If you scale rapidly, a herd of new database instances will be needed.

What does this mean for your data? Well, you may not know how many database instances you have running at a point in time. One approach to manage this is to use a database as a service, which can run all those instances for you. However, this can lock you in.

Instead, it's important to look at the move to cloud native as a whole. Rather than blindly relying on providers which could lock you in, you need to clarify your open source approach and head out on the path that suits you best. That way, you can keep your cloud herd under your control.

voice technology Mozilla gets chatty

Mozilla Common Voice adds 16 new languages to its roster.

he Mozilla Common Voice project (https://commonvoice.mozilla.org) is an open-source initiative that aims to make voice technology, such as voice assistants found in smart speakers, more inclusive, and it's now reached a major milestone. According to a new blog (read it at https://bit.ly/lxf281mcv), Mozilla Common Voice now features 16 new languages, including relatively lesser known languages that popular voice assistants such as Alexa and Google Assistant may not support, such as Basaa and Kazakh.

As Hillary Juma, Common Voice Community Manager, points out, while voice assistants and smart speakers are an increasingly useful part of

PROCESSORS

Russian RISC-V designed CPUs

Eight cores, 2GHz and coming to laptops by 2025.

R ostec, a Russian state-backed company, is working with Yadro and Syntacore to create RISC-V processors for desktop PCs, laptops and servers, initially aimed at government and education devices and looking to launch by 2025.

According to a report by **Vedomosti.ru** (and reported by Anandtech at **https://bit.ly/ lxf281russia**), the new processors will use RISC-V cores, and will initially be an eight-core CPU running at 2GHz and using a 12-nanometer process. Syntacore already develops its own RISC-V core, so it's likely the processors will use that, and will utilise a "high-performance 64-bit RISC-V application core", according to documentation published by Anandtech.

It's likely that when devices powered by these CPUs will arrive in 2025, they'll be running Russian government-approved Linux distros. Whether or not they do appear, however is another matter. We're always pleased to see new RISC-V processors, but there have been plenty of Russian CPUs promised over the years that never materialise – let's hope this is different.



Mozilla Common Voice is a fantastic initiative that brings voice assistants to a much wider audience.

many of our lives, huge amounts of people are still missing out. "Neither Amazon's Alexa, Apple's Siri, nor Google Home support a single native African language."

The increase in Mozilla Common Voice's scope has been driven by a partnership with Nvidia, which, as lead engineer Jenny Zhang explains, "helps us tighten the feedback loop between data collection and machine learning teams that actually use the data." This follows a \$3.4 million investment from the Gates Foundation, Giz and FCDO to support work in East Africa.

SOFTWARE

Krita 5.0 beta released

Digital painting application receives a big update.

hile we review *Krita 4* in this very issue, the beta for *Krita 5.0* has been released, bringing some big changes to the opensource digital painting application. As the release announcement (https://bit.ly/lxf281krita5beta) explains, this version has a new resource system which no longer loads all the tools on each startup. Instead, it caches them after startup, which means that first launch will take longer, but subsequent start-ups will be faster. There's also a new MyPaint brush engine, which does mean some compatibility issues when opening art created in earlier versions of the program.



Krita 5.0 Beta is out and comes with a new splash screen painted by concept artist Tyson Tan.

NEWSDESK

Distro watch

What's down the side of the free software sofa?

LIBREELEC 10.0.0

The tenth version of this media-orientated Linux distro has been released. It comes with *Kodi* (Matrix) 19.1, the popular media player, and as the release announcement (https://bit.ly/kf28Ilibreelec) explains, you'll need to manually update to get it. It's available to download for Raspberry Pi, Allwinner, Rockchip, Amlogic, NXP and standard PCs, but as the release announcement mentions, there remains a major bug that can crash the software when switching profiles. Hopefully that'll get sorted ASAP.

VOYAGER LIVE 11

Voyager Live 11 is the latest release of this Debian-based distro with the Gnome desktop. Based on Debian 11 "Bullseye", it comes with Linux kernel 5.10, Gnome 3.38, *Firefox 78.13* and *LibreOffice 70*. Unlike previous versions, this release preserves the internal structure of Debian, and all package updates are from official Debian mirrors to help avoid any security issues. New scripts have been included to make certain tasks, such as switching between PC and tablet modes, easier. Find out more about this release at https://bitly/lsf28lvoyager11.



LibraELEC is a great distro for people building media centre PCs.



Voyager Live is a graphically impressive Debian-based distro, and a new version is now available for download.



SparkyLinux 6.0 updates the Linux kernel and much more, while still supporting older 32-bit devices.



KDE has published its own fork of Qt 5.15, which KaOS makes full use of.

OPINION

VIRTUAL SUMMER



Lubosz Sarnecki is a senior software engineer working at Collabora.

&& This summer, Christoph Haag and I took part in Google Summer of Code (GSoC) as mentors for *xrdesktop*, the Open Source project bringing the Linux desktop to VR on Valve's SteamVR and Monado. Our students, Remco and Manas, were able to finish their projects and submit merge requests. Kudos!

Remco implemented a virtual keyboard with support for ernojis and 56 languages, while Manas implemented loading and rendering of GLTF models, which included improving our renderer and adding rendering techniques like normal mapping, which will enable *xrdesktop* to display controller models and virtual environments.

As our tools of choice, C/ GObject and Vulkan were used in both projects, maintaining our goal of keeping *xrdesktop* low level and thus providing a performant XR experience.

Contributions like these support our vision to provide a fully open source XR stack that enables complete control and independence for end users and product builders.

These features will get into our next *xrdesktop* release v0.16, which will be tagged soon. If you're brave enough you can also try our experimental next branches on our **freedesktop**. org GitLab.

SPARKYLINUX 6.0

A new stable release of SparkyLinux, based on Debian "Bullseye", has been released, and comes with a host of package updates, along with the Linux kernel 5.10.46 LTS. If you use the rolling release (Po Tolo) of Sparky, you should get this update pretty easily. Unlike many modern distros, SparkyLinux 6.0 still supports 32-bit hardware, which is good news for anyone running this on older machines. To find out more about the latest release, head over to https://bit.ly/lxf281sparky6.

KAOS 2021.08

A new version of this KDE-sporting distro has been released as part of its rolling release schedule, and comes with updates to Systemd, *Curl*, *Mesa* and more. Because Qt 5.15, which KaOS uses for many of its applications, only receives closed source updates these days, KDE has published its own 5.15 fork, which KaOS now uses. There's a host of other handy updates with this release, so check out the full release announcement, along with download instructions, at https://bit.ly/lxf281kaoS08.

OPINION

IT'S ABOUT TIME, TOO



Jon Masters has been involved with Linux for more than 22 years.

Last month was the 30th anniversary of the original announcement of Linux. This month it's time for another anniversary in the form of 17 years since the first "PREEMPT__ RT" (Real Time) kernel patch was posted upstream by Ingo Molnar. It might not be all that much longer until – finally – all of the pieces needed to run a Linux kernel with bounded deterministic performance for latency sensitive workloads are finally in upstream kernels.

I first encountered the "RT" kernel patch not long after Ingo first posted it. In fact, Ingo wasn't actually the first to work on Real Time Linux support. A number of others came before, but Ingo's patches were most influenced by work from Sven-Thorsten Dietrich. Sven had taken a disciplined academic approach to determining the pieces of Linux most impactful upon overall latency, and had created a "real-time" patch series for Montavista Linux that implemented many of the ideas Ingo would refine into what would become PREEMPT_RT later on.

In those early days, I enjoyed gaining experience with Real Time through work on embedded Linux systems (years later, I co-authored a book with one of the RT maintainers on embedded Linux), and then a Real Time product that targeted such users as stock trading platforms. Traders and embedded folks building robots have more in common than you would think, since both care about extremely low-latency response. It was fun being a fly on the wall in those early days, but I didn't at the time think that it would take almost two decades to upstream.

KERNEL WATCH

Jon Masters summarises the latest happenings in the Linux kernel, so that you don't have to.

inus Torvalds announced the release of Linux 5.14, which came mere days after the 30th birthday of his very first Linux announcement. The new kernel includes features such as "core scheduling" (a defence against speculation attacks that span multiple threads) and a new io priority "rq-qos" Quality of Service policy that allows I/O bandwidth to be prioritised within different Linux cgroups.

With the release of 5.14 came the opening of the "merge window" for what will become Linux 5.15. Among the many patches pulled thus far is the core real time "preemption" locking code, and support for scheduling on "asymmetric" systems such as those with Arm "big.LITTLE" clusters of "big" (high performance) and "little" (high efficiency) cores where each of those cores doesn't have uniform architectural capabilities. A classic example comes when some Arm cores don't support 32-bit, but a similar situation could even arise on x86 with some cores not supporting AVX512 vector operations. Other more complex constraints are likely in future.

Real Time Linux now a step closer

As an Operating System, the primary job of the Linux kernel is to juggle various platform resources provided by hardware and make those available to application software programs that need to share them in an efficient manner. A key part of this is scheduling time for each application process that wants to run on the (limited number of) CPU cores available. Traditionally, each running application process (known as a "task" within the kernel) was given a fixed "quanta" of CPU time before the next, and those processes with a higher priority got to have their time quanta before lower priority tasks. Periodically, the kernel would interrupt whatever was running on a given CPU to service a piece of hardware.

This works well enough much of the time, but there are times when whatever the kernel is doing must be preempted right now in order to run a more important task. There are a number of challenges to implement this, not the least of which is that the kernel frequently makes use of "spinlocks" to protect what are called "critical sections" of kernel code that should not be interrupted. To get around this, the "PREEMPT_ RT" patch series converts these into preemptable locks called "trmutexes". The fun isn't over there however because real-time systems classically may suffer from a problem known as "Priority Inversion".

Priority Inversion famously struck the NASA Mars Pathfinder program. A rover was on the surface of the red planet when a low-priority maintenance task was able to prevent another low-priority task from making progress. Unfortunately, a high-priority (critical) task was waiting for a resource owned by the second lowpriority task that never made progress and thus never released the resource. Linux solves this using an approach originally motivated from academia: Priority Inheritance. Whenever a highpriority task is holding, the lower-priority task is temporarily "boosted" up in priority to the same as the high-priority task in order to progress.

» ONGOING DEVELOPMENT

Matthew Wilcox has worked hard on a patch series originally called "folios" that aims to "allow filesystems and the page cache to manage memory in larger chunks than **PAGE_SIZE**. Modern microprocessors use a unit of memory management known as a "page" that's typically around 4KB in size. This is the smallest size of memory that can be tracked by a Memory Management Unit, and is also the fundamental size handled by the low-level (page) memory allocator in Linux.

The trouble is that 4KB isn't all that large these days. In fact, if Linux needs to track all memory in units of 4K then the overhead becomes significant. Instead of doing this, there are times when groups of related pages can be managed as a chunk, or "folio" in the original patches. Linus preferred a much more boring name such as "pageset". The choice of (re)name alone is not likely to get the patches accepted into 5.15, however, in spite of developers calling for clarity, due to ongoing debate about the best approach. Listen up, puddins! Learn how Marvel and DC's Joshua Swaby paints heroes with attitude

BEA KILLER



NO

PLUS/ MASTER THE LIGNE CLAIRE COMIC STYLE LEARN THE NEW TOOLS OF COREL PAINTER 2022 SPEEDY COMIC CREATION IN CLIP STUDIO PAINT



TERRY DODSON REVEALS HOW TO DRAW & PAINT THE PERFECT COMIC COVER

FREEOVER **3 HOURS** OF EXCLUSIVE PRO VIDEO TRAINING



WORKSHOP

PAINT THIS COVE

David Nakayama reworks the iconic comic character

on sale now print and digital editions available at magazines direct.com

Digital editions are also available on iOS, Android, and more platforms via



Zinio

∩ ok

Ŏ



Tr

pocketmags

Buckle up! The 'Dude' shares his art insights

Answers

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

Deskreen discombobulation

I was hoping to try *Deskreen* (LXF278) with my Pi400, so that I can use an older computer as a display. It didn't work for me at all. I did find the support pages of *Deskreen*'s GitHub at https://github.com/pavlobu/deskreen/ issues/59 so confusing that I gave up! Can you advise how I can get this to work?

Thomas Rasche

A I'll start with a disclaimer: I don't have access to a Pi400 so this was done on a plain old Pi 4. You don't say which OS you have on the Pi, so I'll assume it is the standard Raspberry Pi OS, the OS formerly known as Raspbian. The GitHub page you refer to is basically a discussion between people who have already got it working rather than a set of instructions for those who haven't, but near the bottom of the page there is a link to an AppImage build for the Raspberry Pi.

An AppImage is an all-in-one software bundle; all dependencies and the executables should be contained in the one file, which you simply run. It is not the most efficient way of running software – the bundle probably contains duplicates of software you already have on your system – but it is a quick and easy way of using it. The AppImage on the main *Deskreen* site is for x86 processors; make sure you get the Arm one from the GitHub link. Download it to your home directory and make it executable, either by rightclicking in the file manager, or by using chmod in a terminal:

\$ chmod +x Deskreen-1.0.1-armv7l. AppImage

This has to be done only once, then you can run it from a terminal any time you want with:

\$./Deskreen-1.0.1-armv7l.AppImage You can also run it from the file manager with a double-click.

Read-only USB

I'd like to know a procedure in order to create a 'read-only' USB drive to be used with Debian, Ubuntu or Arch-based distros. I'd like to be able to reverse the procedure to restore the default read-write setting, as well. Does something similar to *Diskpart* exist under the hood in Linux in order to create readonly USB drives?

Since I enabled the read-only function with Diskpart, why am I allowed to delete



An AppImage file is the easiest way to try new software on your computer. You download it, run it and can delete it if you don't want it.

or create files on the USB drive while on Linux? Is it due to different file systems, ext4 vs. FAT32?

Also, how could I create a password protected and/or hidden folder with Debian, Ubuntu or Arch-based distros? Terminal instructions or GUI packages will do.

While I'm asking questions, may I trust software available at archive. org – such as a Windows 2000 ISO image which seems to be presented by Microsoft – or are they potentially unsafe and dangerous stuff uploaded by users?

A Forcing an ext4 filesystem to be read-only is easy these days, just use the tune2fs command:

\$ sudo tune2fs -O read-only /dev/sdb1 replacing /dev/sdb1 with the device you are using. You can do this with the filesystem mounted or unmounted, but if is mounted it will remain read-write until you unmount and remount it. The tune2fs command sets a flag on the filesystem so that all attempts to mount it, even if you try to force read-write, will mount it read-only. Reversing the process is simply a matter of clearing the read-only flag, like this:

\$ sudo tune2fs -O ^read-only /dev/sdb1 Depending on your shell, you may need to escape the ^ with a backslash – you'll soon know if this is needed.

Diskpart is a Windows tool, so it only controls how Windows handles the drive. There is no equivalent option for FAT filesystems, so if you want to use FAT you will have to force read-only by putting an entry in **/etc/fstab**. Give the filesystem a label with

\$ sudo fatlabel /dev/sdb1 ROSTICK Then add this line to **fstab**:

LABEL=ROSTICK /mnt/rostick vfat ro 0 0 A simple tool for creating an encrypted folder is VeraCrypt (https://www. veracrypt.fr/en/Home.html). It's a GUI program, so it is basically a matter of following the prompts. VeraCrypt can also create a hidden volume within an encrypted folder, with a different password. Entering one password unlocks the folder



ANSWERS

without revealing the existence of the hidden volume, while entering the other password unlocks everything.

Anything you download is potentially unsafe, but stick to files from trusted sources. In the case of Windows software, only Microsoft distributes them legally so anything you download from elsewhere should be considered suspect.

Ancient and modern

Thanks for the *Roundup* on Desktop Virtualisation. I would like to run some older software on newer, faster hardware. Do any of the virtualisation products support running Windows 98 Second Edition? Gary Hedguist

Any of the systems we covered should be able to run Win 98, as they all have support for older hardware. As support for newer systems is added, the older software is not removed from the supported operating systems, so these all go back quite a way.

One of the reasons for using virtualisation is for running old/ancient software that only runs on older version of Windows, so this ability is unlikely to go away as long as supporting the older systems is still technically possible. VirtualBox actually has a profile for installing Windows 98 - and even Windows 3.1! Agemu also has a Windows 9X option. You may need to check the configuration options to see which hardware is emulated. but I know that they have support for Cirrus graphics cards, an old card verv popular in the Win 98 days, although standard VESA hardware should work too. Similarly, something like a SoundBlaster audio card may be a good pick from the choices offered.

You should also choose a compatible emulated network card. That way Windows can download drivers for any other hardware it needs. Thank you for the positive reaction to the Roundup, and also for the opportunity to run a Windows 98 installer again to make sure it works. Please don't ask again...

Schizophrenic GRUB

I feel a bit 'first world problems' with this question, because it's more of an annoyance than an actual problem, but annoying it is.

I run Mint as my main desktop but also Ubuntu Studio on the same machine for real-time music applications. Sometimes when I apply kernel updates for each, the boot selection display gets changed, not just visually, but the order and default system on each will change



Remember the good old days? Well this is Windows 98, so only the old part applies.

as well. Now, I quite like the pretty (almost, anyway) Ubuntu Studio one, and I also like the bare-bones text-style Mint one, so I'm not that bothered which, but I'd like it to be consistent – ideally with Mint as the default one, so that if I power up the machine and walk away or get distracted, it will boot up without me being surprised.

I've tried not applying the kernel update (on Ubuntu Studio) but it keeps nagging that it's there to be installed; I've also tried waiting until both have the kernel update and doing the Ubuntu one first and the Mint one afterwards, but that doesn't always work.

Unfortunately, the ease and simplicity of using Mint has made me a bit soft in my old age, and I'm struggling to get deep in the technicalities of it now. I assume it's some sort of GRUB configuration that I can do to control this, either before or after? Brian Perryman

A The problem is that each distro has its own instance of GRUB installed. When you install a new kernel, the package manager runs *grub-mkconfig* to update the menu with the new kernel. This has the side effect of making that distro's instance of GRUB the default, until you apply a kernel update to the other distro.

The solution is simple: run grubmkconfig again to reset the menu. It doesn't matter which distro you run it from because, as you have found out, it picks up the kernels for both distros and adds them all to the menu. Ubuntu has a script that calls grub-mkconfig and writes its output to the menu, called update-grub. After installing a kernel update, reboot. If the menu is not the one you want, boot into the distro you prefer and run \$ sudo update-grub

There is another way to handle this, at least as far as the appearance is concerned. The file /etc/default/ grub contains the settings used by grub-mkconfig when building the menu, including settings such as themes and colours. You could make sure they have the same settings in both instances of this file – you can easily mount your Ubuntu root partition in Mint or vice versa - to give a consistent look. This will not help with the order, but you can set the default choice with **GRUB DEFAULT** in this file, using either the title or id of your favourite. Look at the existing menu at /boot/grub/grub. cfg for these values.

DVD demands

I would like my boot order to be: first DVD drive, second Ubuntu on SSD. For years I have had no problem in setting the BIOS to do this including on a computer with UEFI. However I cannot get my new computer to do it. If I set boot DVD first, it will start, providing that there is an OS DVD in the drive. If the drive is empty, it puts a message on the screen telling me to put an OS in. Previously it would have moved on to the hard drive and run Ubuntu. I now have to alter UEFI to DVD if I want to load Clonezilla for example, then alter back again for Ubuntu. What am I doing wrong? Bryan Mitchell

A This sounds like it could be a bug in the firmware. The normal behaviour is to try each item in turn and move on to

ANSWERS

the next if it is not suitable. You should only see that message when there are no bootable devices. Assuming you have carefully checked the various options and there is nothing that could interfere with this – I couldn't see anything in the online manual – I would check to see if there is a BIOS update available for your motherboard. The BIOS menu should indicate the current version; check for an update at **asus.com**.

However, it is not necessary to mess with firmware boot priorities like this, as your motherboard has a boot menu. Press F8 when the display first appears and you will be presented with a list of devices to boot from. That way you can set the firmware to boot from hard drive normally and use the boot menu to boot from DVD or other device.

Incidentally, if you regularly boot *Clonezilla*, it may be worth copying the ISO image to a USB stick using *Etcher* (https:// www.balena.io/etcher) or *dd*. USB sticks boot a lot faster than DVDs.

Linux NAS RAID

Your article on the Ultimate Media server in LXF269 finally got me to try to do something with my underused QNAP 451+ box. The OS that QNAP put on it drives me crazy; it is constantly accessing the disks even when there is nothing going on. I did finally get around to making it boot Xubuntu (based on Bionic) and it worked fine, up to a point. I could mount the RAID1 arrays using a combination of mdadm and mount commands. But I could not mount the RAID5 array, which of course is the one that holds all my data.

I get an unknown filesystem type, brdr. All the mdadm display commands work fine, a dump of /proc/mdstat shows the array as being active and so forth. From all I have been able to find, QNAP use

	eroots woomfondes Kansake		
File Edit View Bookmarks Plugins Settings Help			
<pre>[root@vroomfondel ~ 0]% Volume group</pre>	vgdisplay		
VG Name	majikthise		
System ID			
Format	lvm2		
Metadata Areas	1		
Metadata Sequence No	5479		
VG Access	read/write		
VG Status	resizable		
MAX LV	0		
Cur LV	12		
Open LV	10		
Max PV	0		
Cur PV	1		
Act PV	1 manual in much	I	
VG Size	10.91 TiB		
PE Size	4.00 MiB		
Total PE	2860434		
Alloc PE / Size	2445824 / 9.33 TiB		
Free PE / Size	414610 / 1.58 TiB		
VG UUID	M1TDup-VZcX-DSWi-V88C-2Sj8-f0qb-uu03wX		
[root@vroomfondel ~ 0]%			

LVM adds an extra level of control, and complexity, to filesystem handling. The vgdisplay and lvdisplay commands show what you are working with.

» A QUICK REFERENCE TO... ERROR CAPTURE

The hardest part of solving problems with computers is identifying what the actual problem is. When asking for help, either on a forum, in these pages or just a request to Google, it is important to quote exact error messages. If you are running a program from a terminal, these are plain to see, and to copy and paste – but what if you can't see them, or there's too much info?

You can copy the output from a program to a file like this:

\$ someprogram 2>&1 >file.txt

There are two output channels, standard output and standard error. >file.txt 2>&1 redirects error messages to stdout. Another place to look is in the system logs. On a distro running systemd, you can see the system journal with

root@vroomfondel 22:59 12-Aug-2021

\$ sudo journalctl -5m

This shows entries from the last five minutes, while

\$ sudo journalctl -f

shows log entries as they are added.

as if you have a problem while booting, use your phone to take a photo of the screen. I have even used mine to record a video when a boot problem flashed an error too fast to read, then played it back slowly so I could see what was going on. ext4 on their Raid5 arrays, but mount -t ext4 /dev/md1/mnt/md1 gives me the error. I am not sure what flavour of Linux they use. I also tried a vgscan to see if there was some sort of LVM stuff going on, but no luck.

Do you have any advice on how I can re-use that RAID array with Linux without having to copy everything off of it, and rebuild it under some flavour of Ubuntu? Greg Morse

A QNAP does indeed use ext4 on its RAID arrays, but it also uses LVM on recent releases. I haven't come across a 'brdr' filesystem before and a web search showed up nothing relevant. It would be helpful if you could post the actual output, preferably copied and pasted, of any commands you run. If the RAID5 array is /dev/md1 and it is active, cat /proc/mdstat will tell you this; you can see if LVM is in use by running

\$ vgscan

\$ vgdisplay

This will show any volume groups, mounted or not. Once you have the name of the volume group, say vgXX, activate it by running:

\$ vgchange -a y vgXX

Now you can list the logical volumes with <u>lvscan</u>, which will show the name and size of each one. From here it should be obvious which one holds your data (the biggest one).

Then you can mount it using the device node shown by lvscan. Once you have tested this, you can add the mount details to /etc/fstab in the usual way.

GET HELP NOW!

We'd love to try and answer any questions you send to **lxf.answers@futurenet.com**, no matter what the level. We've all been stuck before, so don't be shy. However, we're only human (although many suspect Jonni is a robot), so it's important that you include as much information as you can. If something works on one distro but not another, then 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** lspci >**system.txt** lspci -**y** >**system.txt**



DISCOVER OUR GREAT BOOKAZINES

From travel and history to gaming and photography, you're certain to find something you're passionate about



Follow us on Instagram () @futurebookazines



www.magazinesdirect.com
Magazines, back issues & bookazines.

Mailserver

WRITE TO US

Do you have a burning Linuxrelated issue that you want to discuss? Write to us at Linux Format, Future Publishing, Quay House, The Ambury, Bath BA11UA or email lxf.letters@ futurenet.com

Elementary Archimedes

I've been enjoying your series on emulating early home computers, several of which I was fortunate to be able to play with during my school days. Even for those machines I didn't encounter, the articles give one a good overview of the various emulation options and useful information - practical advice I could use when I recently had to resurrect and check software from an old Macintosh disk, for example.

I notice that you've now covered emulation of two BBC Micros, and this letter is to ask you to please cover a third (if you've not already planned to do so): the Acorn Archimedes, and specifically the A3000. It was a contemporary of, but superior to, the Atari ST and Amiga 500 (which you've also covered). I'd be interested to hear your experts' advice on the easiest way to re-experience the machine which hosted not just the best version of Lemmings, but also the best version of Elite (as well as David Braben's lesserknown masterpiece, Lander/Zarch. If ever you need to train your fine and precision mouse motor skills, this is the game to use!)

For a more serious and relevant reason to (re-) introduce this micro to your readers - what about the fact that its legacy lives on in the billions of smartphones in the world today, not to mention our beloved Raspberry Pi? **Richard Milne**

Neil says...

The old Acorn Archimedes was where Arm first started

I'm not sure how we've left the Archimedes off the list, though I'm sure it got a mention when we covered the BBC Micro... It'll certainly get covered at some point. I tried to convince a writer to emulate an IBM System 360 Mainframe, but I don't think that's going to happen. We did briefly cover Arm

originating from Acorn in our Open Source Processor feature back in LXF264, though the main focus was RISC-V.

30 vears

I was amazed to hear that this is the thirtieth vear of Linux. What an amazing achievement for something that was never intended to be professional.

I suspect that our lives are influenced by Linux every hour of every day – whether using our Android phones, seeing kioskesque systems built with a Raspberry Pi or browsing the web where nginx or Apache2 will be the webservers.

Dear old Tux has transfromed many lives around the globe.

I was bitten by the bug two decades ago now. I can still remember being sat in a fast food restaurant with a friend of mine, during college free periods. My friend would excitedly talk about the culture surrounding free software and why more eyes on the code is a good thing as well as the technicalities, which, at the time, went over my head.

Then followed a Slackware install at home console-only of course – and two of my friends helped me to set it up as a router, so that the broadband could be shared between multiple PCs with our 10Mb/s network.

Following university, I've worked in IT support for over a decade and while the environments I have been in have primarily been Windows environments. I've always been able to use Linux distros when most suitable. This interoperability is brilliant and is largely due to open standards and protocols such as HTML, HTTP and TCP/IP. Michael Holder

Neil savs...

We've heard some great stories of how many people got started with Linux over the years, and most are no more complex than "I tried this wonderful thing" and people were hooked.

Free stuff!

I read about a suggestion in LXF276 to possibly opt out of the cover DVD to save the planet and all that, and as I glance up from reading this in my place of work (yes, I actually go to a place of work from time to time), my colleague has a poster of Tux on the wall,



and it's relatively large – like, year-planner large. To my excitement and horror of missing out I see a *Linux Format* logo in the corner and on the back of this giant Tux poster I see a year planner! I assume I had missed this in a recent edition of *Linux Format*. I didn't take note of the year of this planner but I was informed it was from a time when phones were not able to do absolutely everything and pen and paper was a thing. I have bare walls at my desk where I should put up pointless company stuff, and since seeing my colleague's poster I have been looking for some Linux-like wall art. So how about LXF introduces some poster stuff? *Neil Robertson*

Neil says...

I hear subscribers should be getting a little present in the post along these lines in a couple of issues...



I Who ♥ Linux stickers? You do!

» LETTER OF THE MONTH

Just do it

Like most people, I would think, I just want my IT to work and when it comes to downloading software I want to use a nice GUI to simply click and collect.

I've got to say at this point that I love *Linux Format* and have subscribed for a few years. I also love Nextcloud, but I would imagine that the several articles that you have produced on Nextcloud have put most people off ever using it. Go to this line, type that command etc. etc. Pah, I say.

All I've ever done is gone to the Nextcloud website and downloaded the latest AppImage version. Once downloaded I made that version executable. Then I opened Nextcloud by clicking on the now-executable AppImage.

The Nextcloud provider that I use is Tab Digital as they were the one offering the most free storage. They reside at https:// nd.nl.tab.digital. I then set up an account and followed the prompts. When it comes time to update to the latest version, which you will receive a message about, all you need to do is go back to the Nextcloud website and download again. All of this was done without a whiff of the command line and was on a Debian/ Ubuntu based system, but I can't see why it shouldn't work with any Linux flavour.

John M

Neil says...

Well that's showing the so-called experts! Jonni and the like tend to favour the command line as it works on headless servers as well as desktops. But as you also mention, packaged AppImage (also Snap and Flatpak) simplify things for everyone. As always it's the magic of open source that there's multiple way of doing the same damn thing!







SUBSCRIBE Save money today!

Sign up today and get your EARFUN AIR Wireless Earbuds



- » Quad-mic system for enhanced noise-cancelling calls
- » In-ear auto detection
- » Custom-built cellulose drivers for enhanced audio
- » IPX7 protection with Sweatshield technology
- » Single earbud mode
- » 35-hour total battery life: seven hours playtime plus 28 hours with charging case
- » Fast wireless or USB-C charging



"A nicely built set of headphones that also boast excellent battery life, wireless

charging support and an expansive listen." *What Hi-Fi?* August 2020



SUBSCRIBE NOW! www.magazinesdirect.com/lin/Earbuds Call 0330 333 1113 and quote BM39

YOUR

Don't miss out

subscribe now

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

Turn to page 63

>> CHOOSE YOUR PACKAGE!

LINUX

EDORA 24 & VOYAGER

into Linux today! FORMAT

16.04 Ultimate Boot CD

Everything you need to back up, restore and

UBCO

READY TO RUN

Fedora

DON'T MISS!

Now with 5 years of Linux User & Developer issues

65 Pages of tutorials and features

A

POWFRED



Terms and conditions: Offer closes 31 October 2021. Offer open to new UK subscribers only. Pricing is guaranteed for the first 12 months and we will

TAKE COMPLETE CONTROL OF YOUR LINUX BOX

Got root? Then get the Linux Format Annual and get the most from your distro and its free and open source software today!



REVIEWS

AMD Ryzen 7 5700G

The AMD APU makes a more-than-welcome return for Chris Szewczyk.

SPECS

Socket: AM4 Arch: 64-bit. Zen 3 Cores: Eight Threads: 16 Process: TSMC 7nm Clock: 3.8GHz (boost 4 6GHz) Cache: 512KB 11.4MB12 16MB1.3 Memory: DDR4-3200. two channel GPU: RX Vega 8 Graphics Compute units: Eight GPU cores: 512 PCIe: v3.0 **TDP:** 65W

MD's Ryzen 5000 series of CPUs has been well received, though, like every other manufacturer of CPUs and graphics cards, AMD has been affected by wider semiconductor market issues. Limited supply meant that production priority was given to the high-end chips in both the red team's CPU and GPU stacks. That means we haven't seen any affordable Zen 3-based CPUs – at least not in the retail market.

This doesn't necessarily change today with the introduction of the 'Cezanne' family of APUs, which combines its strong Zen 3 architecture with capable Vega-based integrated graphics, and are the first in the retail market since the release of the Zen+ based 3200G and 3400G in mid-2019.

Here we have a Zen 3-based processor with a 7nm, 180mm2 10.7 billion transistor monolithic die, with eight cores and 16 threads. And by monolithic we mean that it doesn't match the chiplet design of the other Ryzen 5000-series CPUs, instead squeezing everything into one traditional package design.

Sacrifices have been made

The 5700G has a base frequency of 3.8GHz and a boost frequency of 4.6GHz, yet it still has a desirable 65W TDP. But to meet this 65W target, some sacrifices had to be made; it's not simply a down-clocked 5800X with integrated graphics bolted on. You lose half of the L3 cache (16MB versus the 32MB of the 5800X) and you also lose PCIe 4.0 support, meaning you don't get access to the best SSDs around today.

The integrated graphics does count as best-in-class, even though it's still based on the ageing, albeit reworked and well-optimised, Vega architecture. The 5700G's graphics performance is class leading, but as usual, integrated graphics can't provide the horsepower of even a low-end discrete GPU.

This colourised image of the AMD Ryzen 7 5700G's silicon could make great wall art.



Productivity and gaming tasks all take a big step forward over the popular Zen 2 Ryzen 7 3700X. The exception being sequential SSD performance, where the lack of PCIe 4.0 counts against it when put up

> against the rest of the Ryzen 5000 clique. The 5800X, with its higher boost clock, extra TDP headroom, and increased cache is the better performer. That's as you'd expect, and under CPUlimited gaming loads, the extra cache is helpful in terms of gaining higher frame rates.

When it comes to integrated graphics performance, however, there's the 5700G and then there's clear daylight to the competition. This means games from a few years ago, or less-



Great performance, low power requirements – what's not to like?

demanding eSport titles are perfectly playable. Intel's Xe is a step forward over older UHD graphics, but the EU counts are simply too low in its 14nm Rocket Lake processors to be competitive.

Zen 3 processors deliver awesome power efficiency and the 570OG follows that trait. While looping benchmarks, we saw a peak power consumption of 77W, which is impressive. Don't forget that something like a Core i7 1170OK can use three times that figure.

If you're looking for a powerful HTPC or small formfactor machine, a PC for heavy multitasking, or you need something with integrated graphics to tide you over until the GPU market returns to normality, then the 5700G is an awesome APU. It's also going to be a big upgrade for someone coming from a PC that's a few years old. Add a high-end GPU, even a GeForce RTX 3090 or Radeon RX 6900 XT with all the eye candy turned up, and you'd barely lose any real-world performance compared to any other processor we could mention.

You get the strong IPC and multi-threaded prowess that Zen 3 brings to the table, paired with top-end integrated graphics. You can run a monitor or two and enjoy some basic gaming, or even better than that: true gaming with older or eSport titles, all while keeping around a 65W TDP.

PRICE: £329			
			9/10
FEATURES	8/10	EASE OF USE	5/10

A fantastic all-rounder CPU, combining the best-ever integrated graphics while keeping everything under 65W.

» Rating 9/10

XPG Gammix S50

PCIe Gen 4 and TLC memory for a price you can afford? Jeremy Laird would never fall for that old tale...

SPECS

Capacity: 2TB Interface: M.2 2280. PCIe 4.0 x4 Controller: Silicop Motion SM2267 NAND: Micron 96-laver TLC DRAM cache: 1GB Seq. (r/w): 3.900MB/s. 3.200MB/s IOPS (r/w): 490K, 540K **Endurance:** 1.480TB Warranty: Five years

e're well through the early adopter stage of SSD PCIe4 technology, with numerous high-end PCIe 4.0 drives available, along with several value drives. The new ADATA XPG Gammix S50 Lite, sampled here in 2TB configuration and M.2 2280 format, falls into the latter category. But where some fourth-generation (Gen 4) drives achieve a lower price point by using cheap and not always terribly cheerful QLC or quad-level flash memory, ADATA has managed to price TLC or triple-level flash memory into the bargain - in this case, 96-layer Micron chips. How has ADATA pulled off TLC at this kind of price

point? The answer, at least in part, is a more affordable PCIe Gen 4 controller chip. The original Gammix S50 was a high-end PCIe Gen 4 drive with the Phison controller. This 'Lite' model is cheaper and powered by the new Silicon Motion SM2267 controller.

It's the low-cost option from SM's latest PCIe 4.0 controllers. It's fabricated on a cheaper 28nm production node, where fancier controllers are made on 12nm or thereabouts. It's also limited to four memory channels and two ARM Cortex R5 CPU cores. The SM2267's SM sibling, for instance, is on 12nm, has eight memory channels and sports four Cortex R8 cores.

Doing more with less

That said, this new budget PCIe 4.0 controller is faster than SM's previous-generation high-end PCIe 3.0 controller, the SM2263, with 1.200MT/s peak performance to the older chip's 800MT/s. What's more. the increasing density of flash chips means you can achieve large capacities with just four channels, in this case fully 2TB. One other area of arguable cornercutting is cache allocation. The ADATA XPG Gammix S50 Lite gets 1GB of DDR4 cache where you might expect 2GB for a drive with 2TB of capacity. Still, a bit less DRAM is much better than no DRAM at all.

The official performance claims include sequential throughput of 3,900MB/s for reads and 3,200MB/s and for writes, while the 4K random access is pegged at



ADATA's SSD manages to stay cool under pressure.



such as the previous-generation WD Black SN750 or the Kioxia Exceria Plus, with a lower cost. guad-channel PCIe Gen 4 drive.

Rounding out the speeds, feeds and specs is 1.480TB of write endurance, which should be plenty for all but a tiny fringe of ultra-intense users, and a healthy five-year warranty. This 80mm M.2 drive gets a thin, flat heat spreader that claims to reduce temperatures by up to 20 per cent. All told, you're looking at a cost of around three quarters that of a high-end PCIe Gen 4 drive with a more expensive eight-channel controller. So, it's a pretty attractive proposition on paper.

But what about, you know, the performance? Peak performance in the most forgiving benchmark notching up 3.9GB/s reads and 3.2GB/s writes and are very much competitive with a high-end PCIe 3.0 drive.

As for 4K random access, again it depends on the application used. But the broad-brush conclusion is that the ADATA XPG Gammix S50 Lite 2TB returns numbers slightly above expectation, if not at all remarkable, at 71MB/s for reads and in the low- to mid-200s for writes.

Temperatures are very well managed, with a peak of just 53°C in testing. The well-managed temperatures imply that the drop from 1.1GB/s initial internal file copy speed to a fluctuating range between 300MB/s to 500MB/s after around 350GB of data is related to exhausting the SLC cache rather than thermal throttling. However, 3,300MB/s down to 500MB/s is rather lower than we would normally expect for TLC flash, so that aspect is inconclusive.

VERDICT DEVELOPER: ADATA

WEB: www.adata.com/us/solid-state-drives PRICE: £300 (£131 1TB)

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

ADATA's SSD combines PCIe Gen 4 with TLC flash memory to deliver reasonable performance at a tempting price.

>> Rating 8/10

Linux distribution **REVIEWS**

Pop!_0S 21.04

A desktop distro that offers a different way of working... **Neil Bothwick** spends a night on the tiles with Pop!_OS.

IN BRIEF

An Ubuntuderived distribution aimed at those that use their computers professionally. or at least seriously, but still useful for general purpose computing. The options of a tiling window manager and full filesystem encryption show its intended use. but these can be turned off or not used and Pop!_ **OS** becomes another general purpose desktop distro. Pop! OS is created by a hardware supplier. System76, but works with commodity hardware too.

istros based on Ubuntu are ten a penny – well, cheaper than that because they are free, but you know what we mean. Ubuntu itself is pretty good, not to mention rather popular, so any remake of it needs to offer something special. Pop!_OS is slightly unusual in that it is developed by a hardware vendor, System76. While it is the default OS on their hardware, thanks to the joys of the GPL any of us can download and use it. Pop! OS is based on Ubuntu, but is it sufficiently different to stand out from the multitude of Ubuntu respins?

Pop!_OS uses the Ubiquity

installer from Ubuntu. Apart from the colours the installation experience is the largely the same as installing Ubuntu – easy and quick, with one significant addition: there is an option to encrypt the filesystem. This is not the directory-level encryption some distros offer, but full block-device encryption using the kernel's dm-crypt layer, with a passphrase being required early in the boot process to unlock the disk. This gives some indication of the target market for Pop!_OS: those that want to use the computer for more than personal emails and web browsing.

The focus on productivity and creativity continues with the Gnome desktop, as it uses a tiling window manager. This writer has tried tiling window managers in the past and never been a fan, but we persevered for you dear reader. It's a good way of working, as long as you are not too deeply stuck in your old ways, and is well worth devoting some time to getting to know it. Other advantages of the tiling system are that everything can be controlled from the keyboard. It is also possible to stack windows together, like tabs in a web browser – we used this feature in KDE4 and really missed it when it was removed. However, if you don't like it, the tiling function can be switched off with a click on the top menu bar, giving the overlapping windows that we dinosaurs prefer.

The default desktop appearance is a little cartoony for our tastes, but it is clear and easy to find your way around, and we prefer it to the somewhat drab Ubuntu look. It is only a default and a couple of tweaks in the settings made it much more suited to our tastes. This may sound like a minor niggle, but for an OS intended for creative use, it is likely to be in use for extended periods and the ability to easily adjust it to your needs becomes more important with every passing hour.

Pop!_OS has access to all the software made available through Ubuntu's repositories. If that is not



The tiling window manager can take some getting used to, but is worth persevering with. Some people never want to go back to the old way!

enough, the Pop!_OS Shop (or package manager to you and me) also supports Flatpak packages for those occasions you need something they do not provide.

The support section of the web site contains a number of documents, although no single manual. A clue to the target market for Pop!_OS is that there is an article on switching from macOS to Pop!_OS, but no corresponding document for Windows users. One glaring omission for a distro aimed at those that use their computer for more than recreation is that there is no backup program installed by default. Yes, the support site has some information on setting up *DejaDup*, but we would expect this to be standard on a so called professionally used computer!

Despite the initial impressions, this is actually a distro for serious users and keyboard warriors. The tiling window manager and comprehensive keyboard controls take some getting used to, but once you get the hang of them your productivity will improve greatly. This, and the option to encrypt your hard drive when installing, set it apart from Ubuntu. It really is a different distro, not just a paint job.



» Rating 8/10

NixOS 21.05

A complete working distro created from a single configuration file sounds different, and **Neil Bothwick** discovers that it is.

IN BRIEF

NixOS is a unique distribution where everything is controlled from a single configuration file - everything from the disk lavout to the installed packages to the users and their permissions. It even configures the software it installs according to vour needs. On the downside, the learning curve is not so much steep as almost vertical. It is available for multiple platforms. including the Raspberry Pi.

nstalling Linux is usually a matter of: boot from install media, run installer, make choices, drink a hot beverage, reboot computer and start configuring everything to your needs. NixOS turns this completely on its head.

The starting point for NixOS is a configuration file called **configuration.nix**. This is where you specify everything about how you want your system set up, from which partitions to install to and the filesystems to use, to the list of packages to install and the users to create. You have to create the partitioning yourself – Gentoo users will feel at home here – then run

Q Search mor	e than 10 000 options	
usets.usets		Search
Channel: 20.00 21.05	unstable	
	Showing results 1-35 of 35 options.	Sort Best match +
	users.users	
	users.users. <name>.uid</name>	
	users.users. <name>.name</name>	
	users.users. <name>.home</name>	

NixOS is well documented, and it needs to be. Just look at the number of configuration options!

nixos-generate-config to scan your hardware and create a basic configuration, which you then edit.

Once you have done this, run *nixos-install* and sit back. After it has done its stuff, you should have a ready to use system. This is a lot more work than a traditional install, especially the first time, but it has the advantage of being reproducible, and this is where NixOS really finds its place.

If you want to set up a number of information kiosks or office desktops, all with the same configuration, you only need to create the configuration file once then copy it to each of the computers and run the installer. Of course, you won't get it right first time, and even if your needs change, you can edit the configuration and run *nixosrebuild* to apply your changes.

NixOS actually stores all files under /nix, so everything in the usual Linux hierarchy is a symlink to files in /nix. When you run *nixos-rebuild*, it create new files in /nix and updates the symlinks. The old files are still there and you can select previous incarnations from the boot menu, so you are safeguarded against an upgrade breaking something important.

This all sounds very clever, so what is the catch? It is that much of what you know about administering a Linux system does not apply. Want to export a filesystem by adding it to /etc/exports? You can't - that file is readonly, as is most of /etc. Instead, you add the relevant entries to configuration.nix and rebuild.

The problem is that, at least at first, everything takes three times as long to do because you have to search the documents and forums for the way to do it, because of the alien nature of NixOS. This is different from, say, Gentoo, where you have to edit files by hand, but it is all standard stuff. Thankfully, there is detailed documentation, with examples, of all the configuration options, as well as a comprehensive package list, in a searchable format on the website. When it comes to installing packages, the configuration file is not the only route. That installs packages globally, just like other package managers, but there is also the *nix-env* command that can be used to install software and make changes at the individual user level.

As a general-purpose desktop distro, it is hard to recommend NixOS to anyone but an enthusiast. However, as a means to create custom environments, repeatable across multiple machines, NixOS has a lot going for it.

Yes, you could do the same with Debian and a configuration manager like *Ansible* or *Puppet*, but there is a certain elegance to the NixOS way of doing things. It is not so much a distro as a way of creating and installing your own custom distro.

Is this a good distro? That's a really complicated question. For most users the answer has to be no, but if you need to create multiple instances of the same, or a similar, system, and maintain them easily, it may be just what you are looking for. Just be prepared to put it some time getting it all set up first.



Linux distribution **REVIEWS**

EndeavourOS 2021.08

Arch Linux is beginning to rival Ubuntu as the starting point for modern distros. **Neil Bothwick** looks at one such distro.

IN BRIEF

EndeavourOS is a community distro based on Arch Linux, but made easier to get into thanks to the slick graphical installer, with a wide choice of desktops available to pick and choose from when installing. Created following the demise of Antergos. there really is something to suit everyone. Clean looks and helpful documentation round out this worthy distro contender.

ime was that when you installed a Linux distro. you got to choose what was installed. Distros like Ubuntu changed that so the current approach for most distros is to install a standard desktop and set of packages and let you change things later. This is faster but less flexible and has led to the existence of separate versions of distros for different desktops, such as Ubuntu. Kubuntu, Xubuntu, Lubuntu and so on. EndeavourOS gives the best of both worlds by giving you a choice of online or offline installations.

The latter will install the OS from the ISO image, complete with the XFCE desktop and associated

software. Or you can select the online option where you are given a choice of eight desktops: Xfce, KDE, Gnome, i3 (a tiling window manager), MATE, Cinnamon, Budgie and LXQT. This installs the desktop you choose by downloading it during the installation process. It results in a longer time for installation, but you don't need to be there for that and it means you spend less time sorting out software options post-installation. You can also choose to install no desktop at all, useful for a server.

EndeavourOS is based on Arch Linux and uses its *Calamares* graphical installer, which is clear and quick in use. A nice touch the first time you boot the new system is the Welcome app, a window with a set of buttons to do the things you often want or need to do after installing, such as applying any updates, changing the display or wallpaper and installing certain common extra software. It also has links to tutorials and helpful web pages.

We tried both the offline install of Xfce and an online install of the Budgie desktop. The Xfce desktop has been well covered: it is a lightweight and capable, if somewhat old-fashioned looking, desktop. However, EndeavourOS has managed to make it more attractive with some judicious theming. The Budgie desktop is a good compromise between appearance and performance. It is attractive and easy to use, without the resource usage of the likes of a full Gnome desktop, even though it actually uses Gnome components behind the scenes.

We tested the latest release at only a day old, but Arch is a rolling release distro, so release versions are less important. They only really determine your starting point; everything can be upgraded to the latest version from the Arch repositories with no need to reinstall when a new release appears, as it is just an ISO that installs what you already have. It's a matter of personal



 Pick a desktop, any desktop. Then pick just the parts you need or the whole lot... a truly customisable installation.

preference, but I like rolling release distros as you get updates when they happen, not when their release schedules dictate.

EndeavourOS uses Arch's *pacman* package manager. Unfortunately, there is no graphical interface installed by default, so new users will be left looking up *pacman* command line options in order to install software. This really is a glaring omission on an otherwise well thought-out distro.

The EndeavourOS Wiki contains a wealth of articles on using and setting it up. For those that would rather watch than read, there are also a number of video tutorials on there. In the unlikely event that you can't find what you need on their own wiki, you can always turn to the Arch wiki, which is one of the best such resources available for any distro.

All in all, this is a good distro; Arch is a solid base to build on and the choice of desktops is handled well. The documentation is good – it's just a pity that the ease of initial installation is not carried through to installing other software.



» Rating 9/10

REVIEWS Linux distribution

Linux Mint 20.2

Jonni Bidwell is running low on mint metaphors and time, but the promise of some budget mojito mix expedited this review.

IN BRIEF

A long-standing, well regarded, Ubuntu-derived distribution that puts the needs of its users first. Mint is aimed at both home and professional users with wide business use. ell, it's quite the Mint-fest this issue, what with it being on the cover, the disc, Jonni's virtual machines and (as always) Effy's laptop. Why all this attention? For a start we've always liked Mint, and, spoiler alert, this outing does nothing to change that. Next is that there are some nice under-the-bonnet changes with this release that we think you should know about.

SPECS RAM: 2GB Disk: 20GB under-the-bonnet changes with this release that we think you should know about. Before you get excited though, there's not much in terms of groundbreaking features this release; it's more a continuation of the bold course set out a half dozen or so releases ago. In that time we've seen the introduction of

user-friendly system roll-backs.



As usual, there's an exhibition's worth of beautiful, user-contributed contributed backgrounds in fully updated with 20.2.

ever-improving Nvidia Optimus support, HiDPI simplicity and so many nice tweaks and touches to make this list uncomfortably long if we were to continue. We've also seen further distancing from Ubuntu, through the disabling of the Snap Store, and, to remediate this, Linux Mint taking on packaging duties for *Chromium*.

Back in February, lead developer Clement Lefebvre – backed by statistics from search data – railed at users

» OTHER EDITIONS

The main text covers the Cinnamon edition, but the other two flavours (MATE and Xfce) are well worth your attention. Both of these desktops are a little lighter than Cinnamon, but we still wouldn't recommend using them with less than a gigabyte of RAM. In our testing the Xfce edition had the smallest memory footprint, but only by a whisker. For reference, Cinnamon used around 750MB, MATE weighed in at 600MB, and Xfce scored a slimline 550MB. Don't read too much into those numbers though; fire up a web browser on any of them and open couple of popular websites, and you'll soon see your memory usage skyrocket into the gigabytes.

the newest MATE 1.26 isn't included in this Mint. That's a shame, but something to look forward to in Mint 20.3. Mint Xfce includes the latest 4.16 edition of the desktop, released back in December. Xfce 4.18 is probably a long way off, but we're still looking forward to it.

There's another, lesser-known, edition of Mint too – Linux Mint Debian Edition (LMDE). If you have 32-bit hardware this is the only version you can install, but if you want to learn more about Debian (or don't like Ubuntu) then this is well worth checking out too. It is only available in Cinnamon flavour, but it's easy to switch that out for something more lightweight. See the Customise Mint feature for some pointers. There's no new release of LMDE this outing, but now that Debian 11 is out hopefully we won't have to wait too long. not applying updates in a timely fashion, and in particular the reported 30 per cent of users still running Mint 17. That version went EOL nearly three years ago, so anyone running it on a networked device is at serious risk. As it happened, our machine was still running Mint 20 (supported, like the whole rest of the 20.x series, until 2025) but the update tool enables you to jump straight to Mint 20.2. That went without a hitch for us, in part because there's a nice document that outlines the upgrade on the website.

In an attempt to make updates more visible, the Update Manager (which for a long time only gave you a red dot to notify of updates) will now be much more vocal if updates aren't applied for more than two days. This strikes a nice balance, improving user's security (*Firefox* in particular is updated frequently) without being an unnecessary nuisance. Notifications are only shown for security and kernel updates by default, but it's easy to customise this, as well as the frequency of package list refreshes and notifications.

Automatic updates are also possible, though disabled by default because otherwise you might find yourself suddenly having to restart due to a background update (to Systemd, for example) making things inconsistent. A large warning tells you to make sure to have *Timeshift* configured for regular snapshots if you enable auto-updating. It's also possible to turn on auto updates for Flatpaks and all the flavours of Cinnamon Spices (applets, extensions and desklets). Having Flatpaks update in the background shouldn't cause any problems, and when Spices are updated all you'll see is a transient notification and a hint of display flicker.

If you've ever had to rename a whole lot of files, you'll know that it can be tedious. You might also have taken matters into your own hands, sed-ing and grep-ing and regular expressioning through piped directory listings, in

Linux distribution **REVIEWS**



which case we hope you recover soon.

One of many things we like about Pop!_OS is the bulk rename function (available through the *File Manager*). It might not offer metadata-based renaming, but its simple search and replace and incremental numbering covers an awful lot of use cases. So it's nice to see the introduction of *Bulky*, a similar utility in Mint 20.2. Select a bunch of files in *Nemo*, and *Bulky* is available from the right-click menu. It can do search/replace, uniform insertion/removal and also case adjustment. *Bulky* can do regular expressions too – though we still haven't tried to use them ever since that time a stray '/' character renamed every file on our network to what looked like the ASCII shrug gone wrong. Merely shrug we most certainly did not.

Super search

It's also possible to search inside local files from the Nemo file manager, but on older machines you might want to disable this. Indexing large directories can take significant resources. There's a *Sticky Notes* application too, which will help anyone migrating from the paper alternative. We prefer to keep our Mint desktop clean so we can enjoy the fantastic wallpaper selection, though. The *Warpinator* utility, a simple tool for sending files around your network, has been updated. There's also an

1	You can make Mint's startup even more swift by disabling
l	of staggering various services

*			<u>e o</u>
ô	im-launch No description		
6	mintwelcome Linux Mint Welcome Screen		-
-	Print Queue Applet System tray icon for managing print jobs	Delay 30 s	-
'n	SSH Key Agent CNOME Keyring: SSH Agent		0
•	Support for NVIDIA Prime Shows a tray icon when a compatible NVIDIA Optimus graphics card is deter		0
	System Reports Troublesheet problems	Delay 40 s	0
	Update Manager Linux Mirz Update Manager	Delay 20 s	0
*	xapp-sn-watcher A service that provides the org.kde.QatuaNotilierMatcher interface for XApps		-

unofficial Android port (available on F-Droid and the Google place) so you can easily transfer files from your phone, though you may prefer to use KDE Connect.

New in Mint 20.2 is Cinnamon 5.0. As is common for open source projects (and confusing for pretend tech journalists), the major version bump doesn't really imply groundbreaking changes. Cosmetically, excepting some user-facing changes to a few applets, it looks much like it did in 20.1. A few memory-leaking bugs have been fixed in Cinnamon 5.0 so you might notice a decreased memory footprint, and hopefully find it even more responsive than the previous outing.

If you missed the update back in March, you'll notice there's now an option to restart Cinnamon if (despite memory-leak plugs) it still starts guzzling memory. This is possibly not the most elegant solution, but there's anecdotal evidence that in some configurations – none that the Linux Mint team or we could reproduce – Cinnamon still has a RAM addiction.

Overall this is an excellent release and we strongly recommend anyone running Mint 20 or 20.1 to upgrade immediately. Mint makes it easy to roll back if something goes wrong, so you really have no reason not to. And for those still running unsupported editions, please upgrade: Clem is really worried and you don't want to worry Clem, do you?



Griftlands

Management love a tough negotiation but not with the staff, so **Sam Greer** will be getting a quick visit from the LXF shock troops right about now...

SPECS

Minimum OS: 64-bit. Steam OS, Ubuntu 20.04+ CPU: 2.0+ GHz 64-bit Memory: 4GB GPU: Intel HD Graphics 5000, AMD or Nvidia equivalent HDD: 6GB ames have always been looking for ways to make conversations as compelling as the combat. Why can't a tense negotiation feel as gripping as a tough-as-nails boss fight? That's the question *Griftlands* seeks to answer. It makes for a compelling roguelike, experience thanks to great writing and characters, but doesn't quite live up to the promise of its chatty deckbuilding systems.

The t Griftlands are a chaotic place where corrupt law enforcement and criminal syndicates run things against a backdrop that combines science fiction with fantasy. Brutal it may be, but it sure is pretty. You'll get

to love the scratchy, hand-drawn touch to everything. There's a smart economy to where the embellishments are made, which ensures the game can deliver on its scope without any rough edges. Character sprites have lovely gestures and expressions that complement the writing, while the overworld map is simply littered with tidy icons.

Gangland style

Griftlands' factions are written in broad caricature, in line with the game's cartoonish look – which stands out in a good way among the crowd of indie pixel-art based games. However, there's always a serious undercurrent of scepticism about the status quo. Criminal gangs prove virtuous, while the local police prove to be the most ruthless of all. In this way, developer Klei paints a strong



Sharp writing elevates almost every NPC into someone memorable.

picture of (downtown Bath? – Ed) the world you step into as a 'grifter', the game's shorthand for a smooth-talking mercenary. There are three to play as one, each with their own story, but only one is unlocked to start with: Sal, a once-indentured slave now on the hunt for the woman who sold her.

Revenge is hardly a novel motivation, but *Griftlands* shines thanks to fast-paced dialogue that fleshes out its characters effortlessly. What a relief it is to play an RPG that doesn't involve scrolling through paragraph after paragraph of overdramatic guff that happened 1,000 years ago. There's an immediacy to *Griftlands*, in both the world-building and the stakes. Life's short, and only what's in front of you really matters (*sounds like my Monday morning mantra* – *Ed*).

Even the smallest of characters get a chance to make an impression, as you collect an arsenal of subplots en



The overworld is purely functional but it looks lovely.

Deck-building game **REVIEWS**



Body modifications will give your character the upper hand.

route to your overarching goal. The writing can be damned funny too, the humour emerging from Sal being just a tad smarter than most of the other thugs she encounters (but not that much smarter). There's a charm to these cutthroats, who are often acting out of sheer desperation rather than any kind of malice. It's a world of opportunists, as you'd expect, but there are surprising moments of loyalty and kindness too. We found that the allure of *Griftlands'* story was far greater than that of its expanding deck-building.

Fight, talk, fight

There are two modes of play: battles and negotiations. Battles are turn-based affairs where you draw cards for attacks and abilities, and can earn or hire party members. Overall, they're fun but conventional.

The negotiations are far more interesting. There, your cards are used for building arguments against your opponent. You're constantly looking to chip away at your enemy's smaller points, while piling damage onto their core argument. We really felt like we were juggling a lot of plates, in the best way – the pressure was always on to make the most of each turn.

As an abstraction it's captivating, but *Griftlands* never quite bridges the space between the actual dialogue and the card game in which debates take place. Compared with something like *Signs of the Sojourner*, which much more cleverly uses its mechanics to shape and fuel character interactions, *Griftlands* feels a little awkward. Negotiations are disconnected in a way that turns every interaction into a simple binary of success or failure. We wish we could see the impact of our chosen cards on our relationships or the story.

Worse, committing to negotiation decks will punish you during the game's handful of inescapable combat



Don't take your dog off the lead at the local park, as it can cause issues.



encounters. There are side activities to help, but in general, focusing on combat decks is the easier path through the game.

Negotiations are the game's card battling at its best.

Non-player cards

Griftlands does make good on far-reaching consequences. Recurring NPCs pop up in unexpected places, leading to tough spots where you might have to contemplate a double-cross to achieve your goals. These are also randomised for subsequent runs, to an impressively varied degree. This ensures new runs feel reasonably fresh and that softens the blow when you do die.

You can get a boost on your next run, too – but there's only so much of the sting it can take away, and despite the convincing manner in which side-missions are generated, the main story remains the same. Failing close to the finish line is never not going to take the wind out of your sails. There is a Story Mode setting for those who want to focus purely on the writing, but it feels like a sticking plaster on some harsh difficulty spikes. We were coasting through the game for the first two days, and then out of nowhere comes a fight several orders of magnitude tougher than anything preceding it. It might make narrative sense, but it's not very interesting to keep restarting and bumping into walls like these.

Shortcomings aside, Grifflands is another slice of low-key brilliance from developer Klei. Is there a genre it can't do? The balance of narrative and deck-building made for a much more engaging experience than we've often had with card-based titles, even if it feels like more could be done to connect those systems and bolster the storytelling. Still, we're going to remember our adventures with Sal and the little moments of friendship and betrayal throughout each run. It fits right in here at *Linux Format Towers* playing as a grifter; it ain't much but it's a living.



Charm and good storytelling overcome the limited, experimental deck-building experience.

» Rating 8/10

WE COMPARE TONS OF STUFF SO YOU DON'T HAVE TO!

Krita 4.3 » GIMP 2.10.24 » Karbon 3.2.1 » Inkscape 1.1 » MyPaint 2.1.0



Michael Reed

is such a nerd that he uses SSH to connect to his Linux box from his phone as part of a wild night out.

Digital art programs

Donning a big shirt and a beret, **Michael Reed** examines five open source art packages, and discovers how they can help him tap into his creativity.

HOW WE TESTED...

To assess each package, we put it through its paces by carrying out some typical artrelated tasks and also trying to create some images from scratch. The designers of this type of program are constantly running up against the problem of adding lots of features while making a user interface that's comprehensible. On top of this juggling act, the day-to-day ergonomics have to be smooth, with the most commonly used features remaining accessible. That said, the less routinely used features, such as configuration options, should be well organised and intuitive.

All of the programs supported pressure input from a drawing tablet, but they were usable with a mouse, too.

Bear in mind that many of these programs have a lot of options and features, and hopefully they have a good support community and documentation, and this is another area that we investigated. Programs as potentially complicated as these could end up being a dead end without decent support materials.



elcome to the world of art, created digitally. To this end, we have assembled a line-up consisting of five fantastic applications that are all capable of creating art from scratch, and although they're all worth recommending, each of them has its strengths and weaknesses compared to the others. Which is the "best" may depend on what you're trying to accomplish, and your level of experience.

The applications that we're looking at can be divided into two main categories in terms of how they work. Art packages like *GIMP*, *Krita* and *MyPaint* are primarily bitmap editors, so they affect the pixels that make up the image. On the other hand, vector illustration packages such as *Inkscape* and *Karbon* build up drawings from a collection of lines and curves and shapes.

GIMP and Inkscape are the most featurepacked programs of the five, but that complexity comes at a cost in terms of learning curve, and the other three have the virtue of being a bit easier to get into for the beginner. We're assessing these packages as art packages primarily, but we'll also make some notes about how they perform in other areas of image editing, such as photo editing.

The range of painting tools

What features does the application have and what tasks can it handle?

Ithough Krita is well known as an artist's choice, it and GIMP aren't as far apart as you might think when used for art. GIMP features the essential tools such as paintbrush, pencil, airbrush and ink tools, along with the expected plain, gradient and texture fills.

In addition to the expected bitmap-based drawing features, *Krita* also has vector drawing facilities, and although it wouldn't be a match for a dedicated vector drawing program like *Inkscape*, in that department the features are surprisingly extensive. *GIMP* also has some vector drawing facilities, which are welcome, but they're not quite on the same level as those of *Krita*.

MyPaint has a nice feature in the form of the scratchpad docker. This gives a small canvas in the dock area that can be sketched on or used to mix colours using the paint tools. A similar feature is promised for *Krita*, but it's not in the stable release yet. A feature that *MyPaint* does lack is proper handling of reference images, useful for copying from an existing image or tracing over it. Both are common techniques when learning to draw.

When it comes to painting, another feature we liked about *MyPaint* was the infinite canvas. In other words, you can zoom out and pan around an image, and the software automatically extends the bounds of the canvas as you go. It's a fiddly job to manually add extra space to the canvas in *GIMP* or *Krita* when you need it. *MyPaint's* user interface is a friendly one, but as we got further



You can usually get an overview of what an art package can do simply by looking at the main toolbox.

into exploring it and switching features on, we discovered that it's well-equipped when it comes to pure painting and drawing features. This fits in with our overall experience of testing *MyPaint*: a streamlined workflow for those who just want to use it to draw and paint, with advanced features at the ready if they're needed.

As with all areas of its operation, the drawing features of *Inkscape* tend towards the technical side. For example, you can get down to the exact details of editing a curve, but it's a bit fiddly if you just wanted to doodle something. Fellow vector art package *Karbon* is a simpler program, but it covers the basics and it's easier to get started with than *Inkscape*.

VERDICT

GIMP	8/10	MYPAINT	8/10
KRITA	9/10	INKSCAPE	8/10
KARBON	6/10	ALC: NO.	

MyPaint balances tools with ease of use, but Krita is more fully featured.

Brush and pencil options

Brushes are a fundamental part of digital painting.

Rushes are probably *Krita*'s strongest area. For a start, there are masses of them, divided into categories and each with its own colourful, custom icon. The downside is there's a bit of 'settling in' time while you learn to navigate around them all. If you did need a type of brush that *Krita* doesn't come with, then you can probably create it. For instance, you could take an existing brush and edit it so that it, say, darkens and blurs according to pressure setting or stroke speed, and that interface is nicely done.

All of *GIMP*'s brushes are fully configurable, and just clicking the Brush tool makes this apparent, because the left panel (by default) is filled with many options such as size, angle and hardness. If there was a brush type that you couldn't find, it might be possible to approximate it by changing the parameters or loading in a custom brush shape.

The brush facilities in *MyPaint* are extensive. By default, the selection isn't as extensive (or overwhelming) as *Krita*, but if you



The brush facilities of MyPaint are excellent.

right-click one, you can open the brush editor which has dozens upon dozens of subtle options. The tool options dock, with parameters such as brush size and smoothing amount, is neat and simple. As always seems to be the case with *MyPaint*, these areas look simple at first, but have a surprising amount of depth.

Inkscape has good facilities for brush-based art and customisation of brushes, but like most of that program's features, it takes a bit of work to figure out how to do things.

VERDICT

GIMP	7/10	MYPAINT	7/10
KRITA	9/10	INKSCAPE	7/10
KARBON	6/10/		

Krita is the leader here, but it's a bit overwhelming. MyPaint is, as ever, highly accessible to digital art newcomers.

User interface and experience

Is the creative process helped or hindered?

hat we're looking for in an art package is a smooth workflow that doesn't keep bogging us down with technical options. That means a clear interface, and one that doesn't cause us to stop and start when switching between commonly used features.

The user experience of an art package represents a conundrum because the more features that it offers, the more complicated things necessarily become. This means that less fully featured, or more singularly focused packages have a built-in advantage due to their simplicity.

Often, with applications of this sort, you'll arrive at a work pattern that you come to rely on, so configurability is important. Indeed, being able to come up with your own user-interface layouts can overcome basic design shortcomings in some cases.

All of the applications had the same problem when used with a drawing tablet, because it's difficult to hover over icons to see the tooltip help.

Finally, user interface scaling is particularly welcome for this type of creative work.

Krita

The overall *Krita* user interface is optimised for artist use. It's a very 'open' layout: most of the common options are laid out for a smooth workflow when painting and drawing, and menu diving and working with complex pop-up dialogs, for typical tasks, is minimised. However, there's a certain amount of sprawl because of the sheer number of brushes and options.

The interface is configurable, with movable, switchable docks, and most users will probably find themselves taking advantage of this, once they learn their way around the program. Right clicking with the mouse (or stylus button) brings up a circular menu palette, rather than the usual text based pop-up menu, and it can be customised with your favourite brushes. Touches like this make it clear that *Krita* was designed from the ground up as an artist's tool. There's a lot to the interface, but it's enjoyable to use.

8/10 GIMP

7/10

The *GIMP* interface has a reputation for being difficult for beginners to get the hang of, and it's orientated towards image editing and processing rather than the creation of art. However, that doesn't mean that it's not suited to a more creative role, particularly with configuration. Having said that, it's not as streamlined as *Krita* for painting and drawing, even though both programs cover similar ground.

More recent releases of *GIMP* default to a dark, monochromatic look that some might like. Certainly, don't advise someone to 'click the grey icon' as they won't be able to find it. *GIMP* now also defaults to a main toolbox that groups icons together in a way that isn't conducive to a creative workflow. Thankfully, colour icons, separate icons and larger scaling can be enabled in the configuration options, if you know how. There is also a fair bit of configurability of the interface.



Image processing and retouching

Some of the software on test can edit existing images, such as photos.

G *IMP* is the leader in this section, and it's overflowing with plugins that affect either the whole image or selected regions. The downside is that many of the plugins are similar to one another, and it can take quite a lot of experimentation to sift through them all.

Krita features a lower number of image processing effects compared to *GIMP*, but we'd rate the average quality of what's there to be higher. Both *GIMP* and *Krita* can work with the G'MIC image plugin system, so it's unlikely that you would ever run into a genuine roadblock if you needed a particular image processing feature with either program.

In a pinch, you could use *GIMP* for photo-editing tasks because the application can load camera RAW files, and it covers most of the basics that are needed in this role, such as cropping to a given aspect ratio, altering contrast, gamma and white balance and colour grading an image. *Krita* could also be used for photo editing, to an extent, but it's missing some of the frequently needed tools. This area isn't the primary focus of *MyPaint*, and it lacks cloning and healing facilities along with an absence of plugin support and general features to manipulate the overall tone of an existing image. Both *GIMP* and *Krita* are both very well suited to this area (retouching), with *GIMP* slightly ahead.

Inkscape is a vector drawing package and so isn't much use when it comes to the processing of bitmap images, although it can incorporate existing bitmap images. On the plus side, it's perhaps the leading Linux program for tracing bitmaps into vector images.

VERDICT

GIMP	10/10	MYPAINT	5/10
KRITA	8/10	INKSCAPE	3/10
KARBON	2/10		

GIMP can confidently handle most image processing and editing tasks, with Krita not far behind.

Digital art programs **ROUNDUP**

MyPaint

<u>8/10</u>

When first launched, *MyPaint* gives the appearance of a simplistic paint program, but on closer inspection it's because most UI docks are disabled by default. Once you get stuck in and begin enabling the docks, *MyPaint* starts to offer up the more advanced features, making the most of a three-pane interface. On the one hand, people might be dismissive of the program when first running it, but it does also mean that you could simplify the program if you were setting it up for someone else.

We liked the way that the close document dialog tells you how much work time you'll lose if you don't save. Over and over again, the UI presented the most important options rather than boggling us with an overwhelming array. Overall, it looks like a piece of software designed for people who just want to draw rather than something for the tech-heads, but the depth is there if you want it.

INNOVINE HAN

Inkscape

The version 1.1 update of Inkscape introduces a launch dialog offering startup options such as choosing the UI theme and document type. Essentially, Inkscape uses a three-pane interface, but it's a complicated one that rivals Blender in terms of the sheer number of icons and input dialogs on display. There are three main toolbars, and on the right are the parameter dialogs. However, the complexity of the approach makes sense because Inkscape is a vector drawing program, and you'll often want to alter the parameters of strokes and other objects after they are added to the canvas. The icon scheme is colourful and well-defined.

It is a complicated layout, but the flip side is that familiar operations are available with a few clicks of the mouse. Eventually, a workflow starts to establish itself once you've worked out where the needed features are located.

0

C



6/10

Like a lot of the *Calligra* suite (which *Karbon* is a member of), the user interface has a slightly unfinished look. The layout consists of a standard toolbar along the top and a toolbox containing the drawing tools on the left-hand side. On the right hand are the tool options. All of these sections are scalable, and we suspect that many artists will ramp up the size of the icons to make most of the available screen space, particularly as it isn't an especially intricate interface compared to that of other, similar, programs. Having said that, we were able to bring out some layout bugs by going too big.

Karbon keeps things relatively simple compared to the likes of *Inkscape*, but having said that, it doesn't have anywhere near the total number of features that *Inkscape* does. In typical use, you can usually find what you want, and simply get on with creating vector-based art.



Documentation and tutorials

What resources exist to help you learn how to use the software?

f the packages we've looked at, we consider *Krita* to have the best documentation available overall, particularly when considered as, specifically, an art package. The official website contains a full online user manual. The tutorials selection is also excellent. Searching around on YouTube shows that it's well supported in the area of screencast tutorials made by the wider community.

The Inkscape website features a full online manual, a beginner's guide (downloadable) and various illustrated tutorials. Searching around on YouTube revealed a healthy selection of tutorials covering most areas of the software. Due to its inherently technical nature, you'll almost certainly need these resources.

GIMP is one of the most famous and widely used open source applications, and this means that there's a lot of support for it online. The website contains a number of tutorials along with a full online manual. Searching around on YouTube reveals many tutorials for most aspects of the program, but bear in mind that they're not all aimed at explaining the art creation side of *GIMP*.

MyPaint may not have the instant name recognition as a Linux program, but it's well known on other platforms and an online search revealed plenty of video tutorials. It may even have a slight advantage over some of the other programs, such as GIMP, because a large proportion of the tutorials are made by artists, painting in the software. There is an online manual, but it's a bit on the brief side and some missing sections.

Compared to the other programs in this month's Roundup, *Karbon* is let down by an overall lack of documentation – both officially and from the online community.

VERDICT

GIMP	7/10	MYPAINT	8/10
KRITA	9/10	INKSCAPE	8/10
KARBON	3/10		

All of the applications other than Karbon have good documentation. The MyPaint and Krita community support is more art-focused.

Text handling

Often, when working with images, you'll need to work with text, too.

f the programs we're looking at in this month's *Roundup*, *Inkscape* is the best one for working with text. At all times, text objects remain fully editable and you can carry out tasks such as add shadows and gradient fills, and alter the line style and width. Beyond that, you can convert text into a path to make further alterations to the shape. There's also a tool that makes it possible for you to constrain text to the bounds of another object, such as a box.

When you edit text in *GIMP*, it creates a new layer, and that text remains fully editable – a feature we consider essential. Of course, the downside of this arrangement is that you can't decorate the text using other tools. The overall text facilities are basic but good, and offer the expected selection among the fonts installed on the system and some rudimentary features to constrain the text to a box. In previous years, we had found the text facilities of *GIMP* to be a bit buggy, but putting the current stable release through its paces, everything seemed to work as expected.

Krita takes a similar approach, although the text is added to a vector layer rather than a bitmap one. There's a long-standing bug in the program that prevents the successful selection of text font weight in the dialog. The developers are aware of the problem, but it's taking a long time to fix.



Text remains fully editable while you change the fill and line styles in Inkscape, then add drop shadows for extra impact (if you say so-Ed).

MyPaint doesn't include text-handling facilities as such. However, it'll enable you to add a vector layer containing text with a program such as *Inkscape*. You can then use the 'Edit layer in external app...' option to continue editing. It's a basic but workable arrangement.

VENDIOT			
GIMP	7/10	MYPAINT	4/10
KRITA	7/10	INKSCAPE	9/10
KARBON	7/10		

Inkscape has the strongest features in this area, but all the programs on test can add text to images.

Working with lines and shapes

How is vector drawing handled in this month's Roundup candidates?

I of the packages offer the basics when it comes to using predefined shapes such as rectangles and ovals, along with lines and curves. You might, for example, use these tools to draw a cartoon character with a perfectly circular head and eyes by using the Ellipse tool, or you might add a straight horizon by using the Line tool. As we've already mentioned, *Krita* features some fairly decent vector drawing tools, and these exist on a specialised layer and continue to be editable after they're drawn. However, the two dedicated vector drawing packages (*lnkscape* and *Karbon*) are, naturally, more suited to images that are principally constructed from elements of that sort.

If it's illustration-style image creation that you're interested in, for things like comic characters and project logos, *Inkscape* is hard to beat, but learning how to use it can be a challenge. Once you've placed a basic shape such as a rectangle on the canvas, you can begin to alter every attribute of that object such as its colour, the line style and the bevel of the corners. Beyond that, you can edit the shape itself by, for example, adding additional points to an existing line or curve that, itself, can make up part of a larger shape. The downside of this level of control is that we'd doubt that any newcomer could get into detailed editing without searching the documentation first. From the developers point of view, it's



Karbon might have a sparse user interface, but it offers all of the basic vector drawing facilities that you'd expect.

difficult to implement that level of control without increasing the level of complexity of the application and its interface.

Inkscape couldn't be classed as a CAD package as such, but it can be used for some technical drawing tasks along with things involving page layout such as poster creation and adding text to webcomics. *Karbon* covers the basics of vector drawing. It doesn't go into the same level of detailed features that *Inkscape* does, and this is both a benefit and a disadvantage, as you're less likely to have to resort to the documentation when drawing.

VERDICT —

GIMP	7/10	MYPAINT	7/10
KRITA	7/10	INKSCAPE	9/10
KARBON	8/10		

Vector drawing packages are ideal for working with geometrical shapes.

Digital art programs **ROUNDUP**

The Verdict Digital art programs

e've decided to name *Krita* as our top paint package because it balances a huge number of art features with an artist-focused, attractive and configurable user interface. As well as its prestigious facilities in the realm of art creation, it also has a fair few in the area of image processing. Along with being a bitmap-based tool, it even has some fully usable vector drawing facilities.

Inkscape is, at its heart, an illustration program. You can freehand draw and paint with it, but ultimately you're constructing the overall illustration from shapes, lines and curves. At all times, you can inspect and tweak these elements to your heart's content. This may be handy for those without a steady, practiced hand when it comes to drawing, but many people simply prefer this way of working. It can work with bitmap images, but it's limited when it comes to creating bitmap art. This means that there are certain types of art style that it wouldn't be suitable for, such as pixel art or brush-based art that requires a lot of subtle control.

GIMP is one of the best known open source GNU/Linux applications. It would be doing it a disservice to call it a 'Swiss army knife' as that implies that it's not a robust workhorse when carrying out day-to-day activities, which it is. However, as powerful as it is, it wasn't designed from the ground up to be a tool for creating art from scratch. For this reason, although it has a lot of the same tools as *Krita*, the workflow isn't as problem-free when used in this role.

MyPaint is aimed at people who just want to draw and paint using a computer, and it's proved to be a surprisingly deep piece of software, if you want it to be. By adding and removing docks, it can be turned into either a fully featured art package for professionals or a simplified sketching application. It wouldn't be our first choice as an image retouching and processing application though.

Karbon is a vector drawing package, but it doesn't have the same depth of features as *Inkscape*. Some of the user interface is so sparse it makes the program feel unfinished. However, drawing in the package is a mostly pleasant experience, and it actually helps that the features are focused on that task, even though it's less of a generally useful tool than *Inkscape*.



1st Krita

Web: https://krita.org Licence: GPL 3.0

Version: **4.3.0** Enabling artistic creation was obviously the driving principle behind *Krita*, and it's well-equipped for that type of endeavour.

9/10

2nd	Inkscape	8/10
	IIIKJCupc	0/10

Web: https://inkscape.org Licence: GPL 3.0

Version: 1.1 Inkscape is designed for vector-based art and has a huge array of tools for that and other tasks involving text, graphics and page layout.

3rd	GIMP	8/10

Web: www.gimp.org Licence: GPL 3.0

Version: 2.10.24 *GIMP* has more than competent painting and drawing facilities and a lot more besides in the realm of image processing.

4th MyPaint 7/10

Web: http://mypaint.org Licence: GPL 2.0

Version: 2.1.0-alpha+git.90b36db Not as well known as some of the other packages, MyPaint is an enjoyable and comprehensive painting tool.

5th	Karbon	6/10

Web: https://calligra.org/karbon Licence: LGPL-2.0

Version: **3.2.1** *Karbon* contains most of the vector drawing facilities that you'd expect, but it lacks the detailed features of *Inkscape*.

» ALSO CONSIDER

3D modelling tool *Blender* (www.blender.org) has a mode called *Grease Pencil* for drawing 2D line art (in a 3D world). The vector drawing facilities cover the basics, and it's well-suited to jobs like creating a webcomic.

LibreOffice (www.libreoffice.org) or OpenOffice(www. openoffice.org) come with a vector drawing module that's surprisingly feature-packed. It has some business-orientated features such as built-in shapes like connectors and arrows. On the bitmap painting front, *LazPaint* (https://lazpaint. github.io) is worth checking out as it's fairly easy to get into, and it has lots of interesting features. However, we preferred the *MyPaint* interface and broader featureset. *KolourPaint* (https://apps.kde.org/en-gb/kolourpaint) is a simple package that takes its inspiration from *Microsoft Paint*. *Pinta* (www.pinta-project.com) sits between the two, with more of an emphasis on photo retouching-type jobs.

Jonni Bidwell shows you how to truly make it your own.

inux Mint continues to go from strength to strength, as you'll know if you've already had a play with the latest 20.2 release. If not, what are you waiting for? Fire up that there LXFDVD and witness the, er, Mint-ness forthwith. Or check out our thorough review on page 26. See, now you want to install it, don't you? And that is just the beginning of the adventure. One of the things that makes Mint so cool is its configurability. It's often said (by us) that Mint is an ideal beginner's distro, and it turns out

it's also ideal (we say) for beginners to tinker with.

TUX-MNT20

The flagship Cinnamon desktop can be transformed not just with swishy effects and colourful themes, but with all kinds of extensions, applets and desklets (collectively known as 'spices' in Cinnamon parlance). And MATE and Xfce, the desktops featured in other editions of Mint, are equally seasonable. But we can do better than that: why not mix it up and install a whole new desktop environment? We'll show you how to install the outstanding KDE Plasma, we'll even show you how to make it work with the state of the art Wayland display protocol.

If you like things slimline, we'll show you how to go minimal with the featherweight Sway desktop, again powered by Wayland. Sway is based on the i3 window manager, popular with power users and those who cannot abide desktop bloat. We'll have you doing everything in the terminal and tiling like a pro in no time.

And just so no one gets hurt, we'll start with a handy reminder about how you can use Timeshift to easily undo any desktoprelated mishaps.
The joy of tinkering

Sort out roll-backs so you can customise Linux Mint with impunity and immunity (to problems).

urveyors of historic issues of *Linux Format* may be able to correct this, but as far as our research can tell, the first mention of Linux Mint in our magazine came in the Distrowatch column of LXF094, when Mint 2.2 was released. Even back then Mint was notable for its out of the box experience, bundling codecs, Java and Flash plug-ins and wireless firmware, saving users from having to shoehorn those on there using fragile instructions from a random forum post.

That experience remains central to Mint, and though wireless hardware is well supported on most distros (and no one needs Flash any more), it still shines. Right from the Welcome screen in fact, which will invite you to set up backups using *Timeshift*, switch keyboard layouts, or send and receive files from another machine using *Warpinator*. Oh and there are minimise buttons on windows in Cinnamon – a trend fast disappearing on other desktops, but one which makes many a user feel at home.

Other desktops are going full steam ahead with Client Side Decorations (CSD, which allows applications to draw their own titlebars). This might allow programs to make best use of space and provide a coherent interface. Or it might make them look inconsistent, clumsy or other pejorative terms – it depends who you talk to. At any rate, Mint's X-apps are refreshing in their avoidance of the CSD wave, and Mint's huge fanbase suggests that they're still doing all they can to keep users happy.



Making the menu transparent and pasting Sticky Notes are but one way of customising Mint.

Sooner or later though, you're going to want to change things up. It generally starts with changing your desktop background and Cinnamon theme. These are important, but also quite easy – easy enough that you don't need us to tell you how to do them. What we'll be doing is a little more earth-moving. Tectonic stuff like installing whole new desktop environments, swish display managers, maybe even switching to the Wayland display protocol. And while these aren't without risk, Mint's Timeshift program allows you to back up your system files (much like Apple's *Time Machine* or Windows *Restore Points*), affording an easy way to undo any desktop mishaps.

Even if nothing goes wrong, it's handy to be able to roll back to a cleaner system rather than unpick changes manually: see the walkthrough below. If you've already got *Timeshift* set up, take a manual snapshot now before pouring in all the packages over the page. Go on, you know it makes sense!

EASY ROLL-BACKS WITH TIMESHIFT







Start Timeshift

Fire up Timeshift and set it up to take a couple of daily snapshots to a local drive with plenty of space (at least 1GB more than the current filesystem size). Timeshift only backs ups system files by default, so files in your home directory aren't included. There are better tools for backing these up.

2 Take a Snapshot

It might spring into action immediately if the clocks align. But don't worry if it doesn't, just hit the Back Up Now button to take an on-demand snapshot. Timeshift backs up incrementally so only changed files are stored. Once the snapshot completes, add a helpful description to help future you keep track.

📲 Restore a snapshot

If something goes wrong, you can now easily restore a Snapshot by clicking the button. You might want to examine the files within first, which you can do by rightclicking. Even if things go really wrong, and Mint no longer starts, you can use Timeshift from a live medium. Just point it to the **/timeshift** directory.

>>

Tweaking Cinnamon

See how easy it is to make your mark on Linux Mint's flagship desktop environment and beyond!

f you haven't had a nosey around Cinnamon's many settings, you might be pleasantly surprised at how configurable it is. When Mint 20 was released much ado was made about the Mint-Y theme now having fifty shades of colour variations (okay, it was 32), but we haven't found our favourite hue yet. Check out the palette by opening up the main menu and going to Preferences then Themes.

Dark themes are so common even Windows has them (*but not Google Docs–Ed*) now, but Mint has a corresponding dark theme for each variation. You can download whole new themes from the web too; just don't expect them to all be in line with your design preferences. Hidden away in the Settings section of the Themes dialog are some oft- overlooked options for sorollbars, including the option to disable overlavs.

If you want to experiment with Xfce it's easy to install it and all its apps from the Software Manager. To customise the main panel, go to Preferences and then select Applets. Now you can add all kinds of shortcuts and widgets. For example, select the Expo applet and click the + at the bottom to add a shortcut (via a smooth animation) to an expo-style overview of



your workspaces. For even more efficient workspace shifting (at the cost of some panel estate) add the Workspace Switcher applet. If you want to disable Expo, or any other applet, just click the – button. Like themes, third-party applets can be downloaded by visiting the appropriate tab. Downloaded themes come with no guarantees, so they come with an uninstall option in case they annoy you.

There are some extremely pleasant new wallpapers in Mint 20.2, and we recommend to right-click the desktop and choose Change Background if you haven't already perused them. But before you click there, take a look again at that desktop context menu. In particular, have a gander at the Add Desklets option. There aren't many pre-installed desklets, but if you want a digital clock or photo frame on your desktop then you're in luck. If you delve into the Download tab you'll find plenty more, including an analogue clock as well as more productive tools such as the Google Calendar desklet.

The final flavour in Cinnamon's Spices cabinet is extensions. These change the way the whole desktop behaves. Again there aren't many installed by default, but head to the Download tab and it won't take long for 'Wobbly Windows' to catch your eye. Hopefully you have better luck than us with that particular extension. If you're of that pedigree, you'll remember the Desktop Cube extension too, taking your workspace switching to a whole new level. There are other extensions which some may write off as desktop fripperies, such as being able to tweak window decorations, shadows or transparency, but there's no harm trying them out.

By default desktop effects are enabled in Cinnamon, unless your install has fallen back to software rendering – in which case have a look at the *Driver Manager* for possible remedies. These effects aren't the sort of in-your-face, windows catching fire, stunts of the early 'aughts, but have been designed to help users navigate

>> MUTATING MATE AND EXTENDING XFCE

But what about Mint's other flavours, surely they can be customised too? Indeed they can. You'll find a similar arrangement with themes, extensions and effects in both of these. All three of the Mint desktops are ultimately based on GTK3, so the fundamental desktop elements can be themed with a standard GTK theme from the likes of https:// www.gnome-look.org/browse?cat=135. Extract any themes you like into your ~/.themes directory, and they should pop up in the theming options. If you're using Mint MATE or Xfce, but desire to try Cinnamon, then that's easy. The standard desktop is available through the **mint-meta-cinnamon** package, or you can get a minimal edition via cinnamon-core. Conversely, if you are using Cinnamon and want to try Xfce or MATE, hit up the **mint-meta-xfce** and **mint-meta-mate** (rolls off the tongue nicely that one) packages. Again, there are minimal packages too if you only want the core applications. These will all add a session to your login screen, so you can choose your desktop from the menu to the right of your username. Speaking of the login screen, you can tweak that too. Just go to Administration > Login Window from the main menu. Such cosmetic tweaking can happen even earlier in the boot process, GRUB can be themed and so too can the Plymouth splash screen (https://www. gnome-look.org/browse?cat=108). the desktop. It's reassuring (sometimes) to see where applications were called into being from, and where they disappear to when they minimise. Be that as it may, you might want to turn these off, and this you can easily do by from the Effects option in the Preferences menu.

Besides eschewing the Gnome desktop, Mint has made a couple of other choices that fly in the face of desktop Ubuntu. One is the absence of the Snap daemon, which prevents installing packages from Canonical's Snapcraft store. Another is that there's no Wayland support (yet) in either Cinnamon, Xfce or MATE, the three officially supported Mint desktops. You might not care about next-gen, cross-distro packaging systems or banishing creaky old X.org from your machine. Indeed plenty in the Linux Mint community (looking at the forums) seem to share this sentiment. But there's good stuff in the Snap store, so let's look now at how we might enable that now. Plus, Wayland is pretty impressive now, so we'll look at that over the page, once we have a desktop that supports it.

When Mint 20 was released sans Snap support, an immediate consequence was that there was no way to install the *Chromium* web browser, since Ubuntu (20.04 and later) now only packages it as a Snap. We think more people should use *Firefox*, and if our user agent tracking on **linuxformat.com** is anything to go by, it seems they are. But it sets a potentially worrying precedent; if *Chromium* was to be abandoned by Ubuntu's DEB packagers, then perhaps other popular applications might go the same way. We don't think you should worry, as so far there's been no sign of that. Team Mint now packages its own *Chromium* DEB package, so if you're craving a hint of Google in your browsing then fetch it from the Mint App Store or with a good o' fashioned *apt* incantation at the command line.

We've seen that software is also available as Flatpaks, and that this is enables a wealth of software to be installed from outside the Ubuntu (and Mint) repositories. Flatpak is actually enabled out of the box on Linux Mint, and if you look carefully you'll find some Flatpak applications in the *Software Manager*. For example, if you want the latest version of *GIMP*, the Flatpak edition is probably the second one in the search results. You'll also find Flatpaks of *Spotify* and *Steam* so you can queue up your Rush playlist and play *Space Invaders* like it's the olde days.



If you do want Snaps, perhaps to get the latest version of *Blender* or the PyCharm IDE, that's easy to sort out from the Terminal. We first remove the file that prevents the Snap daemon from being installed, and then install it with *apt*: **\$ sudo rm /etc/apt/prefs.d/nosnap.pref**

\$ sudo apt install snapd

Snaps can now be searched and installed from the command line. You can also browse what's on offer at https://snapcraft.io, but for a complete GUI experience

USEFUL DESKTOP EFFECTS "These effects aren't the sort of in-your-face, windows catching fire, stunts of the early 'aughts."

you'll want to install the *Snap Store* too. Appropriately, it's available as a Snap and it can be yours with: \$ sudo snapd install snap-store

One easy and surefire way to turn heads (or whatever is the virtual equivalent) is to bling your terminal with a little transparency. Not only does this look pleasant, but if you arrange your windows correctly, and get into the habit of arranging them as such, it can be pretty useful as well. Endowed with X-ray vision, one can make out both the terminal itself and the web browser or whatever substrata lay below this.

To enable the opulent opacity effect in the Terminal, go to Edit > Preferences then select the current profile (it will be named Unnamed Profile if you haven't given it), untick the setting about system theme transparency, then tick the box above it. Play with the slider to find the optimum opacity. Or, if you'd rather more drastic changes look over the page, in which we install the highfidelity KDE Plasma desktop environment.



We have fond memories of wobbly windows confusing our graphics drivers so its good to see this still lives on.



Installing KDE Plasma

Transform your desktop with the smooth, svelte, sumptuous experience that is KDE Plasma and go complete next-gen with the Wayland too!

here's no official Linux Mint KDE edition these days, but that doesn't mean Mint users should miss out on the wonderful experience that is KDE Plasma. It's modern, but still has a traditional applications menu. It's incredibly polished, but is nowhere near the resource hog it used to be. Oh, and its *Dolphin* file manager is a joy to work with, especially if you're finding *Nemo* a little too simplistic (*like that attempt at Pixar humour – Ed*). Be that as it may, installing a new desktop environment comes with consequences, and it's good to be aware of these before you blame us for ruining your system.

Firstly, there's the disk space hit. The smallest KDE Plasma metapackage provides a minimal desktop, but according to the screenshot it pulls in some 850MB of dependencies in 446 packages. If you go for the full-fat edition, with all the applications from the KDE ecosystem, that will cost you close to 3GB. Next is the duplication of core utilities such as text editors, media players and screenshot tools. These all start to crowd your application menus, and if you use, say, KDE's Dolphin file manager in Cinnamon, it looks a bit odd.

Finally, it's sometimes hard for *apt* to completely uninstall a desktop. It's likewise hard to repair a broken desktop after you attempt to clear out packages manually. So don't do that; instead, try things out in a virtual machine first, or make use of *Timeshift* to restore things to a known good state (as we demonstrated earlier).

Having installed (at least) the Plasma desktop we can opt to change the display manager (which provides the login screen) from the Mint-themed LightDM to



the Qt-powered SDDM (Simple Desktop Display Manager). If you installed from the command line this will be offered to you, and if you didn't you can get to the configuration by running;



The lightest suite of KDE applications weighs in at around 850MB, but it is quite outstanding.



For some reason we found this SDDM theme much more relaxing than the austere default.

\$ sudo dpkg-reconfigure lightdm

Both display managers work with all major Linux desktop environments (and lightweight window managers), so which you choose is a matter of personal preference – or whichever you can find the prettiest themes for. The default SDDM theme is probably not most people's idea of pretty, but once you're logged in it's easy to change this from Administration > Login Screen (SDDM). For some reason, perhaps an attempt at irony, we found ourselves using a Windows 10-like login theme. Never mind that, you've probably just found yourself immersed in the wonderful world of KDE Plasma. Behold the cool Breeze theme, marvel at the polish and feel at home with the knowledge that all your favourite Mint tooling is just a click away.

KDE 4, now largely retired, received occasional criticism for being too configurable. In part this was fair. Every widget (and there were a lot of widgets) could be configured, a clumsy edit mode gave them a handle about which they could be rotated or stretched, and one was sometimes left wondering what the point of all this was. Worse, successive iterations of KDE 4 got very good at hiding all kinds of key options just when you thought you'd got a handle on where they ought to be. That version of KDE also faced criticism for being something of a resource hog, and shipping with all kinds of graphical frippery enabled. Modest machines would probably have been fine with this, but as the graphics driver ecosystem of the era was far more fragile back then, hardware acceleration was not something that one could count upon.

You'll be pleased to hear, then, that KDE 5 (or KDE Plasma 5 as the desktop prefers to be called) is a much sleeker animal. In our tests it did use up a tiny bit more memory than Cinnamon, and slightly more than Xfce and MATE, but what's a hundred or so megabytes

between friends? It makes not one iota of difference once vou start memorvslayer Chromium. Plasma is certainly configurable. but in a way that is not overbearing. Take the default, medium-weight launcher menu (at the bottom-left, as it should be). Right-click it and select Show Alternatives. You will see it can be swapped for a modern, full-screen launcher (sort of like Gnome) or a more classic cascading menu design.

KDE comes with its own graphical application store called Discover (one of few KDE apps not to capitalise on any opportunity for an

unnecessary letter K). You'll find this already pinned to the favourites menu, and you might also prefer it to Mint's native Software Manager. One thing you'll want to do is sort out Flatpak support in Discover. Fire up a terminal (try the Konsole application) and run:

\$ sudo apt install plasma-discover-backend-flatpak

You can now, after restarting Discover, browse FlatHub (or any other Flatpak repos) by adding them via the Settings option at the bottom right. Just click Show Contents to the right of the repo name. Flatpak is a much more decentralised idea than Snaps; anyone can set up their own Flatpak repository, but the only Snap outlet in town is Canonical's Snapcraft. Both forms are potentially risky though, since there's little to stop a scoundrel uploading a rogue Flatpak or Snap. And while both have some sandboxing capabilities we have no compunction to endorse the downloading of random binaries. Popular applications are easy to spot on FlatHub and common Snap packages have a reassuring 'Verified' badge.

We mentioned that Wayland isn't explicitly supported by any of the Mint desktops, but that is changing. In the latest MATE release a great deal of the desktop now works natively with Wayland, so if you switch the Marco window manager for Compton then you're well on your way to display protocol future. Xfce 4.18 plans to introduce support, though that may be a long way off. So it's really desktops, rather than distros, that enable Wayland – and as luck would have it KDE Plasma has support built in to its Kwin window manager. There's just a couple of packages to pull in to bring it to life: \$ sudo apt install plasma-workspace-wayland

Now if you log out, a new session called Plasma (Wayland) will be available from the menu. The Plasma experience on Wayland has come a long way this year. We're told it even works with the proprietary Nvidia driver now.

One thing that might strike you as jarring about Plasma in general is that your session is automatically saved. If you prefer to start each time without all those stray terminals and whatever else you left open, go to Settings > Startup and Shutdown > Desktop Session.



» LXOT

If you're enamoured with the Ot toolkit, but crave a lighter, nimbler desktop, you should look no further than LXOt. That's what powers the current LTS edition of Lubuntu, and it's what could power your new Livingstonseagull-like Mint desktop. Over the page we'll go pretty much as far as we can go without abandoning the GUI altogether. But if you want something a little more user-friendly and less gymnastic keyboard shortcut-orientated, LXQt may well be for you.

LXQt is the spiritual successor to the GTK2-powered LXDE desktop that used to power Lubuntu and Raspbian. Rather than move to GTK3, which at the time was seen as bloated, LXDE teamed up with the RazorOt effort and the result is LXOt. You can install it with a simple sudo apt install lxqt openbox. If you haven't already installed KDE this will pull in around 400MB of dependencies, but if you've already installed Ot et al, the footprint will be much lighter. Don't forget to take a snapshot first though.

When you start LXQt you'll be prompted to choose a window manager. By default it uses openbox, but it can use Cinnamon's Mutter, Xfce's Xfwm4 or, if you really want to make it pretty, Kwin from Plasma. Openbox is by far the lightest, and for non-scientific comparison purposes a plain LXQt/Openbox session occupied around 500MB of our memories.



LXQt might be just be the lightest desktop environment that still provides all the friendly GUI crutches we've come to rely on.

»

Ultralight Mint

Embrace minimality and learn some keyboard gymnastics with the featherweight Sway desktop.

ne of our new favourite Ubuntu-derivatives is Regolith Linux. It's fairly unique in its choice of the ultra-light i3 tiling window manager. Tiling window managers take some getting used to, and also a whole lot of configuration, but Regolith ships with remarkably sane defaults and easy to learn keyboard shortcuts (i3 is very much keyboarddriven, but converts say they never looked back). Also, it still has all of the GNOME infrastructure and applications for managing sessions and settings, so all of your system administration can be carried out with familiar GUI apps. It's well-known that we're big fans of Pop!_OS too, and in particular its COSMIC

WINDOW TILING IS AN ART "To make the most of window tiling, much like an Etch A Sketch you'll want to use a combination of horizontal and vertical arrangements."

> (Computer Operating System Main Interface Components) desktop. This features a tiling mode that, while not having the diminutive resource footprint of i3, offers users a gentle introduction to the joy and efficiency of keyboard shortcuts and mouse gestures in harmony.

> There's no reason we shouldn't have these sorts of things in Mint too: the i3 window manager is in the Ubuntu repositories. But we're going to try something else. Sway is a lightweight window manager inspired heavily by i3, except that it is for Wayland. If you are



The rxvt-unicode terminal isn't that pretty in its default state, but we're sure you can make it so.

familiar with i3 you will quickly get the hang of Sway; most of the default keys are the same, and you can even use

your own i3status scripts. In fact, you should be able to use your **i3config** file without modifications. Sway is in the Ubuntu repos, but it's an old version from January 2020. It would take some work of the compiling variety to get the latest version working, so let's just install the repo version to dip our toes in:

\$ sudo apt install sway

As before, the login screen should now have a Sway session. Dive into it and you should see the rather fetching Sway logo and top bar. Try anything with the mouse (besides moving the pointer) and you'll realise that you're not in Cinnamon any more – nothing reacts to being clicked, double-clicked, dragged or any such thing. Sway is all about keyboard commands: try pressing Super (the brand-independent name for the Windows key) and Enter. A terminal should spring into life, so now you can at least practise your *Bash* scripting for a while.

Now try pressing Super+2; the terminal will disappear, but not really – cast your eye to the top-left and you'll see we've just moved to a new virtual desktop. If you try to open another one with Super+3, you'll see that this doesn't happen. We didn't need a third desktop because we hadn't opened anything on our second, so Sway quietly renamed the previous workspace to 3.

Go back to workspace 1 and hit Super+Enter to open another terminal. Now you can see what tiling window managers are all about. The first terminal, that was occupying the whole desktop, obsequiously squishes over to the left, making room for a new terminal to the right. If you like, you can start any program you want from either of these terminals (notice the focus follows the mouse so you don't need to click in either one).

But Sway also has its own application launcher, after a fashion. Hit Super+D and you'll see some commands in lexicographical order (some beginning with numbers, and several beginning with the letter A, probably). Start typing the first few characters of **firefox** and you'll see this list get rapidly smaller. Press Enter when **firefox** is highlighted in blue to start it.

If you already had a couple of windows open, things by now will probably be getting a little cramped. You could close one of them (either with a Ctrl+D or traditionally with the Close button). But now is also a good time to introduce window resizing. With at least two windows open, hit Super+R to enter resize mode. You'll see this indicated in the top left of the status bar. Now you can use the H and L keys (like Vi) to make the active window (and remember you can change this by hovering the mouse over a new window) wider or narrower. That covers one dimension. Press Esc to exit resize mode and the hit Super+E while focused on the leftmost window. Bam! Vertical windows. Be careful of case sensitivity here, because Super+Shift+E (often abbreviated to Super+E) is the shortcut to exit Sway.

In order to make the most of window tiling, much like an Etch A Sketch you'll want to get the hang of using a combination of horizontal and vertical arrangements. This seems straightforward at first, but there are a few subtleties that are best experienced for yourself. Before, the Super+E shortcut operated on several windows at once and tiled them uniformly. They were grouped together along a common dimension. But if you hit Super+V on one of them, then hit Super+Enter to open another terminal (or open any other application for that matter), it will open in the other orientation. In this way you'll see that Super+E and Super+V act as orientation triggers, and you'll notice that the former highlights the lower edge, giving a hint as to how the next window will spawn.

Sway's default configuration file can be found at /etc/sway/config. but rather than edit this directly, copy it to ~/.config/sway/config to make personal changes. For example, the line:

output * bg /usr/share/...

sets the background on all displays. You can change this, or indeed set a custom background for each display, by modifying this. The output directive actually controls all sorts of things fundamentally related to the display. For example, if you're running Sway on a virtual machine, putting the final touches to the overdue cover feature, you'll probably need to add a line of the form:

output Virtual_1 resolution 1280x720

in order to make your screen grabs have the correct aspect ratio. You can get a list of display names and modes by running:

\$ swaymsg -t get_outputs

Sway and i3 are famed for their low memory footprints. But in order to keep these low, and maintain



a minimal desktop aesthetic, one has to use lighter applications too. This isn't quite the place to show how much stuff you can run from the terminal, but you should check out the *nnn* file manager, the *w3m* web browser, *mpv* the video player and the *ncmpcpp* frontend to the *mpd* music player. What we will demonstrate is how to swap the default Gnome terminal for something a little lighter.

something a little lighter. You might have noticed when we installed Sway that it pulled in a package called **Suckless Tools**. Suckless (https://suckless.org) is as much a state of mind as a software suite, and encompasses a range of ultra lowresource utilities and daemons. These include the dwm window manager, the *suckless* terminal and a few more. We covered them back in LXF254 and they're actually of limited use for us here because most of them don't cater to Wayland.

These get installed primarily for the sake of the *dmenu* program which (through Xwayland) provides the handy Super+D launcher we met earlier. Have fun tweaking! You'll need to use shutdown from the login screen with Sway; a nimble Super+Shift+E will get you there.

you probably shouldn't install the HWE stack unless something is broken, but it's a safer bet than being

seduced by a stock kernel with a bigger number. Desktop Ubuntu now gets the HWE kernels by default, and so it would seem does Pop!_OS, but not Mint. That's okay, because it's easy to install: \$ sudo apt install --install-recommends linuxgeneric-hwe-20.04

If you are using the Nvidia proprietary driver this is not enough, since you'll need corresponding proprietary modules too. These you can get with: \$ sudo apt install --install-recommends linuxmodules-nvidia-NNN-generic-hwe-20.04 where NNN is the version of Nvidia drivers required by your card. You can find this out by running ubuntu-drivers list. Be aware that lots of users have encountered problems with the 5.8 HWE stack (sound, graphics, virtualisation) and the same is likely to be true for this one.





» KERNEL UPDATES

Being a derivative of the Long Term Support release of Ubuntu, Mint by default uses the same 5.4 series kernel featured there. Don't be put off by the fact that the current branch of the kernel is numbered 5.13, since Canonical backports all manner of features and fixes to the Ubuntu kernel. Also the 5.4 kernel is itself a longterm branch (as you'll see from **kernel.org**). It has always been possible, but not recommended, to use mainline kernels in Ubuntu, but a better way is to activate the hardware enablement HWE stack.

This will give you a newer kernel (which recently bumped from 5.8 to 5.11) which has undergone some Canonical patching and testing, as well as a refreshed graphics stack (new versions of X.org and libdrm and what have you). We mentioned the lowvisibility Edge edition of Mint (nothing to do with the web browser) in the DVD pages, but it's definitely not worth reinstalling just to get a newer kernel. Likewise

www.techradar.com/pro/linux





Josh Lowe Creator of coding tool Edublocks.

» CODING FOR EVERYONE

We're fast approaching the tenth birthday of the Raspberry Pi and little did I know when the small single-board computer was released all those years ago, that it would have such a huge influence on my life.

From a young age I've always had an interest in technology. I would try and see how things worked. By learning how the Raspberry Pi worked it led me down the path of learning how to code. My journey with the Raspberry Pi started by attending the local Raspberry Jam in Preston where I met like-minded people interested in technology and the Raspberry Pi. The community is still one of the best things about the Raspberry Pi and through attending community events. I was able to learn new skills that would help spark my interest in coding.

After learning how to code myself by working with likeminded others at events, I decided that I wanted to create something that would help others to learn how to code. So I created *EduBlocks*, a free drag-and-drop code editor that now helps 75,000 people a month learn how to code using text-based programming languages.

As more and more jobs become tech-focused, there's never been a better time for young people to pick up coding skills and get hands-on with products like the Raspberry Pi. Attending your local Raspberry Jam, Code Club or Makerspace is a great way to do this. When you learn how to code, the possibilities are endless and it's opened up a whole new world of opportunities for me.

New scheme to recycle old Raspberry Pis

Putting a stop to the running joke of everyone having a drawer full of unused Pis.

called OKdo Renew.

helps recycle your

Get cool hard vouchers

for your unused

Raspberry Pi kit.

which not only

used Pis but

t a time when we're all supposed to be getting fully on board the idea of green lifestyles, electronic waste is no longer a joking matter, so that question of what to do with unused but perfectly functioning Raspberry Pis has come into sharp focus.

Say hello to **www.okdo.com** which has partnered with Sony – a plucky Japanese electronics company you may have heard of, that actually manufacturers the Raspberry Pi – to launch the first official Raspberry Pi — recycling initiative



CREDIT: www.okdo.com

Argon Pi NAS ONE, NEO, now EON.

Argon Forty, makers of fine Raspberry Pi casings such as the Argon ONE M.2 and the excellent Argon NEO, has announced a Kickstarter campaign for the Argon EON, a four-bay network storage array powered by your choice of Raspberry Pi 4. It'll be live by the time you read this: https://bit.ly/lxf281argon



CREDIT: Argon Forty

rewards you too! So how does the scheme work? Currently only three models of Raspberry Pi are eligible for the scheme: > Raspberry Pi 3 Model B

- > Raspberry Pi 3 Model B+
- > Raspberry Pi 4

If you have any of these lying around, safely package them up and return them to OKdo following the link below. The old Pis are sent back to the Sony Technology Centre in Wales where they were manufactured. When it gets there, it'll be tested, reconditioned, and repackaged, ready to be sold to its new home. OKdo will be offering the refurbished boards at a lower price and they all come with a fresh 12-month warranty.

Out of all of this you'll get a £10 OKdo voucher you can spend on new shiny kit! To find out more, get complete guidelines on returning hardware see:

www.okdo.com/raspberry-pi-renew.

Pi at the movies

'Marvel'ous use of the Pi...

Wonder how cool special effects come together in films? With the help of the Pi of course! A tweet by **@Arturo182** looked at creating the visual effects for Marvel's *Black Widow*. Actress Olga Kurylenko is decked out in combat gear while surrounded by Pis all controlling Nikon DSLR cameras. This incredible rig is constructed with over 200 individual Pis. www.thescantruck.com



Can you spot the Pis? CREDIT: Marvel, The Scan Truck

SparkFun qwiic pHat

Les Pounder wants another slice of Raspberry Pi and he needs it real 'qwiic'! Complaints about bad gags to the usual address...

IN BRIEF

A low-cost and simple option for those wishing to connect I2C electronics to their Raspberry Pi. Compatible with qwiic and Stemma QT components. which are slightly more expensive than their more common counterparts. The 'qwiic' aspects of the board are the four qwiic connections that enable compatible components to be attached via a single four-pin connector. Connections can be chained together to build more complex projects.

lectronics with the Raspberry Pi's GPIO are amazing, but at first it can seem exceptionally daunting. The many wires and connections can scare many away, but once you get the hang of it the workflow becomes second nature. There are ways to simplify the wiring and one of those is the I2C protocol which uses only four wires to form a bus, with devices connected at addresses.

SparkFun's qwiic pHat, designed for the Raspberry Pi Zero form factor but compatible with all 40-pin GPIO, provides four "qwiic" connectors which we can connect compatible devices to. Two of the connectors are mounted to the side of the board; the others are mounted in the centre pointing upwards. The board also has a 5V power terminal breakout for powering 5V components, and a push button connected to GPIOI7. Around the perimeter there are mounting holes for M2.5 mounts and to secure qwiic-compatible components.

There is no specialist SparkFun software to install as the qwiic pHat is really just an I2C breakout board. All we need to do is enable the I2C interface in the Raspberry Pi Configuration tool and we are ready to read devices. We dug around in our box of bits and found an Adafruit MPR121, which we used in a previous CircuitPython tutorial. Adafruit's MPR121 has a Stemma QT connection, and right now you're thinking that we've made a mistake. Not so, as Stemma QT and qwiic are one and the same interface. They both breakout I2C and are pin-compatible; all that is different is the brand name.

Stemma QT and qwiic components can also be chained together, as many of the components have connectors at two sides of the boards. As I2C uses a bus and address, as long as we don't have clashing addresses we can chain multiple devices to one interface.

A quick CircuitPython install and our MPR121 was easily read by our Python code and detected touch inputs with ease. Qwiic/Stemma QT components are a little more expensive than their standard counterparts, but what extra we spend in money we gain in time and ease of use. This makes them ideal both for the classroom and the makerspace.





Designed for the Pi Zero, qwiic pHat works with all 40-pin GPI0 models of Raspberry Pi.

So experienced electronics experts will now be thinking that the qwiic pHAT is useful, but they already know I2C so this board isn't for them. Well, yes and no. If you are comfortable with I2C and can work with the pullup resistors and wiring necessary. go for it. But if you need a little help, or want to neaten up a project, then qwiic pHat is both useful and much safer than some of the wiring in projects we have come across. We can see qwiic pHat being used in education, embedded projects and art installations for makers who want to get projects built and don't have the time to mess around with I2C.

The biggest issue with the qwiic pHat is that we do not have access to the GPIO. This is more by design than anything else. SparkFun's qwiic pHat is designed to make connecting devices simple, so why add the GPIO? A longer GPIO header would've been a nice addition, but if you are buying this, you will most likely understand and accept that decision. If you really need GPIO access, right now the best option is SparkFun's £6 qwiic pHat for Raspberry Pi 400, which has two qwiic connectors and a full 40-pin GPIO. If you need just a single qwiic connector, the £1.20 qwiic SHIM fits directly to the GPIO using friction and provides a 'qwiic; connection for smaller projects.

SparkFun's qwiic pHat is a cost-effective and simple board that does one job exceptionally well.

DEVELOPER S	SnarkEun		
WFB: www.spa	rkfun com	/products/1594	5
DDICE: \$6	intuit.com	/ products/ 1004	0
PRICE: 10			
	0.00		0./10
FEATURES	8/10	EASE OF USE	9/10

» Rating 8/10

SCRATCH Drawing shapes with maths and code

Les Pounder shows how we can use simple maths to draw patterns and shapes, just like a classic childhood toy.

> s a child I had Spirograph, and it was fascinating how coloured pencils and a rotating cog could create such wonderful patterns.



Les Pounder is Associate Editor at Tom's Hardware and a freelance maker for hire. He blogs about his adventures and projects at bigl.es



Scratch 3 should come pre-installed on your Raspberry Pi OS image, but just in case it is missing it can be installed from the main menu, under Preferences > Recommended Software. Scratch 3 is found in the Programming category; tick the box and click Apply to install. Once installed, Scratch 3 can be found in the main menu under Programming.

patterns in Scratch 3.

Open Scratch 3. On first start Scratch may take a little while to open. We're going to assume that you have an understanding of how to code with Scratch, but if not we covered the basics in previous issues. Before we write any new code we need to click the blue folder icon at the bottom left of the screen to load the Extensions menu. From there select Pen and a palette of new blocks is added to our code.

Draw a square

Next we need to go to My Blocks and click Make a Block. With this we can create our own Scratch blocks which will run a series of steps when the block is used. It works in a similar manner to Python's functions. Call the block 'Square' and click OK. A new red block define Square will appear in the coding area, and we will now write the code that will draw a square on the stage. Go to Pen and connect pen down to the define Square block. From Control, drag a repeat 10 and connect it to the previous block. From Operators drag a _/__ block and drop it on top of the 10 of repeat 10.

A square has four sides, all of equal length and each corner has a 90 degree turn. If we use a little maths we can determine that we need to divide 360 by 90 degrees, giving us the four times that the loop will iterate. Add the equation 360 / 90 to the blank spaces. Inside the loop, drag change pen color by 10 to change the pen colour each time the loop iterates, creating a rainbow effect. Then drag set pen size to 1 and connect it to the previous block. From Operators drag pick random 1 to 10 and drop it on top of the 1 in set pen



You too can create a psychedelic masterpiece with just a few blocks of Scratch code

size to 1. This will randomly change the pen thickness as the loop iterates. From Motion drag move 10 steps and place it under the previous block, change the 10 to 50 so that we have a large square. Next drag turn 15 degrees from Motion and place it under the previous block. From Operators drag a _/ _ block and drop it on top of the turn 15 degrees block. This time we need to turn the 360 / 4 degrees, 90 degrees. Type 360 / 4 into the operator block.

Create a circle

Create another new block via My Blocks and this time call your block Circle and click OK to create. A new red block define Circle will appear in the coding area, and our task now is to create a sequence of code that will draw a circle when we use the Circle block. From the Pen palette drag pen down and attach it to the Circle block. This will place the pen on the screen and enable us to draw. From Control drag a repeat 10 and connect it to the previous block. From Operators drag a _/_ block and drop it on top of the 10 of repeat 10.

To draw a circle we need to do a little maths. A circle has 360 degrees, and to draw one circle we need to make a loop that iterates 36 times, each time rotating 10 degrees. To get 36 loops we need to divide 360 by 10 and the output of the equation will drive our loop to iterate. Inside the loop, drag change pen color by 10 to

YOU NEED

> A PC or a Raspberry Pi 3/4/400 > The code for this project can be run online at https://bit. lv/LXF281 scratch.

Drawing shapes **TUTORIALS**

change the pen colour each time the loop iterates, creating a rainbow effect. From Motion drag move 10 steps and place it under the previous block. Next drag turn 15 degrees from Motion and place it under the previous block. From Operators drag a _/__ block and drop it on top of the turn 15 degrees block. This time we need to turn the 360 / 36 degrees, 10 degrees. Type 360/36 into the operator block.

Draw a triangle

This is the trickiest of them all. In maths we are told that the internal angles of a triangle add up to 180 degrees, and that is still correct. But for Scratch we need to use external angles to draw the shape.

Create another new block via My Blocks and this time call your block **Triangle** and click OK to create. A new red block **define Triangle** will appear in the coding area. From the Pen palette drag **pen down** and attach it to the Triangle block. From Control drag a **repeat 10** and connect it to the previous block. From Operators drag a **/** block and drop it on top of the **10** of **repeat 10**.

The maths this time is to work out the external angle. A triangle has three sides, so if we divide 360 / 3 we get 120 degrees. But remember this loop is the number of times it iterates to draw the sides, so our equation will be 360 / 120.

Inside the loop, drag change pen color by 10 to change the pen colour each time the loop iterates, creating a rainbow effect. From Motion drag move 10 steps and place it under the previous block, change the 10 to a 50 for a larger triangle. Next drag turn 15 degrees from Motion and place it under the previous block. From Operators drag a _/__ block and drop it on top of the turn 15 degrees block. This time we need to turn the 360 / 3 degrees, 120 degrees. Type 360 / 3 into the operator block.

The sequence will use each of the shape blocks in a forever loop. From Events drag When Green Flag



The red Triangle block is a function that we can use to run all of the code within it.



The code to create a forever sequence is rather simple, but there is quite a lot of it so make sure you visit our Scratch project to see it all.

Clicked and place it in the coding area. From Pen connect erase all to clear the screen. From Control drag a forever loop and connect to the previous block. Inside the loop, from Pen drag set pen size to 1 and then from Looks drag hide to hide the cat sprite. From Pen drag pen up and then from Motion drag glide 1 secs to random position, change the 1 to 0.1. Now our invisible sprite will move around the stage without leaving any marks.

From Control drag repeat 10 and change 10 to 24. Inside this new loop, from motion drag turn 15 degrees and move 10 steps to move the sprite 24 times in a slow rotation. From My Blocks drag Triangle and connect it to the previous.

Lastly outside the repeat 24 loop, but inside the main loop, drag pen up from Pen, and another glide 0.1 secs to random position. Click the green flag and watch as a spiral of triangles appears on the screen.

Your next challenge is to try and have your Circle and Square blocks repeat the same sequence.

» TURTLE POWER

In this tutorial we drew patterns with code, but this is nothing new. Logo, a programming language from 1967, introduced the concept of 'turtle graphics', where a small cursor (the imaginary turtle) can be controlled using code. We can draw on the screen using simple commands, or use mathematical expressions to create elaborate art.

Turtle also featured a robot that connected to the computer and was programmed using the same commands. Place a pen in the turtle and you could draw the patterns for real. Logo and turtle have been seen in many systems over the years. Python has a built-in version of Turtle, a module that we can import and use in a similar manner. You can read more about Python turtle at https://docs. python.org/3/library/turtle.html.

Turtle is not just limited to the computer. Computer Science Unplugged has an offline version called Kidbots, where children have to write their own code to 'program' one of their friends. It teaches the basics of computer logic, geometry and how to write simple code. On a sunny day why not have a go in the classroom, or in your garden. See https://csunplugged.org/en/topics/kidbots.

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

OUICK TIP

Using maths

we can draw

any shape in Scratch, but

when we get to

20-sided shapes

things get a little

'interesting' and

we often run

out of screen.

So make sure

vour steps are

smaller than 50,

10 is about right.

PICAMERA Pinhole photography with the HQ Camera

Interchangeable lenses enable us to experiment with that historical curiosity, the pinhole camera, **Mike Bedford** shows you how.



Mike Bedford likes to engage in unusual forms of photography, but he has to admit that the pinhole approach is one of the strangest, but no less interesting for that. he pinhole camera might date back to 1856, but it would be wrong to think it predated cameras with lenses. Instead, it seems that the idea of a lens-less camera was always intended as a curiosity, not a serious suggestion for a practical camera. It has been used as such by enthusiasts, but in the main pinhole cameras have been employed as teaching aids. Irrespective of whether you want to learn more about optics and photography by experiment, or take some shots with a rather unique look, if you have a healthy curiosity we're confident that you'll enjoy trying our various hands-on exercises.

A pinhole camera can be nothing more complicated than a cardboard box, as we demonstrate with our first experiment. While it might seem incongruous to mix old and new technology in this way, we then bring the concept up to date by showing you how to use a Raspberry Pi HQ Camera as a pinhole camera. This isn't just a desktop exercise though – we also provide some guidance on how to make that pinhole camera portable so you can use in the great outdoors.

Because pinhole photos aren't always perfect as they come out of the camera, we'll offer some brief thoughts on photo processing. Finally, to round off our look at this primitive form of photography, we also take a look at pinhole photography with a DSLR.

Illuminating in a cardboard box

Right at the start we have to come clean and admit that you're probably not going to take a photo with our first pinhole camera, even though it's quite feasible to do so.

Here we're putting the final touches to the incredibly simple cardboard pinole camera, and real photos can be taken on a barely more sophisticated variant





The first photo on our H0 Camera pinhole camera, before we made it portable, was shot out of the window. The pinhole was imperfect, the sizing was largely guesswork, and dust specs are clearly visible, but it's quite an eye opener given the lack of a lens.

Instead, the real purpose of this experiment is to demonstrate how simple a pinhole camera can be – after all, this one is just a cardboard box – and to give you a better feel for what's happening when we get embroiled in high-tech pinhole cameras.

First of all you need to find a cardboard box. The size isn't critical, but we suggest going not much larger than the one in the photo, bottom left. Paint it matt black, inside and out, and when the paint has dried, glue a piece of white paper to the inside of one of the smaller faces. Now, using a sharp object like a small bradawl, make a small hole about 1-2mm in diameter through the face opposite the one with the white paper. In picking the faces to which you'll stick the white paper and make the pinhole, make sure that one of the other two faces can be opened. Congratulations, you've made your first pinhole camera.

To see it in action, take it into a dark room and find something that emits its own light as the photographic subject, say a candle or a bare lightbulb. Place your camera on a flat surface such as the floor or a table, and place that object close to the box facing the pinhole. Now turn off the lights and take a look at the white paper through the opening in the box, and you should see an upside-down image of the light-emitting object. If you don't see the image, try adjusting the angle of the box to better frame the shot and/or try moving it closer

Pinhole photography TUTORIALS 👹

to or further from the object.

Now, to take a photo, having set up the camera and the object, just remove the piece of tape from the pinhole for the exposure period. Of course, you'd need to prepare your camera in total darkness, or in red light if you're using photographic paper as opposed to film. You'd need to develop it afterwards, and you'd need to experiment with exposure times. We suspect few of you will choose to try that, so let's move on to some more practical ways of taking real pinhole photos.

A Pi-hole camera

First we're going to see how to build an experimental pinhole camera using a RPi HQ Camera, and then we'll move on to consider how to make it more practical for portable use. For the first step, until you're sure everything's working okay we suggest that you use a Raspberry Pi attached to the HQ Camera on a tripod facing outside through a window, and with the usual keyboard, mouse, monitor and power supply attached.

However, we're not going to be using the camera with one of its normal lenses attached. Instead, we'll be using a pinhole, so let's see how to make one. Take the body cap that will have been fitted to the HQ Camera when you first took it out of its box, and drill a hole about 10mm in diameter through the middle. Now take a piece of aluminium cooking foil and make a small pinhole in it, using a sewing needle – see the box below for information on pinhole diameter.

Alternatively, you might do better using thicker foil, such as that used for some takeaway or freezer-food containers. The advantage of the thicker foil is that you can sand it down with very fine sandpaper to remove any burr. Cut out the pinhole so that the piece of foil



completely covers the hole in the body cap and secure it in place using tape, ensuring that no light can leak through the cap except through the pinhole. Finally, screw the body cap and its pinhole onto the HQ Camera, using the C-CS adapter that you'd use with the longer of the two official lenses. This will provide you with a focal length of about 16mm – see the box on page 50 for information on focal lengths.

Now try taking your first photo using the *raspistill* command line utility. The chances are that the result won't be the sharpest photo imaginable – and that's an understatement – so you'll probably have to try a few pinholes until you get one you're happy with. A pinhole will never compete with a lens, but the imperfect look is part of the appeal of pinhole photography to its adherents, and it is quite amazing that you can get an image at all with no lens. For best performance, though, you might want to abandon the idea of making your own hole in a piece of foil and invest in a professionally made pinhole. These are available in a range of sizes for £17

Details will differ depending on what size box you get, but our layout should give you some ideas about how to build a portable H0 Camera-based pinhole device.

QUICK TIP

If you only use your HQ Camera occasionally, unless you keep a lens attached to it, you really ought to keep its body cap attached, to prevent ingress of dust. A body cap converted into a pinhole 'lens' might not be too effective at protecting your camera, though, so we suggest getting hold of a new one

» PINHOLE DIAMETERS

The operation of a pinhole camera is often explained using a diagram like the left-hand one below, with light rays passing through a perfect pinhole. However, an ideal pinhole has zero diameter which would let no light through, so what diameter should a practical pinhole be?

A larger diameter pinhole admits more light, which is good, but as the middle diagram shows, this allows several light rays to pass from points in the object to the sensor, which will reduce sharpness. But if the pinhole is too small, light rays spread out due to edge diffraction (righthand image) and, again, sharpness suffers. Clearly there's an optimal size. Equations differ, but try the calculator at https://tomroelandts.com/tools/ pinhole-size-calculator with the default 550nm wavelength. For a focal length of 16mm, as in our RPi HQ Camera pinhole device, this works out at 0.178mm, so we used 0.2mm. We have to admit that making your own accurately sized pinhole is difficult, which is why we recommended going to pinholesolutions.co.uk. The aperture (f-number) is an important characteristic of a lens. This is the diameter of the lens as a fraction of its focal length and values of between f/2 and f/22 are common, the larger the number representing a smaller aperture. As well as affecting the exposure, the aperture affects the depth of field, with a smaller aperture giving a larger depth of field. The HQ Camera's 16mm focal length, 0.2mm diameter pinhole has an unusually small f/80 aperture, which explains why dust specs on the sensor are focussed enough to be visible.



This would be a nightmare even for Goldilocks, finding a pinhole that's just right.

>> MORE NOVELTY-VALUE THINGS Subscribe now at http://bit.ly/LinuxFormat



QUICK TIP

You might just have a pinhole camera in your house already -I do. The 'camera' is an entrance lobby which, when both doors are closed, only permits light to enter through the keyhole. It's significantly larger than a pinhole but. when the light is right, an inverted image appears on the opposite wall.

from **pinholesolutions.co.uk**. Another issue you'll invariably notice with your first photos is lots of blurred dots which are actually specks of dust on the camera's sensor. Because of a pinhole's extremely large depth of field, those specs are almost in focus while, with an ordinary lens, they'd be so out-of-focus as to be invisible. To reduce this, you might want to carefully brush the sensor (actually, the filter that covers the sensor) with a small, brand new, perfectly clean brush with very soft bristles.

Going portable

Having got a pinhole that you're happy with, it's time to turn the HQ Camera into a portable camera so you're not limited to shooting through your window.

First of all, because you're not going to want to have to connect a keyboard and mouse to your portable camera, you'll need some software that starts when the Raspberry Pi boots and is controlled by one or more pushbuttons. We used just the one pushbutton, and also a piezo electric sounder for audible feedback connected to GPIO pins. The main reason for omitting the LCD screen is that pinholes don't allow much light through, so if you were to use an LCD panel as a viewfinder you'd probably struggle to see the image in even moderately bright daylight.

We've created Python code (here https://textbin. net/lfylfhigwz) for this project available from www. linuxformat.com/archives. We suggest that you make sure the software is working correctly while you still have a keyboard, mouse and monitor attached, before building everything into a box.

Because the code starts when you boot the Raspberry Pi – there are lots of ways to do that, but we'll leave you to read up on it – the code starts by sounding three short bleeps to indicate that it's booted and is waiting for you to take a photo. Pressing the pushbutton momentarily causes a photo to be taken, something we do by calling *raspistill*. As soon as you press the button the sounder bleeps to confirm that, and when the photo has been taken, it sounds again.

» FOCAL LENGTH

The design of a lens defines its focal length, but nothing about a pinhole has such an effect. Instead, the focal length is defined simply as the distance between the pinhole and the image sensor or film. With the RPi HQ Camera plus the C-CS adapter, that distance is about 16mm, which means it'll give a slight telephoto effect.

If you were to omit the C-CS adapter the focal length would be reduced but, given the fact that the problem with specs of dust will be even greater, we don't recommend it. If you decide to try out a DSLR as a pinhole camera, the distance between the body cap and the sensor will differ between cameras but, generally speaking, it'll behave like a so-called standard lens (35mm equivalent focal length of 50mm), that is one which provides a similar field of view as our eves.

If you find a way of separating your pinhole further from the camera's sensor, the focal length will increase so you'll get a telephoto effect. We haven't tried it out, but we have seen reports of good results being achieved by gluing a carboard tube (painted black on the inside) to a body cap and attaching a pinhole to the other end of the tube. Bear in mind that the pinhole size will be different if you chose a longer focal length.



Our camera might not look too impressive – after all, its lens is made of nothing (well, air, actually) – but its simplicity is surely part of the appeal of pinhole photography.

With the default *raspistill* settings there'll be quite a gap between the two bleeps because it shows a preview for five seconds, and you might choose to reduce or eliminate that if you don't add an LCD panel as a viewfinder. The filename of the image is assembled using the date and time and, although the date and time won't be correct if the Raspberry Pi isn't connected to the internet, it'll ensure that filenames are unique.

Because it's not advisable to turn off a Raspberry Pi by just removing the power cable, the code also allows it to be shut down. This is done by pressing the button while the second bleep you hear while taking a photo is sounding, and is confirmed by three short bleeps, indicating that it'll be safe to remove the power after another few seconds.

We're not going to give a blow-by-blow account of how to build your camera as a portable unit (but see the photo), although we'll give a bit of guidance. We recommend using a smallish plastic box of the type sold by electronic component suppliers, into which you screw the Raspberry Pi and the HQ Camera.

You'll need a large hole through which you can screw the pinhole (or a lens if you want) onto the camera. You'll also need a hole for the camera's tripod socket because, with pinhole exposure times being so long, you're going to need to use the camera on a tripod. We recommend using a USB power bank as the power supply, mounting it on the outside of the plastic box, so you don't need to open it up to charge it.

This being the case, you'll need a hole adjacent to the Raspberry Pi's power socket through which you can pass a USB to Micro-USB lead. The push-button could be mounted on the box, but we put it on a flying lead rather like a camera cable release, so you'll need a hole through which you can thread that lead.

If you do use a flying lead, be sure to provide some

Pinhole photography TUTORIALS 🍑

form of strain relief so you don't risk pulling the lead off the Raspberry Pi's GPIO pins. The other thing you'll need to give some thought to is how to get the photos out of the camera without plugging in a monitor, keyboard, mouse and USB drive, which would require you either to use a very much oversized box or to make lot of extra holes in the box.

There are lots of options, but we took the simple approach of cutting a slot in the box adjacent to the microSD card so we could remove and replace it using a pair of tweezers. This then allows you to copy the files onto any Linux PC with a suitable microSD card slot. If you use this method, do make sure the slot is oversized, otherwise you might snap the microSD card – and we speak from experience.

Finally we come to the viewfinder. The simple solution we used is two 'windows', one behind the other, that you look through. The windows are the same aspect ratio as the sensor, and the one further from your eye is larger than the closer one. If the sizes and their separation are correct – you'll need to draw it out on paper or with a vector graphics package to determine this – when you move your eye closer to or further from the nearest window, it'll coincide with the more distant window when you're seeing the same angle of view as the pinhole.

For the PRi HQ Camera 16mm pinhole, that angle is about 22 degrees. The only snag with this approach is that you get a parallax error if you're photographing close-up, which means that you'll see a different scene to the camera.

Image processing

Photo processing, using a photo manipulation package such as *GIMP*, is highly recommended even if you're using a lens on your camera. With just a pinhole, though, it can make even more of a difference because, putting it politely, pinhole photos need all the help they can get.

For a start you might find that your photos are underexposed. We assume it's some sort of issue with the camera's metering with unusually low light levels, and while it would be better to correct that when taking the photo, you won't always get it right. Next up, we discovered that pinhole photos benefit more from using *GIMP*'s Levels and Curves tools than regular shots. Also,



This was one of our best photos taken on the RPi HQ Camera pinhole camera, using a pinhole from pinholesolutions.co.uk. We found that scenes with bold shapes worked better than pure landscapes.



If you can't stand all those dust blobs on your photos with the HO Camera pinhole solution, fitting a DSLR with a pinhole will provide significantly better results, as this shot of an antique door knocker on a panelled door shows.

given the somewhat inaccurate nature of the viewfinder, you might need to crop your results too. And finally, you might choose to use the Clone facility to reduce the impact of all those out-of-focus dust specs.

The DSLR option

Realistically, you'll struggle to totally prevent those blobs on your RPi pinhole photos, which result from the pinhole's extremely large depth of field which causes specs of dust on the sensor to become visible. Depth of field doesn't depend only on the aperture though, it also depends on the focal length – the real focal length, that is, not the 35mm equivalent.

In particular, the smaller the focal length the greater the depth of field, and our pinhole has a focal length of just 16mm (see the box on the opposite page). And it gets worse. Because the RPi HQ Camera has a small 1/2.3-inch sensor, many times smaller than that in a DSLR, any specs of dust on it will appear much larger as a percentage of the sensor size, and hence also as a percentage of the image size.

Bearing in mind that a DSLR will produce better images, even with a lens, if you have a DSLR you might like to try it for your pinhole photography. The procedure is much the same as with the RPi HQ Camera although, as discussed in our look at pinhole size, you'll need a larger pinhole for optimal results. If you can't find your camera's body cap – and we're sure that many people don't keep them because they expect there'll always be a lens on their DSLR – you really ought to buy a new one rather than trying a makeshift solution such as taping the pinhole directly onto the camera body.

Other than this brief bit of advice on DSLR pinhole photography – surely the ultimate bizarre combination of low- and high-tech – we suggest that you just aim to get lots of experience to see what works and what doesn't as you hone your skills.

If you're still not convinced and you want to see what's possible, we suggest you take a look at the incredible pinhole photography of Cameron Gillie, albeit using a large-format film camera, these are impressive results to see at https://thepinholething. com/pinhole-photography.

QUICK TIP

If a pinhole camera doesn't record a photo, but just projects an image like our cardboard box camera, it's really a camera obscura, not a camera. They can employ either a pinhole or a lens. They're often operated as tourist attractions and are quite impressive look out for one

Building a better file server with the Pi

Christian Cawley shows you how to get the most out of your Raspberry Pi-powered Nextcloud server.



Christian Cawley Like Little Jack Horner, Christian Cawley has been playing with Raspberry Pi for years, even sticking in his thumb (drive) and pulling out a plum, and completely losing the allegory. winning your own cloud storage server saves money, allows you to expand storage as necessary, and can be done with a device as small as a Raspberry Pi. Our previous guide to setting up a Nextcloud server on the Raspberry Pi (LXF280) covered everything you need to know about installing and setting up NextCloudPi and the Nextcloud Ubuntu Appliance, as well as the manual installation of Nextcloud on Raspberry Pi OS.

But with the software installed and connected to your network, the job is only half-done. You still need to configure it for your specific use-case, enable external storage to reduce wear and tear on the Pi's SD card, and make the Nextcloud server accessible from beyond your home network.

The most straightforward way to set up external access to Nextcloud and to add external storage is to do it with the NextCloudPi installation.

External storage

The Raspberry Pi 3 and 4 boot from the SD card. While this can be changed in favour of a more reliable USB device, SD is the default option.

This has obvious advantages, from speed and convenience to simplicity. However, there is a key disadvantage, which is the longevity of the card: regular reading and writing of data to the SD card will shorten its lifespan.

Throw into the mix the fact that your NextCloudPi storage installation is limited to the size of the microSD card in your Raspberry Pi and you have an obvious problem. It's just not big enough to compete with Dropbox or OneDrive or Google Drive, and you're limited to the maximum supported capacity of 1TB.

Adding an external disk drive to Nextcloud requires only a device with USB connectivity (unless you're using a case with a SATA adapter) . For faster data transfer, use USB 3.0+ devices, and for the optimum, rely on SSD or NVME/M.2 flash.

First, connect the USB device and then set up an SSH terminal to your Nextcloud and enter sudo mkdir /media/USBstorage

In your computer's web browser visit the NextCloudPi web panel (using the 4443 port, so https://[IP ADDRESS]:4443) and start the wizard if it



Connect a USB drive or full-size external hard disk drive to your Raspberry Pi for storage with Nextcloud - USB 3.0 preferred!

doesn't launch automatically by clicking the wand icon in the toolbar. Select USB Configuration, input Yes, then Continue, and Format USB.

When this completes, select 'Move data to USB'. You now have a Raspberry Pi-powered Nextcloud server that can store as much data as you need without wearing out the microSD card.

Static access

Configuring Nextcloud to be accessed from devices beyond your network (such as your mobile device while shopping) is possible with a static IP address. But before setting this up for external access, it's smart to configure the Raspberry Pi with a static IP address for access within your network – and, more crucially, your router – first.

To confirm the Pi's IP address and how it was assigned, open a terminal or SSH connection and enter ipr|grep default

If you see the letters dhcp the IP address is assigned dynamically by the router. Make a note of the IP address, then check the nameserver or DNS IP: sudo nano /etc/resolv.conf

Make a note of the IP address listed here for the nameserver before proceeding. Setting a static IP address can be achieved by editing the **dhcpcd.conf** file in *nano*:

Improve Nextcloud server TUTORIALS 👹

sudo nano /etc/dhcpcd.conf

Add the following lines to the end of the file, changing the capitalised text with the information you have to hand:

interface NETWORK

static ip_address=STATIC_IP/24

static routers=ROUTER_IP

static domain_name_servers=DNS_IP

Change **NETWORK** to eth0 or wlan0, depending on the network type (Ethernet or Wi-Fi).

For **STATIC_IP** enter the IP address you wish to use for the Pi – it's smart to use the one already assigned to maintain connectivity.

For **ROUTER_IP**, use the gateway IP for the router on your network – this address is usually printed on the router.

Finally, set **DNS_IP** to the IP address you made a note of after checking the **resolv.conf** file. Finish by rebooting with

sudo reboot

Confirm the changes with Ctrl+X, then Y and Enter, and reboot. You can check the static IP address has been applied using

ip r | grep default

If you can't get a static IP address from your ISP – and many of them don't provide one (see boxout) – your best option for accessing the Nextcloud from anywhere is with a dynamic DNS provider.

Access anywhere

Start by accessing your NextCloudPi web panel and enter the credentials you used during setup.

In the web panel, find Networking and select 'letsencrypt' from the menu. Add a working email address when prompted and click the wand icon to launch the setup wizard and select 'External access'. Answer Yes, then wait for setup to complete. When prompted for Port forwarding, choose 'I will do it manually'.

You'll need to forward ports 443 and 80 (for HTTPS and HTTP) to the Raspberry Pi, targeting the routerassigned IP address for the device or its hostname (nextcloudpi). How you do this will differ from router to router. Many routers offer an online help tool, but you can also refer to the device's documentation to find out how to use port forwarding.

In the next screen, sign up to FreeDNS when prompted, select Subdomains and choose a free domain and subdomain, then input your router's public IP address (check by searching Google for 'what is my IP') and Save.

Next, choose Dynamic DNS and copy the Direct URL for your subdomain into a text file. Remove everything before (and including) the ? so you're left with a long alphanumeric string. On the FreeDNS screen in NextCloudPi, complete the Domain field with the subdomain you chose, then paste the alphanumeric hash from the text file into the Update Hash field and click Finish to complete.

With your Nextcloud installation now fully online it should be accessible from the mobile app or any other device through your browser.



Set up cloud storage that you can access across the internet by configuring port forwarding via static IP or dynamic DNS.

Note that by establishing a static IP you may need to create a new connection in whichever version of the Nextcloud app you're using.

Enhance with apps

With the Nextcloud server accessible from anywhere, there are a few things you can do to take it to the next level. Almost all those things can be found in the apps list. These were mentioned in the previous tutorial but there are several that can really enhance your Raspberry Pi cloud server if you're planning on using it for a home file and media server.

In the Featured Apps list, you'll find a list of snaps that are listed by type. In the Multimedia tab you'll find apps for playing media. There is also an app called *Automated Media Conversion* which converts media files using FFmpeg based on rules that you set. It can also convert files as they're uploaded.

Antivirus and anti-ransomware apps can also be installed to protect your cloud data, along with integrations with mainstream cloud solutions. This is useful for easily migrating your cloud storage across rather than downloading to a PC and then uploading to the Nextcloud server.

» CAN YOU GET A STATIC IP?

Internally, a static IP is relatively easy to set up: simply instruct the device what its IP address is (see main article) and go from there.

Beyond your network, it's a lot trickier. Dynamic DNS tools are useful and have largely replaced the need for a static IP from your ISP – but what if that's the route you'd prefer to take? Dynamic DNS solutions can fail, after all, making a static IP address far more reliable. Relying on your router's external IP address is a risk, as you don't know how often your ISP will refresh it.

In the past, a static IP address could be expensive. These days, however, having a static IP address is more affordable. Some domestic ISPs will let you have a static IP address for a monthly fee, others with a one-off setup free. Others still will require you to switch to a business account, which comes with some shortcomings you might not be comfortable with.

The best way to find out whether you can get a static IP address and learn about any associated costs is to chat with your ISP. You might be pleasantly surprised at how little it can cost!

>> GET US TO SERVE YOU BETTER Subscribe now at http://bit.ly/LinuxFormat

Multiple hard disk drives or SSDs can be connected to your Raspberry

OUICK TIP

Pi Nextcloud and they can run as individual devices or in a RAID setup of your preference.

Erem bis caravan in a field in North Wales, David Butland takes you

CLASS

From his caravan in a field in North Wales, **David Rutland** takes you on a touring holiday through the world of virtual private servers.

he future is in the cloud, they say. But whose cloud? Google? Microsoft? AWS? Why not your own? There was once a time when people ran their own servers from home – under the stairs, in the back bedroom or from a lovingly kitted out homelab in the stripped-down carcass of the garage, meaning the vintage BMW restoration project is left to moulder on the drive.

A server is a simple beast at heart. It's a computer (usually running Linux), with an internet connection, which receives requests for resources such as pages and serves them back to the machine that made the original request. Server. It's in the name, innit.

But servers serve far more than web pages these days. There are progressive web apps, content management systems, databases and a whole host (geddit?) of other toys you may want to play with.

Some of these demand serious investment in terms of time, hardware, power

consumption, and if you or your partner are light sleepers, noise becomes an issue, too. In fact a properly specced home server with adequate cooling often sounds like a jet taking off – so it's time to take your server virtually off the premises.

What is a VPS?

If you're determined to run your server away from home in order to avoid unnecessary concerns about power bills, antisocial noise. and so on, there are a few options available: Rent out a dedicated private server in someone else's data centre. This option gives you an actual physical machine, equipped to your exacting specifications, and on which you can run whatever software you want. It's accessible only by you (making it private). Opting for a dedicated private server with an eight core Xeon and 64GB of RAM may get you the best possible performance, but it's far beyond the needs of most home users, and the financial cost will run to thousands of pounds per year.

Part One! Don't miss next issue, subscribe on page 18! Virtual Private Server. As with the dedicated server option, it's yours exclusively for as long as you choose to rent it, but the key thing here is that it's virtual. Your virtual machine will be one of several being run on one physical machine. You will be given a dedicated IP address through which you can log in and install your own software, and many providers offer some degree of pre-configuration through pre-built images. For example, you can opt for a machine which will have Linux Apache, PHP and MariaDB all set up and ready to go.

The cheapest option is to buy space on a non-VPS web host which is already running Linux. As an example we bought a standard, low-cost hosting package from **Godady.com** at a cost of £12 for a year, and while it was technically possible to deploy and use Nextcloud through the web interface, it didn't run well at all. Even the basic Nextcloud text editor (used to write this article) failed to function, save edits or reliably load documents. It also generated repeated warnings that "Your hosting account is reaching or exceeding its resource limits." Basically a standard webhost isn't configured in the same way as a VPS.

What's a VPS for?

So glad you asked! A VPS can be whatever you want it to be, and you can use it for whatever you choose. Generally you would use it for hosting software which has a front-end which can be accessed through a web browser.

The internet isn't what it was a decade ago, and thanks to the wonders of browser development, it's relatively simple to run beautiful applications which replicate, replace or augment services you would normally have to pay for – either with cold hard cash or indirectly with your personal data which is harvested, stored, analysed and sold every time you go online without protection.

We've been a fan of private servers (both dedicated and virtual) for some time, and deploy our own services to help us keep control of our own data, while enjoying all of the amenity the modern web has to offer. This is some of the self-hosted software we use which should give you some idea of what its all about:

Nextcloud This should probably be your starting point, as it has basic versions of almost everything you could ever need. At its heart it's a drop-in replacement for Google Drive and Dropbox. Where it excels is in its apps, which allow you to further expand the software to cover multiple use cases. There are full-fat office suites,



This is what you can expect if you go for £12 per year non-VPS hosting.

Virtual Private Servers IN-DEPTH

Ubuntu	FreeBSD	Fedora	Debian	Cento	os i	Rocky Linux
20.04 (LTS) x64 🗸 🗸	Select version 🤟	Select version 🔍	Select version	Select versi	lon v	Select version 🗠
hoose a pla	n					Help me choose @
			DEDICATED C	PU		
SHARED CPU						
Basic	General Purpos	se CPU-Opti	mized M	emory-Optimized	sto	orage-Optimized
Basic	General Purpos	se CPU-Opti	mized M	emory-Optimized	d Sto	orage-Optimized
Basic Basic esic virtual machines erformance. Nee blog PU options: () R	General Purpos with a mix of memory ar s, web apps and devites egular Intel m 550 w	se CPU-Opti nd compute resources. It environments. remium Intel minivitae ssb	mized M Best for small proj Premium with NVMs	emory-Optimized ects that can handle AMD NEW	d Sto	evels of CPU
SHARED CPU Basic Isic vitual machines Informance. Ike blog PU options: \$5,000 \$3,0007.000	General Purpos with a mix of memory ar s, web apps and devites equilar Intel m 550 w \$10 mo \$0 of5inuu	e CPU-Opti d compute resources. I t environments. remum Intel in rovas son \$15,mo \$0.022max	mized M Best for small proy Premium \$20,mo \$0.030.heer	AMD SSD \$40, 00000	d Sto	evels of CPU
Basic Basic stic vitual machines rformance. Ide blog PU options: \$5,mo \$30,007,mou 108,1000	General Purpos with a mix of memory ar s, web apps and devites equier intel \$50 \$10 mo \$0 005mu 2 08 1000	e CPU-Opti and compute resources. It environments. remum Intel minivities SSD \$155 mm \$0.022 max 2.08 2.0755	mized M Best for small proj Premium stath WMM \$20 (mo 10.030 have 4 GB 2 CPD	AMD Several Second Seco	d Sto	evers of CPU

While it is absolutely within the LFX discretionary budget to spin up a \$2,400 per month monster VPS, we decided that it would be egregious and excessive for the purposes of this feature.

music and video players, sync clients, mapping and location software, recipe managers, news aggregators... the list goes on and on. But while Nextcloud seemingly does everything, it isn't the best at any one thing, and as time goes on you'll be swapping Nextcloud functions for other services which are dedicated to a single task. >> Jellyfin A FOSS project to manage and play your multimedia library. Jellyfin handles videos, music, photographs, ebooks, and audiobooks, meaning you can listen, watch or read at any time and on any device. >> FreshRSS RSS isn't dead and it is still this writer's favourite way of reading digital news and magazines. FreshRSS will pull in RSS feeds from sources you select, and display them for you in a browser or through one of the many available mobile clients. Even better, with a little jiggery-pokery, FreshRSS will pull full articles rather than stubs - meaning that you never need to visit another website again.

>> WHAT ABOUT AWS?

The one-line explanation of a Virtual Private Server is pretty simple. It's a virtual machine running on a physical machine, alongside other people's virtual machines. You are responsible for the virtual machine, its uptime, its security updates and so on. Your VPS provider isn't going to be interested if your VPS gets hacked, it won't provide security or even updates for you. Case in point, my VPS required downloading around 900MB in updates immediately after initialising the VPS on DigitalOcean. Your resources are strictly limited to what you've paid for and if you max out the RAM on your VPS, it will crash.

Amazon Web Services is slightly different. It is virtual cloud computing – a phrase which means a lot more than the three words would suggest. Virtual cloud computing means that you won't have root access to the system, but security and updates are handled on your behalf. It's also infinitely scaleable, and if your computing needs outgrow what you've paid for, Amazon will still allow you to access more and more resources – meaning that your monthly bill can quickly grow out of hand. Thousands of pounds out of hand in some cases, so be careful...

While it is possible to run web apps in a pseudo-VPS using standard web hosting, it isn't worth the trouble unless the resources usage is exceptionally light.

>>

IN-DEPTH Virtual Private Servers

QUICK TIP

Go for the lowest cost or capacity VPS which suits your needs. It's always possible to upgrade to a more powerful VPS later. Photoprism Do you like being able to back up, sync and view your photos automagically, but don't trust Google and Apple not to check out your personals? You need a self-hosted photo manager hosted on a VPS. Photoprism is our choice. Other people may have strong feelings on other software.

The Fediverse The social network you've probably never heard of. It has millions of users and thousands of private servers which federate together to create a distributed social platform which is surprisingly usable. If you want to be part of it without signing up to yet another monolith administered and moderated by other people, you'll need to set up an instance of your own using Pleroma, Mastodon or Soapbox. A VPS is the perfect place to do it!

There are many more potential uses for a VPS, but space necessitates brevity. We'll be coming back to the software I've mentioned (and more!) in coming issues.

Choosing the right VPS

With great power comes great responsibility. So said Spider-Man, or possibly Voltaire. When choosing a VPS, the adage shifts slightly to "With great power comes great cost", so it's important to know roughly what resources you need your remote virtual machine to have, and an idea of what you want to spend on it.

To have an idea of the specs you require, make a list of the services you'd like to try out then check out the project GitHub pages to see the minimum requirements. You'll also want to read about other users' experiences to see how the software performs in the real world. While Nextcloud will run happily (albeit slowly) on a VPS with an ARM processor and 512K of RAM. *Photoprism* demands a minimum of 4GB of RAM in order to index your epic 20GB collection of holiday photos amassed on the beaches of Europe over the last three decades.

It's a good idea to start small with a low-cost VPS and then add in more memory, storage space and processing power as the need arises.

Being the cheapskates we are, we chose a \$5 per month DigitalOcean droplet with an allotted 1GB RAM, a 25GB share of an SSD, and a 1TB monthly data transfer limit. Naturally, this is scaleable and we can upgrade at any time. At the top end of the scale, DigitalOcean offers a VPS with a staggering 256GB RAM, 32 virtual CPUs, 7TB storage and 10TB transfer limit. At \$2,480 per month, this is probably not the VPS you are looking for.

Get up and running

Every machine needs an awesome hostname – even virtual ones. Virtual Private Servers provided by DigitalOcean are called droplets. It's a pun: your computing requirements are a mere droplet in their digital ocean. It hints at their limitless resources and your insignificance as both a

Finalize and create	
How many Droplets?	Choose a hostname
Deploy multiple Droplets with the same configuration.	Give your Droplets an identifying name you will remember them by. Your Droplet name can only contain alphanumeric characters, dashes, and periods.
- 1 Droplet +	the-avesome.lfx-virtual-private-server
Add tags	
Use tags to organize and relate resources. To	ags may contain letters, numbers, colons, dashes, and underscores.
Type tags here	

Swap usage: 0% Processes: 99	IPv4 address for eth1: 10.116.0.2	
updates can be applied of these updates are sta see these additional up	immediately. Andard security updates. pdates run: apt listupgradable	
e list of available upda check for new updates r	ntes is more than a week old. run: sudo apt update	
ne programs included with ne exact distribution ter ndividual files in /usr/s	n the Ubuntu system are free software; rms for each program are described in the share/doc/*/copyright.	
ountu comes with ABSOLUTE oplicable law.	ELY ND WARRANTY, to the extent permitted by	
ootëthe-awesome-lfx-virtu ootëthe-awesome-lfx-virtu w password: etype new password: osswd: password updated s ootëthe-awesome-lfx-virtu	ual-private-server:-# useradd -m david ual-private-server:-# passud david successfully ual-private-server:-#	
A REAL PROPERTY AND A REAL PROPERTY A REAL PROPERTY AND A REAL PROPERTY AND A REAL PROPERTY A REAL PRO		

The first thing you'll want to do is set up a user who isn't root – trashing the system by accident is a very real risk. Then again, it's only a *virtual* system.

human being (*I feel insignificant enough, thanks–Ed*) and as a user of their services However, as this section can be broadly applied to providers other than DigitalOcean, I'll be referring to the droplet simply as a VPS.

After account creation, which involves verifying your identity by way of adding a \$5 balance to your account, you'll be presented with an array of options, including deploying a web app, hosting a website or static site, deploying a container-based app or deploying a database. The option you want, of course, is 'Deploy a Virtual Machine'. This allows you to set up a VPS, from inside which you can do all of the other stuff anyway.

Clicking through to the next screen enables you to configure the specs of your new VPS, and you can choose from a selection of Ubuntu, FreeBSD, Fedora, Debian, CentOS and Rocky Linux. Each of these distros has a range of recent versions available.

'bunt it up

We chose Ubuntu 20.04 LTS because it's what we use at home and because it's the most popular Linux distro, while there is excellent support available from a range of sources. We stayed with the basic single CPU configuration and selected a datacentre in London – a closer datacentre means less latency. The next choice you'll need to make is whether to use SSH keys or simply a username and password. SSH keys are considerably more secure, but the choice, as always, is yours.

Finally, you will be asked to choose a hostname. By default this will be long, complicated and nonsensical. We chose something long, complicated and sensible: the-awesome-lxf-virtual-private-server.

With that, it's time to click Create Droplet or the equivalent button for your VPS provider, and go and make a cup of tea.

Hot brew in hand, you should come back to your machine after ten minutes or so. You'll see an IP address through which you'll be able to SSH into your shiny new VPS. Note that you'll be SSHing as root, eg: root@159.65.216.202, so the first thing you'll want to do is to create a new, non-root user with its own home directory skeleton using the *useradd* command. useradd -m new_username

followed by a new password for your new user using the *passwd* command:

passwd new_username

Grant your new user *sudo* powers so you can actually get things done by typing:

usermod -a -G sudo new_username

After this you should log out as root and log back in as your new user. Now you can start using your VPS just like your regular machine at home, update and upgrade the already installed packages using sudo apt update and sudo apt upgrade and start adding packages which will make your VPS journey possible.

As a bare minimum, you should have server software such as Apache or Nginx (always choose Apache because then you can say "It's a 'patchy' server"). You'll also need PHP and a database such as MariaDB.

You can install Apache by typing sudo apt install apache2. PHP by typing sudo apt install php , and MariaDB by typing sudo apt install mariadb-server then sudo mysqlsecureinstallation and following the prompts.

A domain of your own

Finally, you probably want a domain name through which you can access your services from out in the world. As we already own **lxf.by**, we chose to create a subdomain of that for our first service: **vps1.lxf.by**. If you don't already have a domain name, they can be picked up for a small cost from a provider such as Namecheap, or for free from Freenom.

You need to make sure that the nameservers at your registrar are pointed towards the IP address given to you by the VPS provider. Fortunately, this is a simple as changing the standard registrar nameserver to those of, in our case, DigitalOcean. From within your registrar's website, find the option for nameserver and change it from standard to custom. In the resulting field, you will want to add something like ns1.digitalocean.com. Nour VPS provider will handle the rest.

There is no point in us going into the exact details here as the location of the nameserver settings will vary from registrar to registrar. Most will have a help document available.

After the DNS changes have propagated (you can check this using **dnschecker.org**), you should be ready to go. Typing in the address of your new domain into your browser should result in you seeing the Apache2 Ubuntu Default Page.

Take a sip of your cooling tea and wipe your brow. The hardest part is over. You have a working virtual machine which you can access through an IP address.

Deploying VPS services

There's no point having a VPS if all you're going to do with it is host a static webpage – especially if that static webpage is the most boring and ubiquitous in existence.



This is where you set custom nameservers using Namecheap. The exact location will vary according to your own registrar.



It's tempting to deploy a web app without setting up your VPS first, but this will only lead to tears. Select the middle option on the bottom row instead.

A VPS is all about useful services. Search Google (other search engines are available) for 'Awesome self hosted' and you should end up on a curated GitHub page with a list of highly rated self-hosted FOSS software. Have a browse and see if there's anything you fancy. You'll come across some which require or use Docker and *docker-compose*. These are outside the scope of this article, but you'll find plenty of Docker tutorials in previous issues, the most recent in **LYF272**.

As a super-quick demonstration piece, I chose Sharlii (https://github.com/shaarli/Shaarli) which is a 'personal, minimalist, super-fast, database free, bookmarking service.'

To install it, download the package into your VPS home directory using:

wget <https://github.com/shaarli/Shaarli/releases/ download/v0.11.1/shaarli-v0.11.1-full.zip> Unzip by typing:

unzip shaarli-v0.11.1-full.zip

(Sidenote: you'll need to actually install unzip first with sudo apt install unzip .)

cd into Shaarli, them move the contents of the directory to where it can be seen by Apache. cd Shaarli

sudo mv * /var/www/html/ remove the original Apache index page: sudo rm /var/www/html/index.html

lastly give ownership of the html directory to apache: sudo chown -R www-data:www-data /var/www/html

That's it. You've installed your first web service on your very own VPS. Have a hunt round and find some more – it's addictive, and before you know it, you'll be wanting to upgrade your server.

	Install Shaarli
	It looks like it's the first time you run Shaarli. Please configure it.
Username	
Password	
Shaarli title	
My links	
Language	
	Automatic ~

Congratulations! You've deployed your first web app on your very own VPS. Awesome Self Hosted has plenty of other great suggestions.

When you first boot your VPS, it will contain practically nothing beyond the bare bones of the OS. You will need to install even basic utilities such as unzip. take some time to plan out what you're going to need.

OUICK TIP

TUTORIALS

TLDR

Better man pages

A fondness of reading is a good habit to inculcate, but **Shashank Sharma** believes it requires more than spending quality time with man pages.



Shashank Sharma is a trial lawyer in New Delhi and an avid Arch user. He's had it with superhero movies and TV shows.

here's very little you can achieve with CLIbased tools and utilities without referring to the manual pages. Commonly known as man pages, these serve as an introduction to all the myriad available options for the different commands, tools and utilities. But as useful as the man pages are, they can also overwhelm you with information. For instance, the man page for the **curl** command is over 3,000 lines long. The man page for **pdftk** similarly runs onto over 500 lines. By comparison, the pages for these commands in the *TLDR* utility clock in at fewer than 140 and 20 lines respectively.

The stark contrast in size will have already help explain the name of the project. TL;DR (too long: didn't read) is a popular internet slang, and is often used to provide a summary of long text. According its own description, *TLDR* is a community-driven effort which strives to simplify man pages into easily digestible pieces of information.

» COMPLEMENT MAN PAGES

The man pages are home to all pertinent information about any given command, such as description of the command, syntax, list of switches/arguments, with explanation, some examples, author information, and more.

If you find the bulk of the information is irrelevant to your needs, there are other projects that will help you quickly get started with the specified command. We've already covered *TLDR* in this tutorial, but apart from that you can also try your hands at a couple of additional projects that make working with CLI a breeze.

The bropages project strives to 'just get to the point'. Unlike man pages however, bropages only provides examples on how to use the specified command. As with *TLDR*, this is a community-driven project, with examples contributed by the users. You'll have to install Ruby on your distribution, and then use the *gem* package manager to install *bropages*. Once installed, you can run **bro**<command> for a list of user curated examples. You can also up/down vote the given examples, and even add your own to the list, all from the convenience of the terminal itself.

The *cheat* project similarly provides a cheatsheet for a large number of command-line tools and utilities. The Python project can be install using *pip* on most popular desktop distributions. Once installed, you can then run cheat <command> to view the cheatsheet, which essentially comprises a large number of examples, for the specified command.

immainlegiplyground: - * marialegiplyground:-5 tifr sfip -t ocean sfip Secure filts inanger argumen. Instructive program to copy files between hosts over SSH. For non-tuteractive file transfers, see scp or rsync. - Connect to a remote server and enter an interactive command mode sfip remote_sport remote_unerfament_host - Connect using an alternate pert: - Con

- Transfer remote file to the local system
- get /path/remote_file
- Transfer local file to the renote system: put /path/local_file

You can override the terminal's colours using the tldr <command> -t <theme> command to define the theme for TLDR, whether simple, base16 or ocean.

While man pages remain the canonical and official source of complete information for different command-line tools, utilities and commands, you don't always need to be completely versed in all the different capabilities before using it. You shouldn't think of *TLDR* as a replacement or alternative to man pages, but only as an effort towards providing the most basic information to enable you to start using the tool as soon as possible.

In fact, the contribution guidelines make this point rather succinctly: "Try to keep pages at around 5 examples. Pages can be longer or shorter when appropriate, but don't exceed 8 examples. Remember, it's OK if the page doesn't cover everything: that's what man is for".

Installing and using TLDR

Some popular desktop distributions carry *TLDR* in their software repositories; however, as the *TLDR* pages are being constantly updated by the community, it makes sense to use the official package.

TLDR is a Node.js utility, which means that you must first install Node.js on your distribution before you can use the *npm* package manager to install it. If you're running Debian/Ubuntu or a derivative distribution, you can install *npm* with the **sudo apt install nodejs npm** command. For Fedora, *npm* is now part of the Nodejs package, so the **sudo dnf install nodejs** command is all you need to install it on Fedora. You'll find distribution specific instructions for installing *npm* package manager on the project's website.

With the *npm* package manager installed, you can now run the **sudo npm install -g tldr** command to

Terminal **TUTORIALS**

install *TLDR*. By default, *npm* installs the specified package in the current working directory, called the local mode, so you must use the -g command option if you want to install the package in global mode.

Before you can start using *TLDR* to quickly assimilate different commands, tools and utilities, you must first update its cache with the **tldr --update** command:

\$ tldr --update

✓ Updating...

✓ Creating index...

If you're comfortable working with Snap packages, you can also install *TLDR* with the **sudo snap install tldr** command. With *TLDR* installed, you can now read the page for any command with the **tldr <command>** command:

\$ tldr ssh

ssh

Secure Shell is a protocol used to securely log onto remote systems.

It can be used for logging or executing commands on a remote server.

- Connect to a remote server:

ssh username@remote_host

- Connect to a remote server with a specific identity (private key):

ssh -i path/to/key_file username@remote_host

- Connect to a remote server using a specific port:
- ssh username@remote_host -p 2222
- Run a command on a remote server:

As you can see, *TLDR* provides a quick one-line explanation of what the specified command does, along with a few command options, with their description and example usage. In this way, *TLDR* helps you quickly assimilate the usage without overwhelming you with information. By comparison, the man page of SSH runs close to 1,000 lines.

Skip to the end!

If you don't remember the exact name of the command, but only have a basic idea of what it can do, you can use the tldr -s<search tldr pages using these words>

command. For instance, run the following command if you're looking for a tool that can be used to manipulate or edit PDF files:

\$ tldr -s edit PDF

Searching for: edit PDF
\$ pdfunite
\$ qpdf
\$ pdftk
\$ mutool
\$ ocrmypdf
\$ pdfjoin
\$ pueue-edit
\$ cpdf
\$ pdffonts
\$ xpdf
Run tldr <command/> to see specific pag

The tldr -s<search terms> command presents a list of commands that feature your search terms. You can then run the tldr <command> command to read more about these suggested tools to narrow down the one that you wish to use. Your choice of keywords is quite important as *TLDR* individually matches each specified word, and provides a list of commands that feature it in the description. So, the command **tldr-senable firewall** will produce different results than **configure firewall**.

Unlike man pages, which are often rather lengthy, *TLDR* pages make it easy to assimilate commands. If you're new to the CLI and not all that familiar with the different commands, tools and utilities, you can take advantage of the **tldr -r** command, which opens a random *TLDR* page from the cache.

You can similarly run the tldr -a command for a list of all the pages in the *TLDR* cache. The command tldr -1 can instead be used to list only the platform-specific pages. If you're currently running Linux, but want to list all the OSX specific pages, you can do so with the tldr -1 -osx command.

The *TLDR* pages are stored in the **~/.tldr/cache/** pages directory, and sorted into different operating systems such as Linux, OSX (or rather MacOS as it's now called), Android, SunOS and Windows. *TLDR* also warns you when you try to access the *TLDR* page about a command that's not available on your current operating system:

\$ tldr xcopy

...

Command xcopy does not exist for the host platform. Displaying the page from windows platform xcopy

Copy files and directory trees.

More information: https://docs.microsoft.com/ windows-server/administration/windows-commands/ xcopy.

- Copy the file(s) to the specified destination:
- xcopy path/to/file_or_directory path/to/destination
- List files that will be copied before copying:
- xcopy path/to/file_or_directory path/to/destination /p
 Copy the directory structure only, excluding files:
- xcopy path/to/file_or_directory path/to/destination /t - Include empty directories when copying:
- xcopy path/to/file_or_directory path/to/destination /e - Keep the source ACL in the destination:

TLDR is fairly straightforward to use, and can help take away some of the 'fear of the unknown' that novice users might harbour towards the command-line.

1/1 - + 🗗 🖵							
1: linuxlala@playground: ~ 👻							
<mark>linuxlala@playground:~\$</mark> t Usage: tldr command [optic	ldr -h ins]						
Simplified and community-	Iriven man pages						
Options:							
-vversion	output the version number						
-l,list	List all commands for the chose cache	en pl	atfor	n in	th	e	
-alist-all	List all commands in the cache						
-1,single-column	List single command per line (use w	rith o	ptio	ns		
-rrandom	Show a random command						
-erandom-example	Show a random example						
-f,render [file]	Render a specific markdown [fi	le]					
-mmarkdown	Output in markdown format						
-o,os [type]	Override the operating system	[linu	x. os	x. s	uno		
linux	Override the operating system	with	Linux				
05X	Override the operating system	with	OSX				
sunos	Override the operating system	with	SunOS				
-t,thene [thene]	Color theme (simple, baselo, c	(cean)					
-s,search [keywords]	Search pages using keywords						
-uupdate	Update the local cache						

Ironically, TLDR doesn't have a man page of its own. If you want information on what the tool can do, run the tldr tldr or the tldr -h command.

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

OUICK TIP

The TI DR

TUTORIALS Better copy and paste



Add real power to your clipboard

Nick Peers dons his best hiking gear and goes hunting for a simple clipboard tool to rule them all. CopyQ is happy to oblige.



Nick Peers dreads to think how many items he's copied and pasted to clipboards over the past 30-plus years. It's enough to turn his hair white (from its current grey). he clipboard is one notable part of any operating system that keeps missing out on some love and attention. This brilliant timesaving tool has one fundamental flaw: it can only remember the last item you copied to it.

Thankfully, there's a cottage industry of clipboard enhancement tools, and one of our favourites has to be *CopyQ*. It enables you to store multiple clipboard entries comprising plain text, rich text, HTML snippets and images, organised as you see fit into tabs, and editable via its own built-in editor. These clips can then easily be pasted into any compatible document as you see fit.

As always, you'll find CopyQ available from the Ubuntu Software Centre, but it's a frozen version (3.10.0 in Ubuntu 20.04 for example). Instead, ensure you have the latest version – 4.1.0 at time of writing – by installing it via flatpak or its own dedicated repo:

\$ sudo add-apt-repository ppa:hluk/copyq \$ sudo apt update && sudo apt install copyq

Basic usage

Once installed, open *CopyQ* from the launcher. It'll place itself in the menu bar and the main window will appear, only to vanish as soon as you click away from it. Bring it back again by clicking the menu bar icon to reveal the main *CopyQ* menu, then choosing 'Show/Hide' option. Better still, prevent it disappearing so quickly by opening Preferences from the same menu, then unticking 'Close when unfocused'. While you're here, you can configure other key settings, such as whether to have *CopyQ* automatically start with your PC.

Once done you can leave CopyQ residing silently in the menu bar. The next time you copy any supported item to the clipboard, such as a snippet of text or an image, CopyQ will capture a copy. Verify this by opening the CopyQ menu – you'll see the item is listed twice: as the current Clipboard item, and additionally as an entry in a menu that sits above it, indicating it's now a part of the CopyQ library.

Note: images copied from your image editor can take a little while to transfer across, so wait a few seconds before checking *CopyQ* has successfully captured them. They'll appear in the list labelled <IMAGE> with only a tiny thumbnail to distinguish them from each other. You'll be able to change this – more on that later.



CopyQ sits unobtrusively into your menu bar, ready to leap into action with its time-saving capabilities when you need to invoke it.

Now copy a second item and when you next open the menu you'll see the Clipboard content change to show the item you've just copied to it. You'll also see the new item appear at the top of the list of library items, with your previous item now sitting beneath it.

Pasting items is simple: to use the current clipboard's contents as you've always done, simply press Ctrl+V. However, if you'd rather paste a previously captured clip, just click it from this list and it'll be pasted into your document at the current cursor point without disrupting the clipboard's current content.

Organise your clips

As it stands, *CopyQ* just made copying and pasting multiple clips a whole lot easier, but it's just the start. As your collection of clips grows, things can become unwieldy, but *CopyQ* offers several ways in which you can label, edit and organise them. Let's start by labelling individual clips: click the *CopyQ* menu icon and choose 'Show/Hide'.

The main *CopyQ* window reopens, and this time you'll see each clip that's been captured by *CopyQ* while it's been running is listed here. First, you can reorganise the order of clips simply by dragging and dropping or using keyboard shortcuts (Ctrl+up arrow/down arrow to move up and down one place, or Ctrl+Home/End to move to the beginning/end of the list, respectively).

You can also file clips into separate tabs to organise them – choose Tab>New Tab to get started. You'll be prompted to name your tab – give it a logical name, and note that you can assign it a keyboard shortcut that will enable you to switch to that tab quickly within the main

QUICK TIP

Copy0 comes with comprehensive documentation. Read it at https://copyq. readthedocs. io/en/latest/ index.html, where you'll be able to discover everything the program is capable of. *CopyQ* window using the & operator. For example, 'Te&xt' would assign Alt+x to the Text tab.

Once created, the main *CopyQ* window becomes a multi-tabbed one. Right-click a tab to rename or remove it, and assign it an icon to make it stand out visually. There are two things to consider: first, the main *CopyQ* menu only displays clips from a single tab at any time – by default this is the currently selected one, but you can specify a specific tab to use via Preferences>Tray (where you can also change the number of clips shown – five by default).

By default, new clips are automatically stored in the clipboard tab, but you can change this too – via Preferences>History>Tab for storing clipboard. If you'd like to move clips to a new tab, just locate the clip and drag it on to your target tab. It'll be copied here rather than moved, so once you've verified it's been copied, simply select the original and press Delete to remove it.

Edit clips

As things stand, clips can only be identified by their content – which is fine in the main window, but can be tricky from the *CopyQ* menu. To add a title, right-click its entry and choose Edit Notes. Simply enter a title and press the Save button. Adding a title is particularly helpful as far as images are concerned, because it displays a descriptive title next to the cryptic <IMAGE> tag in the *CopyQ* menu.

You can also assign tags to clips – right-click the clip and you'll see options for adding, removing and clearing tags (to add multiple tags in one go, separate them using commas). Tags will then appear above the title. Aside from quick visual clues, they can also be used to filter the tags shown: choose Edit>Find to bring up the search box, then enter a tag to apply the filter. The search box can also be used to search for keywords within clips (these are highlighted in the search results).

You can also edit text clips – for example, to correct mistakes or update the text – by right-clicking the text and choosing Edit. You'll see the same editor used for titles appear, but this time you can use the style options to format your text if you wish. If the built-in editor is too basic for your needs, open *CopyQ*'s preferences to the History section and input your choice of editor into the 'External editor command' box.



As your clip collection grows, make use of CopyQ's organisational tools. Tabs are particularly useful in bringing order to your items.

You can link *CopyQ* to an external image editor too, enabling you to edit image clips simply by right-clicking them and choosing Edit – set this up under 'Items > Images' (note to *GIMP* users – *CopyQ* behaves much better with the beta version 2.99 – install it via flatpak and then enter the following into the 'Image editor command' box:

flatpak run org.gimp.GIMP

Pasting items

You've built a clip library, organised it into tabs, but you'd like to make it easier to access your clips. One option is to create a keyboard shortcut that enables you to bring up the *CopyQ* menu bar at the current cursor point in your chosen application.

To do this, open CopyQ's preferences from its menu bar icon to the Shortcuts>Global tab section. Click + next to 'Show the tray menu' to create the shortcut. Choose something unique that won't clash with another application, such as Ctrl+Alt+Super+C, and click OK.

Now switch to your text editor and press the keyboard combo – you'll see the menu bar menu pops up at the cursor point from which you can select your item to paste in. This works fine when your clips are all handily stored under a single tab; if not, set the keyboard shortcut under 'Show main window under mouse cursor' instead to pop up the main *CopyQ* window at your cursor point instead. Note that any commands you set up in this way will also be added to the main *CopyQ* menu. To remove them from the menu without losing the keyboard shortcut, check out the Advanced tips and tricks box (*below*).

» ADVANCED TIPS AND TRICKS

We've covered the basics of *CopyQ* in this tutorial, but there's more. *CopyQ* can be run from the command line. Type **copyq** help into the Terminal window to get a full rundown of available controls (and their limits).

 $CopyQ \ enables \ you \ to \ store \ compatible \ files \ -, png \ (images). \ txt \ (plain \ text) \ and \ . \ html \ (formatted \ text) \ by \ default \ - \ in \ a \ named \ tab, \ keeping \ those \ items \ in \ sync \ as \ the \ files \ themselves \ change. \ To \ set \ this \ up, \ create \ a \ new \ tab \ in \ the \ main \ window, \ then \ open \ Preferences \ to \ tems-Synchronize. \ Type \ the \ tab \ name \ into \ the \ approxpression \ tab \ synchronize. \ Type \ the \ tab \ name \ into \ the \ approxpression \ synchronize. \ the \ the \ the \ the \ the \ tab \ tab$

One final tip: press F6 to customise what commands appear on the *CopyQ* menu, plus add your own. Tick the 'Show Advanced' button after creating a new command to access a Command window where you can set up your own commands. The following example from the *CopyQ* documentation (see the top tip) enables you to automatically filter captured images into a specific tab: Name-Move Images to Other Tab Input-image/png Automatic-true Remove-true Icon=\xf03e Tab=&Images



You can turn files into clipboard content by syncing their folder into CopyQ.

BACK ISSUES

BACK ISSUES >> MISSED ONE?

ISSUE 280 September 2021

Product code: LXFDB0280

In the magazine

As Linux hits

30 years, we show how the kernel became a worldwide phenomenon. We take a look at terminal browsers, plus show how you can make your mark in desktop publishing, build a Pi NextCloud server and run Linux distros from yesteryear!

CELEBRATE

OF LINUX!

DVD highlights

Pop!_OS 21.04 (64-bit), and classic Linux distros (32-bit, unsurprisingly).

ISSUE 277 July 2021

Product code: LXFDB0277

In the magazine Discover what's

new in the latest version of Ubuntu, grab a slice of network-attached storage, code a game in Scratch, emulate the Dragon 32 and set up your own streaming server with *Jellyfin*, Plus we look back at Prestel, the pre-internet data service!

DVD highlights

Ubuntu 21.04 (64-bit) and MX Linux 19.4 (32-bit).

ISSUE 279 Summer 2021

Product code: LXFDB0279

In the magazine Find out how to

video to the world. Elsewhere, we compare five office suites, diagnose and solve Linux problems, emulate the Acorn Electron, set up a virtual network, design circuit boards and manipulate date with Pandas.

DVD highlights

Rescue kit (CloneZilla, System Rescue and Rescuezilla, plus Zorin OS Lite.

ISSUE 276 June 2021

Product code: LXFDB0276

In the magazine We show you

how to set up your own server for your photos, either at home or in the cloud. Learn how to emulate the Commodore PET, manage headless servers with *Cockpit* and render objects in *Blender*. Plus get coding and graph metrics in Python today!

DVD highlights

Manjaro 21 (64-bit) and Tails 4.17 (32-bit).



ISSUE 278 August 2021

Product code: LXFDB0278

In the magazine Want a faster.

want a faster, better server? We show you how to set up some of the best server distros around. We also cover desktop virtualisation, running a mobile second screen, emulating the Altair 8800, multitasking in Python and web-app security.

ÍIND

DVD highlights

Bodhi 6.0 and Lubuntu 21.04 (both 64-bit) and AntiX 19.4 (32-bit).

ISSUE 275 May 2021

Product code: LXFDB0275



In the magazine Get more from

your Raspberry Pi with our great selection of projects. A parents' guide to coding, ana introduction to 3D tool *Blender* and a tour of video conferencing tool *Jitsi Meet*. Plus code a first-person shooter, audio software reviewed and photo effects.

DVD highlights

Mageia 8 (64-bit) and AntiX 19.3 (32-bit).

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

READING IN THE USA?

UK readers turn to p18

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!



>USA From \$132 For 13 issues

REST OF THE WORLD From \$137 For 13 issues ≫EUROPE
From €100
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

TUTORIALS Atari 800 Emulation



Going back to 8-bit with the Atari 800

Les Pounder goes back to a time before E.T. cartridges were buried in the desert and 48KB was considered lots of RAM.



Les Pounder is Associate Editor at Tom's Hardware and a freelance maker for hire. He blogs about his adventures and projects at bigl.es. he first exposure to an 8-bit Atari I had was the 65XE, owned by my mate Scott. Sitting down to play *Millipede*, we had a thoroughly good time. Atari's 8-bit range of home computers dominated the US market but in the UK the home computer wars were fought between Sinclair, Amstrad and Commodore.

Released in late 1979, the Atari 800 and other Atari 8-bit machines of the era were originally developed as a second generation of Atari's popular VCS consoles. But Atari president Ray Kassar wanted to challenge Apple, who had the most popular home computers at the time.

The Atari 800 was marketed as a "timeless computer", driven by how well it could be expanded. The Atari 800's MOS 6502 CPU ran at 1.79 MHz (NTSC) and 1.77 MHz (PAL) and had 16KB, later upgraded to 48KB of RAM. Later the 800XL came with 64KB of RAM. Of the 8-bit Ataris of this era, the 800 and 800XL were the premium models, featuring chunky, clunky keys that could withstand the hammering of an eight-year-old gamer. The beige plastic aesthetic screamed typewriter meets *Blade Runner*, which was a good thing. The 800 had two cartridge slots, one of which is often used for the bundled Atari BASIC cartridge (the 800XL had this built in).

In order to emulate the 800 we chose *atari800*, available via the Ubuntu repositories. Installation is a



The Numen demo is an excellent showcase for clever programming on low-spec hardware. The Atari 800 looks like something from the set of Blade Runner, and it is a joy to behold. A chunky, industrial aesthetic that pleases.

breeze. Open a terminal and first update the repositories, before installing the emulator. \$ sudo apt update

\$ sudo apt install atari800

ATARI

To start the emulator, in a terminal type the following. \$ atari800

You will notice that the emulator starts in a self-test mode; this caught us out, but it is a useful tool nonetheless. Pressing F3 (the Select key on an Atari 800XL) switches between the tests. We chose the audio test, then pressed F4 (Start key) to run the test. We were presented with a simple audio test which played a tune similar to that played by Captain Picard in the *Star Trek The Next Generation episode* 'The Inner Light'. To stop the test and boot the Atari 800 normally, press F5.

To skip the self-test completely we can start the emulator with an argument:

\$ atari800 -basic

The *atari800* emulator is full of rich features, but they are a little hidden. To access the menu press F1 and you will see a plethora of options. In the course of this feature the main sections we will use are Disk Management, Cartridge Management, System Settings and Emulator Configuration. Using the arrow keys, navigate down to System Settings and press Enter. Here we can change our Atari model, from the Atari 400 all the way to a 130XE. The best overall compatibility comes from the Atari 800XL. We can also change

Atari 800 Emulation TUTORIALS

between a PAL and NTSC video system, not a big thing now, but NTSC machines run a little quicker than their European cousins. Emulator Settings, also found in the main menu, is where we can set our emulator to boot into the BASIC interpreter, run a tape on boot (press F4 and Shift+F5 to load) and also set our machine into Turbo mode. Turbo mode makes the BASIC interpreter a little twitchy so perhaps only use it if you are booting directly into a game.

Intro Atari BASIC

Atari BASIC, released in 1979 and created by Paul Laughton and Kathleen O'Brien, was a little different to other versions of BASIC in this era. It wasn't based upon Microsoft's BASIC and it had a few keywords that were specific to Atari BASIC. Originally distributed as an 8KB ROM cartridge, BASIC was soon added as a built-in option for the 600XL and 800XL. There were three releases of Atari BASIC, each fixing bugs and introducing new features.

Revision B was meant to fix all of the bugs in Revision A, including a pesky bug that added 16 bytes to a program every time it was saved and loaded into memory – not ideal when the max RAM was 48KB. All versions of BASIC were provided on cartridge, but some machines, notably the 800XL, XE and XEGS models had it built in.

To determine what version of BASIC we have is pretty simple. In the BASIC interpreter we can type this command to return a value:

PRINT PEEK(43234)

Revision A is 162, Revision B 96 and Revision C 234. It looks like our emulator is running Revision C.

Okay, let's flex a little BASIC muscle. We've done this a few times on many different machines but we start as always with the ol' 10 PRINT project. Each line of BASIC code for a project will start with a number, 10, 20, 30 and so on. This tells the interpreter the sequence of code; it jumps from one line to the next in ascending order. But why do we do this? Quite simply if we make a mistake and miss out a line of code we can insert another line of CODE without messing up the original code. Let's do the 10 PRINT project to illustrate this. **10 PRINT "HELLO WORLD" 20 GOTO 10**

If we **RUN** this code it will print **HELLO WORLD** over and over again. Press F7 on the keyboard to break the running code. But what if we want to add another line? We can insert a new line between 10 and 20. Logically this would be 11, giving us many more options



In the 1980s and 90s, Archer Maclean was a god for home computer fans. International Karate was a particular favourite for the humble C64.

to expand or correct the code. But we are going to use 15 as this is just a simple test. 10 PRINT "HELLO WORLD"

15 PRINT "LXF ROOLZ" 20 GOTO 10 Now RUN this new code and you will see alternating incode full LOWOPLD and LXE ROOLZ on the

lines of HELLO WORLD and LXF ROOLZ on the screen. Press F7 to stop the code.

Making music

As always, we take our BASIC projects a little further and in the past we have made apps to calculate maths problems, controlled turtles, and drawn psychedelic patterns in pixels. This time we shall create a simple progression of three notes.

Start the Atari 800 emulator in BASIC mode.

\$ atari800 -basic

The first line is a variable called **DURATION** and as you may have guessed it will control the length at which a note is played. We've chosen 45 'clock ticks', about half a second.

10 DURATION = 45

Line 20 sees the first note to be played. Atari BASIC has no means to control the duration for which a note is played so we shall hack a FOR loop to play it for the duration of 45 ticks. In the FOR loop we use the SOUND command to control voice 0, with a note of 121,

OUICK TIP

The default emulator window size is tiny! Change this by using the win height and width commands when starting the emulator. Just type the following for a 720P display atari800 -winheight 720 -win-width 1280

» ONLINE EMULATION

If you just want to play games and not have to worry about setting up an emulator, you're in luck as there are many just like us. The emulator *jsA8E* at **https://bit.ly/lxf281atari**, from way back in 2014, runs the Atari 800 in your browser using JavaScript. It seems that online emulators written in JavaScript for retro computers is a big thing. In the course of this series we have encountered online emulators for all systems, even the 486 PC! This emulator takes you straight into a game, and we can recommend *Donkey Kong*, it feels arcade perfect and plays like a dream.

We found a great resource for online emulators at https://bit.ly/lxf281js. Here we found many different emulators with one thing in common, they are all written in JavaScript. The *jsA8E* emulator is there along with Javatari.js for the Atari 2600, and even a great Atari ST emulator in the form of *EstyJS*.

The power of JavaScript as both a general-purpose language for the web and a serious emulation development platform is evident by the sheer number of emulators on offer – preserving the history of computing from the comfort of our browser window.

TUTORIALS Atari 800 Emulation

QUICK TIP

Using the System Settings menu we can easily upgrade our humble Atari 800 with more RAM From the main menu we can add additional disk drives. cartridges and tapes. Press F1 to access the main menu and then choose the peripherals that vou need.

tone of 10 and loudness of 15. Feel free to tinker with the note – musicians amongst us will soon work out the tones to notes.

20 FOR I = 1 TO DURATION : SOUND 0, 121, 10, 15 : NEXT I

Line 30, and here we add a short pause before playing the next note. We have to introduce a pause in this manner because Atari BASIC, like other versions of BASIC, does not have a SLEEP/PAUSE/WAIT keyword. 30 FOR I=1 TO 10: NEXT I

The next sound plays a slightly different note which shows a progression. Again we are controlling voice 0 and the loudness remains the same, 15.

40 FOR I = 1 TO DURATION : SOUND 0, 96, 10, 15 : NEXT I

Line 50 is another pause.

50 FOR I = 1 TO 10 : NEXT I

Line 60 sees the final note in the progression. 60 FOR I = 1 TO DURATION : SOUND 0, 81, 10, 15 : NEXT A final pause. 70 FOR I = 1 TO 10 : NEXT I

Lastly we turn off any audio. 80 SOUND 0,0,0,0

When ready type **RUN** and press Enter to run the code. You should hear a short three-note progression. To run the three-note progression continuously we need to put it into a loop and for that we add line 90, which instructs the code to go back to line 20 and play the first note.

	Atari 800 Emulator, Version 4.1.0	0
HELLO_WORLD		
LXF ROOLZ		
I XE DOOL7		
HELLO WORLD		
LXF ROOLZ		
HELLO WORLD		
LAF RUULZ		
LXF ROOLZ		
HELLO WORLD		
LXF ROOLZ		
HELLO WORLD		
HELLO WORLD		
LXF ROOLZ		
HELLO_WORLD		
LXF ROOLZ		
IXF ROOL7		
HELLO WORLD		
LXF ROOLZ		
HELLO WORLD		
LAF RUULZ		

No 8-bit machine would be complete without some sort of BASIC, and Atari BASIC is a little different to most, but we enjoyed our time with it.

» RESOURCES USED IN THE FEATURE

When writing features we go to multiple sources and for the Atari 800 we went to the Atari BASIC wikipedia page https:// en.wikipedia.org/wiki/Atari_BASIC to get a handle on this version of BASIC. From here we learnt about Atari BASIC's limitations with strings, requiring a string to be created with a maximum length (a Dimension) so that we don't use too

much memory. Remember that some Atari 8-bit machines have just 16KB, so every character counts. The keyword list was especially helpful in researching what projects we could create, and finding the **SOUND** keyword sent us off into a world of noise.

The Atari 800 emulator itself has an excellent help page at https://bit.ly/

90 GOTO 10

Type **RUN** and press Enter to run. The three notes will play indefinitely, but you will soon grow weary of them. To stop the running code press F7 and this will cause the currently playing note to hang, to quickly stop it type this and press Enter. **SOUND 0.0.0.0**

Playing games

The Atari 800 had three ways to load games. First there are the cartridge ports, we can connect a cassette tape drive, or we can use floppy disks. UK users will be familiar with cassettes and cartridges. Tapes were the cheapest and most prevalent way to sell and, ahem, 'redistribute' games. In our tests we found that tape images often failed to load, and we had greater success with disk images.

To load a disk image, you will need your .ATR disk image in a convenient location. In the emulator open the main menu, Fl. Press the down key until you reach Disk Management and press Enter. There are eight possible drives; most games come on one disk so press Enter on Dl. Navigate to your chosen game – note that the emulator will default to the directory from where you opened the emulator. Press Enter to select the disk then press ESCAPE to return to the READY screen.

Press F2+Shift+F5 to cold-restart the Atari 800 with your game in the drive and it will auto-boot. We tested this with Archer Maclean's awesome *International Karate*, and we may have played it for an hour, just to make sure it worked.

Loading a game from cartridge follows a similar process. First we press F1 to load the menu, then navigate to Cartridge Management. Select Cartridge: None and press Enter. Select your cartridge image and press Enter. You will then be asked what type of cartridge it is, and here is where a little trial and error is required. You will see repeated errors until you select the correct cartridge type. For some games you will have to go into the System Settings menu and change your machine to match the supported systems for that cartridge. The best gaming experience comes from using disks – they are fast and work well.

The demo scene

You might expect that the Atari 800 would be too pitiful of a machine for a demo scene. How wrong you'd be! The Atari 800 may not be the most powerful machine, but the demo scene provides plenty of examples illustrating what clever coding can achieve with limited hardware. The pinnacle of the Atari 800 demo scene is

Ixf281help which enables you to learn how to select options in the self-test tool, reboot and make the Atari 800 ready to run a tape/disk/cartridge on boot. It also enables you to up the resolution of the application, so that screenshots were possible. Knowing the correct switches to add when calling the emulator also helps a lot.

Atari 800 Emulation TUTORIALS

Atari 800 Emulator, Version 4.1.0 — 🗆 🏮	Fand but
	cost you
System Settings	sett us to Atar
Machine:Atari 800XL (64 KB)OS version:BB01 Rev. 2 (auto)BASIC:built inBASIC:built inBASIC:Rev. C (auto)BASIC:Rev. C (auto)BASIC:Rev. C (auto)XEGS game:NoneRAM size:64 KBVideo system:PALMosaic RAM:N/AAxlon RAMDisk:N/AAxlon RAMDisk page \$0F shadow:N/A1200XL keyboard LEDs:No1200XL F1-F4 keys:No	• Atar
Keyboard: integrated/attached MapRAM: No	

Fancy a new Atari but can't cover the cost? Try before you buy! System settings enables us to swap 8-bit Ataris at will.

Taquart (TQA) Numen. The demo requires 320KB or more of RAM. To enable this we pressed F1, navigated to System Settings, pressed Enter and then set the RAM size to 320KB (Rambo). The demo starts with a walk through a wooded area, as we approach two guards. Then we enter a new world, a *Doom*-type maze. The demo then proceeds to wow us with all the demo scene tricks in the book, accompanied by a thumping 8-bit tune. 3D worlds, lighting effects, line art and fluid animation all from one 5.25-inch floppy disk and 320KB of RAM. Take a look for yourself at www.youtube.com/ watch?v=J2Abi02Tlwk.

The Atari 800 in 2021

Just like other'80s computers, the Atari 8-bit range has seen interest from nostalgic gamers and coders eager to reminisce about long tape-loading times (we jest – as any C64 or Spectrum user can attest, we waited a long time for games!).

For around £50 we have SIO2SD from https:// lotharek.pl. This device emulates an 5.25-inch floppy drive and we can load an SD card with hundreds of games and applications. We can also save and load BASIC applications directly from the card, preserving our own projects for use with emulators and other models of Atari.

If you would like to take advantage of the Atari 800's built-in cartridge ports then the ATARIMAX flashcart (www.atarimax.com) and software enable you to create your own cartridge-based compilations, ready to go in seconds.

The Rolls-Royce of modern loaders is the SDrive-MAX, an Arduino-powered floppy drive and cassette emulator that enables the user to load and save data with relative ease. The SDrive came onto the scene in 2008 and it has been updated over the years to the MAX variant. For around £37 the SDrive-MAX looks to be the best option for overall ease of use and compatibility at a reasonable price. But what about original Atari hardware? Well, looking around online auctions we can see that Atari 800XL are quite sought-after machines. We saw one for sale at £175, and others with bids approaching £80 with over a week left. External drives were going for $\pounds70$, cassette and floppy games hitting between $\pounds5$ and $\pounds10$.

Before you go off to spend your money, remember this. Old hardware needs careful attention; your dream machine may have been in an attic for 20 years, or kept mint in a box in a collection. Check with the seller before you buy, and make a note of any issues for your YouTube research. A good reference is Adrian's Digital Basement. Adrian has repaired and spoken about 8-bit Ataris and would be a good resource to learn from.

The Atari 800 legacy

The industrial/sci-fi design of the Atari 800 made it a joy to look at. The games and community around it have blossomed and while it may not have had the success of the Commodore, Sinclair and Amstrad, it has something more important: the love of a community who still to this day supports and nurtures this ancient but humble home computer.



QUICK TIP

if your BASIC code won't stop running, fear not! Stop your code dead by pressing F7. If you are making music remember that you will also need to turn off the audio using SOUND 0,0,0,0 otherwise you will get complaints from the neighbours.

Swapping disks is easy; just open the menu, pick a disk and the emulator mounts it ready for use. All we need to do is reboot the Atari and our game loads.

>> WE CREATED ATARI ST FORMAT! Subscribe now at http://bit.ly/LinuxFormat

CERTIFICATES

Understanding the power of certificates

Stuart Burns walks through what certificates are, setting up a certificate server in your local network and how to use it.

ertificates are a critical component of the internet today, so in this tutorial we'll cover

how they can be used and how to make your own

Certificate Authority (CA) to create them. We'll be using a fully updated Ubuntu 20.04 LTS server to create a

certificate server for your local network as well as a root

creating and installing an SSL certificate for an example

these hosts should have static IP addresses. If you want

certificate for distribution with Firefox. We'll also cover

webserver (again, Ubuntu 20.04 LTS server). Both of

to play along you should be able to ping both servers using their FQDN (Fully Qualified Domain Name).

what certificates are, why they are so critical,



Stuart Burns is a Fortune 500 systems administrator who has a passion for open source software, Linux and privacy.

A dead cert

So what is a certificate? A certificate is essentially verifiable proof that an item is what it purports to be. For example, a passport is proof of identity. We put faith in it because we trust that the Certificate Authority (the Passport Office) has vetted and verified the holder's identity, and believe they are who they say they are and that the document (passport) is real.

The same principle applies to verifying identity on the internet. The difference is that the certificate is digital in nature and the certificates are digitally signed by the issuing authority using public key encryption. This means, all things being equal, that they cannot be easily forged or changed. Essentially a certificate provides assurance about the integrity of the item in question. On top of that, a webserver uses this certificate in setting up encryption between the website and the client.

There are several commercial certificate authorities including GoDaddy, Visa, DigiTrust and others. These vendors pay to include their public certificate in a browser – for example, *Firefox*. This means that any website signed by these signing authorities is automatically trusted. To see the range of certificates authorities your browser trusts, in *Firefox* for example, go to Preferences, search for 'Certificates' and then select View Certificates. You will certainly have heard of some of them but not others.

If you want to remove a certificate authority, just select the item in question, click 'Distrust this certificate' and it will be removed from the list. Exercise sysadmin@template:-/signing
 Q = - a
 sadmin@template:-/signing
 Q = - a
 sadmin@template:-/signing
 var about to be asked to enter information that will be incorporated
 to your certificate request.
 to your actificate request.
 you enter '.', the field will be left blank.
 ...
 you enter '.', the field will be left blank.
 ...
 mury have (2) Letter code) [Au]:GB
 as a province have (field have) [Same-State]:North West
 cality have (eq. city)]]:Anonhester
 ganization have (eq. server FQDW or VOB nane) []:webserver.itburns.locat
 atl Address []:staart@meelymintedmedia.com
 ease enter the following 'oxtra' attributes
 be your (reflicate
 equest
 optional conpany name []:
 actificate request
 aothol conpany name []:
 aothol conpany nam

Creating a certificate request for the webserver. The request can then be sent to the CA for signing.

caution when removing a certificate, however, as it may prevent certain sites from working.

Without a doubt, the easiest way to install a basic certificate server for a home network is to use the easyrsa package included in the Ubuntu repository. *Easy-rsa* will do all the heavy lifting for the Certificate Authority creation. There should be a local DNS server configured so that we can properly resolve the servers in question and experiment with our certificates. Without further ado, install the *easy-rsa* package:

sudo apt-get install easy-rsa -y

Create a directory and link the *easy-rsa* to use our configuration:

cd && mkdir easy-rsa && ln -s /usr/share/easy-rsa/* ~/ easy-rsa

All the above has put in place the requirements to be able to build a root CA. The CA has several properties that must be configured in the root certificate, such as Country, City, Organisation name, email and others. These all need to be configured before the CA creation process is started. Out of the box, the **vars.example** (located in the **easy-rsa** folder) will need to be renamed to just **vars**:

cp easy-rsa/vars.example easy-rsa/vars

It is really important that the name is **vars** otherwise the deployment wont work as expected. Open the file: **vi easy-rsa/vars**

To initialise the PKI infrastructure, use the following

QUICK TIP

If you want a simple way to use certificates on an internet-based webserver and enable HTTPS encryption, there is a very easy route: using certbot. This can be installed following the instructions at https://certbot. eff.org.

Creating certificates **TUTORIALS**

command. It initialises the infrastructure for first use. cd ~/easy-rsa

./easyrsa init-pki

Within the file there are several items that need to be uncommented and configured, substituting appropriate values as in this example:

set_var EASYRSA_REQ_COUNTRY "GB" set_var EASYRSA_REQ_PROVINCE "London" set_var EASYRSA_REQ_CITY "London" set_var EASYRSA_REQ_ORG "My super local CA company" set_var EASYRSA_REQ_EMAIL "me@myemail. local" set_var EASYRSA_REQ_OU "My Organizational Unit" set_var EASYRSA_REQ_OU "My super local CA" set_var EASYRSA_NS_COMMENT "My super local CA"

set var EASYRSA CERT EXPIRE 365

Most of the settings are obvious. Once edited, save the file. These can be overridden on the command line if required (more information on easy-rsa can be found at https://easy-rsa.readthedocs.io/en/latest). Now is the time to run the process to build the root certificate server, using the command _/easyrsa build-ca. Note that there are several other options within the file but to keep it straightforward we are only including a subset of most often used. Other changeable values include the length of time the certificate is valid for, encryption schemes and other advanced settings.

The only question asked is the key phrase. This is the password that will be needed to create and revoke certificates for servers. Make sure it is secure and write it down somewhere safe, just in case. If the passphrase is forgotten the only option is to redo everything from the start! It is possible to bypass the password using **nopass** as an option specified on the **build-ca** command, but this is not recommended for obvious reasons. Note that there seems to be an issue with using *easy-rsa* complaining it cannot find the **.rnd** file. If this error is thrown, use the following command from within the **easy-rsa** folder. This should fix this issue: **dd if-/dev/urandom of-pki/.rnd bs=256 count=1**

A matter of trust

It is important to understand how the certificate system works. At the top level of the tree is the root Certificate Authority (CA). As part of the *easy-rsa* CA setup it outputs a digitally signed root certificate (**ca.crt**) and the private key for the root ca (**ca.key** located in the **pki/private** folder). This signed certificate can then installed in web browsers. Once it is installed it signifies

		Certificate M	anager		
		Downloading	Certificate		8
You have bee	n asked to trust a ne	ew Certificate Au	thority (CA).		
Do you want Trust this (Trust this (Trust this (Before trusti procedures (i View	to trust "ca" for the CA to identify web si CA to identify email ng this CA for any pu f available). Examine CA certif	following purpos ites. users. urpose, you shou icate	ies? Id examine its ce	rtificate and its p	olicy and
				Cancel	
View	Edit Trust	Import	Export	Delete or Dis	itrust
<u>v</u> iew	Edit Trust	I <u>m</u> port	E <u>x</u> port	Delete or Dis	trust

that the browser trusts our certificate authority and that by extension, any other certificates that we digitally sign. To install the root certificate is simple enough. Copy (by whatever means) the **ca.crt** file from the **pki** folder. Open *Firefox* and navigate to the settings option. Within the settings, click the 'Privacy and security' portion and scroll to the bottom where there is a Certificates section. Select 'View Certificates', clicking Import from the buttons at the bottom of the dialogue and select the file in question. When asked about how to trust this certificate select only 'Trust this CA to identify web sites'.

This will need to be repeated for all clients that need to accept certificates signed by our new certificate server. You have a fully working CA server. At this moment our browser trusts us, but we trust no-one.

In order to start trusting the computers and devices on the network, a CSR or certificate signing request is needed. A CSR is a package of data formatted for submission to the Certificate Authority and if we chose to accept this request (as the CA owner), we sign the request with the CA private key and then send back the certificate for the approved server. This is the important bit of the CA's role.

If the CA signs the request it is essentially extending its trust to the certificate owner. This is why CAs need to be extremely careful about validating and verifying the requests. To press home that point, not all certificates are created equal.

There are DV certs (to validate a domain), OV (used to validate organisations) and lastly EV (Extended Validation). Each one increases both the scrutiny, time and cost involved, but for example, financial Adding in the root certificate for the signing authority created. Without this the new certificate won't work

QUICK TIP

If the CA is set up and something is amiss, it is possible to recreate the entire certificate authority from scratch, just re-run the ./ easyrsa init-pki command. It will ask to confirm that you want to reset it all by typing YES at the prompt.

» NOT ALL CERTIFICATE AUTHORITIES ARE EQUAL

As we have discussed, certificate authorities can sign many different types of certificates. However, a root CA should only ever really be used to create subordinate CAs. While this adds more complexity, the reason for this is that should one sub CA be compromised for whatever reason, it can be replaced or removed as needed without having to update all browsers and apps globally.

These secondary CAs are referred to as intermediate CAs. For example, *Firefox* trusts the intermediate CAs, which in turn trusts the root CAs. These intermediate certificates are able to sign certificate requests but, when configured properly, won't be able to create further root certificate authorities.

This is known as the chain of trust and is usually provided when a certificate is created. In our example walkthrough we are not using intermediate CAs so there is no chain of trust to be entered, as the immediate parent is the root CA.

TUTORIALS Creating certificates

organisations would very much prefer to use EV because the level of checking and verification is that much more thorough and in-depth.

This tutorial is purely focusing on domain validation. In order to create the CSR there needs to be a copy of OpenSSL. It's probably already installed but just in case: sudo apt-get install openssl -y

Creating a certificate signing request consists of three separate steps. First, create a new private key. mkdir signing && cd signing

openssl genrsa -out webserver.key

After that we create the CSR using our newly created key (be sure to backup that key somewhere safe): openssl req -new -addext "subjectAltName = webserver.key -out webserver.req

It is prudent to mention that there is more than one name for the server. For example **webserver** and **webserver.mysite.local** are two different entities as far as certificates are concerned.

The above command will ask a series of questions that need to be filled in. The important property is the Common Name. This is the FQDN (Fully Qualified Domain Name). If you get this wrong, it won't work! It is

has not been cryptographically verified. Please be sure it came from a trusted source or that you have verified the request checksum with the sender.

Request subject, to be signed as a server certificate for 1080 days:

subject=	
countryName	= GB
stateOrProvinceNa	me = North West
localityName	= Manchester
organizationName	= Newly Minted Media Ltd
organizationalUni	tName = IT
commonName	= webserver.itburns.local
emailAddress	= stuart@newlymintedmedia.com
Type the word 'yes' t Confirm request det	co continue, or any other input to abort. cails: ves
Using configuration t Enter pass phrase for Check that the reques	rom /home/sysadmin/easy-rsa/pki/safessl-easyrsa.cnf //home/sysadmin/easy-rsa/pki/private/ca.key: st matches the signature
The Subjectic Diction	wiched Name is as fellows
countryName	+DDINTARI F. CR
stateOrProvinceName	:ASN 1 12:'North West'
localityName	:ASN 1 12: Manchester'
organizationName	ASN 1 12: Newly Minted Media Ltd'
organizational UnitNam	ACT ASN 1 12: 'TT'
commonName	:ASN 1 12: webserver ithurns local'
emailAddress	:TASSTRING: 'stuart@newlymintedmedia.com'
Certificate is to be	certified until Aug 11 14:13:17 2024 GMT (1080 days)
Write out database wi	th 1 new entries
Data Base Updated	

Certificate created at: /home/sysadmin/easy-rsa/pki/issued/webserver.crt

The process of signing the request so that it can be incorporated into the webserver for use with HTTPS communication.

possible to pass all these parameters to the *openSSL* command directly rather than manually filling it in.

If you look at the request in a file editor it will appear to be nonsensical data, apart from the certificate start and certificate end statements. To view the information it contains (to verify it is correct) use the following: opensal reg -noout -text -in webserver.reg

Now that the certificate signing request file has been generated (the **webserver.req** file) the next step is to copy the request to the CA server. Again, do this by any means but it's just as easy to use the built-in *scp* command. Shown below is an example of copying the file with *scp*:

scp ~/signing/webserver.req sysadmin@10.0.0.217:/ home/sysadmin/

Sign here please

All that remains is to import and sign the request with our CA key. It should be obvious that this key should be kept secure. It is acceptable for a local installation to just 'hide' it, but for commercial CAs the security concern is much greater, given the havoc that could be released if it was lost.

The following should be done on the certificate server. The first step is to import the request into the database that *easy-rsa* created earlier. At the risk of stating the obvious, below, the import request takes the format of a filename and the second variable is a name that gets recorded in the CA database. The following needs to be done on the CA server.

cd ~/easy-rsa

./easyrsa import-req /home/sysadmin/webserver.req webserver

At this point the request has been imported. All that remains is to sign the request (approve it in other words). Do this by using the following command. Don't conflate the **server** and **webserver** items, the first **server** identifies what type of signing is being performed, in our case server signing.

./easyrsa sign-req server webserver

All being well, the newly created certificate will have been created and placed in the **/issued** folder of the *easy-rsa* **CA** folder.

That fully signed request can now be passed back for inclusion in the webserver. To copy it back to the home folder on the webserver use *scp*:

scp ~/home/easy-rsa/pki/issued/webserver.crt
sysadmin@webserver:/home/sysadmin/

» TYPICAL CERTIFICATE-RELATED ERRORS

Certificates have a start and end date. When the browser connects to a web page or even if an **apt-get update** connects to a digitally signed update repository, it checks the local machine time against the certificate validity date and if the time and date are out by more than a small margin the connection will drop and you will see an error. In these situations the first thing to do, once you have checked the remote server is actually available, is to check the date and time on the local machine in question. Assuming the time and date are correct it could also be that the certificate in question has really expired. In this scenario there is nothing you can do until the certificate is renewed.

From the implementer side, some of the typical errors include failing to include Subject Alternate Names. When creating a certificate there should be three items that are included in the request. The first is the FQDN, or in other words the full web address; the second is the hostname (for when someone just types in https://webserver) and lastly the IP address, depending on the situation. This last one isn't always needed and may actually cause errors depending on your configuration. For example, if someone were to include the IP address in a shared hosting environment it wouldn't work as expected. In short, be sure what you need in your certificate.

Creating certificates **TUTORIALS**



Creating the initial CA configuration and doing the bulk of the work with easy-rsa.

Adding it into an example Apache server is next. Install Apache sudo apt-get install apache2 -y . It is important to note that any private key that is used with Apache (or any other web server software) needs to have no password. The reason for this is that if Apache tries to load that certificate it will wait until the password is provided, so if the server is rebooted, Apache will sit there and wait until the password is entered and the webserver won't serve web pages.

Create a folder for the certificates and copy across the files required:

sudo mkdir /etc/apache2/certs

sudo cp ~/signing/webserver.key /etc/apache2/certs/ sudo cp ~/webserver.crt /etc/apache2/certs/

To enable HTTPS for our webserver there is some background work that needs to be done. Firstly, enable SSL support for Apache:

sudo a2enmod ssl

sudo systemctl restart apache2

Assuming Apache2 restarted successfully, the configuration file needs to be edited to reflect our change. Use nano to edit /etc/apache2/sitesenabled/000-default.conf sudo nano /etc/apache2/ sites-enabled/000-default.conf and modify it so it looks similar to the one below. Make a backup first, just in case there is a need to go back:

<VirtualHost *:80>

ServerAdmin webmaster@localhost ServerName webserver DocumentRoot /var/www/html ErrorLog \${APACHE_LOG_DIR}/error.log

CustomLog \${APACHE_LOG_DIR}/access.log

combined </VirtualHost>

sysadmin@ca:-/easy-ris5 ./easyrsa inport-req /hone/sysadmin/w ./easyrsa: 97: ./vars: set_Var: not found Note: using Easy-RSA configuration from: ./vars

sing SSL: openSSl OpenSSL 1.1.1f 31 Mar 2020

he request has been successfully imported with a short name of: webserve ou may now use this name to perform signing operations on this request.

ou may now use this name to perform signing operations on this reque ysadmingca:-/easy-rss\$ []

Importing the signing request into the CA Database. Once imported the data can be viewed and manipulated if required.

<VirtualHost *:443>

ServerAdmin webmaster@localhost

ServerName webserver

DocumentRoot /var/www/html

ErrorLog \${APACHE_LOG_DIR}/error.log

CustomLog \${APACHE_LOG_DIR}/access.log

combined

SSLEngine on SSLCertificateFile /etc/apache2/certs/webserver.

crt

SSLCertificateKeyFile /etc/apache2/certs/ webserver.key

</VirtualHost>

Create the folder **certs** and copy both the **webserver**. **key** file and the **webserver.crt** file into it. Restart the Apache2 service using:

sudo systemctl restart apache2

Now, using *Firefox* on one of the local computers that has the root certificate installed, browse to the webserver using the FQDN, substituting HTTP for HTTPS. If everything has worked it should now be possible to browse using HTTPS and seeing the little padlock icon indicating a secure connection.

At this point we have covered the basics of creating our CA and creating certificates and trusting them. In the real world certificates can be lost or compromised. That is where a CRL or certificate revocation list comes in. To revoke a certificate is straightforward. Implement it by using the following command (again, on the certificate server):

./easyrsa revoke webserver

The CRL works as part of the connection process. When a client is offered a certificate it checks the certificate against the CRL list that the CA maintains. If the certificate is revoked it will throw an error in the browser warning of this.

In summary, all being well you should now be able to create a certificate authority, install the root certificate, create and sign SSL certificates, install them in an Apache server and also revoke them as needed.



» IMPROVE YOUR LINUX SKILLS Subscribe now at http://bit.ly/LinuxFormat

OUICK TIP

TUTORIALS Explore mapping data



Create, edit and use custom mapping data

If you like Google Maps you're going to love QGIS. **Mike Bedford** shows you how to get started and introduces some of its functionality.



Mike Bedford is a big fan of QGIS. However, he admits to still loving good oldfashioned paper OS maps while he's walking on the hills. B Republic discovered a mammoth's tusk that has been engraved with a geometrical design. The pattern was subsequently identified as representing the hills, rivers, valleys and routes of the surrounding area. It was dated around 250,000 BC, making it probably the earliest map ever discovered. In the library of Hereford Cathedral is the Mappa

ack in 1962, archaeologists in the Czech

Mundi. Created around 1300, it's the largest surviving medieval map of the world. The UK's Ordnance Survey was set up in 1791 to help address the perceived risk arising from the Jacobite rising, and of invasion by France, and it went on to become one of the world's foremost national mapping agencies. And 1999 saw California-based ESRI release its *ArcGIS* graphical information system, which would become the world's most used GIS (geographic information system).

These few notable dates in the history of cartography illustrate 27,000 years of change and it's clear that the rate of change has accelerated significantly in recent years. This isn't surprising since digital technology has had the same effect on so many other areas, but it's been suggested that a GIS is so different to its predecessors that it's a totally new entity. Professionals point to support of multiple layers and an ability to carry out analyses of geographically based data as just some of the things you can do with a GIS but not with an ordinary map. And if you're wondering where Google Maps fits in, while we're not going to get embroiled in the debate over whether or not it's a GIS, we would agree with the sentiment that a fully blown GIS is like Google Maps on steroids.





From street maps through topographic maps to aerial imagery, any number of options are available as base maps.

ESRI's well-respected *ArcGIS* isn't cheap at £1.192 per year, including VAT, for professional use. This reduces to £139 for non-commercial use, but that's still a lot unless you have a fairly serious application. Our subject here, therefore, is the free open source *QGIS* which is widely used professionally and is considered to be on par with many professional GIS offerings.

We're going to provide a hands-on introduction to using QG/S but we're certainly not going to be looking at all its features, in fact we won't come close. After all, if you can think of something you believe you ought to be able to do with mapping data, it's a fair bet that QG/S can do it, either natively or via one of its many plug-ins. However, we trust that our tentative first steps will be enough to excite you about the possibilities of QG/S and start you off on your own voyage of discovery.

Orientate yourself

If your main experience with digital mapping has involved Google Maps, the first thing you'll notice when you start up QG/S for the first time is that you don't see a map at all. Instead, unless you open a project that someone else has produced, you'll need to create a new project but, even having done that, you'll be faced with a blank canvas to which you'll need to add data as layers.

Usually you'll want to start a project by adding a base map so, let's see how to do that, but first create a new project at Project>New. Now, in the Browser panel at the top-left, find and expand the XYZ Tiles entry and

QUICK TIP

When you work on a project, you'll end up with several files including, for example, files associated with the layers. For this reason it's good practice to use a new folder for each project and keep all of the project files together.

> Adding OpenStreetMap data as a base map couldn't be much simpler.
Explore mapping data TUTORIALS

double-click OpenStreetMap. This will add the open source OpenStreetMap as a new layer. The map will appear in the main window and you'll also be able to see it listed in the Layers panel at the bottom-left.

Although OpenStreetMap is the only base map shown by default, there are lots of others you can use, and you might like to familiarise yourself with what's on offer. You'll have to search out the URLs yourself, but as an example, right-click XYZ Tiles in the Browser panel and select New connection... In the Connection Details dialog, enter https://tile.opentopomap. org/%7Bz%7D/%7Bz%7D/%7By%7D.png as the URL,

give it a name such as OpenTopoMap and click OK.

OpenTopoMap will appear below OpenStreetMap under XYZ Tiles and, if you double-click it, OpenTopoMap will appear as a second layer. In fact, it'll appear to have replaced OpenStreetMap in the main window, but that's only because it's the top-most layer: it appears above OpenStreetMap in the list, and it's fully opaque. We'll look at opacity later but, for now, note that you can swap to seeing the OpenStreetMap again by clicking the tick box next to OpenTopoMap so it's not displayed, or by dragging OpenStreetMap to the top of the list.

This method of adding base maps isn't limited to maps in the normal sense of the word. You can also add satellite or aerial imagery using the same method. To add Google imagery, for example, use https://mt1. google.com/vt/lyrs=s&x=%7Bx%7D&y=%7By%7D&z=%7Bz %7D as the URL.

Elevation data

So far, we've seen how to add various layers to a QG/S project, although all of those layers were essentially images of one type or another. But since a GIS is concerned with far more than just images, or maps if you prefer, let's take a look at how other types of data can be added.

To start we're going to add some elevation data, but first a few words on types and sources. The types you'll mostly find are DSMs (digital surface models) and DTMs (digital terrain models). DSMs include objects such as trees, buildings and cars, as well as the actual ground, while in a DTM the former types of object are



filtered out so it represents only the topography. In the UK, elevation data has been produced by the Environment Agency at a resolution of 1m or 2m, and is distributed freely online by DEFRA. Over the Pond, the USGS makes elevation data available for the US at resolutions ranging from 10m to 30m, but with the aim of migrating to 1m. They also provide Space Shuttlederived data for most of the world at 30m resolution. Commonly elevation data is provided as a GeoTIFF file, which is a geo-referenced variant of the familiar TIFF image format, although there are other formats, and *QGIS* supports most of them.

Once you've downloaded the elevation data it's time to import it into a QG/S project, ideally one into which you've already added a base map. Select Add Layer from the Layers menu and then choose Add Raster Layer... In the dialog, select the elevation data file by clicking the three dots against Raster Dataset(s) and clicking Add.

Because the elevation data might cover a smaller area than the base map, it's possible that you won't see the new layer on the map, even though it appears in the Layers panel. If that happens, just right-click the name of the layer containing the elevation data and select Zoom to Layer. Your elevation data will often appear as a rectangle, because you'll have downloaded a so-called Freely available elevation data can be added to a project and formatted in so many ways.

» QGIS OR MGIS? TIME TO LOOK BEYOND EARTH...

If you were wondering what QGIS stands for, the Q is for Quantum – or at least it was until 2013 but we guess now Q just means Q – and the G stands for Geographic which includes the "geo" suffix that relates to the Earth. As we proved in our experiments, though, it might be equally appropriate to call it QMIS, since we used it to display Martian data. Indeed, such systems have been used with data representing a wide range of non-terrestrial astronomical objects.

Interested in the wider Solar System? Then you'll probably want to investigate the various planetary base maps and digital elevation models that are freely available. However, you're going to need

to delve into the subjects of CRSs that's coordinate reference systems although it's a subject you really need to be familiar with whatever your use of OGIS. Needless to say, data representing anywhere other than the Earth won't use any of the CRSs devised for terrestrial use. And we have to admit that we struggled with one particular Martian base map which loaded, but wouldn't display. However, we did do better with elevation data so that might be a good place to start. Search for "Mars MGS MOLA_DEM_mosaic_global_463m" and you'll find that it can be downloaded from https://astrogeology.usgs.gov. And if the moon is your thing, you might be

interested in https://lunaserv.lroc.asu. edu, which hosts data from the Lunar Reconnaissance Orbiter Camera and provides information on how to use it in QG/S.



From Io to Titan, maps and elevation data is available for lots of heavenly bodies and visualised in QGIS.

TUTORIALS Explore mapping data

QUICK TIP

With v3 DOIS now offers a 3D visualisation feature This is commonly seen in online mapping applications, but considered a novelty by some GOIS professionals, . The 3D view can be manipulated by zooming, panning, rotating, tilting and so on.

tile, although it's possible it might be less than a complete rectangle because coverage isn't always 100 per cent.

Initially the elevation data might not look too impressive. It's just in shades of grey with white representing the highest elevation in the tile, and black representing the lowest. In passing, this means that if you have two elevation layers, they'll be scaled differently, but this can be solved by merging them using Raster>Miscellaneous>Build Virtual Raster...

The uninspiring greyscale image is just a start and *QGIS* has plenty of tricks up its sleeve. For a start let's just change the colours: right-click the layer and select Properties. Now, in the Layer Properties dialog, ensure that Symbology is selected on the left, under Render type choose Singleband pseudocolor instead of Singleband gray, and select one of the options for Color ramp. You might not find that any of the Color ramps are to your liking, and none correspond to the common colouring of altitudes on topographic maps, but you'll notice that you can edit each of the colours.

Moving away from just changing colours, Hillshade is another interesting option for Render type. It'll revert the colours to greyscales in the process, but there's a solution to that. As before, use the Layer properties to select your preferred colour scheme. Next, with the elevation layer selected, select Analysis>Hillshade... from the Raster menu, before accepting all the defaults in the dialog and clicking Run. Unlike last time, the latter won't be converted to a hill-shaded version, but a new hill-shaded layer will be created.

As always, the new layer will be at the top of the list, so drag the coloured elevation layer to the top in the Layers panel and, of course, it will obscure the hillshaded layer. However, adjusting the transparency of the pseudo colour layer will give the best of both worlds. This option is available in Layer Properties.

And as a final way of processing elevation data, a few words on contours. Depending on your choice of base map, you might already be able to see contours. However, if you're using a base map without contours, you can generate these from elevation data at Raster>Extraction>Contours....

As this example of water depth during a flood illustrates, how any geotagged data can be used in QGIS.

Beyond conventional maps

So we've seen how QGIS handles data that's basically an image, or in other words a map, and we've seen that it can process elevation data in so many ways. But this





Points can be added to a project and data fields can be defined for subsequent display or analysis.

is far from the limit of its data-handling capabilities. *QGIS* can work with any type off data so long as it's georeferenced. So, for example, this could enable you to see the variation in rainfall or hours of sunshine across a region, or it could provide a means of visualising the correspondence between the geological classification of rocks and surface features.

A recent project involved showing how QG/S can operate with other types of data, and via a brief return to elevation data, how it can carry out calculations on that data. The exercise involved using rainfall and elevation data to simulate areas at risk from flooding. The simulation software could input raw elevation data, so no need for QG/S there, except for the fact that elevation data for the area of interest was incomplete. In particular, there were holes in the preferred 1m resolution data, but 2m data was available for the area.

QGIS was used, therefore, to compile a file of elevation data using 1m resolution where present, and filling the gaps with 2m data where it wasn't. This was achieved by loading both sets of elevation data as separate layers and then using QGIS' raster calculator (Raster>Raster Calculator) to combine them. This was done using the statement ("1m_dtm@1" = 0) * "2m_dtm@1" having first ensured that the **no data** value in the 1m elevation data file is 0. It generates a new raster layer with values defined by the expression, and it's fairly clear how it achieves the desired effect.

This is only presented as an aside, but the ability to carry out any arithmetic operation on raster data is a powerful tool. *QGIS* then came into play again when the simulation software had crunched the input data and created a geo-referenced output file of water depths. By importing this to *QGIS*, and viewing it, partially transparent, over a base map, at-risk areas could be clearly identified.

Lines, points, photos

Informative pushpin labels are a familiar sight on online maps, and *QGIS* doesn't disappoint here. This sort of information can be created elsewhere and imported as a new layer, or it can be created within *QGIS*, as we're about to see. As always, you need at least a base map in your project before proceeding.

In the Layer menu, select Create Layer>New Shapefile Layer... In the dialog, give the layer a name and choose a type, which must be Point, Multipoint, Line or Polygon. Now start adding fields, and there can be as many as you want. In our example, we're creating a layer

Explore mapping data TUTORIALS

that shows places at which particular minerals were found. Our first field, therefore, is the type of mineral, so we give it the name "type", and specify that the field is defined as text. Having filled in the first field, click Add to Fields List and you'll see it appear in the list below. We also added a second field, namely the date on which the mineral was found, in just the same way. When you've defined all the necessary fields, just click OK and you'll see that the layer appears in the Layers panel, although nothing will appear on the map, because there are no points defined yet.

To add points, make sure the new layer is selected and click the Toggle Editing icon (yellow pencil) in the Data Source Manager toolbar. Once you're in editing mode, click the Add Point Feature icon (three green dots and a crosshair), and the cursor will change to a crosshair. Click it on the map at the place of interest and the Feature Attribute dialog will appear. You might not want to bother with the "id", but fill in the various other fields and click OK. A coloured dot will appear on the map. Now click the Toggle Editing toolbar again to turn off editing and see how a user can interrogate that point. Click the Identify Features icon (the "i" on a blue circle with an arrow cursor) in the Project toolbar. Click a coloured dot and a window will appear providing data on all the fields defined for the layer.

You might find it surprising that some of that data doesn't appear alongside the point without interrogating it, but since there's no limit to how many fields you can assign, it's fairly obvious why that doesn't happen by default. However, you can choose to add labels for all points on a layer, which you'll do by clicking the Layer Labelling Options icon ("abc" on a yellow arrow) on the Data Source Manager toolbar. In the Layer Styling dialog, ensure the Labels tab is selected (the "abc" on a yellow arrow again), choose Single labels in the top box and choose the field you want the point to be labelled with against Value. You'll see the point labelled on the map. You might want to fine-tune the options under "Placement". Also, if the size, shape and colour of the point aren't to your liking, this is something else you can adjust in the Layer Styling dialog, this time on the Symbology tab (the paintbrush).

QG/S also enables you to show photos at the location they were taken, but only if they're geolocated. So, if they were taken on a phone you'll be okay, although if you took them on a camera they might not be. Bizarrely, given that QG/S doesn't seem to be short on features,



As an alternative to points, places of interest can be shown on a map using geo-tagged photos. we could find no way to import non-geotagged photos. That's not a show-stopper, though – it just means you've got to add the geotags separately. There are plenty of standalone applications and web-based utilities to do that (for example, **www.tool.geoimgr.com**), although many limit how many photos you can process a day without paying a fee.

So, with a folder full of geotagged photos at the ready, here's how to import them into a *QGIS* project. In the Processing menu, select Toolbox>Vector creation>Import geotagged photos. Against Input folder select the folder containing the photos you want to add and click Run. A new layer will appear in the Layers panel and you'll notice some small markers appear on the map, at their correct locations, just like the ones we added previously.

However much you zoom in, though, they won't appear as photos unless you make some changes. Bring up the Layer Styling dialog and ensure the Symbology tab is selected. Single Symbol will be shown and below that the marker type will be shown (expand it if not already expanded to show Simple Marker). Click Simple marker and then, against Symbol Layer Type, choose Raster Image Marker instead of Simple Marker. Lower down there's an unnamed box where you can define a particular raster image, but instead we want it to pick up the image at that point. So, click the icon at the right, which looks like a filing cabinet and a couple of arrows. to allow the Data Defined Override option to be selected. and against "Field type: string", choose "Photo (string)". Photos will replace the blobs, but they will be very small so specify a new size against Width.

QG/S is a very powerful tool and what we've presented here is only a taster to what it can achieve. Experimentation is one of the best ways of learning, so what are you waiting for?

» ANALYSIS OPPORTUNITIES

Although the ability to analyse geographical data is considered a key feature of a GIS, we've largely ignored this feature. As you grow in your confidence with QG/S, though, depending on your application you might want to delve into analysis, and here the sky's the limit.

For a start, there's lots that can be done with elevation data. Mobile phone companies use GIS software to visualise the coverage that would be achievable from base stations at given positions. This sort of exercise can be used to decide where to locate base stations and forms the basis of published coverage maps. But it goes much beyond that, a fact that should be obvious when we think about the Shapefile layers that we created and populated with points. There can be any number of such layers, and the points in each of those layers are defined by their geographical location plus data in several fields.

We can, therefore, start to think about a *QGIS* project as a geotagged database and the potential starts to become clear. A common approach is to carry out some sort of analysis and create a new layer from the results that can be viewed on the base map. The *QGIS* documentation shows an example of an exercise that involved analysing the cause of a cholera outbreak in London in 1854. And while that exercised used fairly simple analytical tools, *QGIS* can also be interrogated using SQL queries.

>> MAP OUR SHAMBOLIC SUB'S SYSTEM Subscribe at http://bit.ly/LinuxFormat

With some actions that create new layers, those layers will be temporary scratch layers. The method we described for creating a layer with photos is one such case. These layers will be lost when you close a project **Calthough QGIS** will warn you), but you can prevent this by right-clicking the layer and selecting Make Permanent.

OUICK TIP

TUTORIALS Steam gaming



Get better Steam and Proton gaming

Michael Reed looks at what it takes to run a large variety of games under Steam, including those designed to run on Windows.



Michael Reed used to think he was doing well if he could game in 2, 4 or 8 colours. It was all simpler back then ..

f you're at all interested in gaming on PCs, you've probably come across Steam, Valve's platform for distributing, updating and running games. Steam makes it possible to purchase a game. install it over the internet and then run it from the Steam interface.

Ah, but that brings us back to the age-old Linux gaming conundrum of support, as not every PC game is designed to run on anything other than Microsoft Windows. That said, there are plenty of Steam games that will run on Linux, and quite often, a Windows Steam game can be convinced to run on Linux even though some 'fettling' by the user may be required. This approach is officially supported by Steam using a system called Proton.



THE STEAM INTERFACE

Main menu

Easily overlooked, this is the menu bar. It's mostly concerned with options to configure Steam itself

Steam Sections

2 The tabs that take you to different parts of Steam. Each entry is also a pulldown menu if you hover over it.

Game library controls

3 Three filter icons to flick between: installed games, owned games and native Linux games.

Search

The search box. If you need this, you may have too many games! Game library

5 The list of games in your library. Left-click to be taken to the game's page. Right-click for the properties menu.

Run a game

6 A button that changes between play, install and stream functions.

Guilt-o-meter

7 Total play time. Keep this covered when your boss, partner or parents are present.

Getting all of these things running, and then possibly optimising the results, is what we'll be looking at in this tutorial.

Steam powered

Steam itself is installed through a custom program called the Steam Installer. This makes sense because Steam updates itself and the games you install with it without relying on Linux's own update systems. These days, the installer is in the official repositories of many Linux distributions such as Ubuntu and Fedora and their derivatives. As Steam is proprietary software you may have to enable a specific repository, such as 'Multiverse' for Ubuntu or 'Nonfree' on Fedora. Having done this, you can, for example, install Steam on Ubuntu as simply as typing sudo apt install steam-installer

If you're running a DEB-based distribution, but you can't find the Steam installer in the official repositories. you can obtain the installer directly from the Steam website (steampowered.com) and install it with the usual dpkg -i [name of archive] as the super user.

There is also another, non-official, way of installing Steam that all Linux users might find interesting, and that is installation via the Flatpak system. This offers a few advantages of its own in terms of privacy and sandboxing, as Steam is a system unto itself once it's on your system. A standard Steam installation is probably safe to use in the vast majority of cases, but if you are concerned (or if the other installation methods don't work), give the Flatpak variant a look (search for 'Steam' on https://flathub.org).

Games in Steam

Let's get started with an example that we can use to explore some of the features of Steam. Use the search feature at the top of the Store page to search for OpenTTD, a free management game, and click it to go to the product page for that game. Underneath the preview images and movies we find a bar with some pertinent information. The first point of interest is the price, which in this case is 'Free' (yay!). There are also three icons that indicate that the game is compatible with Microsoft Windows, macOS and Linux. Actually, that final icon, a piston arm over a wheel, indicates that the game has native support for typical desktop Linux distributions in

Steam gaming **TUTORIALS**

addition to SteamOS, Valve's custom, game-orientated Linux distribution. Valve sells specialised devices that run SteamOS, including the recently announced (and rather tasty-looking) Steam Deck hybrid console. Developments like these are good news for gaming on desktop Linux because it incentivises Valve to keep up the support.

It's worth scrolling further down the product page to determine what the recommended specifications are and make sure that your machine meets them. In the case of *OpenTTD*, they're fairly light.

Installation options

As the game's free, we could begin the installation procedure, but there are a few points to make about installation options in advance. It's quite common to end up with a system that has more than one hard drive installed, and for this reason, you might prefer to install games to a drive other than the system drive.

To begin to configure this, the option that you're looking for is called Steam Library Folders, and it's located in Steam > Library > Steam Library Folders. From here, you can add an alternative installation folder located anywhere in the Linux filesystem that you have permission to access. The great thing about doing this is that the default game installation location is still available as an option in the installation dialogue. So, you can make decisions about installation location



Getting ready to play Firewatch on a laptop without 3D support, thanks to Steam Remote Play streaming.



based on available hard drive space and performance requirements, particularly useful if you're rocking an SSD as the system drive. Note that you can install to other types of storage such as network drives or USB flash drives, but the performance tends to be poor in the case of most games.

To carry out the installation, click on 'Install' or (confusingly, in the case of free games) 'Play', near the top of the product page. Complete a test of the system by clicking on 'Play' once the installation has completed to check that the game runs properly.

As we've already said, Steam handles its own updates, both for itself and the games that you have used it to install. The updates for games are often huge, but you can monitor them, pause them or cancel them on the Downloads page (hover over Library and click Downloads). This is fortunate because, as well as consuming bandwidth, the updates are compressed and encrypted and this means that they use a lot of CPU time while they are downloading, in our experience keeping two cores fully occupied for the entire duration. Handily, you can exit Steam at any time, and updates will resume next time it is launched.

Windows games

But what if you want to run a game that doesn't offer a Linux version? In such cases, there is a good chance Enabling Steam Play within the Steam settings dialog so that we can run both supported and (officially) unsupported Windows games on Linux. Naughty!

QUICK TIP

Most Linux users like to know what's running on their system, but when you close the main window, Steam doesn't actually stop running, even though it's a fairly substantial application. Close it by leftclicking the icon in the control panel area of your desktop environment or by selecting (Steam > Exit) in the top menu.

» RUNNING GAMES REMOTELY

Steam can run games remotely between two Steam-equipped setups. The way it works is that the remote computer runs the game and compresses the video and audio. It then streams this video and audio to the other computer while the client sends back mouse, keyboard and controller input. The two computers that you use for this don't even have to be of the same architecture, and there's an Android app too. We tested it out between a Windows computer with a GNU/Linux one over a home network, and found that it worked very well indeed.

The performance-limiting factors are the speed of the host that's running the

game and the latency and bandwidth of the network, rather than the power of the client computer. So you could have a tricked-out gaming computer with a decent graphics card hosting the game, paired with a low-powered Linux box as the client.

Over gigabit Ethernet, it was possible to run an action-orientated game with no noticeable input lag at all while running the game at 1080p. There was some slight degradation of graphics quality as video compression was in use. However, we were able to alleviate this to an extent by increasing the graphics balance setting (Steam > Settings > Remote Play > Client Streaming Options To This Computer: Beautiful).



Streaming a game from a Windows PC onto a Linux PC. Hold down Escape for the streaming menu.

TUTORIALS Steam gaming

QUICK TIP

You can choose a specific version of Proton - worth trying if you're not getting quite the same results as reported on ProtonDB. Do this by rightclicking the game name in the Library tab. From here, go to Properties > Compatibility > Force the use of a specific Steam Play compatibility tool and trv matching the version number in the ProtonDB entry.

that it can be made to run using a system called Proton. An official part of Steam on Linux, Proton makes use of both *Wine* and DXVK, layers that translate Windows software and graphics calls into native Linux ones. However, there are some caveats. Firstly, compared with how the game runs natively on Windows, performance and stability might be an issue. Secondly, it does raise some ethical questions as, arguably, running the Windows version under Linux might reduce pressure on the developers to make a Linux native version of a game.

Getting back to the first point, the best way to get answers about the viability of running a particular game is to search for it on the ProtonDB website (www. protondb.com) to see what luck other Linux users had with it. The database itself is well organised and includes information about when the testing was carried out and what the specific hardware and software configuration used was.

The level of success of running the game on Linux is rated from 'Borked' to 'Platinum'. So, if you find a recent, highly encouraging report about the game running well on a similar set-up to your own, the odds are good. In addition, ProtonDB reports contain information on tweaks you can use to make stubborn games run properly. More on that in a moment.

To run Windows versions of games on Linux, visit the

» GET THE BIG PICTURE

Steam can offer a gaming experience that is ergonomically closer to console gaming through what it calls Big Picture Mode. This pushes Steam into a full-screen mode with controller support. We tried it out and we liked what it had to offer.

Big Picture Mode is invoked by clicking on a small icon at the top of the screen, and the first thing you'll notice is that every element of the user interface is jumbo-sized for better legibility on typical TV setups. You navigate around this interface using the mouse or a controller – all from the comfort of the sofa.

However, we found that the hardware support wasn't quite as slick as the interface undoubtedly is. For example, there is no clear way of specifying which screen should be used for the actual game playing. The most common workaround is to temporarily make the target screen your primary display, or set up display mirroring. Similarly, there was some work needed to get the controller working properly in the games themselves. You may have to experiment here to find a well-supported setup. Top marks for the interface, but room for improvement for ease of hardware configuration was our verdict.



The Big Picture interface is suited to hassle-free navigation from a distance.



Success! Due to his diligence and commitment to his work, the author put in a few hours of play time in Fallout 4 to make sure it was completely functional.

Steam Play settings (Steam > Settings > Steam Play). There are two options here to be enabled: 'Enable Steam Play for supported titles' enables approved games to run. These are games that Valve has tested and knows to work properly. Ticking 'Enable Steam Play for other titles' enables you to attempt to run all Windows games that you have access to through Steam, even if Valve has not tested the game. Again, examining the entry for the game on ProtonDB will give you the best idea of what to realistically expect.

Troubleshooting and tweaking

If you look through ProtonDB, you'll notice quite a lot of references to *Winetricks* (https://wiki.winehq.org/ Winetricks). So what is it? Winetricks is a helper script that installs various Windows components that are needed by games. The script pulls through the component and makes it visible to the *Wine* subsystem, and it's worth knowing about, even if you're using *Wine* outside of Steam. For example, if you typed winetricks corefonts it would download the Microsoft Corefonts, a set of fonts often needed by Windows applications, and install them to a location that *Wine*, and therefore Windows applications, can see.

Things get a bit more complicated if you want to apply Winetricks fixes to a specific Proton game, but not by much. Thankfully, there is a helper script for the helper script called Protontricks (https://github.com/ Matoking/protontricks) for just this purpose. If you have a full Python set-up (see your distribution documentation to install Python) you should be able install it with pipx install protontricks.

To apply Protontricks to a specific game, you need to discover the specific game ID of that game, using Protontricks itself. For example typing **protontricks -s** fallout found *Fallout* and *Fallout* 4 installed on our Linux gaming machine. This might seem like a lot of work, but we've got to admit that it's pretty cool, being able to add all of the resources that a game might need on a game by game basis.

This is necessary as different games might need different versions of a resource. You might even find that this flexibility makes some older games more viable on Linux than on native Windows, and this will be increasingly so as time goes by. The syntax is **protontricks [game ID] [Winetricks command]**, but most of the time, you'll simply be cutting and pasting the needed commands from the ProtonDB entry.

Moving on from Protontricks, some games need special launch options to be added to the game. The

Steam gaming **TUTORIALS**

OUICK TIP

SteamOS icon (a

piston arm over

a wheel) when

vou're looking

various parts

of the Steam

application. In

addition, you

can often filter

by platform on a

given area, even if you have to

scroll around to

find the filters.

games on

for Linux native

Look for the



Examining the compatibility reports for a Metro 2033 on ProtonDB. It looks promising.

dialogue for doing this is accessed by right-clicking the game name in the Library section in Steam, selecting 'General' and then entering the launch options that you've found on the relevant ProtonDB entry into the box. Most such options will end with the string **%command%** at the end. If you add more than one launch option at once, remove the extra instances of **%command%** as it should occur only once, at the end of the overall option string.

If you have to alter the INI files and such of the game, you can browse to the game folder by right-clicking the game name, clicking 'Properties...' and then 'Browse...' in the 'Local Files' section.

If any of that seems a bit confusing, don't worry; we're going to go through an example of getting an awkward game to run properly using these techniques.

An apocalyptic adventure

It's time to take a trip into the forbidden wastelands of lawlessness and chaos. No, we're not talking about rebooting into Windows. We're going to take a look at what it takes to get post-apocalyptic first person role playing game *Fallout 4* running under Linux. Don't worry if you're not a fan of that particular game (heathen that you are) as we've chosen this one because it provides a good example of getting something a bit tricky to run properly.

We had already purchased the game and played it quite a bit under Windows, and as it was already present in the Steam library, we installed it by selecting it and clicking the install button.

Remember, as this game isn't approved by Valve, you have to enable installation in the Steam Play settings, as detailed earlier on. This led to a 39GB download that took ages and made the CPU fan spin quite fast while it was doing so. Steam and Proton added all of the components that the game needed to run, so we didn't have to make any use of Protontricks in this instance.

At that point, all we had to do was to click on 'Play' to launch the game, and it did launch as expected, but there were some problems. Firstly, the sound was crackly and distorted for the introductory movies, and entirely absent for the game itself.

Secondly, the mouse kept moving in a crazy way, making the game unplayable as well as silent. Still, early signs looked promising for what is quite a technically demanding game. Leaving the game to find some fixes, we hit upon another problem, because it had hung rather than exiting cleanly.

Last things first, the hang-on-exit problem was solved, in the short term, by process-killing the running game. We did this by pressing Alt+F3 to open a virtual terminal. From here, we typed top to view all of the user processes on the system in Top, seeing Fallout4.exe as the first item. killall -9 Fallout4.exe killed the process, and pressing Alt+F7 got us back to the regular desktop. On subsequent runs of the game, we learned to close down cleanly by exiting to the main menu (rather than straight to the desktop) and quitting the game from there. Little workarounds like this are fairly common when you're running a Windows game under Proton.

Searching ProtonDB, it turned out that the sound problems we encountered are very common for Fallout 4. The solution was to cut and paste WINEDLLOVERRI DES"xaudio2_7=n,b" PULSE_LATENCY_MSEC=90 %command% into the launch options of the game. A quick jump back into the game confirmed that this had largely fixed the sound, but it took regressing to an earlier version of Proton, rather than the default experimental build, to make it work perfectly.

The mouse problem was a little bit more complicated to track down. Sure enough the suggestion, repeated over and over in ProtonDB, to add the line **bBackgroundMouse-1** to an INI file turned out to be the right one, but we initially had difficulty locating the file because it was duplicated in multiple locations, and it was impossible to figure out if what we were doing was having any effect.

Eventually, we discovered that we needed to alter the **Fallout4.ini** file located in the **compdata** folder, within the simulated **My Documents** Proton folder, rather than the identical file within the game directory itself.

At this point we customized a character and ventured out into the apocalyptic wasteland, experiencing performance that was similar to running the same game on the same hardware via the Windows partition. Success!



Download in progress. Keep an eye on transfers like these because they can consume a lot of bandwidth, CPU time and storage.

>> LET US PAY GAMES WITH YOU! Subscribe now at http://bit.ly/LinuxFormat



techradar

The home of technology

techradar.com

THE BEST NEW OPEN SOURCE SOFTWARE ON THE PLANET

HotPicks PDF Mix Tool » DroidCam » Fractal » Polyfoto ImprovedTube >> LSD >> Pam_usb >> Billiards » Blues Brothers » EasyEffects » Latte Dock

PDF TOOL

PDF Mix Tool Version: 1.0.2 Web: https://gitlab. com/scarpetta/pdfmixtool

ver wondered why it's not possible to edit a PDF document? That's because historically PDF was invented as a target format for viewing and printing and it never implied modifications. Of course, you can use Okular and Evince to leave some annotations, or open a PDF in Inkscape to disassemble it to separate objects, but none of the above resembles direct editing. Moreover, if you need to shuffle pages or extract parts from a PDF, you'll definitely need some sort of professional software.

Enter the gorgeous PDF Mix Tool, which we have just discovered. This application solves a lot of issues regarding page management as well as manipulating several PDF files at once. To be exact, PDF Mix Tool lets you merge files, extract pages (such as odd/even, or a custom selection), rotate them, change layout and even convert your PDF file into a booklet for two-sided printing. PDF Mix Tool is a versatile PDF machine with a very intuitive interface. It is powered by Qpdf, a contentpreserving PDF transformation software known to process PDF documents very accurately.

The main window has two areas: the left one with available actions, and the right area with the current action's details. All you need to do is to open one or several PDF files, set up processing options and hit the Save or Generate button (depending on the action). It is quite easy to do common PDF editing tasks in PDF Mix Tool, whether you need to split a document apart, merge files, crop page borders, or collate pages in a custom way. Combining the available actions enables even more complex things with your PDF files. Of course, the primary goal of the application is to ease pre-print routines if you want to have your PDF files on paper. However, nothing prevents you from shuffling pages of electronic magazines, creating collections of your favourite sections, or maybe splitting large PDF files.

with the ease and comfort provided - × by PDF Mix Tool.)-E aftere Miteen an.



Choose what to do with 1 your document(s)

PDF Mix Tool has a handy sidebar with lots of frequently used actions for PDF files. Merge, extract, delete, rotate and much more.

Provide the source document 2

This top panel differs across available actions, but normally it allows you to browse for one or several input files. No worries, everything is completely non-destructive.

Manage profiles

3 Choose Menu > Multipage Profiles to tell PDF Mix Tool how to handle multi-page PDF documents. The default profile plays well with classic A4 documents with portrait orientation.

Review document information

4 Jump to this action to quickly find out who created and when the PDF document was created. In many cases you can also find out the PDF generating software, such as Inkscape, Scribus or InDesign.

Manage page layout

5 Things like margins, page size and alignment can be easily changed in PDF Mix Tool in order to help you unify different input files, or prepare a document for booklet printing for example.





Edit PDE files

EXPLORING THE PDF MIX TOOL INTERFACE...

CAMERA APPLICATION DroidCam

Version: 1.8 Web: https://github.com/ dev47apps/droidcam

D roidCam is a spectacular piece of software that covers several categories in one go: it is hardware based, money-saving and fun. It's a software camera application designed to turn your mobile device into a network-attached IP camera. You need to have a smartphone, a Linux machine and Wi-Fi to connect both parts with each other, but it seems like almost everyone qualifies for these requirements.

Despite the name, which suggests that this application requires an Android device, you can use *DroidCam* with iOS, or between Linux machines without any mobile devices at all. The 'money-saving' idea implies that you don't need to buy an expensive IP camera to set up a home surveillance system – simply use your phone for that.

DroidCam provides an easy to use GTK interface that lets you quickly select the connection mode and get video in a few mouse clicks. It normally goes like this: install the application on your mobile phone using the Play Store or App Store, make sure the phone is connected to the same network as your Linux machine.



and finally provide the phone's IP address to the *DroidCam* application in Linux and press Connect. If everything was set up correctly (and no firewall is blocking port 4747), you should see the picture from the phone's camera.

The Linux client requires the **v4l2loopback** kernel module to be able to create a **/dev/videoX** device and receive the video stream. Make sure you have that module (**sudo modprobe v4l2loopback**) beforehand. Also, it is possible to test *DroidCam* without network connection at all using the USB connection. *DroidCam* supports any video input that works with the V4L subsystem in Linux, including various remote streams and even flatbed scanners. However, there is still more sense in using *DroidCam* for something more sane, like broadcasting or remote video control.

MATRIX CLIENT

Fractal

Version: 4.4 **Web:** https://gitlab.gnome.org/GNOME/fractal

e've not been paying enough attention to secure instant messaging of late, which is something that needs to be fixed. For the sake of open source purity and user liberation we need to discard many popular options including *Telegram* and *WhatsApp*, and look towards open IM options. Basically, the choice gets narrowed down to IRC, Jabber and Matrix, of which we believe the last one enjoys the rapidly expanding user base and growing popularity.

Here we have *Fractal*, a beautiful and very userfriendly desktop client for using Matrix networks. This application has been written in Rust and GTK and it sports a very clean and intuitive interface, with a list of chat rooms on the left and the current room contents on the right. If you haven't used Matrix before.you need to know that it is a de-centralised protocol for federated real-time communication, ideal for collaborative work, team chats and so on.

The Matrix communication model could not be more different to the client-server approach that we're used to. There is no direct way to simply 'send a message'.



Fractal is one of the best Matrix clients, with everything you need for group chats.

Breathe new life

into an outdated smartphone using video surveillance

powered by DroidCam

Instead, every member of a room constantly syncs their local room states and histories between each other to keep it all consistent and self-sovereign. In order to start using *Fractal*, you need to self-register by creating a local account on your home server, or, alternatively, register online at **app.element.io/#/register** and become able to join rooms on **matrix.org**. Naturally, a public server will require its own registration.

Once your Matrix identity is set up, hit the + button in *Fractal* and choose to join an existing room. The syntax looks like **#gnome:gnome.org**, where **#** is a room prefix. The registered user normally has their nickname prefixed with @ (@user:gnome.org). You can then chat with other people, share images and other files, invite other people to rooms or accept invitations.

Polyfoto

Polyfoto

Version: GIT **Web:** https://github.com/ shanedrabing/polyfoto

he photo mosaic is a fun and a technically advanced example of image manipulation. In most cases it can be created using commercial image editing software, or using online services – with limitations for size, number of images and so on. We've also seen a lengthy online tutorial on creating photo mosaic using *Adobe Lightroom*, where you ended up with a grid of images superimposed over the target image using the Soft Light filter.

Fancy doing a mosaic with a purely open source software, with no fuss or time-consuming steps? That's a good reason to try *Polyfoto*. This is a small Python-based utility which works like a charm: fast and predictable. As for prerequisites, only Python 3 is required, as the *Polyfoto* install script takes care of the rest. Go to the project's directory and use the following command to install:

\$ sudo python3 setup.py install

The setup script pulls all required dependencies, mainly the Python bindings to OpenCV. Once that's done, choose a source image to play with and provide a



directory stocked with other images that you wish to become tiles. You'll definitely need lots of such images, and that could be shots from your photo library, or maybe icons taken somewhere from **/usr/share/icons** (that's the easiest path). Anyhow, when all images are in their places, you can create a mosaic using the following command:

\$ python polyfoto.py -f /path/to/input.png -d /path/ to/source/images -o /path/to/output.png -n 64

The last parameter defines the number of rows. It makes good sense to increase it if you need more details, because the more rows you have the larger the mosaic resolution. The quality of mosaic also depends on the number and variety of source images. *Polyfoto* will play best if you have a lot of images with different colours and lighting.

Frame By Frame

•

approximates each source image to its prevailing colour and use it as a tile for the resulting mosaic.

BROWSER EXTENSION

ImprovedTube

Version: 3.232 Web: https://github.com/ code4charity/YouTube-Extension

ere we have another good reason to write about YouTube and open source, but this time it has nothing to do with YouTube-dl – the pervasive back-end that powers dozens of GUIs for video downloading. Instead, we'll keep watching YouTube videos using a web browser, but with a little helper called *ImprovedTube*.

This is a super-cool browser extension that adds a bunch of customisation options to the **youtube.com** website layout. *Chromium, Firefox* and all their scions are supported. The project description advertises *ImprovedTube* as having over 80 features that affect the looks and the behavior of YouTube. We were curious to give it a whirl, and it was really awesome!

When the extension is installed, its icon is greyed out until you open YouTube. Click the icon to reveal the menu and change something under any of the nine available categories. Dozens of useful settings reside there, and here are few examples. You can adjust the website layout, for example to make the YouTube logo smaller, tell the search bar to auto-hide and only pop up on mouse hover, apply a different

theme (for a darker background), collapse and hide comments and more. Some handy quality-of-life features include forced theatre mode, explicitly defined video quality or subtitles mode, channel trailers autoplay switch, reversed playlist processing and whatnot.

With ImprovedTube, you can add certain videos to the blacklist, gather statistics with the built-in analyser and even define custom shortcuts. The abundance of tweaks is breathtaking! People who want to concentrate on videos can remove or hide all the unneeded stuff, while those who need the "social" side of YouTube – likes, comments, subscriptions or stream chats – will also be able to customise their set up accordingly. Userdefined shortcuts alone is a killer feature. Make YouTube look and play the way you like it with the help of this handy toolkit.

IN O # 🖑 🖌 🖬 🗿

COMMAND LINE TOOL

LSD Version: 0.20.1 Web: https://github.

com/Peltoche/Isd

S ometimes we wonder if there is a way to bring our readers more happiness during these challenging times, especially with a help of open source software. Maybe that could be a nice addition to the terminal, something fanciful and fun, yet be of good practical use? We recently discovered *LSD* (short for *LSDeluxe*, in case you were wondering) and had a really inspiring time using it in a Linux terminal.

If you didn't know, *LSD* is a colourful replacement for *Is*, the GNU/Coreutils command for listing files. *LSD* is not a revolutionary next-gen thingy, but it still hopes to become your preferred *Is* equivalen, for a couple of reasons. One is that *LSD* is about colours.

It's not just about the looks of your files listings, but also their readability. Administering a server with lots of directories, each populated with a wealth of various files, definitely makes it a lot of a challenge to quickly locate certain files. *LSD* offers a good colour convention and helps you distinguish files from directories, objects with different permissions from one another, quickly locate empty directories and so on.

autorogenetation autorogenetation-

List your files and directories with this enhanced version of the traditional 'Is' command.

Secondly, LSD sports good integration with nerdy fonts and makes use of extra symbols and icons. In order to unleash the full power of LSD, you'll need to use some sort of 'nerd' or 'awesome' flavour of your favourite terminal font.

Chances are that such fonts are already available in the package manager of your Linux distribution. Other bells, whistles and terminal power-ups (such as *Powerline*) are welcome but not required to make *LSD* work properly.

Test your LSD with something like this:

\$ lsd . -l -h

Notice that each item in the list now has its own symbol reflecting its type. Thanks to this nice feature, *LSD* makes it visually easier to tell media files apart from code, directories from files etc.

LSD is written in Rust and seems to be very robust. If you like it, feel free to set appropriate aliases in your ~/.bashrc to always use *lsd* when *ls* is invoked.

SECURITY TOOL

Pam_usb

Version: 0.7.3 **Web:** https://github. com/mcdope/pam_usb

ecurity is a crucial part of computing, and even if you are only administering your localhost, it's always good to have strong passwords everywhere. Some people take extra measures to protect their sensitive data and use extra authentication methods, including hardware keys. A hardware key usually resembles a USB dongle and provides a trusted way of limiting access to software and hardware assets. The good news is that you can set up a similar access barrier and try it at home or at office without buying any security dongles.

Pam_usb enables ordinary USB thumb drives to become hardware keys to authenticate users in Linux. The software consists of the PAM module that you need to build from source, and user-defined settings with registered devices and unlock rules. A USB thumb drive is detected via its UUID – which is not that secure as a cryptographic key since you can fake a UUID, but on the other hand nobody knows the UUID of your USB drive in advance.

First of all, make sure you have PAM and Udisks



Use your flash USB drive instead of the root password to make the login policy more convenient.

development packages installed before building the pam_usb module with the *make* command. After doing **sudo make install** make sure that **pam_usb.so** has landed where it should (that differs across Linux distros). For instance, the software was tailoured for Ubuntu-like distros and didn't know that PAM modules in, say, Fedora live in /usr/lib64/security and not in /lib/x86_64-linux-gnu/security.

After that add the following in your PAM configuration by editing the **common-auth** or **system-auth** file under **/etc/pam.d/**: auth sufficient pam_usb.so.

Now you are ready to register your USB drive – see the 'Getting started' guide. When all is said and done, you will be able to login via su <user name> without a password if the corresponding USB drive is connected.

Games HOTPICKS

Billiards

Version: GIT Web: https://github.com/ tailuge/billiards

he world of open source games lacks some categories, and one of those is billiards. We rarely find a decent billiards simulator with freely available source code, and it's easy to guess why. Lots of maths work, complex game physics and many real-world billiards features that are hard to implement, like rolling friction, resilient cushions, true-to-life collision behaviour are all involved. Therefore we were extremely lucky to find a high-quality billiards implementation written in pure TypeScript and available for everyone at GitHub.

It's a classic one-person project showcasing coding skills and various challenges that are specific to ball movement. Don't be scared of those rough balls that look more like polygonal spheres! They move, roll, twist and bounce very much like real-world billiards balls.

The game is certainly a WIP thing, yet it already allows playing and training. *Billiards* is a 3D single-player game where you can aim and shoot with a finely controlled cue stick that can handle top, back and side spins, with full control over the kick power. After



shooting, the game automatically switches to the top view to better show you how the ball moves, and after that takes you back to aiming.

Use – and + keys to change the view by moving the camera along the vertical axis, and use the left and right arrow keys to move the aiming position along the horizontal axis. Hold Shift to enable more precise control. Aiming is implemented very well and the game is already a very decent trainer for practising your billiards.

No rules or penalties are implemented in the game so far, but given the constant shooting with the white cue ball, *Billiards* will likely grow into a classic eightball simulator. The live demo is available at http://tailuge. github.io/billiards/dist, but it's also possible to run a local instance of the game using *npm* (see **Readme. md** for detailed instructions). *Billiards* requires a WebGL-capable web browser. Billiards is a very promising simulator of its kind. Aim and shoot just like in real life (well, kind of)!

ARCADE GAMES

Brothers

Version: GIT Web: https://github.com/cyxx/blues

B lues Brothers is an open source project that lets you play the early 1990s games *Blue Brothers*, *Blues Brothers: Jukebox Adventures* and *Prehistorik 2* natively on Linux. Build using the make command and find the following three executables in the main **sources** directory: **blues**, **bbja** and **pre2**. Copy the respective file to the directory with the game and launch it – that's enough to start playing right away.

You can also inspect the **--help** section of the launcher to find some great extra options. For instance, the default game window is very small due to very low original resolution, and you may want to pass the **--fullscreen** option in order to fix it. Also, you can instantly jump into any level, or enable cheats and get unlimited lives and energy.

Our favourite game is definitely *Prehistorik 2*, where you lead a caveman through wild lands and fight hostile animals. Even realising how small the game's textures



are and how limited the colour palette, it's still amazing to see the amount of artwork stuffed into the tiny 1MB. *Prehistorik 2* includes over a dozen levels with different landscapes, complex portals and hidden platforms, four weapons, air-frame power-up, bosses, wind and snow effects and a lot more.

We haven't advanced in *Blues Brothers* and its sequel far enough, but these two games are no less impressive. There are small differences between the Blues engine experience and the original DOS experience, but we haven't encountered any issues that would spoil the general impression.

If you are a fan of those Titus games, do give the Blues engine a try!

Worms, dinos and extinct birds are ready to attack our hungry prehistoric hero

EQUALISER

EasyEffects Version: 6.0.3 Web: https://github.

com/wwmm/easyeffects

n LXF228 we reviewed *PulseEffects*, a powerful standalone audio equaliser for any Linux system that ran Pulseaudio. As time passed by we felt like it was necessary to revisit this outstanding software due to recent changes in the Linux audio subsystem. As you may know, Pulseaudio is already getting to the state of being legacy software, while the new Pipewire sound server is going to supersede it.

Pipewire is another intermediate layer between your media applications and the hardware, yet it is less bloated, less buggy (hopefully), and promises to better handle audio sessions for the sake of more efficient management. Also, Pipewire was designed with a powerful security model that makes interacting with audio and video devices from containerised applications easy, with supporting Flatpak applications being the primary goal. As such, the *PulseEffects* project was renamed to *EasyEffects* to reflect that its no longer a Pulseaudio-related thing.

We tested *EasyEffects* using the latest code, which required us to set up some bleeding-edge versions of

 Prod
 2
 4 Rayer
 4 rayer
 2 rayer
 2 rayer

 AddPage +

Gtkmm and GTK 4 dependencies taken from GNOME's GitLab servers. The application itself is an evolutionary improvement over the older *PulseEffects*, and it works and plays like a charm.

The whole idea behind a system-wide media stream control centre is that it puts you in control over everything that is played on your Linux system. *EasyEffects* replaces both Pulseaudio plug-ins and *Gstreamer* filters and provides the single point for altering audio properties. The application has a solid number of filters including equaliser, convolver, gain control, limiter, noise reduction and much more.

Control input and output levels, fine-tune harmonics, change stereo scene – these are only a few of the possible usage scenarios for *EasyEffects*. It feels like a very pro-level tool, which will be landing in mainstream Linux distros pretty soon.

DESKTOP ENHANCEMENT

Latte Dock

Version: 0.10 Web: https://invent.kde. org/plasma/latte-dock

his all began with envy of the catchy parabolic zoom effect that was first implemented in the Apple OS X dock. Although we have several implementations of that eye candy for Linux these days, it is *Latte Dock* that became the most mature and advanced dock panel.

The reality is that it's already much more than a dock; it is a complex and versatile system for creating and managing custom panels of very different sorts and types. More than that, the fresh *Latte Dock 0.10* release is no longer KDE-only software, as it boasts a desktop-agnostic widget explorer and happily runs anywhere outside of KDE Plasma without issues.

Latte Dock particularly shines with desktop customisation. The application enables setting both MacOS-like docks and classic panels, place it along any screen edge, add any combination of widgets and adjust their placement with paddings and spacers. It takes very little time to mimic the Ubuntu layout with its vertical left-side panel, or set up macOS Big Sur's recognisable floating dock that seems to be slightly



Version 0.10 was long-awaited and brings a great number of new dock and panel customisations.

Manage and adjust

audio streams like a pro with the help of EasyEffects.

pulled off from the screen edge, or copy the new Windows 11 bottom panel with centered icons, and so on... There is really a lot of room for creativity!

Latte Dock provides a parallel system of desktop panels in KDE Plasma, which means that you can impose a Latte panel over the traditional Plasma panel. For instance, you can fill the convenient Plasma panel with the Panon widget (LXF276) and then run a semi-transparent and blurred Latte panel above it for the sake of pulsating panel effect.

The new *Latte 0.10* has some cool improvements and lets you place several panels along the same screen edge, play with new hiding options, backgrounds, panel corner rounding, better move and align widgets and plenty of more advanced customisation tricks. By the way, it still zooms the dock icons very smoothly. Don't miss it!



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**

CODING ACADEMY

DAST

Build a dynamic app security pipeline

Dynamic Analysis Security Testing takes centre stage in the third instalment of our Web Application Security series with **Tim Armstrong**.



Tim Armstrong is a former Lead Engineer turned Developer Advocate specialising in networking. software development. and security. You can find him on Twitter as @omatachyru or via his website at www.plain textnerds.com. he battle between developers and malicious hackers is one that developers have been losing. A lot of the time, it comes down to mentality and company priorities. Hackers, like burglars, only need to find a single open window or unlocked door to get in. You wouldn't check that you've locked your door only once every few months, yet this is the exact approach many companies take to security.

Dynamic Analysis Security Testing (DAST) is perhaps the most overlooked stage of any security pipeline, frequently relegated to a check-up every six months by an outside consultancy that does an automated scan with *Burp Suite* or *Zed Attack Proxy (ZAP)* and provides you with a (hopefully short) report and an invoice in the range of £3,000-30,000, mostly depending on the scope. In most cases, the consultants don't go further than the automated scan because at that point they already have enough to write a multi-page report.

But here's the thing: when malicious actors (aka hackers) attack your web app, site or API, they aren't checking if your code is neatly formatted, they're essentially doing dynamic analysis. They're looking for a place where you've not validated the input, an endpoint that you've forgotten to protect, cookie slack, a vulnerable login system, leaked credentials and hundreds of other things that are very difficult to detect statically. If you're relying on a spot test every six months then odds are you've got security holes that you're not aware of.

Building DAST into your CI/CD only takes a few minutes and gives you effectively that same information that you'd get from a pen-test where all they did was run an automated scanner. The main difference is that instead of it only occurring every six months, the scan happens every time someone merges a PR to the main branch – meaning you find out about the vulnerability when it gets merged. Ultimately this means that when you do bring in the external consultants for the sixmonth check-up, you actually get your money's worth!

In this tutorial, you'll be adding DAST to the GitLab CI/CD pipeline that you've built over the course of this series. If you haven't read the earlier instalments yet it's a good idea to check those out first, but if you just want

👻 Create New App	Initial Scan Environment Configure initial scan environment. You can add more later.	
App Details	Development	
Environment	The emolectrone your app is naming in Host	
	http://localhost:8000	
YAML		

Adding an app in StackHawk is pretty straightforward, with a clean workflow and UI.

to dive in at this point, then you can pick up a copy of the progress so far at https://gitlab.com/ plaintextnerds/web-app-security-tutorial2-lxf280.

Too many acronyms

When it comes to DAST there are a growing number of solutions in the market. Most of them use hosted scanners that run on a periodic schedule, so either you need to expose a test environment to the internet and have the app scan that, or you need to set it up to scan production – which is a bit late (and could lead to instability of the production environment).

The goal here is, as it was when you added static analysis and software composition analysis to the pipeline, that you know about any vulnerabilities before the code goes live so that you can fix them before it gets deployed.

What you need, then, is something that can run inside the pipeline to test the service on commit. So what are your options? Well, you could build your own solution around an open-source tool like ZAP, or you could pick up an off-the-shelf solution. Unfortunately, there aren't many DAST solutions that you can build directly into your Cl pipeline, with the leaders in this space being StackHawk and GitLab. Both based their scanners on ZAP, meaning that they can run in a Docker container in your pipeline (or even locally).

StackHawk's scanner, *HawkScan*, is a little more advanced than GitLab's version and has support for multiple authentication methods and makes it easy to customise the scanner. While StackHawk holds the lead in scanning capabilities, GitLab is ahead in pricing (for

Security testing CODING ACADEMY

proprietary/closed-source software) as it's included in GitLab Ultimate, and obviously, it's also ahead in its integration with the rest of the GitLab platform.

What really separates these two solutions from others in this space, however, is their dedication to open-source, as both companies have decided to make their solutions free to open-source projects. Which, considering that they are both built on *ZAP* (which itself is an open-source solution), should perhaps not be a surprise. Unfortunately, however, it is not that common, as even though a number of their competitors are also likely based on *ZAP*, the vast majority of them do not extend the same offer.

Street – er, StackHawk

Kicking things off with StackHawk, the first thing you'll want to do is register an account at **http://stackhawk. com** so that you can set up your app and get an API key-pair for your *HawkScan* instance with which to push the scan results.

When you open a Developer account you get to use it for free for one app, so you can follow along even if you're not working on an open-source project. If you are working on an open-source project, make sure you contact the StackHawk team to unlock that free upgrade for your own team (they also help out start-ups with special deals).

With an account set up, make sure you're on the Applications dashboard and hit the 'Add an App' button. This will open up a modal dialogue box where you can give it a name, configure the environment type, set the hostname, and generate the Application ID and stackhawk.yml config file.

Next, you need to generate your API key. Do this by clicking your profile picture at the bottom left and then going to Settings. From here go to API Keys and create a new key and copy it into your clipboard. Then, head over to the Settings page for the CI/CD in your GitLab project and add it as a new variable with the key **HAWK_API_SECRET** and with both the 'Protect variable' and 'Mask variable' boxes ticked. While you're on this page you'll also want to add the **app_id** as **HAWK_APP_ID**; however, you don't need to tick the boxes for this one.

Next, you want to edit the **.gitlab-ci.yml** to add a new stage and of course, the new job. To do this add the line **- dynamic-analysis**

. . . .

<

Something you might notice is that this job is quite different from those defined in previous instalments of this series; this is because *HawkScan* needs to run in a DinD (Docker in Docker) environment. But how does it



know how to run your app so that it can test it? The answer to that question of course is that it doesn't, so that's what you'll need to define next.

To run the app in the CI/CD pipeline, you'll need it running in a Docker container. This means you need to define a build stage that makes a Docker image from the source code of the merge request. To do this you'll need a Dockerfile to build and a stage in the pipeline that will build it.

When it comes to building a Dockerfile for a Django project you only really need it to be a handful of lines long, like this:

FROM python:3.9 WORKDIR /usr/src/app COPY requirements.txt / RUN pip install -r requirements.txt COPY ./i_am_vulnerable/..

EXPOSE 8000 CMD ["python", "manage.py", "runserver",

"0.0.0.0:8000"] Place this **Dockerfile** in the project's **src** directory

next to the requirements.txt.

To build this Docker image in the CI/CD pipeline you'll need to add a new stage called **build** to the list of stages, placing it directly before the **dynamic-analysis** stage. Then you'll need to add a build job, that will look something like this:

» DEFENCE IN DEPTH

Modern security practice involves accepting that, at some point, you will be breached. That doesn't mean that you give up; it means that you plan for the eventuality. Once you have been breached you need to detect and prevent lateral movement through the system so that you know what has been leaked.

If you've been hit by ransomware, you'd better hope that you have off-site backups. If your administrator credentials for your hosting provider get stolen, you'd probably wish you had your infrastructure defined as code so that it can be rebuilt quickly. If your database gets leaked, then you'd want to be certain that you've encrypted sensitive data so that your customers don't get exposed.

This kind of planning for disaster is essential in today's world, and it's not going to change any time soon. Laws won't stop criminals and to put salt in the wound, at the moment, governments around the world seem to be single-minded in making it easier for them (with proposals like backdoored encryption, and online 'real identity' verification being promoted by politicians).

Make sure you download the stackhawk.yml file and make a note of the App ID as you'll need it later.

QUICK TIP

CI/CD pipelines can be optimised by adding rules that dictate when a stage or job is to be run, which can reduce costs and improve the developer experience.

>>

CODING ACADEMY Security testing

build-docker:

stage: build

image: docker:20 services:

- docker:20-dind

script:

- cd src

- docker login --username \$CI_REGISTRY_USER --password \$CI_REGISTRY_PASSWORD \$CI_ REGISTRY

- >

docker build

--tag \$CI_REGISTRY_IMAGE:\$CI_COMMIT_ SHORT_SHA

- docker push \$CI_REGISTRY_IMAGE:\$CI_ COMMIT_SHORT_SHA

So what does this do? Looking through the script section of this job, you can see that it will login to the GitLab Docker repository for the project, then build the **Dockerfile** found in the root of the project directory tagging it with the current commit tag, after which it pushes the container up to the GitLab Docker repository for the project.

This is all good to go, so you'll need to make it available to HawkScan. To do this you'll want to replace the script section of the HawkScan job with the following code:

- docker login --username \$CI_REGISTRY_USER --password \$CI_REGISTRY_PASSWORD \$CI_ REGISTRY

- docker run --name djangoapp -d \$CI_REGISTRY_ IMAGE:\$CI_COMMIT_SHORT_SHA

-1

docker run --link djangoapp -v \$(pwd):/hawk:rw -t \ -e API_KEY="hawk.\${HAWK_API_ID}.\${HAWK_API_ SECRET}" \

- -e NO_COLOR=true $\$
- stackhawk/hawkscan

When you've finished, the first few lines of your .gitlab-ci.yml should look something like this:

stages:

- static-analysis - composition-analysis - build - dynamic-analysis hawkscan: stage: dynamic-analysis
- image: docker:20
- services:
 - docker:20-dind
 - before_script:

Variables

Once you've

generated your StackHawk API

secret and APP

ID you need to

put them where

HawkScan can find them.

varia	Jies				cond
Variables	store information, like pas	words and secret keys, that you can	i us e in job script		
Variables					
• Рл • Из	etected: Only exposed to sked: Hidden in job logs. N	protected branches or tags. fust match masking requirements. L			
	HAWK_API_SECRET				
	HAWK_APP_ID				
Add va	Reveal values				

Linux-Format-Demo				
arquintad Da Ducatian aq 00, 2021 at 0054 057 49 axts	Nonitiras Vere 9.11.6			
4 5 Nation	8	0 August	0 Mart Arrest	 0 Nata Persita

From the scan results view you can see how many of the issues are new, and how many still remain since you last triaged the results.

- docker pull stackhawk/hawkscan

variables:

APP_HOSTNAME: \$CI_REGISTRY_IMAGE script:

- docker login --username \$CI_REGISTRY_USER --password \$CI_REGISTRY_PASSWORD \$CI_ REGISTRY

REGISTRY

- docker run --name djangoapp -d \$CI_REGISTRY_ IMAGE:\$CI_COMMIT_SHORT_SHA

-|

docker run --link djangoapp -v \$(pwd):/hawk:rw -t \ -e API_KEY="hawk.\${HAWK_API_ID}.\${HAWK_

API_SECRET}" \

- -e NO_COLOR=true \ stackhawk/hawkscan
- stacknawk/nawkscar
- build-docker: stage: build
- image: docker:20

services:

- docker:20-dind
- script:

- cd src

- docker login --username \$CI_REGISTRY_USER -password \$CI_REGISTRY_PASSWORD \$CI_
- REGISTRY

docker build

--tag \$CI_REGISTRY_IMAGE:\$CI_COMMIT_ SHORT_SHA

- docker push \$CI_REGISTRY_IMAGE:\$CI_ COMMIT_SHORT_SHA

Then you'll need to add the **stackhawk.yml** file that was generated at the start to the root of the project. Finally, you'll need to edit a couple of lines in that file, so change the host field of http://djangoapp:8000 to http://djangoapp:8000/bad_sql and uncomment the antiCsrfParam field and set it to be csrfmiddlewaretoken. Then git add, git commit and git push those changes up to the GitLab project.

This will trigger the pipeline which, in addition to running the SAST and SCA stages defined in previous instalments, will now build a Docker image, push that to the GitLab project's Docker Registry and then run *HawkScan* against the image, posting the result up to your StackHawk account when it's complete.

Scan results

Results are broken down into three main categories. High category should be fixed immediately as they pose an immediate danger to business continuity; Medium

Security testing CODING ACADEMY

	+ findings	Want	https://diangoapp:443/1/vote/	No.
			Exidence Activity	
HTTP Only Site MERLIN				
and the second se			Request	
and good in the last of the little of the				
Deatabaret Historychals convOreitShiCheadineetier Layer Frofaction, Died, Sheetind			PET https://djangang.ckk/l/womar Aft toor Agent: Moztla/5.8 bindow Af D Program mortache	
			kaferen: http://djacgoop.comm/// kests djacgoop.comm Caskins cirfickon-skitck/reasion/sering	

The finding details view provides a very useful overview of the issue, the evidence, and the paths where it was detected so that you can tackle it efficiently.

are generally issues with known exploit paths, but might not be a direct risk to business. Low category tend to be informational leaks (such as server versions) that could make an attacker's job easier.

Clicking one of the findings brings up a summary of what the specific vulnerability is, including some notes on how a malicious hacker might abuse it.

Next to the findings is a complete list of the paths scanned, so you know if *HawkScan* was able to find and scan a particular path in your application. This can be really helpful for ensuring that you have full coverage of the application. Web crawlers are rarely perfect, and the one utilised by *HawkScan* is no different; however, if the application you're scanning has a GraphQL or OpenAPI/ Swagger schema (or even a SOAP descriptor) then *HawkScan* won't need to use a crawler and should hit 100 per cent of the paths every time.

It should be noted that just because a vulnerability isn't detected doesn't mean that it isn't there. At the time of writing the SQL Injection vulnerability in the *bad_sql_practices* Django app that is used as the base example for this series was not detected by *HawkScan* (or any of its competitors).

Triage and false positives

Once we know of a vulnerability we need to triage it. The first step to a good triage process is to discard known false positives. When testing developer environments, like the one configured in this tutorial, it's common to exclude things like TLS/SSL certificates. It's no surprise then that HawkScan detected it as an 'HTTP Only site'.

Because of this the StackHawk UI wishes to inform you of the dangers of HTTP Only websites. However, as this is a development-grade deployment this is not actually a concern. To mark this finding as a false positive, open up the finding, and at the top right of the page you'll see a button labelled Validate and a dropdown called Actions. From the drop-down you can select 'False Positive' and provide a description as to why it should be ignored.

Once you've filtered out known false positives, the next step of triage is to ensure that you have tickets in your project management for all of the remaining risks. Any high-risk vulnerabilities should be expected to break sprint and receive immediate attention, as failure to do so would mean knowingly leaving the door wide open to attackers.

Medium-risk vulnerabilities are commonly scheduled into the next sprint. This isn't ideal as, while they aren't

» MAKING A POINT EFFECTIVELY

If you're struggling to advocate for fixing a particular vulnerability, one of the best tools is learning how to demonstrate the risk. If you can show the vulnerability in action and the damage that can be done, you will find the conversation shifts very quickly to how quickly you can fix it. If it doesn't, you might need to speak to someone higher up in the organisation. If even that doesn't help then the organisation as a whole has a serious culture problem, and will likely have a big breach in the future.

Demonstrating vulnerabilities is a far more effective way to get time to fix vulnerabilities than just discussing them because it makes things far more tangible for someone who isn't as capable or knowledgeable as you are.

This is one of the main reasons that hacker conferences like Black Hat and DEF CON are important, as they disseminate both defence and attack strategies, giving a more rounded exposure to the field. After all, without knowing how to exploit something, understanding the potential risk and work out how to protect it is very difficult, and expressing the risk to others effectively is even harder.

normally as big of a problem as the high-risk ones, when coupled with other vulnerabilities they can be just as bad as a high-risk vulnerability. However, if you attempt to treat everything as a sprint-breaking priority then management might start to ignore you as if you were crying wolf – ultimately leading to the opposite of the desired outcome.

Making a plan

If you've followed this series thus far, you now have a pipeline containing SAST, SCA and DAST. Once all three stages do not find any vulnerabilities, you should be in good shape moving forward.

However, it's important to remember that this just means there are no vulnerabilities found – not that they aren't there.

This is why it's important to take time to make a security breach response plan and to build your defence in depth. When developing a complex stage to a CI/CD pipeline, it can be a good idea to comment out the existing stages – otherwise, you can find yourself spending a lot of time waiting.

5 jobs for main in 8 minute	es and 57 seconds (queued for 1 seconds)	nd)	
3			
• e0004373 🔓			
No related merge request	s found,		
No related merge request	s found,		
No related merge request	s found,	N III A	Recently and the
 No related merge request peline Needs Jobs 50 T Static-analysis 	s found. cots 8 Composition-analysis	Build	Dynamic-analysis
th No related merge request ipeline Needs Jobs (\$ 1 Static-analysis	Composition analysis	Build	Dynamic-analysis

A completed pipeline can take up to 15 minutes to run (or longer in more advanced projects), so combining stages can sometimes be necessary.

>>> DISCOVER FLAWS IN OUR SYSTEM Subscribe now at http://bit.ly/LinuxFormat

CODING ACADEMY Pseudo-3D effects

Classic pseudo-3D racing road effects

With a bag of 10p pieces in hand, **Andrew Smith** whisks us back to the classic arcade days to recreate pseudo-3D racing games.



Andrew Smith is a software developer @ NHS Digital, has a Bachelors degree in Software Engineering and an MSc in Computer Networks. ue the budget wibbly-wobbly 'going back in time' special effects, as for this coding tutorial we're going to look at some the old-school techniques used in some of the classic racing video games such as *Road Rash* (1991), *Outrun* (1986) and *Pole Position* (1982). Designed to work smoothly on low-powered hardware, these are smart visual tricks that create a fake 3D effect.

What is known as pseudo-3D techniques were used to create a simulated 3D racing effect. The games would often be played by a single player or two players against computer opponents. The Pseudo 3D Road project (created by Ray Tomely) that we will be looking at, even though not a full video game, is a selection of examples of pseudo-3D techniques demonstrating ways to generate the 3D effect. You will see that once the project has been downloaded, these programming techniques are located in different folders that each demonstrate a pseudo-3D effect.

Getting started

To get started we will need a few things: Python, PyGame and the Pseudo 3D road project. To install Python, open a terminal window (Ctrl+Alt+T) and type sudo apt-get python3 followed by sudo apt-get install pip3. Then install the PyGame module by typing pip3 install pygame.

Finally grab a copy of the Pseudo 3D road project from https://raytomely.itch.io/pseudo-3d-roadcollection by clicking the Download button on the page. Once downloaded, extract the contents of the pseudo_3d_road_collection.rar file into an accessible location on your system.

As an example, the whole project has been put into a folder called **PythonProjects** which was created before downloading the project. The source code and project can be retrieved from the **LXF281** DVD. This tutorial will focus on the source code located in the folder called simple_road. If you're not already in that folder, type cd simple_road to get into it.

To edit and view the source code you can either use a default text editor installed on your flavour of Linux (Ubuntu for example) or you could use something more



The absolute classic racing game Pole Position by AtariSoft.

specific such as Notepad++, PyCharm or VS Code. The choice is entirely up to you. For this tutorial, we will be using gedit to view and edit the source files. When using this method to view/edit source files, it maybe helpful to open up two console windows where one will be used for editing/viewing source files and the other is a terminal window for executing the PyGame code.

Within the **simple_road** folder, the source code file that we want is **simple_road_curve_segment_demo.py**. Before we look at the source code, navigate into the **simple_road** folder and type the following to execute the script.

\$ python3 ./simple_road_curve_segment_demo.py

On successful execution you should get the output as seen in Figure 2 (page 94). The script executed has only been implemented with a forward control using the Up cursor key, and you will need to close the program with a mouse to end the program. As you hold down the Up cursor key, you will see the generated road go from being straight to having a bend in it and then back to being straight again.

Getting curvy

To view/edit this source file type the following into another terminal window:

\$ gedit ./simple_road_curve_segment_demo.py You will see that it's commented throughout to indicate what the variables are used for. As with any

Pseudo-3D effects CODING ACADEMY

other Python/PyGame script, at the top of the source code the libraries needed to run the script are declared. Also notice that continuing on from this, global variables for the script are declared to define the colours for the other parts of the road. See the following below:

import pygame,sys

from pygame.locals import *

BLACK=pygame.color.THECOLORS["black"] WHITE=pygame.color.THECOLORS["white"] RED=pygame.color.THECOLORS["red"] GREEN=pygame.color.THECOLORS["green"] BLUE=pygame.color.THECOLORS["blue"] YELLOW=pygame.color.THECOLORS["yellow"] SCREEN_WIDTH=640 SCREEN_HEIGHT=480

HALF_SCREEN_HEIGHT=int(SCREEN_HEIGHT/2)

The screen resolution settings **SCREEN_WIDTH** and **SCREEN_HEIGHT** are set to relatively low values in regards to today's device screen capabilities. Please feel free to change these for a higher resolution screen setting if you're not happy with them. In addition to changing this, you will need to add two lines of code to ensure that the images loaded will always meet the screen resolution specified.

Add the following lines just after the images **light_** road.png and dark_road.png have been loaded, as shown below.

light_road=pygame.image.load('light_road.png'). convert()

light_road = pygame.transform.scale(light_road, (SCREEN_WIDTH, SCREEN_HEIGHT))

dark_road=pygame.image.load('dark_road.png').
convert()

dark_road = pygame.transform.scale(dark_road, (SCREEN_WIDTH, SCREEN_HEIGHT))

You will also need to alter the following line of code in the script. Change

bottom_segment={'position':240,'dx':0}

bottom_segment={'position':SCREEN_ HEIGHT/2,'dx':0}

If the **SCREEN_WIDTH** and **SCREEN_HEIGHT** are changed and the above changes are not done, the display will look odd when running the program.

The main operation of the script is carried out in a defined **main()** function which is shown in part below:

Intervendel-Luburtu-01:~/PythonProjects/pseudo_3droads ls IndreweddelLuburtu-01:~/PythonProjects/pseudo_3droads ls Indorexetsd_road projected_road readme.txt simple_road z_map_ indreweddelLuburtu-01:~/PythonProjects/pseudo_3droads ls -1 indrexetsd_road readrewed woog Aug 30 2020 projected road incxnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road ner.m.r. 1 andrew andrew 4006 Aug 30 2020 projected road incxnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incxnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 1 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x 2 andrew andrew 4006 Aug 30 2020 projected road incrnwr-x

The final folder structure once the project has been downloaded

def main(): pygame.init() #Open Pygame window screen = pygame.display.set_mode((640, 480),) #add RESIZABLE or FULLSCREEN #Title pygame.display.set caption("simple road") #font font=pygame.font.SysFont('Arial',30) #images light_road=pygame.image.load('light_road.png'). convert() dark_road=pygame.image.load('dark_road.png'). convert() light_strip=pygame.Surface((SCREEN_WIDTH,1)). convert() dark strip=pygame.Surface((SCREEN WIDTH.1)). convert() light strip.fill(light road.get at((0,0))) dark strip.fill(dark road.get at((0,0))) #variables texture_position=0 #this is used to draw the road As with all PyGame scripts, PyGame is initialised with a pygame.init() function call as can be seen from the above code. Also notice that when the screen display is set up, where there is the line: screen = pygame.display.set_mode((640, 480),) #add **RESIZABLE or FULLSCREEN** it misses out using the variables SCREEN WIDTH and SCREEN HEIGHT declared above. Replace the value of 640 with SCREEN WIDTH and

Replace the value of 640 with SCREEN_WIDTH and the value of 480 with SCREEN_HEIGHT so that the screen resolution can be adjusted to suit your device.

Notice that there are just two images loaded to create the intended effect and the rest of the effect is created by use of internal colour schemes to create alternating strips.

» CREATING THE ROAD

The road is created with a combination of two images loaded near the beginning of the script, **light_road.png** and **dark_ road.png**, and also two colour strips created with light and dark tones for the grass area outside of the road. The colour of the grass is picked from the light and dark images respectively for the grass colour. Let's look at the code that generates the graphics for the road.

 screen.blit(light_road,(curve_ value,i+HALF_SCREEN_ HEIGHT),(0,i,SCREEN_WIDTH,1))

else:

screen.blit(dark_ strip,(0,i+HALF_SCREEN_HEIGHT)) screen.blit(dark_road,(curve_ value,i+HALF_SCREEN_

HEIGHT),(0,i,SCREEN_WIDTH,1)) As can be seen, when rendering the

As can be seen, when reindering the road, an if/else statement is used to determine whether to render a dark part of the road or a light part of the road. The grass outside of the road is rendered at the same time the road is. What is also worth pointing out is that when the strips are drawn, they stretch from one side of the window to the other; the road is overlaid on top of the strips to give the illusion the grass is at each side of the road. To help create the curve effect in the road (left-hand bend), only a part of the image for light and dark road are redrawn with an offset x-position by curve_value. Each image segment which is rendered on the screen is incremented slightly more than the last to give the road a left-bend.

>>

CODING ACADEMY Pseudo-3D effects

In every PyGame program there needs to be a main loop that renders the graphics used in the program and that also controls input for the game via keyboard and/ or mouse depending on the PyGame application being developed. Let us have a look at the main loop

implemented in simple_road_curve_segment_demo.

while True:
#loop speed limitation
#30 frames per second is enough
pygame.time.Clock().tick(30)
for event in pygame.event.get(): #wait for events
if event.type == QUIT:
pygame.quit()
sys.exit()
#Movement controls
keys = pygame.key.get_pressed()
if keys[K_UP]:
As can be seen from the code above the main while

As can be seen from the code above, the main **while** loop ends on a Boolean condition to identify when the program has finished. In the case of the above, the loop ends when a False condition is raised. Continuing on from this, the frame rate is declared at 30 frames a second and further on from this the input/keyboard control is set up.

Finally after doing all this we come to the final part of the loop as shown below:

pygame.display.flip()

From the above we update the contents of the entire display. The main focus of the render processing is done in the PyGame main while loop as shown below: while True:

Setup / Control code

#Movement controls

- keys = pygame.key.get_pressed()
 if keys[K_UP]:
- road_pos+=road_acceleration
- if road_pos>=texture_position_threshold: road_pos=0
- top_segment['position']+=curve_speed

#if we reach the curve's end we invert it's incrementation to exit it

- if top_segment['position']>=curve_map_lenght: top_segment['position']=0
- bottom_segment['dx']=top_segment['dx']

top_segment['dx']-=0.01 #+0.01 to exit a left curve and -0.01 to exit a right curve

top_segment['dx']*=-1

Figure 2: This is what a successful execution of program should look like. As can be seen from the above code segment, when the Up cursor key is pressed, the road position is increased as per the **road_acceleration value** (currently set to 80) and when the value of **road_pos** exceeds the value of **texture_position_threshold**



(currently set to 300), the value of **road_pos** is set back to 0 and the process is repeated. If you think of a conveyer belt or treadmill mechanism where the track used is on a continuous cycle or loop, this is a very similar idea here for the pseudo-3D effect. The road is continually re-generated after a certain point is reached.

To help control when the bend on the road appears, there is the following line of code after that which deals with the Up cursor key and after the initial road processing code:

curve_map_index+=curve_increment

Initially, the value of **curve_map_index** is set to -1, however the value is incremented by **curve_increment** (currently set to 2) on each cycle that the player has the Up cursor key pressed.

As the Up cursor key is pressed, you should start to see the road go into what looks like a left bend and then back to being a straight road again. Let's look at the code that controls this.

- if curve_map_index>=curve_map_lenght:
- curve_map_index=curve_map_lenght curve_increment*=-1

#if we exit, we invert its incrimintation to enter again #we invert the curve's direction to change the way elif curve_map_index<-1:

curve_increment*=-1

curve_direction*=-1

The point at which the left-bend appears and goes back to being a straight road again is controlled by an **if else** statement, or in Python **if elif**, as the Up cursor key is pressed. The first part of the **if** statement deals with the situation where if the curve has reached its maximum curvature it starts to make the road look straight again. The second part of the statement (**elif**) deals with the case of making the road bend to the left.

For those that are new to incrementing variable values in Python, it may be worth pointing out that curve_increment *-1 is equivalent of writing curve_increment *-1. This is again the same with the variable curve_direction = curve_direction *-1.

Graphics rendering

To help learn about how this program works further, you may find it useful to play around with the following variables:

road_pos=0 #remembers our position on the road road_acceleration=40 #the speed we traverse the road texture_position_acceleration=8 #strip "stretch" value texture_position_threshold=300 #strip division value half_texture_position_threshold=int(texture_position_ threshold/2) #define drawing light or dark road

To save time scrolling through the code in getting to what you want, use the IDE search facility usually brought up by pressing Ctrl+F and then type the name of the variable or function you are looking for.

The two variables that might first be of interest to play around with values first could be **road_acceleration** and **texture_position_acceleration**.

Change the values of each variable and then run the program as shown before to see the effect this has on the running of the program. You may want to do this a number of times to get used to the effect of using different values.

Pseudo-3D effects CODING ACADEMY



Our retro road running the basic straight road effect.

In every PyGame program there has to be a part that is written to render all the graphics so after user input they can be updated. This is done with a **for** loop situated in the second half of the **while** loop as seen in part below:

for i in range(HALF_SCREEN_HEIGHT-1,-1,-1):

if top_segment['position'] < i:

dx = bottom_segment['dx']

else:

dx = top_segment['dx']

ddx += dx

current_x += ddx

curve_map[i] = current_x

curve_value = curve_map[i]

if texture_position<half_texture_position_threshold: screen.blit(light_strip,(0,i+HALF_SCREEN_

HEIGHT))

screen.blit(light_road,(curve_value,i+HALF_ SCREEN_HEIGHT),(0,i,SCREEN_WIDTH,1))

else:

screen.blit(dark_strip,(0,i+HALF_SCREEN_ HEIGHT))

screen.blit(dark_road,(curve_value,i+HALF_ SCREEN_HEIGHT),(0,i,SCREEN_WIDTH,1))

Going further

The rendering of the graphics is by default set to take place in the bottom half of the screen and not the top half, hence the use of **HALF_SCREEN_HEIGHT**. The above code renders both the dark and light sections of the road onto the bottom half of the screen. The top half of the screen is kept clear for the blue sky effect.

After you have gained more confidence with the code shown, you may want to add further functionality to the existing code. As you're currently forced to close the program by closing the window with a mouse, it may be a good idea to end the program when the Escape key is pressed instead.

The colour of the sky is currently a plain blue. A possible addition here would be to create an image of a white cloud in an image editor such as *Gimp* and then create an animation of clouds moving across the top half of the screen.

» GENERATING HILLS

An example of how to create a hill effect is given in the Python script file **simple_hill_road.py**. From viewing the code, you will see that new variables have been added to create the hill effect. hill_ position, hill_velocity, hill_acceleration and hill_sharpness are the main ones which have been added. The code that deals with the Up arrow key event is the same but the rendering process is slightly different even though similar in some ways to how the left-bend was created in **simple_road_curve_segment_demo.py**.

Again, you will see that a for loop is used for rendering the graphics to the bottom half of the screen. The hill effect for the road is created with the following code.

hill_position+=hill_velocity*hill_sharpness

To increase how steep the hill is, increase the value of hill_sharpness which is currently set to the value of 4. To make the road appear flat again, set the value to 0. The position of the hill is continually tracked as well to determine whether to regenerate a light part of the road or a dark part of the road.

There is currently recorded a current position of the hill and old position which the difference between the current hill position and the old hill position is continually worked out.

As can be seen in the script, the acceleration value is set to a constant 80 per Up cursor key press. Those of you who from a physics background may want to create a more realistic acceleration algorithm where the rate of acceleration may gradually increase and decrease.

Even though it may involve some planning and some re-construction of the existing code, since components of a race track have been demonstrated – straight roads, bends and hills – it maybe a good exercise to build a small race track.

In other source code examples, images of vehicles have been added. It maybe a good idea to take the existing image(s) and modify them in an image editor to create an image of a vehicle turning left or turning right and implementing into the script, writing code that will handle the left and right cursor keys. Instead of scrolling through the code line by line looking for what you want, instead press Ctrl+F which will usually bring up a search facility and type in a variable or function name.

With a bit of maths the basic road can be made to bend.



» DRIVING YOU ROUND THE BEND... Subscribe now at http://bit.ly/LinuxFormat

Un the disc



Discover the highlights from this month's packed DVD!

» START HERE

USING THE LXFDVD

Using Linux for the first time can be very confusing. It'll most likely be unlike anything that you've operated before, especially if you're used to Microsoft Windows or Apple macOS.

Generally our DVDs are designed to be run directly, which is to say that when you first power on your PC (or Mac, see below) it should 'boot' from the DVD - so before Windows or macOS even starts to load - with Linux running directly from the DVD. This trick is known as a Live Disc. It enables you to try out the various versions of Linux without having to install or change anything on your PC. Just remove the DVD, restart your PC and it'll be exactly as you left it.

While many systems will boot from a DVD when it finds one, many will not. See below for the standard process for enabling booting from a DVD on various desktops and laptop PCs.

The alternative option is to locate the ISO file on the DVD and write this to your own USB thumb drive and attempt to run that. We recommend using Etcher from https://balena.io/etcher that's available for Windows, macOS and Linux, Good luck!

BOOT THE DISC

Many PCs should boot automatically if they're turned on with a disc in the drive. If not, many offer an early Boot Menu accessed by tapping a key while powering up from cold: F9 (HP), F12 (Dell, Lenovo), F8 (Amibios) or F11 (Award BIOS). Alternatively, use the BIOS/UEFI to adjust the boot order to start with the optical drive. Again, this is accessed by tapping a key during power up, usually Del but sometimes F1 or F2.

Some new UEFI PCs require access via Windows: holding Shift select its Restart option. If you're still having problems using the DVD then visit www.linuxformat. com/dvdsupport

Mac owners: Hold the C key while powering on your system to boot from the disc.

FRESHLY MUDDLED MIN SPECS: 2GB RAM, 20GB DISK

Linux Mint 20.2 "Uma" 64-bit

id your olfactory senses notice a whiff of Mint as you pried the DVD from its containment? Perhaps you should seek help - it's supposed to be odourless. The latest version of Linux Mint, however, will decongest your desktop senses with some refreshing new features. Besides the new selection of wallpaper. the main focus in this outing is the thorny issue of updates. We all know we should do them, especially if they concern security. But there's an awful lot of inertia around clicking that little Update Manager notification in the taskbar. In an attempt to reduce this, updates that have gone unnoticed will now be a little more vocal.

Master of the Mint Clement Lefebvre made it clear his team don't want to make updates become the nuisance they are in certain operating systems originating from Redmond, WA. But at the same time they expressed concern at search data (via the custom Yahoo search that ships with Mint's Firefox package) which shows that a considerable faction of users don't update

packages in a timely manner. Worse, a concerning proportion are still running EOL versions of Mint. As you'll see in our feature and corresponding review, the 'nag' frequencies can all be customised, so we're sure it can be made less annoving, if annov you it does.

You can also turn on automatic updates - either the whole gamut or just those which fix security issues. A nice touch is that if you opt to update Flatpak packages in the background, unused runtimes will be deleted. This is handy because one's ~/.local/share/flatpak directory can, if left unchecked, easily grow to gargantuan proportions if you get over-excited with the huge array of software available through FlatHub.

If you have very new hardware and run into problems then you might consider taking a peek at the lesser-known Mint Edge edition. This one features a newer kernel, so it hasn't undergone the thorough testing of the 5.4 kernel used by the Ubuntu LTS and Mint. It's hidden away at https:// www.linuxmint.com/edition.php?id=291



Easy backups and the Mint IRC channel should ensure Mint-related mishaps are rare

> IMPORTANT NOTICE!

DEFECTIVE DISCS: For basic help on running the disc or in the unlikely event of your Linux Format coverdisc being in any way defective, please visit our support site at www.linuxformat.com/dvdsupport. Unfortunately, we're unable to offer advice on using the applications, your hardware or the operating system itself.



FREE AS IN AWESOME MIN SPECS: 4GB RAM, 15GB DISK

Elementary OS 6.0 "Odin"

ell that was a surprise. Here we were thinking there'd be no room for Elementary OS on the DVD, but it turns out it just fitted. With all of 4MB to spare. So excited were we to achieve this that we forgot we can't ISO-boot Elementary without making a new initramfs. And said filesystem would occupy significantly more than 4MB. So you'll have to write it to a USB stick (or DVD if you really must) yourself. As always, our DVD support pages have all the answers.

Before you do that though, we'd like to take a minute to remind you that Elementary OS should be considered a paid product. The price is 'what you can', which includes \$0, but please don't think our giving it away on the cover means you shouldn't pay something. The Elementary OS team relies primarily on payments from the download screen (and the app store) to keep their efforts funded, so if you'd like it do get yourself to https://elementary.io and send 'em some dollars.

We won't be including Elementary OS on the digital download this issue, since that would make it too easy for people to download for free. Digital subscribers can download this directly (as they can any distro), and indeed pay for it. We hope you understand the rationale here – please do write in with your feelings. But not if you don't understand 'free as in beer' vs 'free as in speech'. It's really not that hard.

We've been waiting a long time for this release, since not long after Ubuntu 20.04 came out in fact. Let's just say it's well worth the wait.



Elementary is quite, quite stunning. And much better behaved than the Macs at Future Towers.

Elementary OS has always been stunning; people say it looks like macOS but we say it looks much better. Apart from the new dark theme, you can add your own accent colours now, or even have one automatically chosen based on the current wallpaper.

Applications in Elementary OS's bespoke AppCenter are all packaged as Flatpaks, with strict security settings enforced. In fact, noncurated apps, which are loaded through the Sideload helper, are subject to these restrictions too. Go to Settings > Applications > Permissions to see what apps can do, and if it concerns you, to revoke some of those permissions. Besides security, efficiency is an important focus this release. Thanks to touchpad gestures everyday tasks become a little less awkward. Carefullydirected three finger swipes can summon the Multitasking View or switch workspaces.



THE LXF LIBRARY

Advanced Bash Scripting Guide Go further with shell scripting.

Bash Guide for Beginners Get to grips with the basics of Bash scripting.

Bourne Shell Scripting First steps in shell scripting.

The Cathedral and the Bazaar Eric S. Raymond's classic text explains the advantages of open development.

The Debian Book Essential guide for sysadmins.

Dive Into Python Everything you need to know.

Introduction to Linux A handy guide full of pointers for new Linux users.

Linux Kernel in a Nutshell An introduction to the kernel written by master hacker Greg Kroah-Hartman.

The Linux System Administrator's Guide Take control of your system.

Tools Summary Overview of GNU tools.

GNU Emacs Manual Six hundred pages of essential information!

Producing Open Source Software Everything you need to know.

Programming from the Ground Up Take your first steps.

 Image: Section of the section of t

Unlike Jonni's accent, elementary's backgroundbased accentuation makes things clear and harmonious.

» NEW TO LINUX? START HERE...

Never used a Linux before? Here are some handy resources: • Read our quick-install guide http://bit.ly/LXFinstall • Looking for an answer? https://askubuntu.com • Want to delve more deeply? https://linuxjourney.com



Linux comes baked into Windows 11, so get the most out of Microsoft's latest by using Linux!

Build the LXF server

We kick off a new series looking at being a sysadmin running and growing your own Linux server. Let's get things stood up!

Start your own distro

We step you through how you can author your very own bespoke distro, pick your own apps, theme, settings and more!

Easy wireless file sharing

Need to quickly fling a few files around your wireless network? We have the perfect, simple tool to make that happen.

Get a grip on GPS

We explore the GPS system and how you can utilise it to tell accurate time, amongst other useful things.

Contents of future issues subject to change - as we might be too busy launching Windows Format! Haha, no.



Future Publishing Limited, Quay House, The Ambury, Bath, BAI 1UA **Email** linuxformat@futurenet.com FDITORIAL

Editor Neil "not in the office" Mohr Desktop editor Jonni Bidwell Art editor Efrain Hernandez-Mendoza Operations editor Cliff "on holiday" Hope Group editor in chief Graham Barlow Senior art editor Jo Gulliver Editorial contributors

Editorial contrustors Tim Armstrong, Michael Bedford, Neil Bothwick, Stuart Burns, Christian Cawley, Matthew Hanson, Jon Masters, Nick Peers, Les Pounder, Michael Reed, Ed Ricketts, David Rutland, Andrew Smith, Shashank Sharma, Alexander Tolstoy Cover illustration magictorch.com

Raspberry Pi is a trademark of the Raspberry Pi Foundation. Tux credit: Larry Ewing (lewing@isc.tamu.edu) and The GIMP. ADVERTISING

Commercial sales director Clare Dove Senior advertising manager Lara Jaggon

Head of commercial – Technology Dave Randall dave.randall@futurenet.com Account director Andrew Tilbury andrew.tilbury@futurenet.com

INTERNATIONAL LICENSING

inux 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.

Head of Print Licensing Rachel Shaw 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

CIRCUI ATION

Head of newstrade Tim Mathers PRODUCTION AND DISTRIBUTION

Head of production UK & US Mark Constance Production project manager Clare Scott Senior ad production manager Jo Crosby Digital editions controller Jason Hudson

THE MANAGEMENT

Chief audience and ecommerce officer Aaron Asadi MD, tech specialist Keith Walker Head of art & design Rodney Dive

Commercial finance director Dan Jotcham Printed by Wyndeham Peterborough, Storey's Bar Road, Peterborough, Cambridgeshire, PEI 5YS Distributed by Marketforce, 5 Churchill Place, Canary Wharf, London, E145HU www.marketforce.co.uk Tel: 0203 787 9001

The second secon

We are committed to only using magazine paper which is ediform responsibly managed, certified forestry and chicoine-ranufacture. The paper rith imagazine was sourced and uced from sustanable managed forests, conforming to strict ormental and socieconomic standards. The manufacturing rmill and printer hold full FSC and PEFC certification and citation.



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 ontent and advertising solutions for passionate onsumers online, on tablet & smartphone and in print.







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.





PAPER POWER

60% of the energy used to produce paper and paper packaging in Europe comes from renewable sources.

Discover the story of paper www.lovepaper.org



Source: Confederation of European Paper Industries (CEPI), 2018 CEPI represents 92% of European pulp and paper production

9006