



LINUX PRO

MAGAZINE

ISSUE 248 – JULY 2021

Brain Tools

Stay sharp with these free tools for study and play



COVID Crunch

Explore pandemic data with Gawk and Bash

Z-Wave

Raspberry Pi as a home assistant

LINUXVOICE


Voice and video chats with Tox

Zotero personal research assistant

10

FREE SOFTWARE GEMS!



The  **GNOME**[™] Conference

GUADEC

July 21–25, 2021 | Online

GUADEC is GNOME's main annual event. Held since 2000, it brings together Free Software enthusiasts and professionals from all over the world.

Join us at GUADEC 2021 to hear about the latest technical developments, learn new skills and tools, attend talks, and participate in workshops.

Learn More
GUADEC.ORG



PROFITS AND VALUES

Dear Reader,

The interface of politics with tech is always complicated – even when the stated values of the country and the company are similar. But in a world of global corporations and diverging values, the interplay is almost impossible to follow.

For instance, skirmishes of tech companies with China have been in the news recently. I should start by saying, few western social media companies even operate in China. Google, Facebook, Twitter, and other social media giants are all blocked, having fallen out with authorities for various reasons – including for refusing to comply with censorship demands.

LinkedIn remains in operation behind China’s “great fire-wall,” but the Chinese government is leaning heavily on them to step up their censorship. According to the *New York Times*, China’s Internet regulator scolded LinkedIn in March for failing to enforce censorship rules [1]. As a punishment, the Chinese government suspended new LinkedIn sign-ups in China for 30 days and ordered the company to perform a “self evaluation” and submit a report about it.

Fast forward to May, and LinkedIn appeared to have learned their lesson. British journalist Peter Humphrey, a critic of the Chinese government, claimed that LinkedIn censored his reference to the Chinese government as a “repressive dictatorship” and his assessment of Chinese government media as a “propaganda machine” [2]. Humphrey apparently reported that he received a notification from LinkedIn that his comments were removed for “bullying and harassment.” His account was then placed in a restricted state because of “behavior that appears to violate” the terms of service.

After a public outcry, LinkedIn restored Humphrey’s account and some (but not all) of the content. A spokesperson for LinkedIn told Bloomberg, “Our team has reviewed the action, based on our appeals process, and found it was an error” [3].

Info

- [1] “China Punishes Microsoft’s LinkedIn Over Lax Censorship”: <https://www.nytimes.com/2021/03/18/technology/china-linkedin-censorship.html>
- [2] “LinkedIn Accused of Censoring for Beijing”: <https://reclaimthenet.org/linkedin-accused-of-censoring-for-beijing/>
- [3] “Microsoft’s LinkedIn Accused by Noted China Critic of Censorship”: <https://www.bloomberg.com/news/articles/2021-05-12/microsoft-s-linkedin-accused-by-noted-china-critic-of-censorship>
- [4] “Censorship, Surveillance, and Profits: A Hard Bargain for Apple in China”: <https://www.nytimes.com/2021/05/17/technology/apple-china-censorship-data.html>
- [5] Amnesty International China Overview: <https://www.amnesty.org/en/countries/asia-and-the-pacific/china/>

There is a lot to learn in this kind of a story, although you can never quite see to the bottom of it. One question is, why is LinkedIn still in China when companies like Google and Facebook have already left. One (perhaps naive) answer is that the other companies have better values than LinkedIn. This interpretation is certainly something to consider, or at least, it is something you can’t really rule out, but other factors are also at play. For instance, LinkedIn is owned by Microsoft, which is more like an old-school IT company than a social media company and has some irons in the fire that complicate the equation. I should also note that some observers believe the presence of LinkedIn in China creates a tension for censors that would not be present if they disappeared. In this example, for instance, LinkedIn went along with the censorship at first, but under pressure from the international community, they backed down. That can only happen if companies with an international presence are part of the China scene and therefore part of the conversation.

But it is also important to remember that this devil’s dance of tech versus censorship is not just a social media problem. Another recent report in the *New York Times* details Apple’s long history of spouting liberal values in the west while caving to government demands in China [4]. According to the report, “...thousands of apps have disappeared from Apple’s Chinese App Store over the past several years, more than previously known, including foreign news outlets, gay dating services, and encrypted messaging apps. It has also blocked tools for organizing pro-democracy protests and skirting Internet restrictions, as well as apps about the Dalai Lama.”

What should we do about all of this? Of course, there are many layers and levels of possible involvement, depending on your time and interest. If you want to get actively involved, Amnesty International is a good place to start [5]. But one thing we can all do is keep talking about it and calling it out when we see it so that others will see it. As the LinkedIn example illustrates, these companies really do care what people think, and shining light on an issue can sometimes lead to a better outcome.



Joe Casad,
Editor in Chief



ON THE COVER

36 COVID Crunch

Create custom scripts to access online pandemic data.

62 Lakka

Torpedo the past with this retro gaming distro.

88 Zotero

Use this personal research assistant to create better and more thorough reports.

56 Z-Wave

Build a smart home with a Rasp Pi brain.

74 Tox

Peer-to-peer messaging with free tools.

NEWS

08 News

- Armbian 21.05 Now Available
- StarLabs has Released Another Linux Laptop
- Ubuntu 21.04 Adds Support for Active Directory
- Gnome 40 Available on openSUSE
- Kernel 5.13 Merges Apple M1 Hardware Support
- KDE Launches the Qt 5 Patch Collection

11 Kernel News

- Want vs. Need
- Hiding System Resources ... from the System

COVER STORIES

14 Anki

Anki brings a virtual flash card box to the desktop. Thanks to a useful collection of add-ons, you can adapt this helpful learning tool to suit your needs.

20 Dia

If you need to make the occasional diagram, Dia gives you professional results with a minimal learning curve.

24 gbrainy

Develop your logical thinking skills with this collection of brain teasers.

28 Manim

Create video sequences to present mathematical problems in an engaging way.

REVIEW

32 Distro Walk – Zorin OS

We interview one of the creators of Zorin OS, a Debian derivative with a minimal yet friendly desktop.

IN-DEPTH

36 Statistics with gawk

With very little overhead, you can access statistics on the spread of COVID-19 using gawk scripts and simple shell commands.

42 Command Line – Pandoc

Convert files from one markup format to another at the command line.

46 Nix and NixOS

This package manager and native distro offer a promising new approach to the challenge of managing packages.

52 Programming Snapshot – SymPy

Whether he's filling the bathtub with water or routing electricity through resistors, Mike Schilli juggles mathematical formulas with Python.

55 Charly's Column – User Groups

Charly investigates the three most frequently asked questions about user groups.



Cover Stories

Sometimes you want the computer to think for you, and sometimes you want the computer to *make you think*. This month we present a selection of free Linux tools for learning and thinking. Flash through virtual flashcards with Anki, create mathematical presentations with Manim, and, to keep your mind sharp, take a turn with the gbrainy brain teaser collection.

MakerSpace

56 Z-Wave

Use your Rasp Pi as a powerful smart home control center.

62 Lakka

This Pi-based distro comes with everything you need to play retro games. Just bring your Raspberry Pi and, ideally, a simple gamepad.

LINUXVOICE

67 Welcome

This month in Linux Voice.

69 Doghouse – Distributions

Which distribution is the best choice for you? maddog offers some suggestions.

70 Micro

Micro brings the ease and intuitiveness of a graphical editor to the Linux terminal.

74 Tox

The Tox protocol uses file-sharing techniques for messaging and audio-video chats.

82 FOSSPicks

This month Graham looks at Vital, Shortwave, doggo, Bonzomatic, Ultimate Plumber, Open Brush, and more!

88 Tutorial – Zotero

Use this reference management tool to organize citations and create bibliographies.



Ubuntu 21.04 and Fedora 34 Workstation

Two Terrific Distros on a Double-Sided DVD!



Ubuntu 21.04
64-bit

This month's DVD features the latest versions of two major distributions. First up is Hirsute Hippo, otherwise known as Ubuntu 21.04. The Ubuntu team chose a hybrid desktop environment of Gnome 3.38 and Gnome 40, so much of the interest in this release is in security enhancing features.

Ubuntu tried using Wayland as the default display server in the 17.10 release, only to encounter difficulties and return to the X Window System. Now, with Ubuntu 21.04, Wayland is once more the default. With better performance and high security, Wayland is a long-awaited change, although recording and screen capture apps may not work. Fortunately, switching to X Window is a minor task. Other security features include a power mode option for laptops and a recovery key option for encrypted disks. Probably the most noticeable change is the default permissions for home directories, from 755 (the user can read, write, and execute, while the group and others can read and execute, but not write) to 750 (the user can read, write, and execute; the group can read and execute, but not write; and others cannot read, write, or execute.). This change hides home directories from others; in case others manage to log in, it prevents them from running binary files.

Like all Ubuntu releases, Hirsute Hippo is a solid choice for new users, with plenty for them to discover as their expertise increases.

Defective discs will be replaced.

Please send an email to subs@linux-magazine.com.

Although this Linux Magazine disc has been tested and is to the best of our knowledge free of malicious software and defects, Linux Magazine cannot be held responsible and is not liable for any disruption, loss, or damage to data and computer systems related to the use of this disc.



Fedora 34 Workstation
64-bit

In contrast to Ubuntu, Fedora has used Wayland for several years. This release includes a bundle of enhancements to make Wayland practical with KDE's Plasma. Another major difference from Ubuntu 21.04 is that Fedora 34 Workstation uses the full Gnome 40 for its desktop environment.


As usual with a Fedora release, this one also incorporates cutting edge technology. The default Btrfs filesystem improves the speed and reduces the size of compressed files by using zstd. Even more importantly, the long-awaited PipeWire replaces the temperamental PulseAudio for managing sound. If, like many users, you have struggled to get PulseAudio to work on your system, the introduction of PipeWire alone may make Fedora 34 Workstation worth your time.

Like all Fedora releases, this latest one consists of only free software, as well as some first glimpses of new applications that will later find their way into Red Hat Enterprise Linux and, eventually, other distributions as well. Beginners may struggle with Fedora, but experienced users should find it a snapshot of the latest in Linux.

THINK LIKE THE EXPERTS

Linux Shell Handbook 2021 Edition

This new edition is packed with the most important utilities for configuring and troubleshooting systems.



Here's a look at some of what you'll find inside:

- Customizing Bash
- Regular Expressions
- Systemd
- Bash Scripting
- Networking Tools
- And much more!

ORDER ONLINE:

shop.linuxnewmedia.com/specials

NEWS

Updates on technologies, trends, and tools

THIS MONTH'S NEWS

- 08 • Armbian 21.05 Now Available
- Star Labs Has Released Another Linux Laptop
- 09 • Ubuntu 21.04 Adds Support for Active Directory Plus Other Major Changes
- Gnome 40 Available on openSUSE
- More Online
- 10 • Apple M1 Hardware Support to be Merged into Linux Kernel 5.13
- KDE Launches the Qt 5 Patch Collection

Armbian 21.05 Now Available

For anyone that makes use of ARM-based architecture, you have a friend in Armbian, a Debian-based Linux distribution, purpose-built for ARM and embedded devices (including development boards and Linux phones). Armbian is a simple-to-use, lightweight, secure Linux distribution, optimized for ARM-based hardware. This new release includes the 5.11 kernel, which now includes support for the Orange Pi R1 Plus as well as improved NVidia Jetson Nano support.

But Armbian isn't limited to only single-board computers and other embedded devices. This distribution can also run on ARM-based laptops such as the Pinebook Pro (<https://www.pine64.org/pinebook-pro/>). Other supported devices include Banana Pi BPI-M3, Firefly, ODROID N2+, ODROID-XU4, ROCKPro64, NanoPi K2, NanoPi M4V2, and Tinker Board computers.

Although Armbian defaults to the Xfce desktop, Armbian 21.05 also brings support for the Gnome desktop environment. You'll also find USB-C DisplayPort and eDP outputs enabled for the NanoPC-T4 board, as well as the usual host of bug fixes, patches, and updates.


Download a version of Armbian for all supported devices from the official download page (<https://www.armbian.com/download/>) and enjoy Linux on your ARM-based hardware of choice. You'll find support available for over 130 different devices.

Star Labs Has Released Another Linux Laptop

Star Labs (<https://earth.starlabs.systems/>) is not just a research facility in the world of *The Flash*, it's also a company that sells Linux laptops. Up until now, Star Labs only had one device for sale: the 11" Star Lite Mk III, which offered an 11.6" screen, 1.1GHz quad-core Intel Pentium N5000 CPU, a 240GB SSD drive, 8GB of LPDDR4 onboard memory, and up to seven hours of battery life.

The Star Lite was spec'd to be a low-end machine and started at \$478. At the moment, the Star Lite is unavailable for purchase. However, the company will soon release the StarBook (<https://starlabs.systems/pages/starbook#europe>), which can be spec'd quite a bit beefier than the original, with a 14" matte IPS display, an 11th Gen Intel CPU (either i3-1110G4 or i7-1165G7), up to a 1TB NVMe SSD drive, up to 64GB of DDR4 memory, a full-sized backlit keyboard, and up to 11 hours of battery life. The base-model StarBook can be purchased now at \$929 and a fully spec'd model will run approximately \$1,852.

Although the StarBook won't compete with flagship Linux laptops (such as the System76 Galago Pro or Serval WS), it should serve as a quality mobile Linux



machine for on-the-go users. When you go to purchase your StarBook, you can select your choice of Linux distribution, from Ubuntu 20.04.2, elementary OS 5.1.7, Linux Mint 20.1, Manjaro 21.0, MX Linux 19.4, or Zorin OS 15.3 (Core or Ultimate).

Pre-order your StarBook now at <https://starlabs.systems/products/starbook#>.

Ubuntu 21.04 Adds Support for Active Directory Plus Other Major Changes

In a move that should surprise no one, Canonical has made it considerably easier for admins to join Ubuntu desktop machines to Active Directory domains and use Group Policy to set password requirements, user access controls, and even tweak desktop environment settings (such as login screen backgrounds and required applications).

Canonical has even made it possible for the integration of a Ubuntu Desktop into an existing Active Directory domain to be an automated and effortless process, with the help of the System Security Services Daemon (SSSD).

Active Directory isn't the only new feature that should be considered a major step forward for Canonical's desktop. The developers have finally shifted over to the Wayland graphics server by default. This change brings considerably faster performance over its predecessor, X.Org.

Finally, one long-rumored feature is the private home directory. Before Ubuntu 21.04, any user could view the contents of another user's home directory (but not make changes). Now, all home directories are private, so the permissions shift from 755 to 750. This particular feature will only be implemented in clean installs and not upgrades. Making this shift is important. According to Ubuntu's Security Tech Lead, Alex Murray (<https://discourse.ubuntu.com/t/private-home-directories-for-ubuntu-21-04-onwards/19533?u=d0od>), "...a lot of things have changed in the last 14 years, not least of which that Ubuntu has a significant customer and user-base in the public cloud and server space." Murray continues, "For these users, there is generally 1 admin account and perhaps a number of less privileged worker accounts, and so world-readable home directories now present more like a foot-gun than a feature – in this case, if a worker account is compromised, an attacker could now more easily access sensitive data from the other worker accounts or the admin account."

Download the ISO of Ubuntu Desktop 21.04 from the official Ubuntu Downloads page (<https://ubuntu.com/download>).

Gnome 40 Available on openSUSE

Gnome 40 is the latest iteration of the vaunted desktop but has yet to reach the majority of Linux distributions. Fortunately, those anxious to give the desktop a try need to look no further than openSUSE Tumbleweed, which is the rolling-re-

lease distribution that includes the latest and greatest software updates. Although Fedora 34 will also default to Gnome 40, that release is still in beta.

What's the hype about? Although the latest release of Gnome isn't exactly mind-blowing, it does deliver a much more efficient workflow, thanks to a horizontal flow within the Activities overview. The workspaces are now at the top of the overview, making it much easier for users to drag application windows to a specific space. Gnome 40 also makes it easier for you to open the Applications launcher and then open an app directly to the workspace you want to use. Again,



MORE ONLINE

Linux Magazine

www.linux-magazine.com

ADMIN HPC

<http://www.admin-magazine.com/HPC/>

Preload Trick

- Jeff Layton

A topic system that administrators learn as they gain experience, called the "LD_PRELOAD Trick," can help fix misbehaving applications, upgrade applications, and even improve application performance.

Rethinking RAID (on Linux)

- Petros Koutoupis

Often, you find yourself attempting to eke out a bit more performance from the computer system you are attempting either to build or recycle, usually with limited funds at your disposal.

ADMIN Online

<http://www.admin-magazine.com/>

Kafka: Scaling producers and consumers

- Jesse Yates

A guide to 10x scaling in Kafka with real-world metrics for high throughput, low latency, and cross-geographic data movement.

Protect privileged accounts in AD

- Evgenij Smirnov

Granular protection for highly privileged accounts is granted by the Protected Users group in Active Directory and Kerberos authentication policies.

Harden your Apache web server

- Christoph Mitasch

Cyberattacks don't stop at the time-honored Apache HTTP server, but a smart configuration, timely updates, and carefully considered security strategies can keep it from going under.

thanks to the horizontal workflow, this is a significant improvement over the previous iterations of the GNOME desktop.

Other features found in openSUSE Tumbleweed include the 5.11 kernel, Audacity 3.0, and the latest FWUPD (the firmware updater tool).

Read the official announcement from openSUSE (<https://news.opensuse.org/2021/04/16/gnome-40-kde-frameworks-plasma-update-in-tw/>), and then download a copy of openSUSE Tumbleweed (<https://get.opensuse.org/tumbleweed>).

Apple M1 Hardware Support to be Merged into Linux Kernel 5.13

Hector Martin has merged the initial support for Apple M1 hardware into the Linux system on a chip (SoC) tree. Martin is the founder of Asahi Linux (<https://asahilinux.org/about/>), a project to port Linux to Apple Silicon Macs. The project was started in 2020, using the M1 Mac Mini, MacBook Air, and MacBook Pro hardware. The Asahi goal is “not just to make Linux run on these machines but to polish it to the point where it can be used as a daily OS.”

Now that M1 support has been merged into the tree, it should make it into the Linux kernel for the 5.13 release (which should come out sometime this summer). That does not mean, however, you’ll be able to run Linux on Apple Silicon this summer. In fact, at the moment there is no timetable for full support. The reason for this is that porting Linux to Apple Silicon is a daunting task. Because Apple doesn’t release any documentation for the M1 hardware, everything must be reverse-engineered and drivers must then be written.

But as of April 8, 2021, the `arm/apple-m1` branch has been merged into `Linux-next` (the holding area for code expected for the next kernel merge window. To view the code that has been merged, take a look at this SOC commit (<https://git.kernel.org/pub/scm/linux/kernel/git/soc/soc.git/commit/?h=for-next&id=0d5fe4b31785b732b71e764b55cda5c8d6e3bbbff>). Although the Asahi Linux environment will now boot on the M1 hardware, it only provides serial and frame buffer console access. In other words, there’s a long way to go. And, according to Martin, “we absolutely do not recommend buying M1 hardware for that purpose unless and until the Asahi project gets much, much farther down the road than it has managed so far.”

KDE Launches the Qt 5 Patch Collection

At the end of 2020, Qt 6 was released to serve as the next-generation Qt application framework. This new iteration has made it possible to deliver more modern software and KDE has every plan to fully adopt this new release for the entire software stack.

However, KDE still very much relies on Qt 5 for both desktop and applications. With KDE’s goal of migrating to Qt 6, they had to do something to ensure nothing falls by the wayside. To that end, KDE has decided (until Qt 6 adoption is finalized), to maintain a collection of patches for the Qt 5.15 release. These patches will include both security and standard fixes to make sure KDE continues to remain secure and stable.

In regards to this, Aleix Pol, KDE e.V. President, said, “To transition to great future technologies like Qt 6 we need to have the peace of mind that our current users are catered for. With this patch collection, we gain the flexibility we need to stabilize the status quo. This way we can continue collaborating with Qt and deliver great solutions for our users.”

As for Qt 6, the plan is to have support sometime in 2021.

To find out more about the KDE Qt 5 Patch Collection, read the official initiative (<https://community.kde.org/Qt5PatchCollection>). To find out where KDE stands with Qt 6, check out the Phabricator (<https://phabricator.kde.org/project/board/310>).



**Get the latest news
in your inbox every
two weeks**

**Subscribe FREE
to Linux Update
bit.ly/Linux-Update**



Zack's Kernel News



Chronicler Zack Brown reports on the latest news, views, dilemmas, and developments within the Linux kernel community.

By Zack Brown

Want vs. Need

Ryan Houdek wanted to enhance compatibility layers in Linux. A compatibility layer is used when you have a piece of software that was compiled to run on a different system, and you want it to run on yours. Maybe the software expects a certain system file to exist, or certain opcodes at the CPU level, or certain system calls. A compatibility layer will provide those things so the software can run. A lot of cloud service companies like Google and Amazon use Linux's compatibility layers to make one piece of hardware look like a whole bunch of other pieces of hardware.

So compatibility layers are not new in Linux, but Ryan wanted to run old software compiled for 32-bit CPUs on 64-bit systems and offered up a general justification for compatibility layers. One of his main points was that "Not all software is open source or easy to convert to 64-bit," and that a lot of gaming software fell into this category.

Ryan pointed to various attempts in the Linux world to work around these problems, such as Qemu, a generic CPU emulator. But the problem with such attempts, he said, was not emulating the CPU, it was emulating various system resources such as memory handling and input/output controls.

He posted a patch to address the whole issue in what he felt was a more comprehensive way, by exposing compatibility system calls to user space. System calls are the library of functions the kernel provides so that user space can use the hardware and other resources on a given system. Ryan wanted to create a new set of system calls that would behave the way older software expected.

This was always going to be an iffy proposition.

Steven Price, for example, remarked, "Running 32-bit processes under a compatibility layer is a fine goal, but it's not clear why the entire 32-bit compat syscall layer is needed for this." And he added, "QEMU's user mode emulation already achieves this in many cases without any changes to the kernel."

Steven went on to say, "Clearly there are limitations, but in general I can happily 'chroot' into a distro filesystem using an otherwise incompatible architecture using a qemu-xxx-static binary."

So exposing an entirely new set of system calls didn't appeal to Steven, though he agreed that memory and input/output handling were serious issues when it came to any sort of compatibility layer.

In particular, Steven agreed that "ioctls are a mess."

Input/output controls (ioctls) are a nightmarish fantasy of one of the outer gods, possibly Nyarlathotep. Ioctls exist in the nether region between what you need the hardware to do, and what system calls are able to provide.

The `ioctl()` system call is intended to be extended by device driver writers, so that the inputs to `ioctl()` can be relevant to any particular device driver. In this way, the kernel doesn't need to be loaded up with new system calls every time a new device comes out on the market. Instead, the single `ioctl()` system call can take all of that malignant energy unto itself, growing darkly beneath the surface for all time. If you asked a kernel developer about documenting all the behaviors of `ioctl()`, they would begin to laugh, cry, and explode simultaneously. Try it and see. Or don't. They have suffered enough.

Steven remarked, "ioctls are going to be a problem whatever you do, and I don't think there is much option other than having a list of known ioctls and translating them in user space."

Steven also agreed that memory handling was difficult to manage, in terms of converting between 32- and 64-bit systems. But he also mentioned, "I've seen examples of `MAP_32BIT` being (ab)used to do this, but it actually restricts to 31 bits and it's not even available on `arm64`. Here I think you'd be better off focusing on coming up with a new (generic) way of restricting the addresses that the kernel will pick."

In any event, Steven said, even exposing a full system call compatibility

layer would not save anyone from having to do “a load of fixups in user space due to the differing alignment/ semantics of the architectures.” And he pointed out, “You are already going to have to have an allow-list of ioctls that are handled because any unknown ioctl is likely to blow up in strange ways due to the likes of structure alignment differences.”

Meanwhile, Mark Brown pointed out that a certain amount of compatibility layering did already exist in Linux. He said, “this has been deployed on Debian for a long time – you can install any combination of Debian architectures on a single system and it will use qemu to run binaries that can’t be supported natively by the hardware.”

In response to that, Catalin Marinas remarked, “The only downside I think is that for some syscalls it’s not that efficient. Those using struct iovec come to mind, qemu probably duplicates the user structures, having to copy them in both directions.”

Regardless, Catalin opposed Ryan’s patch, saying “Those binary translation tools need to explore the user-only options first and come up with some perf numbers to justify the proposal.” The perf tool is for performance analysis of the Linux kernel. It’s common to see developers asking for the perf numbers for a given patch, to make sure it doesn’t slow things down too much.

Arnd Bergmann also spoke in favor of Qemu, saying, “qemu does a nice job at translating the interfaces for many combinations of host and target architectures at a decent speed, and is improving at both the compatibility and the performance over time.”

But he also acknowledged:

“The ioctl emulation in qemu is limited in multiple ways:

- *It needs to duplicate the kernel’s compat emulation for every single command it wants to handle, and will always lag behind what gets merged into the kernel and what drivers a particular distro ships.*
- *Some ioctl commands cannot be emulated in user space because the compat code relies on tracking device state in the kernel.*
- *In some cases, emulation can be expensive, both for runtime overhead and for code complexity.”*

Arnd opposed Ryan’s patch as well, or at least was not convinced it was needed. He thought it might be better to try to address the ioctl insanity on its own, rather than the entire system call layer. And once ioctls had been brought under control, Arnd felt, the rest of the system call layer would not pose many more problems.

David Light also had issues with Ryan’s patch. As Catalin pointed out already, performance would be a significant question. But David said, “I don’t think the problem is only the performance. The difficulty is knowing when structures need changing. A typical example is driver ioctl requests. Any user space adaption layer would have to know which actual driver has been opened and what internal structures it has. [...] It is much easier to get it right in the code that knows about the actual structures.”

Mark Rutland also made the point that exposing a full set of compatibility system calls might be more than Ryan actually needed. For example, all Ryan really needed, Mark R. said, was “being able to limit the range of mmap() and friends.” He went on, “I think that for this series x86 emulation is a red herring. ABIs differ across 32-bit arm and 32-bit x86 (and they have distinct arch-specific syscalls), and so those need distinct compatibility layers. If you’re wanting to accelerate x86 emulation, I don’t think this is the right approach.”

Mark R. went on to say that, in fact, a more targeted approach could actually benefit more projects. He said, “For example, having variants with an explicit address mask would also benefit JITs which want to turn VA bits into additional tag bits.”

Finally, Mark R. voiced his opposition to Ryan’s patch, adding, “However, I do think that we can make emulation easier by extending some syscalls (e.g. mmap()), and that this would be worthwhile regardless of emulation.” He concluded, “I appreciate that people have 32-bit applications, and want to run those, but I’m not convinced that this requires these specific changes. Especially considering that the cross-architecture cases you mention are not addressed by this, and need syscall emulation in userspace; that implies that in general userspace needs to handle conversion of semantics, and

what the kernel needs to do is provide the primitives onto which userspace can map things in order to get the desired semantics (which is not the same as blindly exposing compat syscalls).”

Amid all the voices expressing concern, Amanieu d’Antras said he liked Ryan’s patch and thought it should be accepted.

First of all, Amanieu felt that speed and efficiency were not as important as others had suggested, and that the main point was correct emulation. And the bottom line, he said, was that user space simply “does not have the information or the capabilities needed to ensure that the 32-bit syscall ABI is correctly emulated.”

Amanieu also pointed out that while exporting a compatibility system call layer would not solve all the problems, it would allow emulators to take up the various slack of “memory management, signal handling, /proc emulation, ptrace emulation, etc.” With that done, the emulator could pass the resulting system call through to the user code, where it would properly do its thing.

He listed a bunch of technical requirements that he felt couldn’t be done in user space, saying, “these issues are all solved by exposing compat syscalls to 64-bit processes and ensuring is_compat_task/in_compat_syscall is true for the duration of that syscall.”

Amanieu said something similar was already done in Linux, specifically “on x86, syscalls made with int 0x80 are treated as 32-bit syscalls even if they come from a 64-bit process.”

Mark R. sympathized with the point that “there are cases where the emulator cannot do the right thing due to lack of a mechanism.” But rather than try to compensate for this by exposing more mechanisms from the kernel, Mark R. said, “where the emulator does not have knowledge, I don’t think that it can safely invoke the syscall.”

He went on to say that there were numerous cases where the kernel would be unable to determine what the correct behavior might be. And that, “the kernel cannot possibly do something that is correct in this case.” He concluded, “Maybe your emulator avoids these, but that’s no justification for the kernel to expose such broken behaviour.”

Amanieu completely disagreed with Mark R.’s take on the situation, saying,

“I disagree that any broken behavior is exposed here.”

At that point, the conversation came to an end.

It’s doubtful that anything short of absolute necessity would lead Linus Torvalds to accept a full-on system call compatibility layer into the kernel. It would lock Linux into emulating ancient hardware essentially forever.

But Ryan and his supporters have some valid points as well. Aside from everything else, it’s fundamentally good to be able to run compiled binaries where the source code has been lost, and where the only other way to run that code would require a chronometric displacement aperture of human proportions, in which case we would have other problems.

Ultimately, there seems to be plenty of room for compromise – I would guess Linus would want the kernel to export anything that was truly needed by user space emulators, and the emulator people will be satisfied with that since it will let them do their main thing.

Hiding System Resources ... from the System

Mike Rapoport recently submitted a patch to implement secret memory areas in Linux. As he explained, “The pages in that mapping will be marked as not present in the direct map and will be present only in the page table of the owning mm.” And he went on, “such secret mappings are useful for environments where a hostile tenant is trying to trick the kernel into giving them access to other tenants’ mappings.”

Part of Mike’s idea was to hide the secret memory in a file – or at least to make it accessible via a file descriptor, as if it were an ordinary file. He explained:

“Hiding secret memory mappings behind an anonymous file allows usage of the page cache for tracking pages allocated for the ‘secret’ mappings as well as using address_space_operations for e.g. page migration callbacks.

“The anonymous file may be also used implicitly, like hugetlb files, to implement mmap(MAP_SECRET) and use the secret memory areas with ‘native’ mm ABIs in the future.”

In Mike’s vision, the feature would be disabled by default and would require a

boot-time command-line argument to activate.

David Hildenbrand had some questions. Among other things, he said, “secret” memory allocations would be invisible to the various memory management features in the kernel. This would mean that blocks of memory would be sitting immobile in RAM, while other blocks could be moved around and reorganized as needed. So a lot of Linux features, such as process migration and other cool things, might be tripped up by “secret” memory.

Mike clarified, saying that actually, “secret” memory (which he called `secretmem`) would only be allocated from a region of RAM that was never moved around anyway. So it wasn’t that `secretmem` was especially unmovable; it was that the kernel already had unmovable memory, and `secretmem` used that.

This didn’t satisfy David, and the two – along with Michal Hocko – embarked on a technical discussion of different types of memory and the conditions under which they were used. As Michal put it:

“A lot of unevictable memory is a concern regardless of CMA/ZONE_MOVABLE. As I’ve said it is quite easy to land at the similar situation even with tmpfs/MAP_ANON|MAP_SHARED on swapless system. Neither of the two is really uncommon. It would be even worse that those would be allowed to consume both CMA/ZONE_MOVABLE. One has to be very careful when relying on CMA or movable zones.”

Throughout the discussion, everyone also made reference to various aspects of memory handling that needed to be documented somewhere in the kernel. So, the lack of that documentation may have made the conversation more difficult.

Another element that may have made the conversation difficult was the fact that other areas of the kernel have to deal with similar issues – hot plugging, for example, as Michal pointed out.

There was also the question of how users could analyze their own systems. As David said at one point, “With plenty of `secretmem`, looking at `/proc/meminfo` Total vs. Free can be a big lie of how your system behaves.”

David also pointed out that at least in its current form, `secretmem` might allow processes to hog memory that

should stay available to other processes. As he put it, “`Secretmem` gives user space the option to allocate a lot of `GFP_{HIGH}USER` memory. If I am not wrong, ‘`ulimit -a`’ tells me that each application on F33 can allocate 16 GiB (!) of `secretmem`. Which other ways do you know where random user space can do something similar? I’d be curious what other scenarios there are where user space can easily allocate a lot of unmovable memory.”

And at one point, David gave some of his own context for the discussion:

“I am constantly trying to fight for making more stuff MOVABLE instead of going into the other direction (e.g., because it’s easier to implement, which feels like the wrong direction).

“Maybe I am the only person that really cares about ZONE_MOVABLE these days :) I can’t stop such new stuff from popping up, so at least I want it to be documented.”

Though Michal confirmed that “MOVABLE zone is certainly an important thing to keep working. And there is still quite a lot of work on the way. But as I’ve said this is more of an outlier than a norm. On the other hand movable zone is kinda hard requirement for a lot of application[s] and it is to be expected that many features will be less than 100% compatible. Some usecases even impossible.”

The technical discussion continued for quite awhile, with James Bottomley joining in at one point, with some virtual machine developer perspective.

Ultimately the discussion ended inconclusively. It’s clear that there are many perspectives to consider and many stakeholders within the kernel, as well as a general lack of documentation and a lack of clarity on the behavior of various parts of the system.

`secretmem` is also a security feature, but not an essential one. It doesn’t seem to plug any particular hole, but is intended more as camouflage to guard against the possibility of an attack. As such, it’s not the sort of feature that Linus Torvalds tends to leap towards. Linus seems to be more focused on closing known security holes, rather than on reducing the surface area of potential attacks. So for now, it’s unclear if `secretmem` could make it into the kernel, either in this form or another. ■■■

Virtual flash cards

Card Game

Anki brings a virtual flash card box to the desktop. Thanks to a useful collection of add-ons, you can adapt Anki to suit your needs, making it one of the most efficient learning tools. *By Erik Bärwaldt*

The flash card is still one of the most effective methods for anchoring learning material in long-term memory. But instead of working with a physical index card box and decks of flash cards, today's students use computer programs. Anki [1], considered to be one of the most innovative learning programs for electronic flash cards, goes far beyond everyday needs and even offers a web-based counterpart.

Anki includes several convenient features that you won't get when shuffling through an old-school stack of note cards, such as built-in statistics to chart your progress. The program also lets you rate your success with individual cards, keeping the most difficult at the front of the deck for additional repetitions. The Anki: flashcard system also comes with a large number of add-ons that let you adapt the software to meet a wide range of learning needs. For learning foreign languages, Anki offers the ability to use different character sets, such as Chinese or Russian. Anki uses LaTeX to display formulas, making it ideal for higher mathematics. In addition, Anki can also handle multimedia content; unlike conventional flash cards, Anki can integrate videos, audio files, and images as learning content.

Installation

Anki is available in the repositories of all the common distributions. You can also download a tarball for 64-bit systems from the Anki website, which is recommended if your distribution's repos offer an older Anki version. To set up Anki, unpack the roughly 150MB tarball, which will add about 460MB of data to your hard disk. Then change to the newly created directory and install the package using:

```
sudo make install
```

Then start Anki by either selecting the entry in the start menu or entering the `anki` command in a terminal window.

If you want to compile Anki manually, the source code is available on the project's GitHub page [2]. You will

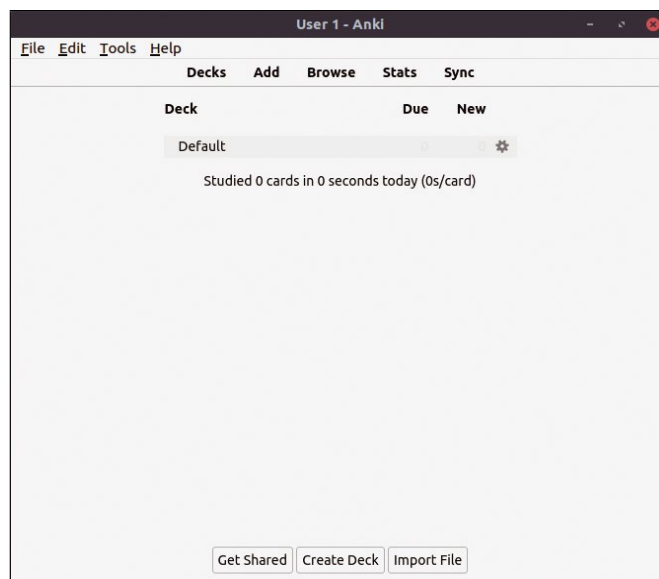
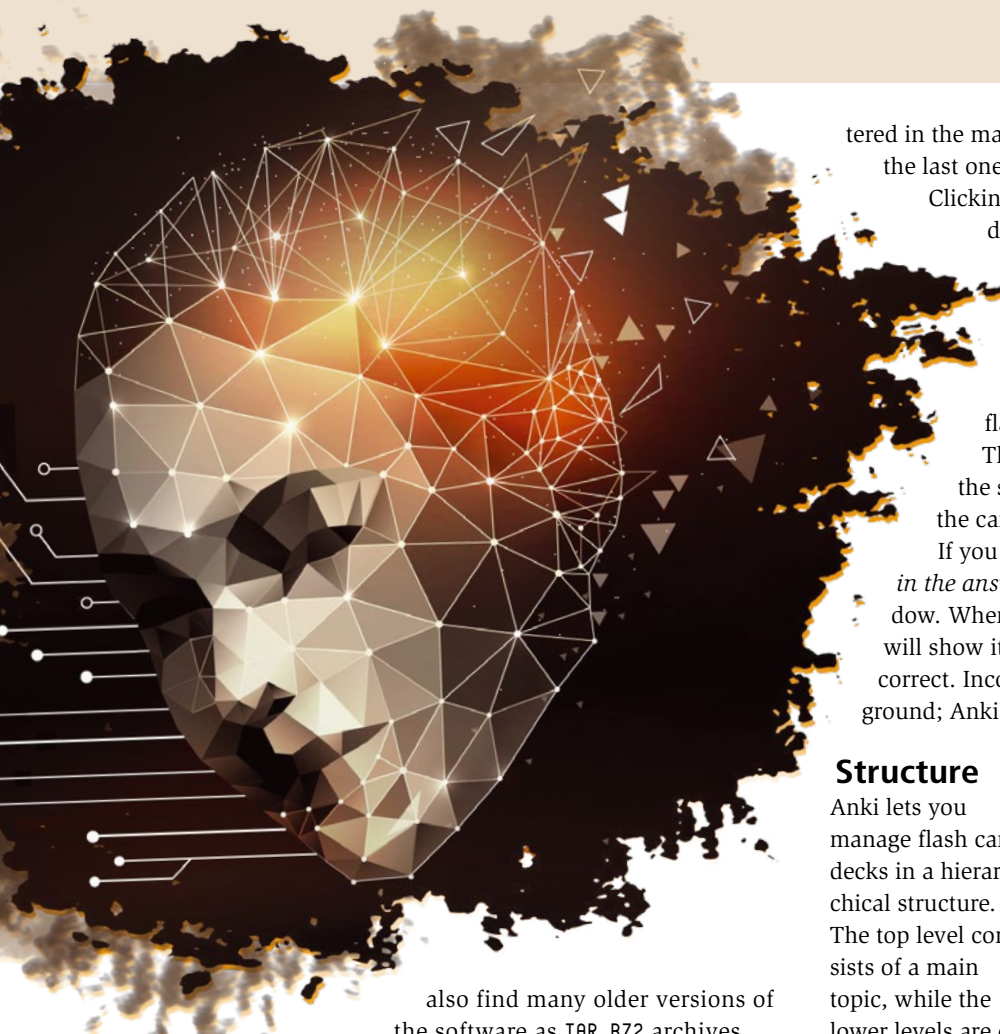


Figure 1: Anki's main window where you start adding virtual flash card decks.



also find many older versions of the software as TAR.BZ2 archives, which can be useful if the latest version of Anki cannot be installed on older Linux systems due to missing dependencies.

Operation

Upon initial start-up, Anki opens a language-localized window, so you don't have to make any adjustments (Figure 1).

To get started, you need to create a deck of virtual flash cards with the corresponding content on their front and back sides. Click the *Create Deck* button bottom center. Enter a name for the deck in the window and confirm it.

After clicking on *Add* below the menubar, you can start entering the contents for each card. The input window consists of two fields for free text input for the front and back of the flash cards; it also allows formatting of the entered data via a small buttonbar above the input fields. You can format the text input, define a background color, link a multimedia file to the index card, or play an audio recording, which is useful for practicing the correct pronunciation of foreign language terms. In addition, the software lets you to enter short fill-in-the blank text. You can also insert chemical or mathematical formulas as needed (Figure 2) using the hamburger menu on the right side of the buttonbar.

The *Type* selection field lets you specify the form in which the entered terms should be queried. You choose whether Anki queries the cards on both sides or only the front or back side at any given time. Once you have entered the front and back text, click on the *Add* button to save your card.

After completing the input of a deck of cards, Anki displays the name of the deck with the number of flash cards you en-

tered in the main window. If other decks already exist, only the last one to be activated is shown in the main window.

Clicking *Decks* in the top buttonbar opens the deck display, where all the existing decks are shown, including the decks that need to be reviewed and the new decks to be learned.

Click on a deck to activate it. Next, click *Study Now* to the right of the deck display to start querying. The program switches to a flash card view where the first term appears. Think about the correct answer and then press the space bar to reveal the answer on the back of the card (Figure 3).

If you want to type in your answer, select *Basic (type in the answer)* in the *Type* field in the Add dialog window. When you type in your answer on a card, Anki will show it with a green background if your answer is correct. Incorrect answers are shown with a red background; Anki then displays the correct answer below.

Structure

Anki lets you manage flash card decks in a hierarchical structure. The top level consists of a main topic, while the lower levels are divided into sub-topics. To do this, you create the individual sub-topics as separate decks and then drag and drop them onto the main deck.

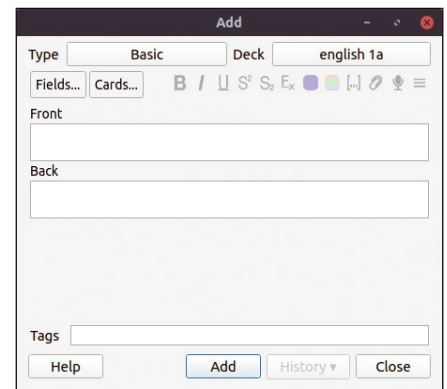


Figure 2: Anki's input dialog offers various options for formatting questions for your flash cards, including fill-in-the-blank text and multimedia files.

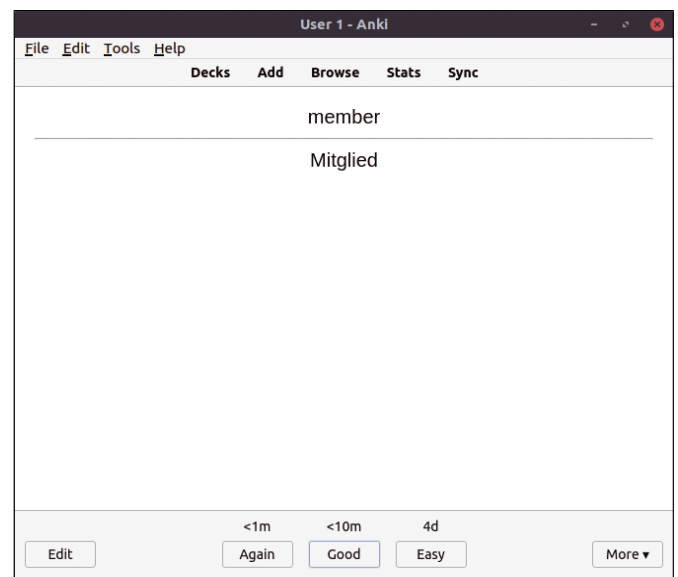


Figure 3: The front and rear sides of each flashcard appear in the program window.

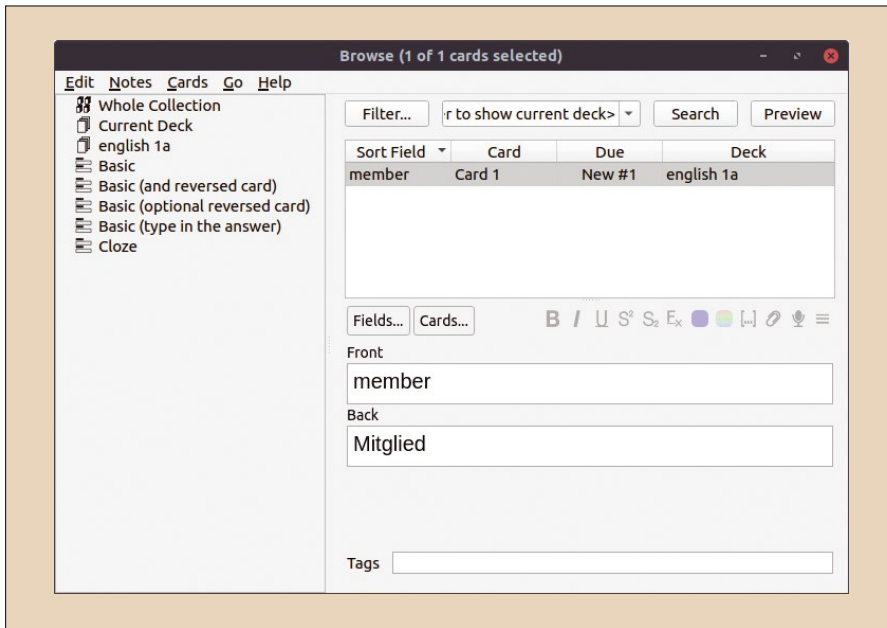


Figure 4: You pause your exercises via *Browse*.

Alternatively, when creating the decks, you can define sub-topics by first specifying the name of the main deck when assigning names followed by two colons (::) and then the name of the sub-deck. Anki then displays each of the sub-decks indented below the main deck.

You should define your deck structure at the beginning of data input. When creating a hierarchical structure, Anki hides sub-decks entered in the main deck. The sub-decks can only be viewed via *Browse* in the main window's top buttonbar.

However, after creating a tree structure, new flash cards can still be created in the main deck. To do this, select the desired target deck in the *Decks* selection field of the Add dialog window. Anki adds the flash card (front and back) when distributing the flash card to multiple sub-decks and then displays the total number of new flash cards in the main deck's overview.

You can also create an extended structure when entering the learning content in the individual decks. Especially for complex topics, it is advisable to use keywords or tags. The *Tags* field lets you enter several keywords separated by a space. These keywords can also be targeted selection criteria later on when learning the content on the individual cards.

Hands On

After you select a deck to study, Anki displays three categories, each containing several cards. *New* denotes the first level in which all newly entered cards are located. In *Learning*, Anki groups all the cards that are in an ongoing initial learning process. The cards that you have already learned but now want to review for long-term memory appear in the *To Review* group. Anki takes into account all the different types of cards and prompts for responses in a mixed order. Anki's algorithm assumes a daily learning speed of 100 cards, including review.

If you plan to take a break in studying, it is recommended to pause the decks. To do this, click the *Browse* link in the vertical bar on the left in the new window (Figure 4). On the right,

you will see a list of all the cards in this deck. You can now select the desired cards and then right click on the selected deck. From the context menu, select the option *Toggle Suspend* or press Ctrl + J.

After closing the window, you will no longer find any cards designated for study in the selected deck in the deck overview. To un-pause, repeat the same steps as for pausing. Afterwards, the number of cards to be learned appears again in the deck overview.

Spaced Repetition

Anki's learning method uses spaced repetition, where memorization after several repetitions spread over a period of time anchors a topic in long-term memory (see the "The Leitner System" box). To implement this concept, the Anki developers introduced a rating system consisting of four levels that is applied individually for each card.

On the first level, click *Again* if you are unable to answer the question (or had great difficulty doing so); Anki will call this card again within the next 10 minutes. If you are able to answer the question, click the *Good* button. If this is the first time *Good* has been selected for a flash card, Anki transfers the card to the second level, which causes the card to be called up again after three days. If you click *Good* again for the same card on subsequent review, Anki moves the card to the third level and calls it again after five days. If you select the fourth level, *Easy*, Anki extends the recall interval to six days.

As learning outcomes improve, the call intervals also lengthen, ensuring a lasting learning effect through regular spaced repetition.

The Leitner System

The Leitner learning system, developed in 1973 by Sebastian Leitner, implements a form of spaced repetition using index cards separated into a series of boxes. The card index boxes let users review the cards at regular intervals, with the review intervals increasing from box to box. You decide which level Anki assigns to each card by pressing the *Again*, *Good*, and *Easy* buttons at the bottom of each flash card.

In daily learning, the system ensures that a different number of cards is retrieved from each box. While new cards are initially located in the first box, cards that have already been queried and successfully answered move to the back of the next box. Cards that are not answered or answered incorrectly find their way back into the first box. This ensures that terms that are difficult to learn are initially queried more frequently, while cards that have already been answered correctly over a longer period of time (and therefore successfully learned by the user) are queried less frequently.

For lasting learning success, you need to use the Leitner system at regular intervals, querying a maximum of 100 terms a day that you want to learn permanently.



Statistics

To help you keep track of not only your study load, but also your learning and memory performance, Anki offers a statistics feature that uses bar and pie charts to visualize your progress.

To generate the statistics report, click on *Stats* in the top toolbar in the main window; you then see all the information in a separate window (Figure 5). In this window, you can narrow the statistics to a specific deck or a deck collection as well as a time period (either the last 12 months or the entire learning history).

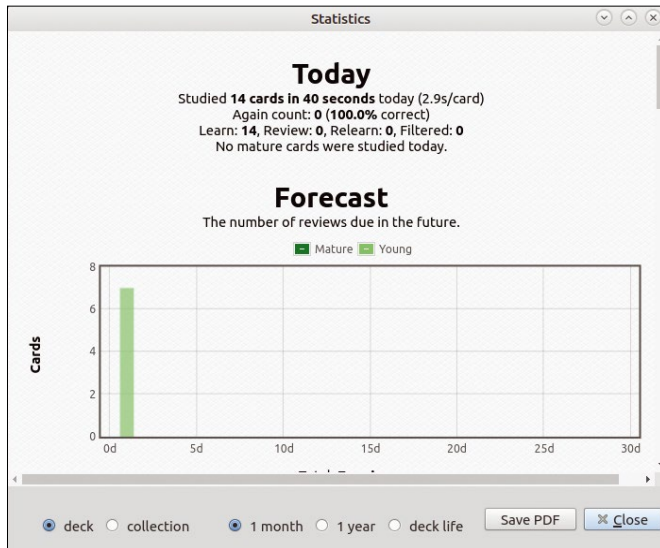


Figure 5: Comprehensive statistics in Anki document the user's learning success.

In addition to these statistics, Anki displays the most relevant data over several different time intervals in a colorful, grid-style calendar overview. To save the report permanently, use the *Save PDF* button at the bottom. After requesting the storage path and file name, Anki creates a corresponding PDF file that contains all the details, making it quite extensive.

Add-Ons

Anki can be extended quite easily with additional modules. A dedicated web page [3] lists numerous add-ons (Figure 6). These add-ons serve not only to visually customize the program interface and also extend the software with additional functions.

To install an add-on, select *Tools | Add-ons* in the menubar of the main program window. In the dialog that opens, you first see an empty segment on the left, where the integrated add-ons will appear later. On the right, there are several buttons for managing the extensions.

On the left side of the dialog, click *Get Add-ons*. Another small dialog will open where you click *Browse Add-ons* at the bottom. The application then opens the web page for the available add-ons. After clicking on one of the listed add-ons, you see a detailed description, including a specific code that you will find below the description. To install the add-on, enter this code in the corresponding field of the Install Add-on window. After clicking *OK*, Anki installs the add-on and adds it to the list view. To enable the add-on, restart Anki.

Import and Export

Anki lets you import data from various other learning programs and export your own data to other formats. Use *Import*

Title	Ratings	Modified	Anki
AnkiConnect	158	2021-02-22	
Button Colours (Good, Again)	146	2020-02-08	2.1.16~
Speed Focus Mode (auto-alert, auto-reveal, auto-fail)	174	2019-06-01	
Searching, PDF Reading & Note-Taking in Add Dialog	140	2021-02-16	<=2.1.40
Edit Field During Review	107	2020-10-06	2.1.33~
Batch Editing	106	2019-06-02	
AwesomeTTS (Google Cloud Text-to-Speech) [unofficial]	217	2020-10-19	2.1.35~
Mini Format Pack	128	2018-07-24	
Custom Background Image and Gear Icon	152	2020-12-27	2.1.35~
No Distractions Full Screen: Clean review interface + Tablet/Touch support	82	2020-08-13	2.1.29~
Symbols As You Type	78	2020-04-30	2.1.24~
Replay buttons on card	191	2017-11-19	2.1.16~
The KING of Button Add-ons	71	2020-04-17	2.1.24~
Advanced Review Bottom Bar	70	2020-12-05	2.1.35~
Anki Simulator	69	2021-02-21	2.1.38~
Customize Sidebar	45	2020-01-15	2.1.18~
Dancing Baloney: Allows changing wallpapers (with live preview)	64	2020-03-26	<=2.1.35
Advanced Copy Fields	40	2019-05-03	
Leaderboard	79	2020-12-09	2.1.34~
Progress Graphs and Stats for Learned and Matured Cards	60	2020-03-28	2.1.23~

Figure 6: Anki can be extended quite easily with a variety of add-ons.

and *Export* in the File menu for this purpose. Besides some third-party formats, Anki also has the ability to import media files and CSV files created with LibreOffice or OpenOffice. Since these need to meet certain requirements, such as a uniform character set, it is worth taking a look at the documentation before importing data.

AnkiWeb

AnkiWeb [4], a web-based variant of Anki, makes it possible to learn content and maintain the database independently of an individual computer.

However, AnkiWeb, which is primarily intended as a supplement to the locally installed program, only lets you create and learn with text-based flashcards. Different deck hierarchies and multimedia content are not supported by AnkiWeb. You also need to log in to the appropriate page first to use AnkiWeb.

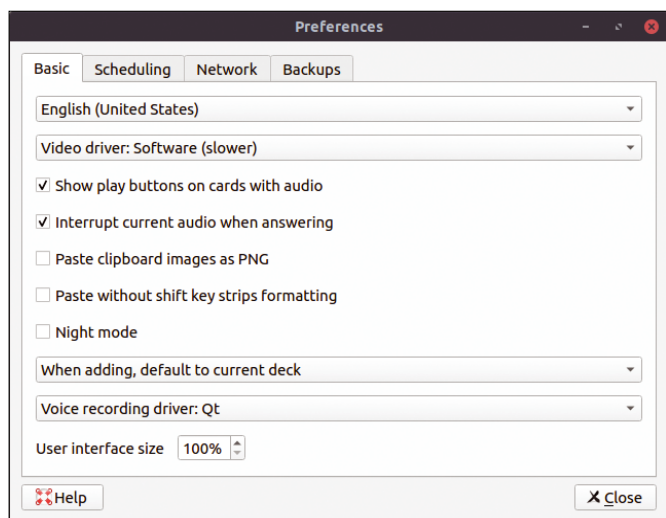


Figure 7: Anki's Preferences dialog is very clear considering its range of functions.

The main goal of AnkiWeb is to back up existing flash card decks in the cloud and sync them across multiple computers. Anki does this by synchronizing the data files between the individual computers every time the application is opened and closed.

Profiles

If more than one person uses the locally installed version of Anki, a separate profile can be created for each via *File | Switch Profile*. To do this, just click the *Add* button and enter a profile name. To select a profile, double-click on the profile name.

To configure a profile, select *Tools | Preferences...* when the profile is active. You can configure each profile independently, with Anki saving the adjustments along with the profile (Figure 7).

Conclusions

Anki currently offers probably the most comprehensive program for learning terms and content based on the Leitner flash card system. The software appears to be very sophisticated in terms of functionality and, unlike some other programs based on the Leitner system, is not limited to simple index cards with individual terms and short texts. It also lets you create complex learning processes including multimedia content. Consequently, Anki is recommended not only for students, but for anyone interested in quick and sustainable learning. ■■■

Info

[1] Anki: <https://apps.ankiweb.net>

[2] Anki on GitHub: <https://github.com/ankitects/anki>

[3] Anki Add-ons: <https://ankiweb.net/shared/addons/>

[4] AnkiWeb: <https://ankiweb.net/about>

Author

Erik Bärwaldt is a self-employed IT-admin and technical author living in Scarborough (United Kingdom). He writes for several IT-magazines.

Public Money

Public Code



Modernising Public Infrastructure with Free Software



Free Software Foundation Europe

Learn More: <https://publiccode.eu/>

Create diagrams with Dia

Perfect Presentation

If you need to make the occasional diagram, Dia gives you professional results with a minimal learning curve. *By Erik Bärwaldt*

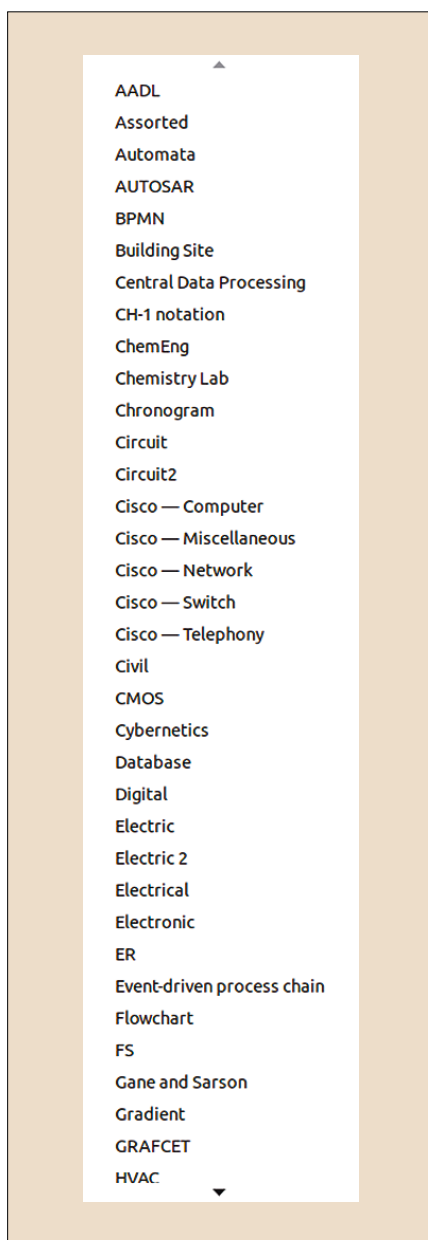


Figure 1: Dia offers additional graphical elements for multiple disciplines.

Diagrams make it easier to understand complex issues. Creating such graphics using LibreOffice Draw or Gimp can quickly test your patience, as these programs offer an extremely large number of functions, which requires a massive learning curve.

For the occasional diagram, what you really need is an easier-to-use solution that gives usable results fast. With Dia Diagram Editor [1], professional-looking graphics and diagrams can be designed quickly with hardly any training.

Installation

Dia, which is based on the GTK+ toolkit, can be found in popular Linux distribution's repositories and set up with the respective package manager. You will typically get the latest available version since the current version, Dia v0.97, dates back to 2013. You can also find other (including older) versions on SourceForge, including various BSD derivatives, macOS, and Windows. Once installed, you will find Dia in the Applications section of your desktop's main menu.

Getting Started

When you launch Dia, a clear-cut program window opens. On the left, you will find a toolbar with numerous graphical elements that can be integrated into the diagrams. At the top below the menubar, a buttonbar allows quick access to important functions. The workspace, or canvas in Dia parlance, on the right is divided into a grid of boxes that allows precise positioning of graphic elements.

You can start using Dia without any further configuration. To create your first diagram element, use the object icons located in the lower half of toolbar on the left. If you don't see the object you need, click on the selection field in the middle of the toolbar. From the context menu that opens, select *Other sheets* to access additional object collections from a wide variety of disciplines (Figure 1), such as civil engineering, chemical engineering, electrical engineering, electronics, IT, lighting, hydraulics, and circuits.

To create a diagram, drag the desired object from the object collection to the canvas. The object will now have small green square handles at each corner and in the center of each side. You can change the size and aspect ratio of an object by clicking on a handle and then dragging it while holding down the left mouse button. You can reposition the object by clicking inside the object, holding down the left mouse button, and dragging the entire object to the desired location.

If you've placed several objects on the canvas, you can connect them with lines or geometric shapes using the drawing tools located at the top of the toolbar on the left. Click the desired tool and then click on the starting point for your line (or shape). While holding down the left mouse button, drag the mouse to the endpoint, and the corresponding line or geometric shape will now appear on the canvas. In this way,



you can also integrate freehand symbols into the diagram and connect the individual objects with lines.

Additional Libraries

If you can't find the object you need, you can extend Dia with additional object collections with a separate repository that offers still more object collections [2].

You can integrate these collections into your existing application using the Diashapes tool. You can download the Diashapes tool for 32- and 64-bit systems from the project's website [3]. After installing the tool, integrate the additional libraries into your Dia installation (Figure 2).

Labels

Dia supports free labeling of diagrams. To do this, select the *T* icon from toolbox (top left). Click on the canvas where you want to place label and then start inputting text. Double-clicking on the text opens a small dialog window where you can format the text, including font, size, and fill color. You can reposition your label on the canvas by dragging the text while holding down the left mouse button.

Images

Aside from the existing objects, Dia also lets you integrate image files (in common formats such as JPG, PNG, or SVG). Drag the desired image from your file manager directly to Dia's canvas. Dia

first positions the image in the top left corner, but you can drag it to the appropriate location. Image files can also be modified like other objects using the small green handles (Figure 3).

Point of View

With very detailed graphics such as electronic circuits, you will not be able to see all the individual elements that are displayed on the canvas on a conventional monitor. To remedy this, you can use the zoom in and zoom out icons located in the buttonbar at the top to change the view size. Alternatively, you can use the View menu.

For accurate positioning and adjustment of objects, Dia superimposes a grid on the entire canvas by default, with rulers at the edges. If you don't want to see the grid or rulers, go to the View menu and uncheck *Display grid* and *Display rulers*. From the View menu, you also can select *AntiAliased* to smooth out edges on enlarged bitmaps. Clicking on the magnifying glass icon in the toolbar enables magnification. Unlike the zoom function, this tool only enlarges the area that you select on the desktop after enabling the function.

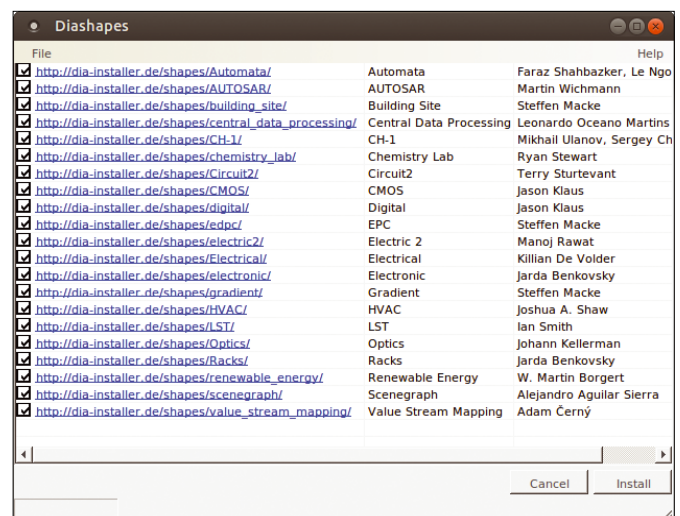


Figure 2: The Diashapes tool lets you extend your Dia object toolbox.

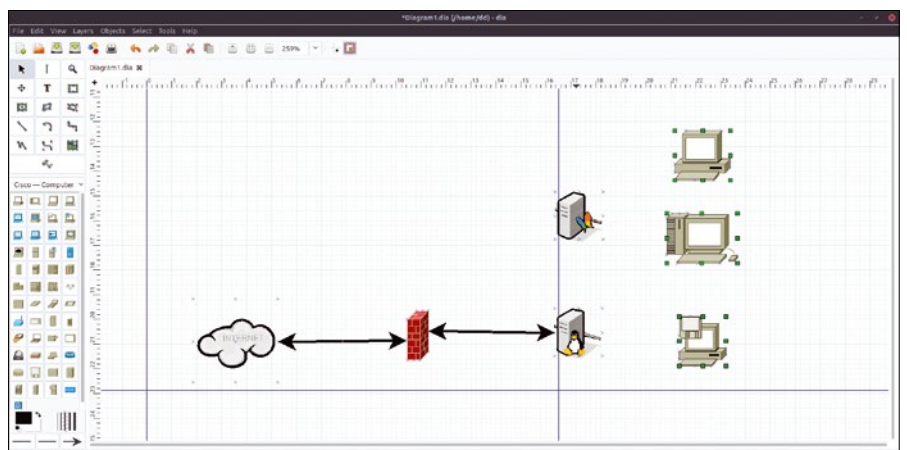


Figure 3: Dia lets you use images in your diagrams.

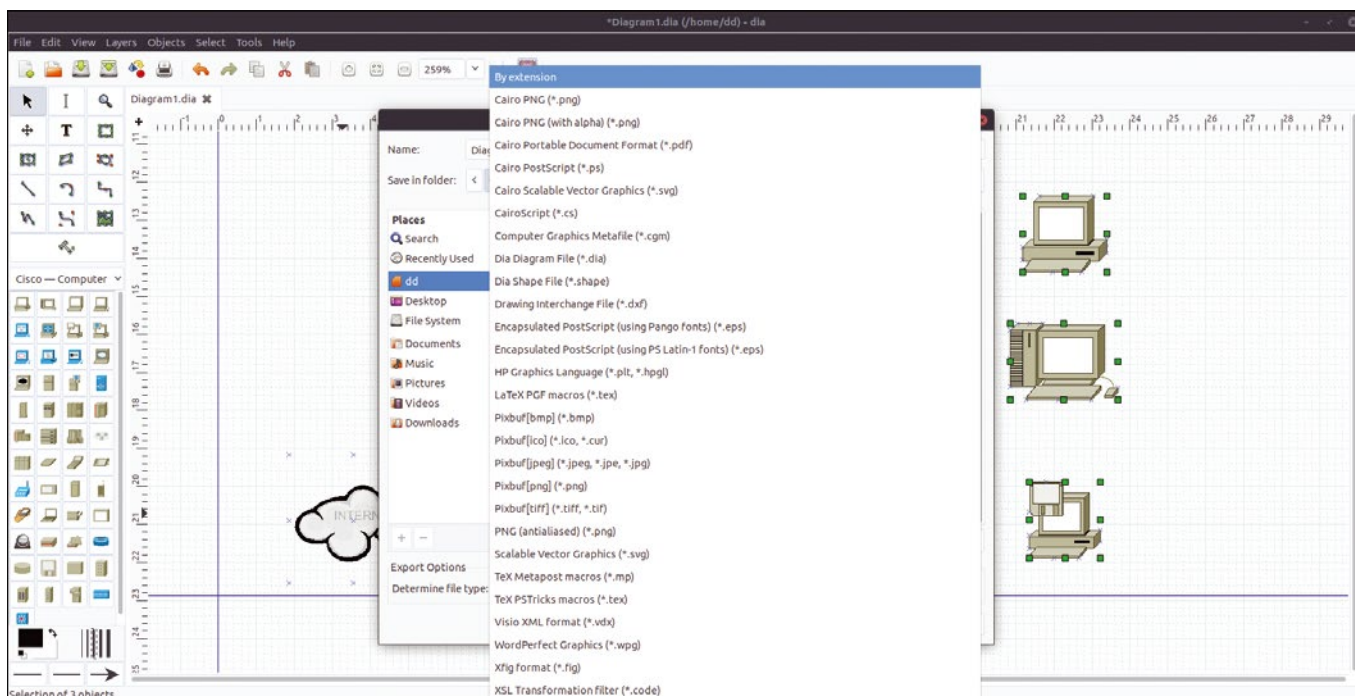


Figure 4: Dia exports diagrams to numerous formats without any problems.

Dia lets you to keep several diagrams open simultaneously for editing. While the toolbar and buttonbar remain the same, the canvas has a tabs along the top that let you toggle between the canvases.

Export

Dia has very powerful export filters, which means that you can integrate diagrams into other applications using a standardized image format. *File | Export...* lets you open a small file manager for this purpose. In the selection field at the bottom, you

specify the desired target format (Figure 4). The diagram can be converted to PDF, PNG, JPEG, BMP, and CGM, as well as several PostScript variants. In addition, the DXF format used in AutoCAD CAD software is also available.

Printing

While the canvas onscreen always appears in landscape format, the finished diagram does not have to print or be saved in that format. The *File | Page Settings...* dialog lets you modify paper size, print orientation (landscape or portrait), and scale (Figure 5). The changes made here only apply to the current diagram. If you want to change the default settings for all new drawings, you can modify the settings in the *File | Settings* dialog. Among other things, you can specify the default paper size and orientation.

Conclusions

Dia supports an intuitive work approach and focuses on the essentials without requiring extensive training. You can quickly create diagrams using numerous ready-made objects and add text, geometric elements, connecting lines, or even images. The powerful export function saves the finished product in many common formats for further processing in other applications. If you only need to make the occasional diagram, Dia offers a real alternative to heavyweights like Gimp or LibreOffice Draw. ■■■

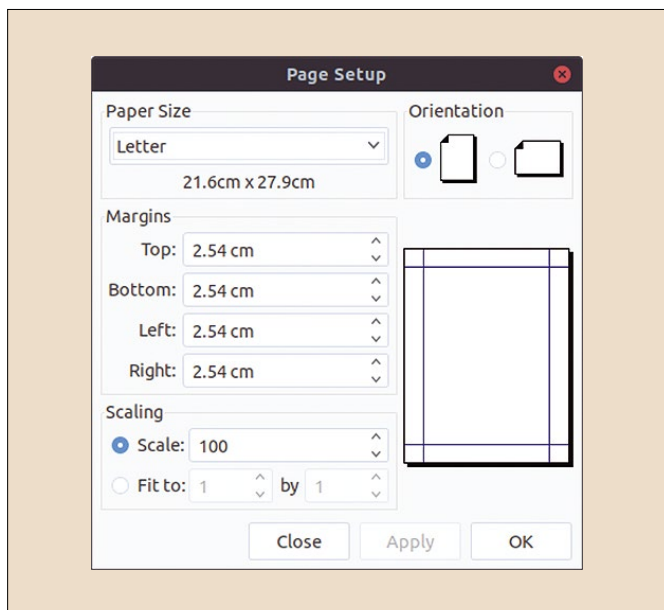



Figure 5: The Page Setup dialog for printing is intuitive.

Info

- [1] Dia: <https://sourceforge.net/projects/dia-installer/>
- [2] Object libraries: <http://dia-installer.de/shapes/index.html>
- [3] Diashapes: <http://dia-installer.de/diashapes/index.html.en>

REAL SOLUTIONS *for* REAL NETWORKS



ADMIN is your source for technical solutions to real-world problems.




Improve your admin skills with practical articles on:

- Security
- Cloud computing
- DevOps
- HPC
- Storage and more!



SUBSCRIBE NOW!
shop.linuxnewmedia.com

Check out our full catalog:
shop.linuxnewmedia.com



A brain teaser game

Brain Trainer

Gbrainy can help you develop your logical thinking skills with a collection of brain teasers.

By Erik Bärwaldt

Not only are brain teasers fun, they also promote logical thinking skills. The PC has long since replaced the good old puzzle book, making memory training far more efficient with more flexible tasks. Gbrainy, a puzzle game originally developed for the Gnome desktop, is a fun way to train your memory.

Gbrainy includes numerous logic games from different areas. The application, which can be used across all platforms, trains both your memory and your mental arithmetic skills. The games are divided into categories (Figure 1), with three difficulty levels for each category, which lets users advance to a higher difficulty level after practice. The program also keeps a success log for each user, and if so desired, a help function is available for solving tasks.

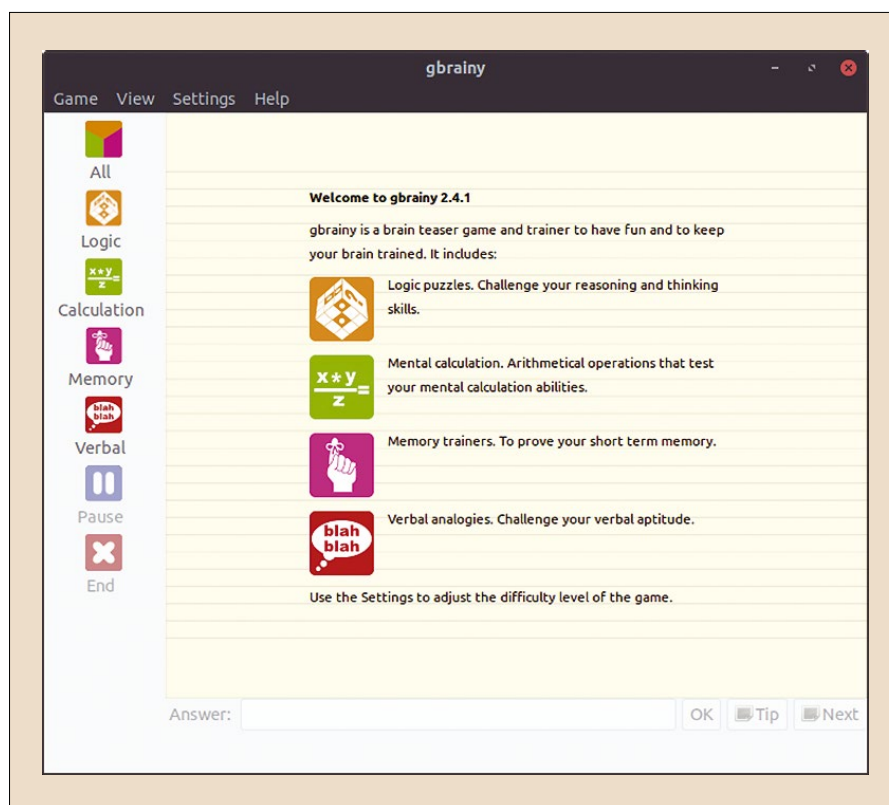


Figure 1: Gbrainy's program window gives you all the information you need to get started.



On the Record

Gbrainy can be found in the software repositories of virtually all major distributions. Some BSD derivatives already have gbrainy in their software archives, but beware of older software versions.

An overview of the existing precompiled packages is maintained by the developers on the project's website [1]. If you want to build the program yourself, the source code is also available for download as a tarball from the website.

Setup

When you install gbrainy with a precompiled package, a launcher will appear in your desktop environment's application menu. Clicking on the launcher takes you to an intuitive program window (refer to Figure 1) with basic information about the game and the different options. On the left, a vertical toolbar lists the individual game categories. At the bottom of the window, you'll find an input field.

To configure basic options, go to *Settings | Preferences*. Under *General Settings*, you can choose to skip games that use colors if you have color-blind players or automatically search for extensions at program startup.

The Preferences dialog is also where you select the difficulty level: *Easy*, *Medium*, or *Master*. By default, gbrainy is set to *Medium*. For memory games, you can also set the amount of time in seconds for users to memorize the task.

Under *Player's Game Session History*, you can manage when a game session is saved and the maximum number of records stored in the game history. By default, gbrainy only saves a game session after five games have been solved, with a maximum of 20 records stored in the game history. Both values can be adjusted.

Getting Started

The gbrainy games are groups of quick little puzzles consisting of a single question (Figure 2). Depending on the category, you might be asked to pick a number that matches a pattern, solve an analogy (ice is to slippery as glue is to _____?) or memorize an image then answer a question about it. The overall feeling is that you are taking an IQ test or college entrance exam (Figure 3).

When you are ready to play, click on the game category button in the toolbar on the left, and the first exercise will appear in the main window on the right. Once an exercise is active,

gbrainy grays out the game category buttons and activates the *Pause* and *End* buttons. Below the main window, you will find the Answer field where you enter your solution to the exercise, along with three buttons: *OK*, *Tip* (for additional help when available), and *Next*.

During game play, a status bar at the bottom shows the number of games played, your score, the elapsed time, and the current game title. After entering your answer in the Answer field, click *OK*, and gbrainy displays the correct solution above the input line. A green check mark and a bell indicate that your answer was correct. To move on to the next exercise, click *Next*.

Click *Pause* in the toolbar on the left if you need to take a break in the current game. While paused, the display area remains empty, so you cannot continue working on the solution. Clicking on *Resume* takes you back to the current exercise. When you finish the current game, gbrainy presents a graph with statistics about the game as well as hints on how to improve your performance (Figure 4).

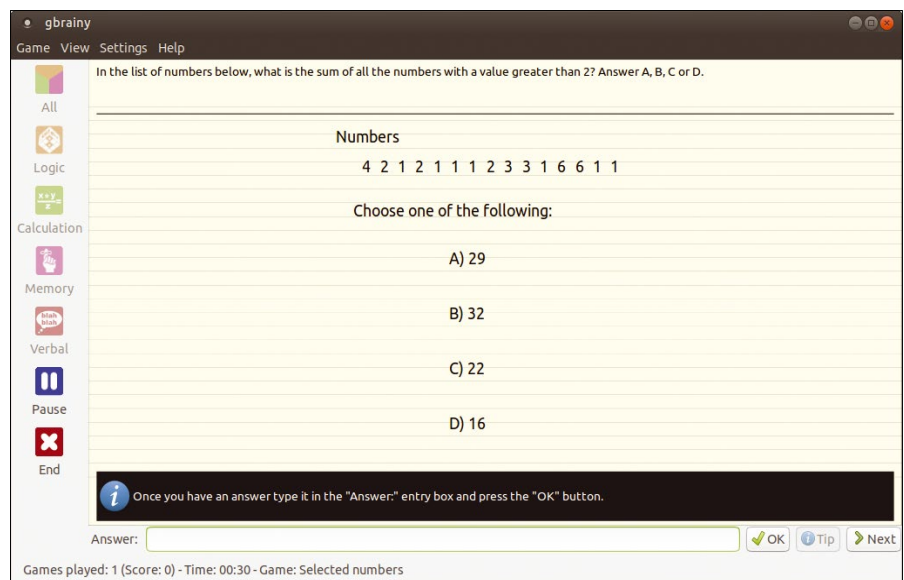


Figure 2: Gbrainy questions are puzzles designed to be solved in your head.

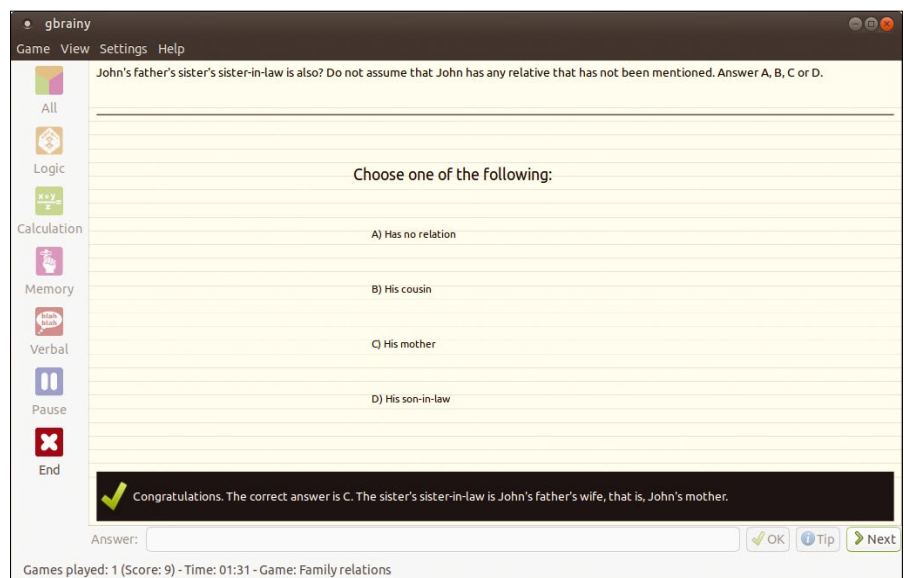


Figure 3: Gbrainy feels like an IQ test or college entrance exam.

On Paper

For users who want to solve tasks on paper, gbrainy offers an export function. Navigating to *Game | Export Games to PDF for Off-line Playing...* opens a window where you can select which game types, as well as the number of games, to be exported (Figure 5). If you only have a monochrome printer, you can choose to skip games that use color. You also can specify the desired difficulty level of the tasks to be exported.

After selecting your options, specify the path and file name for the document to be exported in the *Output File* field and press *Save*. Your saved document will also contain the answers for the selected exercises on a separate sheet, making export useful for school and training scenarios.

Special Features

If you want to customize the game selection, go to *Game | Custom Game*. By default, all games are selected here.

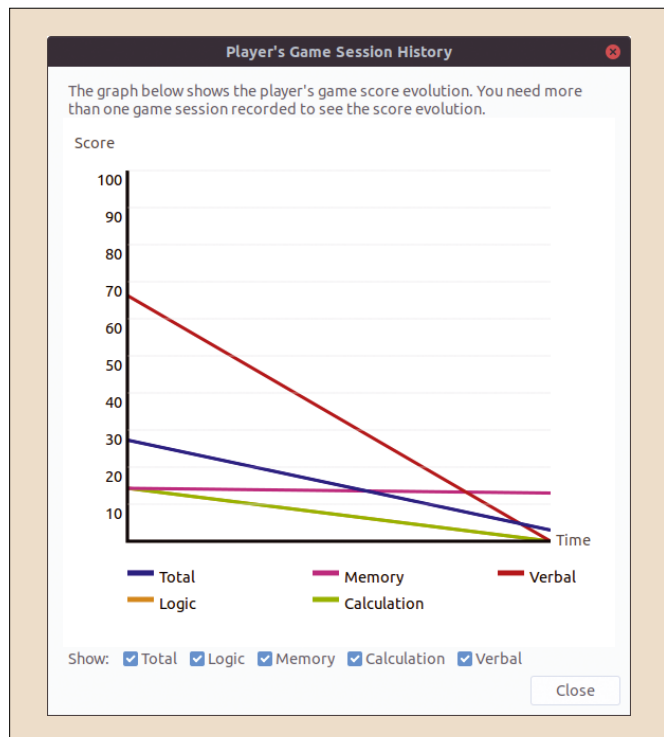


Figure 4: The statistics provide immediate information about your most recent game session.

Checking the box to the right of the game category removes the task from the collection. In the preview to the right of the list, you can see the currently selected task in the main window, which makes the selection easier.

Once your selection is made, click the *Start* button to immediately start the defined game collection. Using the *Game | New Game* menu, you can also enable each of the four game categories individually, so that only questions from an individual category are used in a given session.

Conclusions

Gbrainy lets you practice logical thinking skills. It is equally suitable for students and adults. With the ability to adjust learning levels and select different game categories, you can easily adapt the brain teasers to meet your needs. With the PDF export function, you can print out questions, along with answers, for offline use, making gbrainy suitable for home schooling. ■■■

Info

[1] gbrainy: <https://wiki.gnome.org/Apps/gbrainy>

Author

Erik Bärwaldt is a self-employed IT-admin and technical author living in Scarborough (United Kingdom). He writes for several IT-magazines.

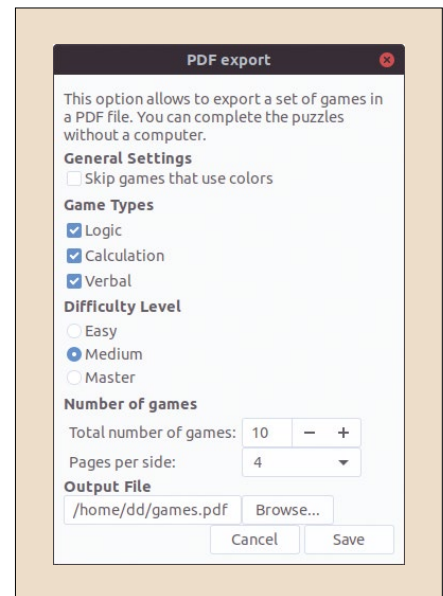



Figure 5: Using the PDF export option, you can create documents for printing.

SHOP THE SHOP
shop.linuxnewmedia.com

GET PRODUCTIVE WITH 101 LINUX HACKS



Improve your Linux skills with this cool collection of inspirational tricks and shortcuts for Linux geeks.

- Undelete lost files
- Cure the caps lock disease
- Run C one-liners in the shell
- Disable your webcam and mic
- And more!

ORDER ONLINE:
shop.linuxnewmedia.com/specials

Illustrating mathematical concepts in video with Manim

Explainer

Manim lets you program video sequences with a few lines of Python code to present mathematical problems in an engaging and scientifically accurate way. *By Frank Hofmann*

A precise description of a complex problem is sometimes difficult to achieve. Plain text will work eventually, but it often requires a great deal of imagination on the part of the reader. In a conventional classroom setting, an instructor might attempt to illustrate a topic by drawing on a chalkboard or whiteboard during a lecture, but it is fair to ask whether the power of the computer might offer a better option than a professor talking and writing at the same time. Presentation tools such as Impress or PowerPoint evolved as a way to illustrate key concepts in advance, but these tools typically have limited support for mathematical expressions, and they tend to present information in a static, slide show format. Movies and video sequences, on the other hand, can follow a spoken presentation organically in a way that often enhances learning and builds understanding.

However, shooting, trimming, and post-editing video footage can take time, even if you are experienced with video production. Manim [1] is a computer graphics tool that speeds up the task of building mathematical relationships into videos (Figure 1), letting you create a graphic video image with a simple Python script.

Manim, which is based on Python, combines the Cairo graphics library, the FFmpeg recording program, the SoX sound tool, and the LaTeX word processing system. SoX and LaTeX let you create optional acoustic effects and output mathematical formulas. The finished video all lands in an MP4 file.

The name Manim is an abbreviation of the words “mathematical animations,” which describes the basic idea behind the project. The software was created by US mathematician and developer Grant Sanderson, who uses it for his mathematical

Author

Frank Hofmann works on the road – preferably from Berlin, Geneva, and Cape Town – as a developer, trainer, and author. He is also the coauthor of the Debian package management book (<http://www.dpmb.org>).

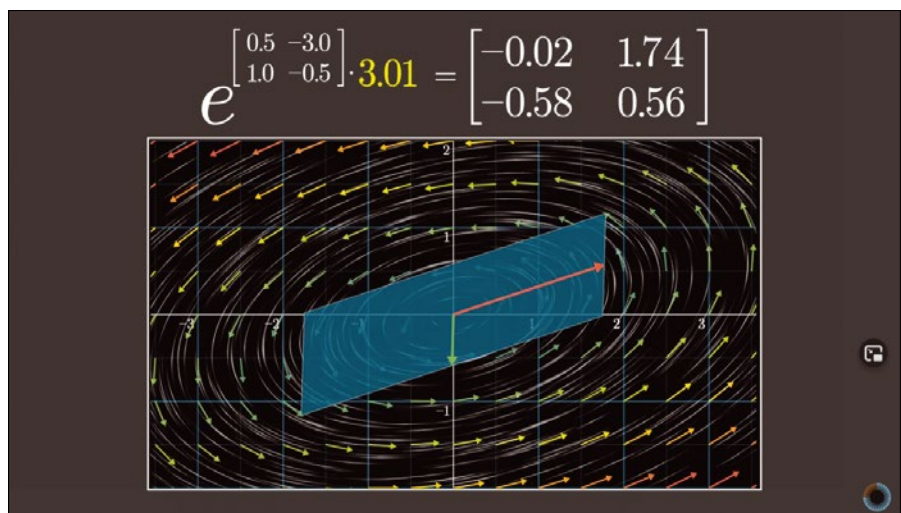


Figure 1: Graphic images can help the reader makes sense of mathematical expressions. © Grant Sanderson



Installation

Currently you won't find Manim packages for the popular Linux distributions, so you will need to go to a bit of manual work to compile and install it. As for dependencies, Manim requires Python version 3.7 or newer and the Python package manager pip with some development libraries, FFmpeg, SoX, Git, and LaTeX. LaTeX (which alone requires around 2.5GB of disk space) is optional.

If you need to install any of these dependencies, retrieve them with your favorite package manager (Listing 1, line 1).

For LaTeX, you need either the *texlive-full* package or the three other packages, *texlive-latex-base*, *texlive-latex-extra*, and *texlive-fonts-extra*. The full installation is handled by the call from line 2 of Listing 1; line 3 installs the three packages.

Line 4 gets the Manim library from the GitHub repository. The command creates the `manim/` folder in the current directory, to which you then change (line 5). The command in line 6 adds other libraries that might still be missing. Depending on your hardware, this step could take some time, as it involves picking up the current source code for the libraries and compiling it for your system.

Use the command from line 7 in Listing 1 to create your first video. The code refers to the sample script, `example_scenes.py`. The first parameter specified in the call designates the Manim library as a local Python file. This is followed by the sample script

explainer videos on his science channel 3Blue1Brown [2]. The videos at the 3Blue-

1Brown site offer insights into how to use Manim for creating engaging and informative presentations on technical topics. Each video is a moving tapestry, with equations, graphics, geometric shapes, and occasional cartoon characters emerging and disappearing to illustrate concepts described in a spoken narrative (Figure 2).

A community variant of Manim is also available. The community version is based on Sanderson's code but maintained separately by the community [3]. Both variants of Manim are on GitHub [4].

Manim consists of two parts: the program code for creating the visualization effects for the videos on 3Blue1Brown and the general code. The code for the 3Blue1Brown videos has a proprietary license (rights are held by 3Blue1Brown). The general code is under an MIT free software license. I used the original variant on a Debian 10.6 system when writing this article.

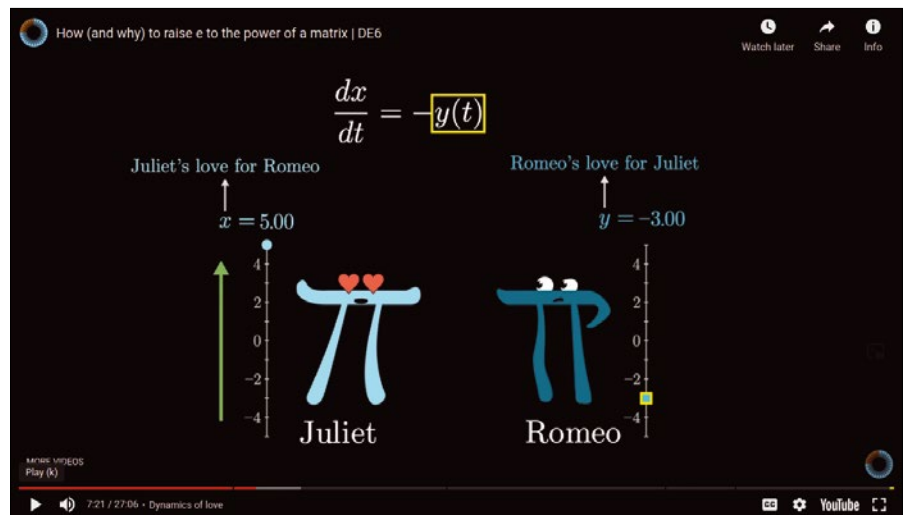


Figure 2: Manim is capable of illustrating some extremely complicated concepts. © Grant Sanderson

Listing 1: Installation and Initial Startup

```
01 # apt-get install python3.7 python3-pip ffmpeg sox libcairo2-dev libjpeg-dev libgif-dev git
02 # apt-get install texlive-full
03 # apt-get install texlive-latex-base texlive-latex-extra texlive-fonts-extra
04 $ git clone https://github.com/3b1b/manim.git
05 $ cd manim/
06 $ python3 -m pip install -r requirements.txt
07 $ python3 manim.py example_scenes.py SquareToCircle -lp
```

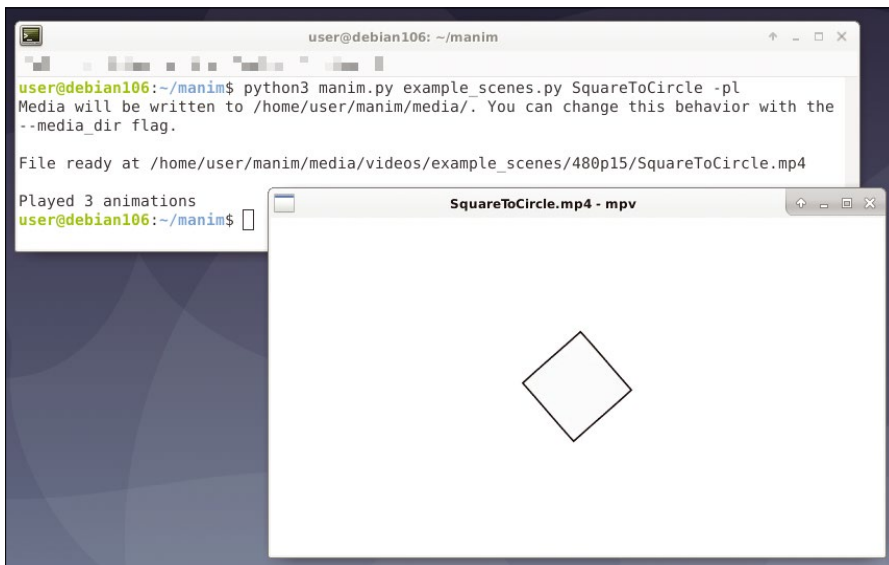


Figure 3: Rendering and playback of the sample scenes.

with the animations and, as the third parameter, the name of the class in the sample script with the animations you want Manim to execute. At the same time, it serves as the name of the output file, which is automatically given an extension of `.mp4`.

Listing 2: Sample Script

```
01 class SquareToCircle(Scene):
02     def construct(self):
03         circle = Circle()
04         square = Square()
05         square.flip(RIGHT)
06         square.rotate(-3 * TAU / 8)
07         circle.set_fill(PINK, opacity=0.5)
08
09         self.play>ShowCreation(square)
10         self.play(Transform(square, circle))
11         self.play(FadeOut(square))
```

Table 1: Geometric Objects

Object	Class
Arc	Arc(), ArcBetweenPoints()
Arrow	Arrow(), CurvedArrow(), CurvedDoubleArrow(), DoubleArrow()
Circle	Circle()
Dot	Dot(), SmallDot()
Ellipse	Ellipse()
Line	Line(), DashedLine(), TangentLine()
Polygon	Polygon(), RegularPolygon()
Triangle	Triangle()
Rectangle	Rectangle(), Square(), RoundedRectangle()

Table 2: Supported Animations

Animation	Method
Fade in and out	FadeIn(), FadeOut(), FadeInFrom(), FadeOutAndShift()
Grow	GrowFromPoint(), GrowFromCenter(), GrowFromEdge()
Transform and rotate	ClockwiseTransform(), CounterclockwiseTransform(), FadeToColor()
Rotate	Rotate()

Two switches follow at the end: `-l` (“lower”) causes the video to be rendered in the slightly lower quality of 480p15 instead of 1440p60; `-p` (“pre-view”) gives you immediate playback of the generated video file (Figure 3).

By default, Manim saves the generated files in `./media/videos/`. You can adapt this path to suit your needs via the `MEDIA_DIR` environment variable.

Understanding Animations

The source code for the video from Figure 3 contains just a few lines. The code consists of the Python class `SquareToCircle()`, which derives from the internal Python class `Scene()` (Listing 2, line 1). `SquareToCircle()` contains only one method named `construct()` (line 2).

Lines 3 and 4 define a Manim object named `circle` of the type `Circle()` and a square of type `Square()`. Lines 5 and 6 first initialize a clockwise rotation of the square via the `rotate()` method and then a rotation by the specified angle. The code from line 7 colors the circle pink, with an opacity of 50 percent.

Lines 9 to 11 use the `play()` method for three animations. Calling `ShowCreation(square)` creates the square, `Transform(square, circle)` renders the gradual transformation of the square into a pink circle, and `FadeOut(square)` slowly fades out the square.

Table 1 lists some of the geometric objects that Manim currently supports. This list is based on the analysis of the source code, since the documentation does not currently provide this information. Table 2 summarizes some of the animations that Manim supports. The existing documentation provides a reference.

Manim also offers the ability to change the camera position. You need this feature if you want to highlight parts of the overall image by zooming in and out, or visualize movement along an axis.

Manim in Practice

The gaps in the documentation make it a little difficult to get started with Manim. Blog posts [5] that explain the procedure for specific use cases [6] are helpful.

In Listing 3, lines 3 to 6 define a mathematical formula in LaTeX notation. Lines 7 and 8 define two rectangles and assign them to the second and fourth components of the formula. First, line 9 outputs the formula in full. Line 10 adds the first rectangle to it, after which the `wait()` method adds



a pause. Line 12 replaces rectangle 1 with rectangle 2. Manim adds motion to it: The rectangle slides to the right to the end of the formula. Line 13 specifies a short pause before the video ends.

To see the effect, save the class in your Manim directory as a Python script `box.py`. To render the video (Figure 4), run the command from the first line of Listing 4.

The last example generates a graph of the sine function and plots it in a Cartesian coordinate system (Listing 5). A data structure named `CONFIG` specifies the appearance of the coordinate system. The *Y* axis runs from 0 to 100 (lines 3 and 4). The *X* axis has a vertical line, and the *Y* axis has a horizontal line every 10 units for ease of reading (lines 5 and 6). The *Y* axis shows the corresponding values.

Listing 3: `box.py`

```
01 class MoveFrameBox(Scene):
02     def construct(self):
03         text=MathTex(
04             "\\frac{d}{dx}f(x)g(x)=", "f(x)\\frac{d}{dx}g(x)", "+",
05             "g(x)\\frac{d}{dx}f(x)"
06         )
07         framebox1 = SurroundingRectangle(text[1], buff = .1)
08         framebox2 = SurroundingRectangle(text[3], buff = .1)
09         self.play(Write(text))
10         self.play(ShowCreation(framebox1),)
11         self.wait()
12         self.play(ReplacementTransform(framebox1, framebox2),)
13         self.wait()
```

Listing 4: Rendering the Examples

```
$ python3 manim.py box.py MoveFrameBox -pl
$ python3 manim.py function.py SinusPlot -pl
```

Listing 5: `function.py`

```
01 class SinusPlot(GraphScene):
02     CONFIG = {
03         "y_min": 0,
04         "y_max": 100,
05         "y_axis_config": {"tick_frequency": 10},
06         "y_labeled_nums": np.arange(0, 100, 10)
07     }
08
09     def construct(self):
10         self.setup_axes()
11         dot = Dot().move_to(self.coords_to_point(PI / 2, 20))
12         func_graph = self.get_graph(lambda x: 20 * np.sin(x))
13         self.add(dot, func_graph)
```

The code starting in line 9 first outputs the two axes using the `setup_axes()` method (line 10). Line 11 creates a dot, which it places at position $(\pi/2, 20)$ using the `move_to()` method. Line 12 uses `get_graph()` to define the instruction for the sine function; line 13 outputs the dot and the sine function.

Add this class to your Manim directory as a Python script, named `function.py` this time. Trigger rendering of the video with the command from the last line of Listing 4 (Figure 5).

Conclusions

Manim proves to be a perfect tool for presenting mathematical relationships as videos. The tool is comparatively easy to use and delivers results that would require considerably more effort using conventional methods. ■■■

$$\frac{d}{dx}f(x)g(x) = \boxed{f(x)\frac{d}{dx}g(x)} + g(x)\frac{d}{dx}f(x)$$

Figure 4: Formula highlighting a component.




Figure 5: Function plot with the sine function.

Info

- [1] Manim: <https://github.com/3b1b/manim>
- [2] 3Blue1Brown: <https://www.3blue1brown.com/>
- [3] Manim community edition: <https://docs.manim.community/en/v0.1.0/index.html>
- [4] Manim on GitHub: <https://github.com/ManimCommunity/manim/>
- [5] "Creating Math Animations in Python with Manim": <https://gilberttanner.com/blog/creating-math-animations-in-python-with-manim>
- [6] "Getting Started Animating with Manim and Python 3.7": <https://talkingphysics.wordpress.com/2019/01/08/getting-started-animating-with-manim-and-python-3-7/>



Designing for newcomers

Zorin OS

Artyom and Kyrill Zorin set out to design a user-friendly operating system, resulting in Zorin OS, a favorite among new users. Bruce talks to Artyom about Zorin OS and where the project is headed.

By Bruce Byfield

First released in 2009 by brothers Artyom and Kyrill Zorin, Zorin OS [1] quickly became known as a Debian-derivative with a minimalist desktop interface based on Gnome. Today, Zorin continues to tweak its desktop, but has also become known for its appeal to new users – particularly for its Zorin Appearance app, which in some editions emulates the design of other operating systems, including Windows 10 in Zorin 16 [2]. In addition, Zorin 16 also supports a wide variety of installation methods, ranging from apt and Wine to Snap and Flatpak.

In a break from this column’s normal format, I interviewed Artyom to find out more about Zorin OS, including its past, present, and future.

Linux Magazine (LM): How did you and Kyrill get involved with free software?

Artyom Zorin (AZ): I first came across Linux and free and open source software back in the summer of 2008, when my brother Kyrill and I stumbled upon videos of Ubuntu. We found the concept of an operating system other than Windows and macOS intriguing, so we decided to give it a try. Immediately, we were amazed at how much faster and more customizable it was, not to mention the security and privacy advantages. We knew that Linux was much better than Windows and macOS in a lot of ways. But after showing it to our father, it became clear as to why it didn’t have the level of popularity that the incumbents had: It just wasn’t user-friendly enough.

We noticed that there wasn’t a strong enough focus on design simplicity in the Linux world. The unfamiliar user interfaces that most distros came with presented roadblocks to regular computer

users like my father when trying to make the switch away from Windows. We saw this as the biggest challenge preventing Linux from gaining a wider user base.

Back then, I was 12 years old and my brother was 14, so we had no experience and education in computer science or programming. However, we knew about a national school science fair here in Ireland and thought it would be interesting to enter with a project to try and solve this problem with Linux and learn how to do it along the way. That project was Zorin OS.

LM: How is the Zorin project organized? How do the company and community interact?

AZ: We’re always rethinking how we organize ourselves as a company and as an open source project so we can deliver a greater product and better execute our mission.

We see the current incarnation of “Zorin, the organization” as a mix between a traditional community-focused open source project and a business. Over the years, we’ve learned that there are pros and cons to both approaches, and, as a project that’s trying to do something different in the Linux world, we need to innovate how we run ourselves.

In practical terms, this means that we pay very close attention to our users’ feedback. We’ve made it a policy for everyone on our development team to help users and Zorin OS Ultimate customers with their technical support queries. That helps us to find what new features users are looking for, what’s missing in Zorin OS, and how we can improve the system further. However, we accept the fact that we can’t implement every idea and suggestion from every user and can’t cater to every kind of user out there. Every decision we make is first and foremost guided by our mission: to bring the power of Linux to newcomers by making the entire experience easier and more accessible. If a new feature suggestion or product change doesn’t serve this mission, it generally doesn’t make it past the drawing board.

Like other open source projects, we receive contributions from other members of the community, both indirectly and as direct code commits to our sub-projects.

The inverse is also true, in that we make patches to existing third-party software in order to improve Zorin OS’s user experience and submit many of them upstream so the entire Linux ecosystem benefits.

LM: What is Zorin’s design philosophy, and which features reflect it?

AZ: From the very beginning, our main goal was to make Zorin OS as easy and frictionless as possible. If you normally use Windows, you should be able to move to Zorin OS without having to learn anything new. The start menu and taskbar are right where you remember them. We’ve also made it possible to change the desktop layout with a click using Zorin Appearance, so if you’re used to macOS or another operating system you can feel right at home. Of course, you still have all the power of an Ubuntu-based Linux distribution under the hood.

We also pay close attention to crafting a beautiful visual design language. That way, new users aren’t only attracted to Zorin OS and Linux because of the technical and practical advantages, but also because they want to use a system that’s visually pleasing and inviting.

Rather than trying to load as many features and tweaks as we can, we think carefully about whether a new feature is intuitive and fits into our larger cohesive

vision of the desktop. Our thought process always starts with defining the user experience we want to provide and then we work our way back to the technology.

LM: Tell us about the Zorin Appearance option (Figure 1) for emulating other OSs.

AZ: The configurable desktop layout in Zorin Appearance is one of the main features that helps make Zorin OS so friendly to new users. It lets you select one of the preset layouts, and the desktop rearranges itself to resemble that interface with one click.

From the early days of Zorin OS, we’ve heard from countless users that this was what has helped them to get comfortable with using Linux. The ability to simply choose the interface you’re most familiar with helps you keep your muscle memory and makes you feel more at home in the new Linux environment. This goes beyond being just another customization option, it helps the operating system to adapt to you and not the other way around.

By default, Zorin OS’s desktop layout is similar to that of Windows – which most computer users out there are familiar with. You can use Zorin Appearance’s desktop layout feature to make it look and feel like macOS, Windows XP, Ubuntu, GNOME Shell, or a touch layout that works great on computers with touchscreens.

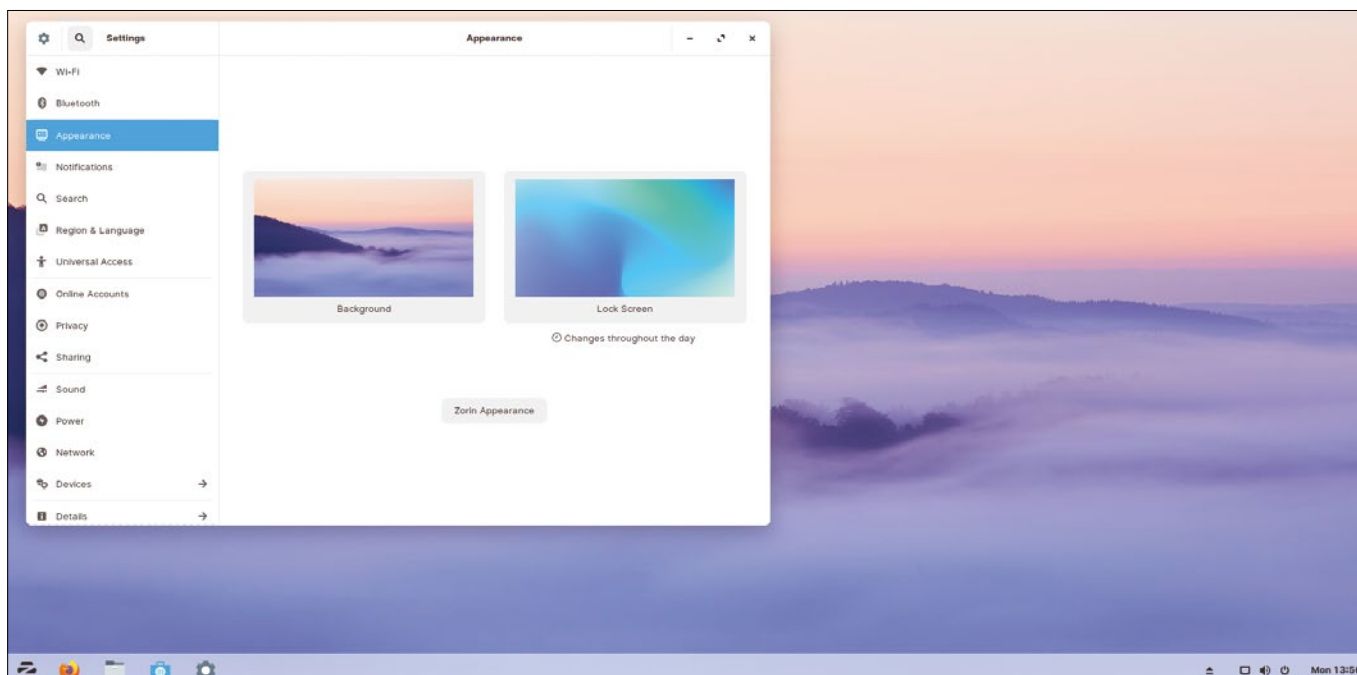


Figure 1: With Zorin Appearance, you can configure your desktop layout to emulate the operating system with which you are most familiar. Commercial versions have more emulations.

One of the new improvements in Zorin OS 16 is the introduction of a Windows 10X-like desktop layout option. It has a modern and streamlined UI that adapts well to computers with touchpads, mice, or touchscreens. I've personally been using this layout as my daily desktop for the past few months, and I love how clean and elegant it looks. Its simplicity makes me feel like the interface gets out of the way, so I can focus on what I'm working on better.

LM: Does Zorin OS offer any special support for WINE and PlayOnLinux?

AZ: We've made it super simple to get Windows apps running in Zorin OS compared to most other Linux distros. Simply double-click on an .exe installer, and you'll get guided through installing Wine and PlayOnLinux in a couple of clicks. We think it's the best way to let anyone easily run their Windows apps while also keeping the operating system safe and secure out of the box.

Zorin OS also has a built-in database of Windows installer files for apps that are already natively available on Linux. If the system detects that the .exe you double-clicked on is for a known app in this database, it will automatically suggest you to install the recommended native package in the Software store.

Some Windows apps and games have even been packaged up in containerized Snap and Flatpak packages, preconfigured with Wine as a runtime. Because Zorin OS 16 now comes with both the Flathub and Snap stores preloaded, all you need to do is search for the app – like Notepad++ or TrackMania, for example – in the Zorin OS Software store to install it. You don't even need to know anything about Wine or PlayOnLinux – it just works like a native app.

LM: How many users does Zorin OS have?

AZ: Since the release of Zorin OS 15 – our latest major version – in June 2019, it's been downloaded over 2.5 million times around the world. Over 67 percent of these downloads came from Windows and macOS, reflecting [the success of] our mission to bring the power of Linux to people who've never had access to it before.

LM: What are your future plans for Zorin OS?

AZ: We've heard many success stories from businesses, schools, governments, and other organizations that made the switch from Windows to Zorin OS. However, the one downside that kept popping up in these cases was that deploying the operating system and managing computers was initially challenging and time-consuming. As we looked around to see if there were any tools to help these users, we found that there weren't any that were sufficiently simple and suited for Linux desktop workstations. It became clear to us that the difficulty in managing and maintaining a Windows-to-Linux desktop migration was one of the strongest reasons why more organizations weren't moving to Linux.

To remedy this, we are planning to launch a new remote management tool later this year called Zorin Grid. Its goal is to make it possible to manage a fleet of tens, hundreds, or even thousands of workstations as easily as one, with little to no training. We've taken the lessons we've learned from the businesses and organizations we've heard from and are designing Zorin Grid to address the roadblocks they experienced in order to make the transi-

tion as seamless and effortless as possible. Over the long term, we're planning to make Zorin Grid work with many other Linux distributions – not only Zorin OS – to make the entire Linux ecosystem more viable for use in the workplace.

We believe that this could be the missing piece in the puzzle that could allow lots of businesses, schools, and organizations to make the leap to Linux, and avail of all the practical, security, and cost-related advantages it has.

Another exciting thing on our roadmap is our plan to port Zorin OS 16 to ARM-based devices like the Raspberry Pi. One major goal that's motivating us to focus on ARM devices – and especially ultra-low-cost ones like the RPi4 – is to make computers accessible and useful in the developing world. So far, the Raspberry Pi (and the Linux distributions made for it) have generally been aimed at computer tinkerers and beginner programmers, and it has become a tremendous platform for these use cases. By porting Zorin OS Lite and Education Lite to the Raspberry Pi 4 and 400, we're aiming to make it into a great general-purpose computing platform and educational resource for people who've never had access to a computer before.

Zorin Core is available as a free download. Three commercial versions are also available: Zorin Ultimate, Lite, and Education. ■■■

Info

[1] Zorin OS: <https://zorinos.com/>


[2] Zorin 16 (latest features):

<https://blog.zorin.com/2021/04/15/introducing-zorin-os-16-test-the-beta-today/>



What?!

I can get my
issues
SOONER?



Available anywhere, anytime!

Sign up for a digital subscription and enjoy the latest articles on trending topics, reviews, cool projects and more...

Subscribe to the PDF edition: shop.linuxnewmedia.com/digisub

Now available on ZINIO: bit.ly/Linux-Pro-ZINIO



Editing statistical data with gawk

Stats Don't Lie

With very little overhead, you can access statistics on the spread of COVID-19 using gawk scripts and simple shell commands. *By Goran Mladenovic*

Open a paper, turn on the radio or TV, and you are confronted with statistics concerning the spread of COVID-19. Sometimes, it's hard to wrap your mind around all the data, and even then, you are never sure if the media is taking the data out of context. For instance, a report of 30,000 new infections in a country doesn't provide any information about the number of tests or the positivity rate.

If you want to drill down into the data concerning the spread of COVID-19, you can use gawk scripts and simple shell commands to process, evaluate, and present these statistics in a more meaningful way.

A Few One-Liners

The statistical data on the spread of COVID-19, usually collected daily and organized in timelines, is freely available

on the Internet. Our World in Data (OWID) [1] has been recording the data since the beginning of the crisis. On its website, OWID makes the data available for download in various formats, including CSV format, which is well suited for further processing.

To get started, you should download the CSV file from OWID to see what it contains. Looking at the extensive `owid-covid-data.csv` file with a simple `less` or `cat` command may result in eye strain when trying to read the more detailed data for individual countries. Even if you open the file with an editor or OpenOffice Writer, it is still quite difficult to get a useful overview. A better option is to use shell one-liners to extract the information you need.

To find out the column titles in the CSV file, you can use the one-liner shown in line 1 of Listing 1. The `head` command at the beginning of the pipe

outputs the first line of the file. Then, `sed` converts all commas to line feeds. Next, `cat -n` numbers the output, and `pr` distributes the whole thing in columns across the screen, resulting in a useful overview (Figure 1).

Looking at Figure 1, you can assume that the first four columns (`iso_code`, `continent`, `location`, `date`) will not change over time, but you can't be so confident about the order of the other column titles. New key data is added all the time, as knowledge about the virus evolves, treatments are discovered, and vaccinations are introduced. For more information on column title descriptions (as well as information on the data and data sources), visit OWID's GitHub page [2].

You may also want to know which countries provide data. The one-liner from line 2 of Listing 1 reveals that OWID has data for over 200 countries

Listing 1: Useful One-Liners

```
01 $ head -n 1 owid-covid-data.csv | sed 's/,/\n/g' | cat -n | pr -T --columns=2 --width=90
02 $ cut -f 3 -d "," owid-covid-data.csv | sort | uniq | sed -r '/World|International/d' | cat -n | pr -T
    --columns=2 --width=120
03 $ grep -i germany owid-covid-data.csv | cut -f 4-8 -d "," | column -t -s ","
```

```

dd@ubuntu:~$ head -n 1 owid-covid-data.csv | sed 's/,/\n/g' | cat -n | pr -T --columns=2 --width=90
1 iso_code 31 new_tests_smoothed_per_thousand
2 continent 32 positive_rate
3 location 33 tests_per_case
4 date 34 tests_units
5 total_cases 35 total_vaccinations
6 new_cases 36 people_vaccinated
7 new_cases_smoothed 37 people_fully_vaccinated
8 total_deaths 38 new_vaccinations
9 new_deaths 39 new_vaccinations_smoothed
10 new_deaths_smoothed 40 total_vaccinations_per_hundred
11 total_cases_per_million 41 people_vaccinated_per_hundred
12 new_cases_per_million 42 people_fully_vaccinated_per_hundred
13 new_cases_smoothed_per_million 43 new_vaccinations_smoothed_per_millio
14 total_deaths_per_million 44 stringency_index
15 new_deaths_per_million 45 population
16 new_deaths_smoothed_per_million 46 population_density
17 reproduction_rate 47 median_age
18 icu_patients 48 aged_65_older
19 icu_patients_per_million 49 aged_70_older
20 hosp_patients 50 gdp_per_capita
21 hosp_patients_per_million 51 extreme_poverty
22 weekly_icu_admissions 52 cardiovasc_death_rate
23 weekly_icu_admissions_per_million 53 diabetes_prevalence
24 weekly_hosp_admissions 54 female_smokers
25 weekly_hosp_admissions_per_million 55 male_smokers
26 new_tests 56 handwashing_facilities
27 total_tests 57 hospital_beds_per_thousand
28 total_tests_per_thousand 58 life_expectancy
29 new_tests_per_thousand 59 human_development_index
30 new_tests_smoothed
dd@ubuntu:~$

```

Figure 1: A lean one-liner delivers a quick overview of the column titles in the CSV file.

```

dd@ubuntu:~$ cut -f 3 -d "," owid-covid-data.csv | sort | uniq | sed -r 'r /world/international/d' | cat -n | pr -T --columns=2 --width=120
142 Niger 183 South Sudan
143 Nigeria 184 Spain
144 North America 185 Sri Lanka
145 Northern Cyprus 186 Sudan
146 North Macedonia 187 Suriname
147 Norway 188 Sweden
148 Oceania 189 Switzerland
149 Oman 190 Syria
150 Pakistan 191 Taiwan
151 Palestine 192 Tajikistan
152 Panama 193 Tanzania
153 Papua New Guinea 194 Thailand
154 Paraguay 195 Timor
155 Peru 196 Togo
156 Philippines 197 Trinidad and Tobago
157 Poland 198 Tunisia
158 Portugal 199 Turkey
159 Qatar 200 Turks and Caicos Islands
160 Romania 201 Uganda
161 Russia 202 Ukraine
162 Rwanda 203 United Arab Emirates
163 Saint Helena 204 United Kingdom
164 Saint Kitts and Nevis 205 United States
165 Saint Lucia 206 Uruguay
166 Saint Vincent and the Grenadines 207 Uzbekistan
167 Samoa 208 Vanuatu
168 San Marino 209 Vatican
169 Sao Tome and Principe 210 Venezuela
170 Saudi Arabia 211 Vietnam
171 Senegal 212 Yemen
172 Serbia 213 Zambia
173 Seychelles 214 Zimbabwe
dd@ubuntu:~$

```

Figure 2: To fit the country output on your screen, adjust the command's parameters.

(Figure 2). Depending on your terminal size, you can adjust the `--columns` or `--width` parameters.

Finally, you may want to filter out a specific country's data. You can do this with a combination of `grep` and `cut`, as shown in line 3 of Listing 1 (Germany in this example). However, this output is missing the column titles. A better way to extract this information is to write a script that takes into account the column sequence – the script needs to be able to

discover the number of columns and which columns contain the data.

Scripted

Listing 2 shows the script to filter out a specific country's data. Like any script, it starts by specifying the interpreter. In lines 2 to 6, you can see basic declarations that are stored in variables for easier access later. In principle, these are fairly simple things like the URL, the file name, two columns titles

needed later, and the first key data (the world population).

The `download()` function (lines 8-12) first checks to see whether the file has already been downloaded and whether it corresponds to the current date. If the answer is no to either of these criteria, the download takes place, overwriting a previously downloaded file.

Next, `country_selection_menu()` (lines 16-21) displays all countries for selection (Figure 3). You then input line by line the selected ISO codes, which are loaded into an array. Do not worry about uppercase/lowercase when entering the ISO codes. Press `Ctrl + D` to confirm your selections.

For column titles, `col_selection_menu()` (23-28) lets you do the same thing using the column title numbers for input. When inputting your column selections, you do not have to input the date column: It is always parsed by the script, because listing the other data without this important column does not make much sense.

The `output()` function (lines 30-60) takes care of the data output. Starting in line 39, the function contains a `gawk` script with Bash commands embedded in it. Use the following format:

```
'GAWK CODE'"BASH CODE"'GAWK CODE'
```

This function is passed an ISO code as a parameter, which is converted to uppercase letters in line 31 and ends up in a variable. With this help, `grep` then retrieves the right name from the list of countries and stores it in a variable (line 32). Line 33 stores the first line of the file (the column titles) and, using `grep` again, the country's data block in a variable so that Bash or `gawk` can work with it. Lines 34 and 35 filter out the current population size and density of the country from the output block.

Line 36 creates a format string for a later `printf` command for `gawk`. The code determines the longest entry in the output block for each column title you select, creating a number string in which each number has a percent sign (%) at the beginning and an `s` appended to the end (e.g., `%10s %20s ...`).

The actual output starts on line 37, which is handled by Bash. The `gawk` script takes over in line 39, and it expects as parameters the preallocated

Listing 2: Data Evaluation (query.sh)

```

01 #!/usr/bin/bash
02 URL=https://covid.ourworldindata.org/data/owid-covid-data.csv
03 FILE=owid-covid-data.csv
04 col_population=`sed lq $FILE | sed 's/,/\n/g' | cat -n | grep "population$" | cut -f 1`
05 col_populationdensity=`sed lq $FILE | sed 's/,/\n/g' | cat -n | grep "population_density$" | cut -f 1`
06 world_population=`grep -i world $FILE | tail -n 1 | cut -f $col_population -d ","`
07
08 download(){
09     test -e $FILE || wget -O $FILE $URL
10     test `ls --full-time $FILE | gawk '{print $6}'` = `date +%Y-%m-%d` ||
11     wget -O $FILE $URL
12 }
13
14 download
15
16 country_selection_menu(){
17     country=`cut -f 1,3 -d "," $FILE | sed 1d | sort | uniq | sed '/International/d;/World/d' | sed 's/^(/;/s,/)/`
18     pr --columns=3 -T <<<$country
19     echo "Enter ISO codes line-by-line and terminate input with Ctrl-D:"
20     readarray -t country_selection
21 }
22
23 col_selection_menu(){
24     sp=`sed lq $FILE | sed 's/,/\n/g' | cat -n`
25     pr -T --columns=2 <<<$sp
26     echo "Enter columns line-by-line and terminate input with Ctrl-D:"
27     readarray -t col_selection
28 }
29
30 output(){
31     searchstring=`tr [:lower:] [:upper:] <<<$1`
32     local country=`grep "$searchstring" <<<$country`
33     outputblock=`cat <(sed lq $FILE) <(grep -E "^"$searchstring $FILE)`
34     population_country=`tail -n 1 <<<$outputblock | cut -f $col_population -d ","`
35     populationdensity_country=`tail -n 1 <<<$outputblock | cut -f $col_populationdensity -d ","`
36     colformat=`for col in 4 ${col_selection[*]};do cut -f $sp -d "," <<<$outputblock | wc -L | sed 's/^\s/g;s/$/s/g' ;done`
37     echo -e "\n"
38     echo `tr [:lower:] [:upper:] <<<"$country`
39     gawk -F "," -v pop=$population_country -v world_pop=$world_population
40         -v density=$populationdensity_country <<<$outputblock '
41     BEGIN{
42         ub1="Population: "pop", percentage of world population: "pop*100/world_pop" %"
43         ub2="Inhabitants per km<+>2<+>: "density
44         stars = gensub(/./, "*", "g", ub1)
45         print stars
46         print ub1
47         print ub2
48         print stars
49     }
50     FNR == 1{
51         printf ""echo ${colformat[*]}`"\n", ""echo 4 ${col_selection[*]} | sed 's/^\$/;/s/ /,$/g'`""
52     }
53     FNR > 1{
54         printf ""echo ${colformat[*]}`"\n", ""echo 4 ${col_selection[*]} | sed 's/^\$/;/s/ /,$/g'`""
55         datasetcounter++
56     }
57     END{
58         print "\nDatasets in total: "datasetcounter

```

Listing 2: Data Evaluation (query.sh) (continued)

```

58     }
59     '
60 }
61
62 country_selection_menu
63 spalten_selection_menu
64
65 test -e Evaluations && : || mkdir Evaluations
66
67 for land in "${country_selection[@]}"
68 do
69     output "$land"
70 done | tee "Evaluations/"`echo ${country_selection[*]} | sed 's/ /_/_/g;s/$/_/'`"date +%a_%d_%b_%Y".txt" | less
71
72 read -p "More queries? (Y/N)" continue
73 test $continue = "y" || test $continue = "Y" && $0 || exit 1

```

variables and the output block that was read out.

The BEGIN block (lines 40-48) works with the variables. In ub1 and ub2, strings are assembled, or literally lumped together, by the code simply concatenating them. In line 43, a general substitution occurs; this simply converts each character in ub1 to a star (*) and stores the result in the stars variable. Lines 44 through 47 then output the variable contents.

Two more blocks take care of outputting the appropriate headers and data

columns (lines 49-55). After querying the special variable FNR, gawk knows which line is currently up and then executes the appropriate block. The variable datasetcounter (line 54) records how many records there are in total; the result is then output in the END block with the appropriate text. You will notice that not all columns for each country are completely filled: Due to the vast amount of data, data glitches are bound to happen, and some countries do not always collect the same metrics as others.

Last but not least, all previously defined functions must also be called. Lines 62 and 63 handle the query of the countries and columns to be displayed. Line 65 checks if the directory ./Evaluations/ already exists and creates it if necessary to be able to store data there later. The calls in lines 67 to 70 then work through all the selected countries and pass the results to the output function.

At the end of the for loop, another tee statement writes the data to the ./evaluations/ directory, with the appropriate

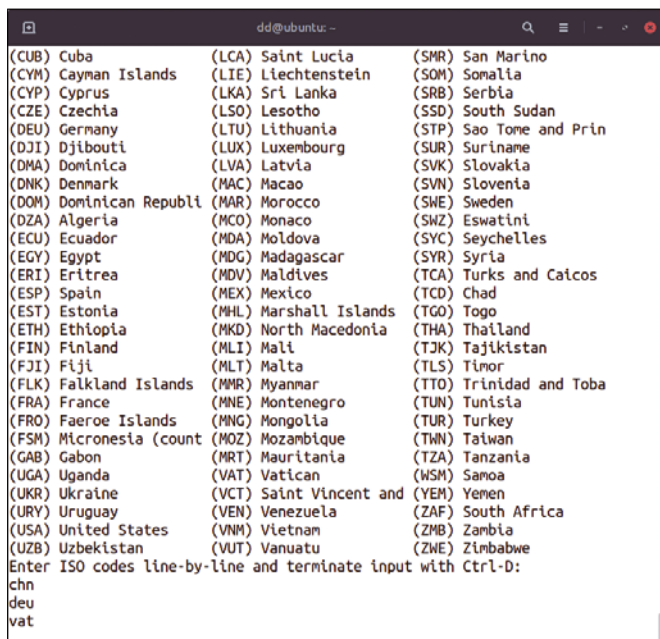


Figure 3: From the country selection menu, you can choose the ISO codes for the countries you want included in your data.

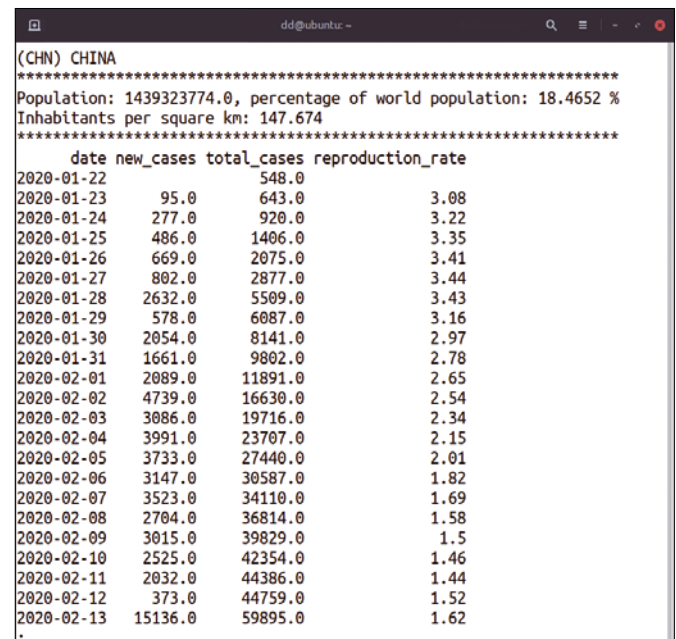


Figure 4: The script lets you track how the pandemic developed in China or any other country. You choose the data you want to see.

date and country specified, and then copies it to standard output, where `less` lets you view the data in the terminal (Figure 4). Finally, the routine in lines 72 and 73 prompts you to start more queries. If you want to, the script restarts; otherwise it says goodbye with an `exit` command.

The script lets you view any key data that you are interested in, giving you a more detailed view than what you might find in the media. In particular, the data for individual countries can be more accurately compared. Keep in mind, this script can not tell you how a reproduction (R) value, incidence, or other magic number was calculated in the original data from OWID.

Making Your Own Calculations

How much money does the pharmaceutical industry make on COVID-19 tests worldwide? A precise answer could require weeks of investigation, but you can come up with a quick estimate using `gawk` and `Bash`. In the `total_tests` column, most countries show the total number of tests. All you have to do is add up the current test numbers for all countries and multiply the result by an estimate of the average test price. You now have a rough idea about the worldwide pharmaceutical industry's gross revenue (or turnover) on tests.

Listing 3 provides a script to calculate the gross revenue for COVID-19

tests. In line 2, the script uses the `download` routine from Listing 2. For this reason, both scripts (Listing 2 and Listing 3) must reside in the directory. Line 3 stores the file name in a variable. However, this can also be left out, since the source statement from line 2 also takes this from the first script (Listing 2).

Line 4 determines the correct column for the test count, and line 6 bundles all the countries into an array. The `for` loop in lines 8 to 12 goes through all countries one by one and creates three columns separated by commas: the country name, the total number of tests, and a total number formatted to the thousandths decimal place. Then a `sort`

Listing 3: Test Revenues (turnover.sh)

```
01 #!/usr/bin/bash
02 source <(sed -n '2,/download$/p' query.sh)
03 FILE=owid-covid-data.csv
04 total_test_col=`sed 's/,/\n/g;lq' $FILE | grep -n total_tests$ | cut -f 1 -d ":"`
05
06 readarray -t all_country <<<`cut -f 3 -d "," $FILE | sort | uniq | sed -r '/World|International/d`
07
08 for land in "${all_country[@]}"
09 do
10   val_country=`grep -F "$land" $FILE | cut -f $total_test_col -d "," | grep -E "[0-9]" | tail -n 1`
11   grep -q "[0-9]" <<<$val_country && echo "$country,$val_country,`(sed -r 's/\. [0-9]*$/' <<<$val_country) |
   numfmt --grouping`"
12 done | sort -t "," -k 2 -n | gawk -F "," '
13   BEGIN{
14     printf "%-25s %20s\n","Country","Tests"
15     price_per_test=180
16   }
17   {
18     printf "%-25s %20s\n", $1, $3
19     tests+=$2
20   }
21   END{
22     "numfmt --grouping "tests |& getline tests_grouped
23     printf "\n%-20s %25s\n","Tests taken: ",tests_grouped
24     printf "%-20s %25s\n","Price per test: ",price_per_test" euros"
25     "numfmt --grouping "tests*price_per_test |& getline megaturnover
26     printf "%-20s %25s\n","Turnover generated:",megaturnover" euros"
27   }
28 '

```



```
dd@ubuntu:~$
```

South Africa	10
Pakistan	10
Hong Kong	11
Belgium	11
Chile	11
Poland	12
Colombia	13
Iran	13
Israel	15
Saudi Arabia	15
Australia	15
Slovakia	19
Denmark	22
Austria	25
Canada	28
Spain	37
United Arab Emirates	39
Turkey	40
Germany	51
Italy	52
Russia	123
Europe	129
United Kingdom	129
China	160
India	254
North America	385
United States	385
Tests taken:	2,295,820,992
Price per test:	180 euros
Turnover generated:	413,247,778,560 euros

```
dd@ubuntu:~$
```

Figure 5: Based on the the global number of COVID-19 tests and the estimated sales, the pharmaceutical industry is unlikely to go bankrupt in the next few years.

command sorts this data by the second column (total number of tests).

Now a `gawk` script takes care of the formatted output. The `BEGIN` block creates a suitable heading and sets the current price for tests. The block outputs the name of the country and the formatted numbers, adding up the numbers from the second column at the same time. The `END` block then displays the gross revenue for test sales (Figure 5) – a number that the pharmaceutical industry probably would not like to be known.

In the `END` block (line 21), you can see how `gawk` handles pipes. Pass a shell command to a `|& getline variable` (lines 22 and 25). `gawk` then catches the shell command output and stores it in the variable, letting it then work with it downstream. In this way, you can implement even more complicated subroutines, but at the cost of more complex handling. Most importantly, you must close pipes at the right place to produce useful output. In this case, however, it is a simple shell statement, which then also self-terminates.

Conclusions

The media rarely highlights data on population density, the average age of the population, life expectancy, and so on, making it difficult to compare outbreaks across countries. For example, Germany has a population density that is 10 times higher than Sweden, making more stringent measures to contain the pandemic necessary (recommending versus mandating masks).

If you are looking for a more nuanced view of the COVID-19 data than what you might find in the media, `gawk` scripts and a few shell commands let you evaluate COVID-19 data using the criteria you choose. ■■■

Info

- [1] OWID: <https://ourworldindata.org/coronavirus-source-data>
- [2] OWID GitHub page: <https://github.com/owid/covid-19-data/blob/master/public/data/owid-covid-codebook.csv>

Shop the Shop

shop.linuxnewmedia.com

Missed an issue?

You're in luck.

Most back issues are still available. Order now before they're gone!

shop.linuxnewmedia.com

GET IT NOW!
SAVE TIME ON DELIVERY WITH OUR ALTERNATIVE PDF EDITIONS

Document conversion from
the command line

Swiss Army Knife



Pandoc lets you convert files from one markup format to another at the command line. *By Bruce Byfield*

A strength of free software is that applications usually have everything users need for a specific purpose, a tendency that is especially strong in apps for KDE and the command line. Pandoc [1], a universal document converter, exemplifies this strength.

First released in 2006 by John MacFarlane, a philosophy professor at the University of California, Berkeley, Pandoc is a Haskell library for converting between text formats, especially those using a markup format (Table 1). In effect, it is an all-in-one replacement for the dozens of scripts that exist in many distributions for the same purpose.

Pandoc is not equipped to precisely convert complicated layout, such as margins and tables, in formats like PDF or Open Document Form (ODF). However, templates can be created for different formats. Sometimes, though, converting content alone is far better than not at all.

Author

Bruce Byfield is a computer journalist and a freelance writer and editor specializing in free and open source software. In addition to his writing projects, he also teaches live and e-learning courses. In his spare time, Bruce writes about Northwest coast art (<http://brucebyfield.wordpress.com>). He is also cofounder of Prentice Pieces, a blog about writing and fantasy at <https://prenticepieces.com/>.

Moreover, in many cases, Pandoc is adequate for simple formats, like articles or essays, especially in a markup language. It also has advanced features for slide shows, citations, and bibliographies.

By default, Pandoc produces a document fragment as standard output (Figure 1). The general output type plus the specific format must be specified, as well as the input source:

```
pandoc -f markdown -t latex pandoc.txt
```

The result is a fragment for the extension specified that can be pasted into another document. To save the output, you must specify a file using the `--output (-o)` option. If you want a complete file, rather than a fragment, add the `--standalone` option. As with many command-line options, saving to a file produces no output unless something goes wrong.

If you do not specify the input and output, Pandoc will attempt to guess them. To ensure formatting, a template file can be specified (see the Templates

section below). Use the `-t` option to list the types of formats supported. If multiple input files are specified, they are concatenated into a single output file with a space between the contents of each input file.

Templates

Each supported format has a default template stored in `/usr/share/pandoc/data/templates/`. Most follow the naming structure `default.FORMAT`. Exceptions include ODT's template, which is named `default.opendocument`, and PDF, which shares the `default.latex` template. In addition, EPUB uses `epub-page.html`, `epub-coverimage.html`, and `epub-titlepage.html`. You can view the default template using the command `pandoc -D FORMAT` (Figure 2).

You can write or download custom templates [2] or modify copies of existing templates [3] if the default template does not meet your needs. Templates consist of fields with fixed values and may include variables that are

```
bb@nanday:~/work/journalism/2021/05-may$ pandoc -f markdown -t latex pandoc.txt
Pandoc Document conversion from the command line

\begin{center}\rule{0.5\linewidth}{\linethickness}\end{center}

Table 1: Supported Source Formats

\begin{verbatim}
Creole
DocBook
EPUB
```

Figure 1: By default, Pandoc writes a fragment to standard output.

replaced by elements of the source file, often automatically. For example, the variable `<title>${title}</title>` is replaced automatically by the source file's title. More advanced users can include if/else or conditional statements. For a full description of custom templates, see Pandoc's man page and user guide [4].

In the end, if content is more important than structure, you can generally use the default templates without tweaking them.

Note that early releases of Pandoc required additional applications to convert to PDF. Several online sources like Wikipedia continue to list this requirement, but it is now obsolete.

Input/Output Options

Instead of templates, you can do some formatting using options. To eliminate any ambiguity in the command structure, you can specify the input format with `--from FORMAT (-f FORMAT)` or `--read FORMAT (-r FORMAT)`, and the output with

`--to FORMAT (-t FORMAT)` or `--write FORMAT (-w FORMAT)`. Similarly, although the default directory for all output to a file is `.pandoc`, you can specify another directory with `--data-dir=DIRECTORY`.

Other options affect the internal formatting. For instance, while the default format is to replace tabs with spaces, `--preserve-tab (-pv)` will override the default. When setting up tabs, you may also use `--tab-stop=NUMBER` to change the default four spaces used for tabs. You can also use `--base-header-level=NUMBER`

Table 1: Supported Formats

(<← = conversion from; → = conversion to; ↔ = conversion from and to)	
Lightweight markup formats	
↔	Markdown (including CommonMark and GitHub-flavored Markdown)
↔	reStructuredText
→	AsciiDoc
↔	Emacs Org Mode
↔	Emacs Muse
↔	Textile
←	txt2tags
HTML formats	
↔	(X)HTML 4
↔	HTML5
Ebooks	
↔	EPUB version 2 or 3
↔	FictionBook 2
Documentation formats	
→	GNU Texinfo
↔	Haddock markup
roff formats	
↔	roff man
→	roff ms
TeX formats	
↔	LaTeX
→	ConTeXt
XML formats	
↔	DocBook version 4 or 5
↔	JATS
→	TEI Simple
Outline formats	
↔	OPML
Bibliography formats	
↔	BibTeX
↔	BibLaTeX
↔	CSL JSON
↔	CSL YAML
Word processor formats	
↔	Microsoft Word DOCX
↔	OpenOffice/LibreOffice ODT
→	OpenDocument XML
→	Microsoft PowerPoint
Interactive notebook formats	
↔	Jupyter notebook (IPYNB)
Page layout formats	
→	InDesign ICML
Wiki markup formats	
↔	MediaWiki markup
↔	DokuWiki markup
←	TikiWiki markup
←	TWiki markup
←	Vimwiki markup
→	XWiki markup
→	ZimWiki markup
↔	Jira wiki markup
Slide show formats	
→	LaTeX Beamer
→	Slidy
→	reveal.js
→	Slideous
→	S5
→	DZSlides
Data formats	
←	CSV tables
Custom formats	
→	Custom writers can be written in Lua
PDF	
→	Via <i>pdflatex</i> , <i>lualatex</i> , <i>xelatex</i> , <i>latexmk</i> , <i>tectonic</i> , <i>wkhtmltopdf</i> , <i>weasyprint</i> , <i>prince</i> , <i>context</i> , or <i>pdfcrow</i>

Table courtesy of the Pandoc website [1]

```
bb@nanday:~$ pandoc -D opendocument
<?xml version="1.0" encoding="utf-8" ?>
<office:document-content xmlns:office="urn:oasis:names:tc:opendocument:xmlns:office:1.0" xmlns:style="urn:oasis:names:tc:opendocument:xmlns:style:1.0" xmlns:text="urn:oasis:names:tc:opendocument:xmlns:text:1.0" xmlns:table="urn:oasis:names:tc:opendocument:xmlns:table:1.0" xmlns:draw="urn:oasis:names:tc:opendocument:xmlns:drawing:1.0" xmlns:fo="urn:oasis:names:tc:opendocument:xmlns:xsl-fo-compatible:1.0" xmlns:xlink="http://www.w3.org/1999/xlink" xmlns:dc="http://purl.org/dc/elements/1.1/" xmlns:meta="urn:oasis:names:tc:opendocument:xmlns:meta:1.0" xmlns:number="urn:oasis:names:tc:opendocument:xmlns:datastyle:1.0" xmlns:svg="urn:oasis:names:tc:opendocument:xmlns:svg-compatible:1.0" xmlns:chart="urn:oasis:names:tc:opendocument:xmlns:chart:1.0" xmlns:dr3d="urn:oasis:names:tc:opendocument:xmlns:dr3d:1.0" xmlns:math="http://www.w3.org/1998/Math/MathML" xmlns:form="urn:oasis:names:tc:opendocument:xmlns:form:1.0" xmlns:script="urn:oasis:names:tc:opendocument:xmlns:script:1.0" xmlns:ooo="http://openoffice.org/2004/office" xmlns:ooow="http://openoffice.org/2004/writer" xmlns:oooc="http://openoffice.org/2004/calc" xmlns:dom="http://www.w3.org/2001/xml-events" xmlns:xforms="http://www.w3.org/2002/xforms" xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" office:version="1.2">
  <office:font-face-decls>
    <style:font-face style:name="Courier New" style:font-family-generic="modern" style:font-pitch="fixed" svg:font-family="'Courier New'" />
  </office:font-face-decls>
  <office:automatic-styles>
    $automatic-styles$
  </office:automatic-styles>
$for(header-includes)$
```

Figure 2: The beginning of the default template for ODF.

to set the first heading level to use and `--smart (-S)` to use typographic characters such as smart quotes and em dashes (instead of two hyphens).

Individual formats also have their own formatting options. For instance, in HTML5, `--section-div` adds `<div>` or `<section>` tags, which can be formatted with CSS style sheets created outside Pandoc. LaTeX, ConTeXt, and DocBook output can use `--chapters` to convert the top-level headings into chapters, while `--no-tex-ligatures` suppresses ligatures in LaTeX or ConTeXt output, which can be convenient with some recent OpenType features. More generally, several options are intended primarily for code, such as the self-explanatory `--no-wrap`, `--columns=NUMBER`, `--no-highlight`, and `--highlight=STYLE` (with options of `pygments`, `kate`, `monochrome`, `espresso`, `zenburn`, `haddock`, and `tango`). Many of these options can reside in a single file that is specified with `--defaults = FILE`, eliminating the need to continually structure a detailed command.

For many output formats, options provide most formats with the exception of spacing options. However, layout can be added via CSS style sheets and linked with `--css=URL`. Some output formats have specific options for style sheets, such as `--reference-odt=FILE` (ODT), `--reference-docx=FILE` (DOCX), and `--epub-style-sheet=FILE` (EPUB). If you regularly convert to such formats, developing a style sheet may be worth the effort. You may even find a style

sheet online that you can use with little or no modification.

Special Uses

Besides routine format conversion, Pandoc has several special uses. For instance, Pandoc supports several slide show applications, including PowerPoint. However, to judge by the available options, its main emphasis is on Beamer, a LaTeX-based presentation application [5]. The markup for a Beamer slide is as simple as starting each one with `##`. To Beamer's own thorough array of features, Pandoc adds options of its own. While converting a file for use in Beamer, Pandoc can define a logo, title graphics, navigation symbols, Beamer theme, and the aspect ratio for slides. Common layouts include slide backgrounds, transitions, and lists in which items are displayed one at a time. There is even an option to add Beamer options to the converted presentation. In addition, Pandoc can convert a Beamer presentation to an article. Pandoc's emphasis on Markdown provides a professional slide show application regardless of the office suite used.

Pandoc also has extensive support for citations and bibliographies. Using the option `--citedoc`, Pandoc can generate citations from a source file and a bibliographic database specified with one `--bibliography=FILE` for each bibliography used. BibLaTeX (.bib), BibTeX (.bibtex), CSL JSON (.json), and CSL YAML (.yaml) are all supported formats.

By default, Pandoc uses the *Chicago Manual of Style* citation style, although other citation formats can also be defined. There is even a `--citation-abbreviations=FILE` option that can define abbreviations for often used titles. The citations and bibliography are kept separate from the Pandoc files, making it easy to update and then generate a new file.

Only an Overview

I've only described some of Pandoc's most useful features. In addition to the features mentioned here, Pandoc also includes support for math equations, EPUB, and other special uses. Its functionality is impressive, but best learned one task at a time to avoid being overwhelmed. Probably, too, your ability to use Pandoc depends on your knowledge of Markdown, style sheets, and other related subjects. But for those with the knowledge to use it docu-ments, Pandoc has the thoroughness that is the hallmark of the best Linux applications. ■■■

Info

- [1] Pandoc: <https://pandoc.org>
- [2] Custom templates: <https://bookdown.org/yihui/rmarkdown/template-pandoc.html>
- [3] Other templates: <https://github.com/jgm/pandoc-templates>
- [4] Pandoc User's Guide: <https://pandoc.org/MANUAL.html>
- [5] Beamer: <https://www.overleaf.com/learn/latex/Beamer>

IT Highlights at a Glance

The collage features several overlapping content pieces:

- ADMIN HPC**: A newsletter snippet with sections for 'HPC Up Close', 'Highlights', and 'Further Reading'. It mentions 'High-Performance Distributed Python' and 'CEIN Sponsors Free Lecture Series on Quantum Computing'.
- ADMIN Update - Hottest Links**: A newsletter snippet with a 'Highlights' section and a 'Most Read Articles' section. It lists articles like 'Certificate Transparency' and 'SMPTel Has Done Choking Blow to CVE-ID'.
- LINUX UPDATE**: A newsletter snippet with sections for 'FEATURED ARTICLES', 'FURTHER READING', and 'MOST READ'. It includes articles like 'Video Conferencing with Jitsi' and '20 Years of Linux Magazine'.
- SC20**: A snippet for the Security Conference 2020, featuring a 'FULL PROGRAM' and 'FULLY VIRTUAL' badge.
- Be Prepared**: A snippet for 'Effective Protection for T8, Z87 Networks' by Genio.
- 20 Years of Linux Magazine**: A snippet celebrating the anniversary of Linux Magazine.

Too busy to wade through press releases and chatty tech news sites? Let us deliver the most relevant news, technical articles, and tool tips – straight to your Inbox. Subscribe today for our excellent newsletters:

ADMIN HPC • ADMIN Update • Linux Update

and keep your finger on the pulse of the IT industry.

ADMIN and HPC: bit.ly/HPC-ADMIN-Update
Linux Update: bit.ly/Linux-Update



Experimental package management with Nix and NixOS

Brand New Package

NixOS and the Nix package manager offer a promising new approach to the challenge of managing packages in Linux. *By Mats TAGE Axelsson*

In most Linux distributions, the files associated with the OS end up in specific places. Most Linux distributions subscribe to the filesystem hierarchy defined in the Linux Standards Base (LSB), which specifies the familiar directory names you are accustomed to if you work in Linux (`/etc`, `/dev`, `/bin`, etc.). On most Linux systems, if you do not have the files in the right places, the applications will not find them. But experimentation is at the heart of Linux, and every rule has an exception. An innovative project called NixOS [1] takes a different approach to system configuration and package management.

On NixOS, all files are in the Nix store. For applications to find them, NixOS links to the correct locations with symlinks. This approach makes all kinds of things possible. For instance, you can have several versions of any library and let the application know which link to use. For development environments, you can even test an application using system files from

different distributions, just using the package manager.

On a NixOS system, everything is reproducible. A good backup of your home directory and a few `*.nix` files is all you need to get back up after a crash.

At the heart of NixOS is the Nix package manager. Nix is a cross-platform package manager that can run on other distros as well, but NixOS was created to provide a native environment for testing and experimenting with Nix.

When you make a change to your system, NixOS creates a new *generation*. A generation contains all the links to the software you want to run. When you come into your favorite boot manager (GRUB?), it will give you the default boot or alternatives. The alternatives contain all versions of the system you have created since you installed. Yes, as you will learn later in this article, saving all previous versions of the system can waste a lot of disk space, so you will want to take the time to retire older generations you are no longer using.

Advantages and Disadvantages

From a practical point of view, changing your configuration is easier on NixOS than on other distributions. Let's say you want to use another window manager but still keep the first one for backup. If you make the install and something goes wrong, or if you change your mind, you can roll back. To roll back, reboot and pick the earlier generation.

You might be missing a few packages, but Nix is popular enough to boast 80 thousand packages in the default repository. You can create your own package if you are capable enough. There is also support for the ApptImage and Flatpak systems, although Snap support seems a little lacking.

NixOS does not have a graphical installer. You need to change the settings for keyboard, locale desktop environment, and similar by changing the `configuration.nix` file. Keep in mind, also, that the Nix package manager does not have a graphical interface, which could be a disadvantage for some users.

Photo by Jess Bailey on Unsplash

When you first try NixOS, all you need to do is partition your drives, or label the existing ones, and run `nixos-generate-config`. Next, you change one file (`configuration.nix`), run `nixos-install`, and you are done.

NixOS lets you put all the packages that you want to use in the `configuration.nix` file. You can then use `configuration.nix` to recreate the same configuration on any other machine.

Nix Package Manager

The real strength of NixOS is the Nix package manager. The package manager has many unique features, such as rollback, support for multiple versions, and the ability to reproduce the install. The Nix package manager is an interesting option if you are developing software and you are getting tired of the multitude of virtual environments.

To fetch the Nix package manager only, you need to use the link available at the project website. The installer will ask you for `sudo` access to create the `/nix` directory as root.

```
sudo mkdir /nix
curl -L https://nixos.org/
nix/install | sh
```

The more cautious of you will download the script and read it through before running it. Bear in mind though, that if you run the install with a non-privileged user, the script will ask for root, so it can create and write to the `/nix` directory.

You can also create a multi-user environment with the installer script. To do this, use the `--daemon` option:

```
sh <(curl -L https://nixos.org/
nix/install) --daemon
```

Notice that you are still piping the script to the shell. Many seasoned administrators will not bat an eyelid over this, except for the security implications of running a

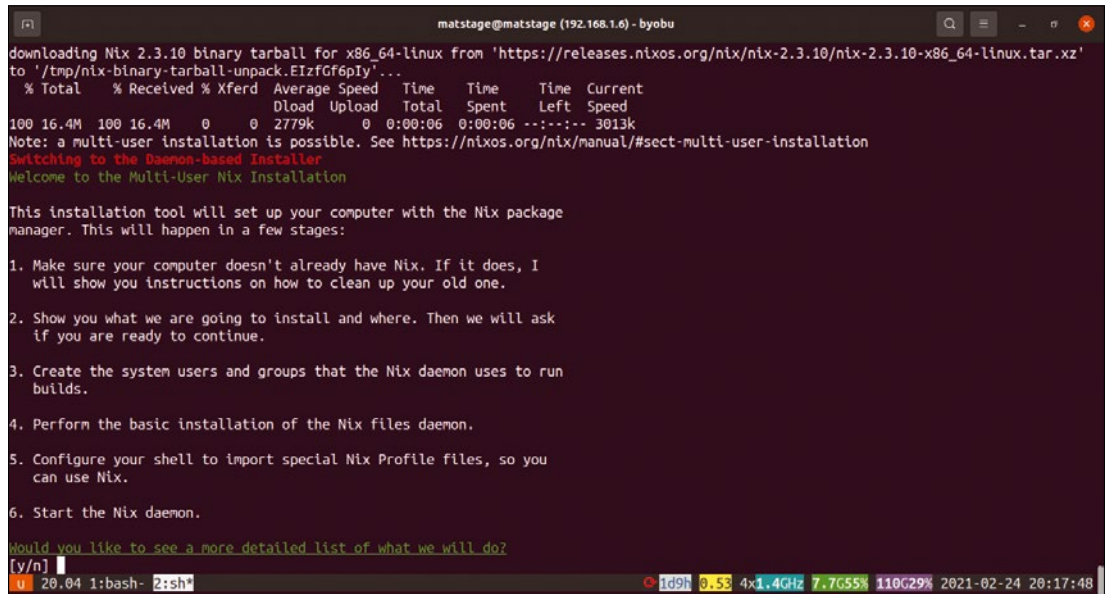


Figure 1: The Nix package manager is very informative. If you pay attention, you will know every detail of your system.

script straight from a URL. You can see the start of the script in Figure 1.

The installer is aware of whether you have already installed Nix before and warns you about it.

The installer has several different configuration options, depending on what you want to do, but the main idea is to support any version of software at any time. One option is to install the Nix package manager using the installer and then pick a package and see if it has a version you like.

```
nix-env -i gimp
```

After having installed Gimp, you use the `which` command to discover the version:

```
which gimp
```

It is a good idea to check through the applications you are already using and see if they are present and maintained in the NixOS repositories. NixOS is still not very common, so some packages may not have the latest available version. If you want to check what software is available before installing, the NixOS website has a great searchable list (Figure 2).

After you have installed and tested a few packages, you should consider the space on your system. Due to the store and the way Nix stores the files, you may have many binary files that are identical. To fix this, search through the Nix store and re-link the files.

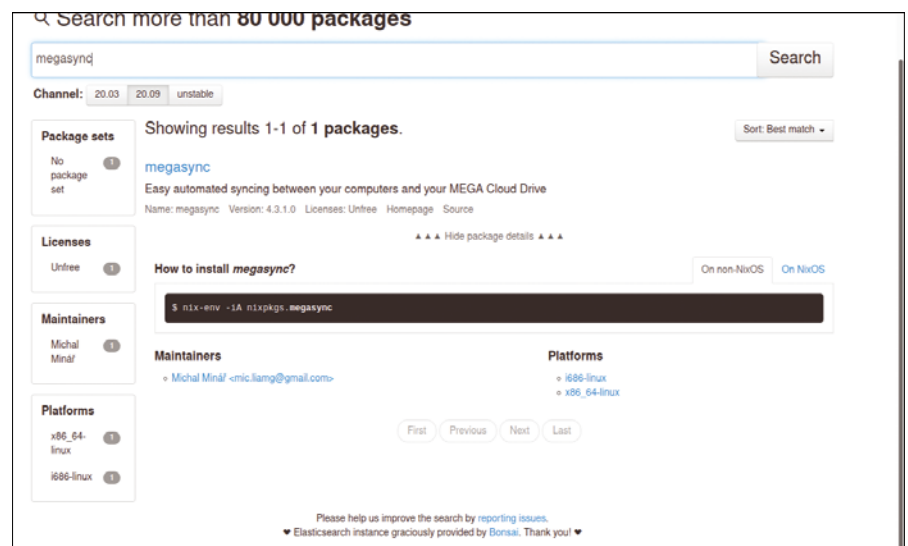


Figure 2: The NixOS website includes a convenient package search feature.

Don't worry – there is a simple command line tool for this task. Check your Nix store using the `nix-store` command:

```
nix-store --optimise
```

Best practice is to have NixOS do this automatically on a regular basis. What the command does is go through all the binaries and compare them. When it finds a match, it creates a new link and erases the, now unnecessary, file.

When you have tried your favorite applications, you can list them using the `nix-env` command again. Any applications that you cannot find in the Nix repositories might be available using AppImage and Flatpak. Support for Snap packages is a bit more dubious. To store your choices in a separate file, redirect the result to a text file:

```
nix-env -q >>Nix-applications.txt
```

Listing 1: Teensy Environment

```
01 shell.nix:
02
03 { system ? builtins.currentSystem }:
04
05 let
06   pkgs = import <nixpkgs> { inherit system; };
07 in
08   pkgs.callPackage ./default.nix {}
09
10
11 default.nix:
12
13 { stdenv, lib, gcc, gcc-arm-embedded-, teensy-loader-cli }:
14
15 stdenv.mkDerivation rec {
16
17   buildInputs = [
18     gcc
19     gcc-arm-embedded-4_7
20     teensy-loader-cli
21
22   ];
23 }
```

```
matstage@matstage:~/git$ python
Command 'python' not found, did you mean:
  command 'python3' from deb python3
  command 'python' from deb python-is-python3

127 matstage@matstage:~/git$ python3
Python 3.8.5 (default, Jul 28 2020, 12:59:40)
[GCC 9.3.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> quit()
matstage@matstage:~/git$ nix-shell -p python3

[nix-shell:~/git]$ python
Python 3.8.7 (default, Dec 21 2020, 17:18:55)
[GCC 10.2.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> quit()

[nix-shell:~/git]$
```

Figure 3: The `nix-shell` command shows that the Python version is different.

You can use this list when you decide to switch your entire system to NixOS.

Dev Environments

Once you install the Nix package manager, you can use it to keep your development environments clean and consistent. With the `nix-shell` command, you can specify exactly what versions you have available. If anyone else is sharing the project, they can use the `default.nix`

and `shell.nix` files to create an environment identical to yours.

To install an individual package, use the `-p` option. If you want to run Python without installing on your system, or a version other than your current system, try the following:

```
nix-shell -p python3
```

You will get a prompt (Figure 3).

Theoretically, you could install your whole environment with the `-p` option, but that would be counterproductive. Instead, you can put all the complicated stuff in the `shell.nix` and `default.nix` files.

Inside these files, you can define your installation. The `nix-shell` command calls the `shell.nix` file first. Listing 1 shows an old example that starts an environment for running the Teensy development board, courtesy of Richard Zetterberg [2].

In Listing 1, the curly brackets at the top contain the environment you need for this shell. In most cases, you will put `stdenv` here. You can see that this file needs `gcc`, `lib`, and the special `gcc` version for ARM devices. This does not tell you how to compile the package, though, which you do with `stdenv.mkDerivation`. In Nix parlance, the word derivation means building something from something else.

Listing 2: Python 3.8

```
01 { pkgs ? <nixpkgs> , ... }:
02
03 with pkgs;
04 mkshell {
05   buildInputs = [
06     python38Packages.flask
07   ];
08 }
```


To make this work, you create the directory and place the files inside. Only one of the files is required, but most people mix in a separate default.nix to separate the parts of the system. In the file, the most important part you want to know about is the curly brackets: {...}. The curly brackets contain the standard environment you want to take from. The <nixpkgs> references the list of files that exist in the repository. For a shell to work, you only need to put those curly brackets in and then add pkgs.mkshell.

The little script file in Listing 2 gives you the ability to start Python 3.8 in a shell. Since the flask packages are also added, you also have those modules available. To find the packages you want, see the NixOS website package page [3]. When you run nix-shell, the install happens on its own. You do not need to install it first.

Preparing Your System

The first issue you'll face when transitioning to NixOS is which applications to enable. Most users will find all the applications they need in the NixOS repositories. (By the way, there are also some alternative repositories.)

Even though the installer is command-line only, it is surprisingly easy to use. You can use any standard partition scheme. The install manual has examples that you can copy if you don't have any special demands. You do have to know how to handle a partition program. They suggest you use partition labels, which is really useful. You'll need to partition your disks, mount the root partition, boot your new system, and run the nixos-generate-configuration command:

```
nixos-generate-configuration 2
--root /mnt
```

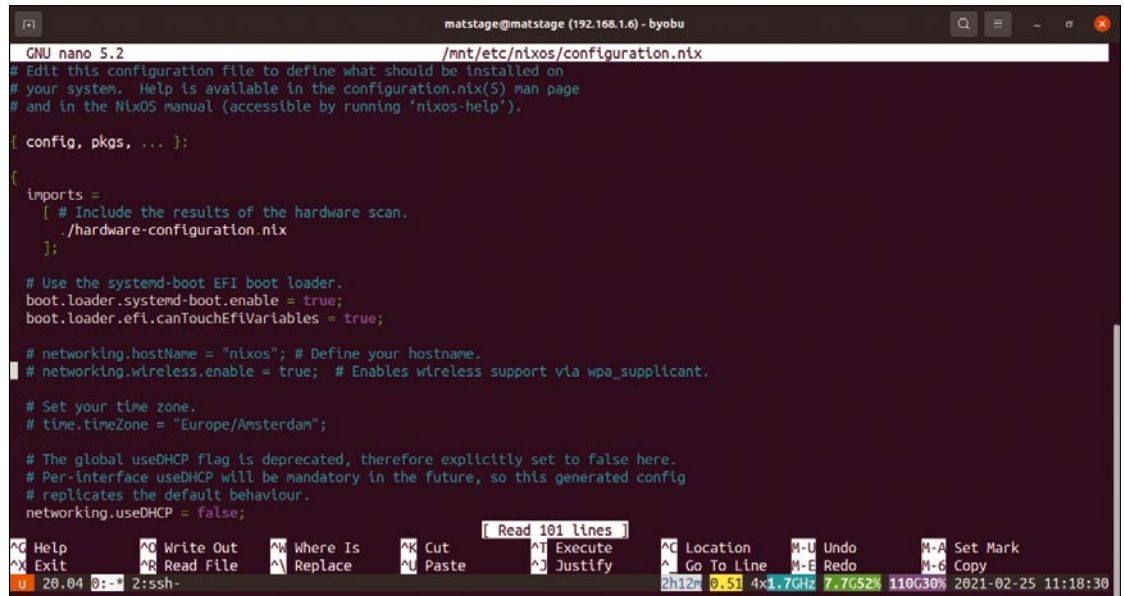


Figure 4: The top of the configuration.nix file points to the hardware file; you can add your own files for special purposes.

Use the manual that is included in the install ISO. The manual looks long and complicated, but it is actually comprehensive and clear. When this command has finished, you will need to edit your configuration.nix file.

configuration.nix

The nixos-generate-configuration command generates two files: configuration.nix and hardware.nix. In most cases, you will do nothing with the hardware file, since the installer has set everything according to the current system you are on. The interesting part is the configuration.nix file. Inside,

you will have all the settings you need and the applications you want on the general system.

Starting from the top, the file has the curly brackets that define your environment (Figure 4). You change this setting for extreme cases; on a desktop, leave it alone. Second, you have an import statement that imports the hardware file. You can use this section if you want to break out some of your configuration. Any Nix file that you reference will be added to your configuration.

Next, there should be a few lines about your boot loader. The Nix installer

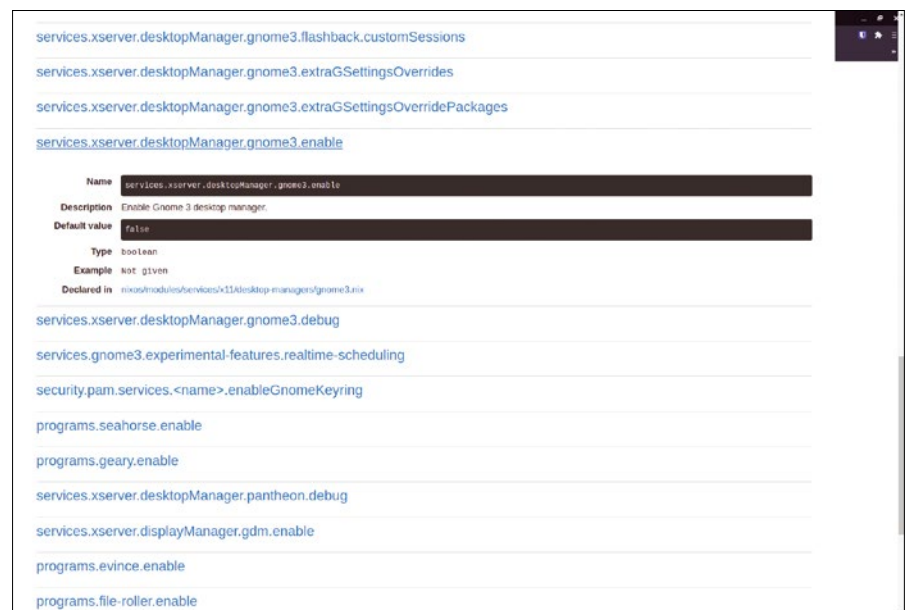


Figure 5: Exploring options and applications.

detects UEFI and BIOS; don't change the boot loader configuration unless you have very particular circumstances. As you go down the file, most options are clear and follow standards for locale and time zone. When you want to choose a desktop system, however, you are left in the dark – the configuration file has nothing for you.

Fear not! The NixOS website has a very nice interface for looking for options and applications [3]. On this page, you can enter in the application that you want. If you search for Gnome, you get a long list of options (Figure 5). Click on the option you want and the *Name* given in the list is the name you put in the file. The Gnome entry needs the setting:

```
services.xserver.desktopManager. ⌘
gnome3.enable = 1;
```

The list of options has more than 10,000 entries! You can use them when you want, but don't let the sheer number hold you back. All options you need, initially, are in the original file. You will only need to add a window or desktop manager. Once you have written the file, you can install. The procedure is one script. At the end of the script, you need to set the root password.

The script compiles all programs that are not available in binary form, so it might take a while. You have options for distributing the compile to other computers. If you have problems, you can also choose the `--show-trace` option to see a diagnostic trace.

After the installation, reboot into your new system.

Updating Applications

At some point a package might require an upgrade. You can use the `--upgrade` action for this job. More surprising is that you can upgrade to the same versions. As the manual states, this might seem useless. However, you can use it to recompile your current applications to newer libraries and also cut out duplicates of libraries. Enter the following command:

```
nix-env --upgrade --leq
```

You can also skip the `leq` parameter, in which case the command will upgrade everything without optimizing the libraries you use. An upgrade leads to another concern; you may have many versions of any file in your system, which means you'll have many generations to deal with.

Generation Handling

Creating many generations can waste large amounts of file space. The utility that handles this is `nix-env`, but you should use the wrapper `nix-collect-garbage`. `nix-collect-garbage` has the standard delete option that takes away all old generations. You can do this manually, but more useful is the `--delete-older-than` option. You should set this as a regular script:

```
nix-collect-garbage ⌘
--delete-older-than 10d
```

Using Nix in the Cloud

Nix is not supported by many cloud companies, but many still provide ISO options. NixOS has a page describing which ones are friendly and which ones you can still get started on. The procedure is similar, though getting the disks created is a bit of a chore when full support is not included. The list of services is available at the project website [4].

Conclusion

Nix is very powerful tool for handling diverse configurations. Many developers appreciate the ease at which they can just run the Nix shell and have all dependencies correctly configured every time. Ordinary users will appreciate the power of rollbacks and the possibility to run old software without bending over backwards. ■■■

Info

- [1] NixOS: <https://nixos.org/>
- [2] Teensy development on NixOS: <https://rztterberg.github.io/teensy-development-on-nixos.html>
- [3] NixOS package search: <https://search.nixos.org>
- [4] NixOS-friendly cloud hosts: https://nixos.wiki/wiki/NixOS_friendly_hosters

Author

Mats Tage Axelsson keeps aging hardware for the challenge of handling crashes. Oh, and he also is a cheapskate.

Hit the ground running with Linux

FREE DVD!

JOIN THE LINUX REVOLUTION!
← ALL THE SOFTWARE YOU NEED!

**GETTING STARTED WITH
LINUX**

• MORE POWERFUL • MORE SECURE • MORE FUN

LEARN HOW TO SET UP A LINUX SYSTEM TO:

- Listen to Music
- Play Games
- Process Photos
- Surf the Web
- and Much More!

2020 Edition

WWW.LINUX-MAGAZINE.COM

Want your friends and colleagues to make the switch to Linux?

This single issue shows beginners how to:

- install Linux
- download and install free software for your Linux system
- play games
- create documents and spreadsheets
- process photos
- play music and videos
- and much more!

ORDER ONLINE: shop.linuxnewmedia.com/specials

Python package simplifies algebraic equations

Quick Fill

Whether he's filling the bathtub with water or routing electricity through resistors – Mike Schilli juggles mathematical formulas with the assistance of the Python SymPy package.

By Mike Schilli

still remember a traumatic event as an elementary school student: A weekend newspaper had set a logic puzzle for kids, the solution of which it promised to publish in the next issue a week later. It involved a bathtub with two taps: one of which filled the tub in 10 minutes, the other in 15. The question was how long would it take to fill the tub if both taps were turned on all the way.

As a little boy, I was absolutely sure that 10 plus 15 equals 25, which is 25 minutes. My father laughed and suggested that couldn't be true, because two taps would fill the bathtub faster than one alone. The next weekend, I was initially triumphant, because there it was – printed in black and white in the following issue – the con-

firmation that 25 minutes was the correct solution.

But disillusionment followed one week later: After receiving angry reader comments, the newspaper had to admit that it had made a mistake, because it does not take 25 minutes, but only 6 minutes, to fill the tub with both taps. I nearly fell off my chair and decided at that point to become a famous columnist peddling logic puzzles.

Faucets and Resistors

It was only much later, in college, during a lecture on circuit technology, that I came across a similar problem, which somehow translated nicely to a solution to the bathtub problem: parallel arrangement of resistors in a circuit. The electrical engineer measures the resistance of a conductor in ohms, and

the higher the value, the more it slows down the flow of electrical current.

Now how do you calculate the equivalent resistance of a parallel circuit with two resistors R_1 and R_2 as shown in Figure 1? The solution involves the currents I_1 and I_2 flowing through each resistor. They add up to the total current I after the two branches are reunited. The measured voltage is the same everywhere, that is V , and according to Ohm's law $V = R \cdot I$. Ergo, with individual currents I_1 and I_2 adding up to I , the result comes out as $V/R_1 + V/R_2$.

If we now replace the two parallel resistors R_1 and R_2 with a substitute resistor R , we also get $I = V/R$ for the total current, so this gets us: $V/R =$

Author

Mike Schilli works as a software engineer in the San Francisco Bay area, California. Each month in his column, which has been running since 1997, he researches practical applications of various programming languages. If you email him at mschilli@perlmeister.com he will gladly answer any questions.




Figure 1: Two resistors R_1 and R_2 in a parallel circuit ...




Figure 2: ... and their equivalent resistance R .

Lead Image © Andrea Danti, 123RF.com

$V/R1 + V/R2$. The voltage V can be canceled out to give $1/R = 1/R1 + 1/R2$ (Figure 2). After transformation, the following result appears:

$$R = (R1 * R2) / (R1 + R2)$$

Today, computer programs perform the computational work instead of the engineer doing it manually. Listing 1 [1] shows the development of the formula using SymPy, a symbolic algebra package in Python. It can be used to define symbols, which the package later leaves intact when evaluating formulas, instead of immediately replacing variables with values and determining the formula's numerical result. With the `simplify()` function, SymPy can simplify expressions using the rules of algebra – just as any mathematically skilled person would do.

Thus, line 3 in Listing 1 defines a whole slew of symbols. For the two resistors of the circuit and their equivalent resistance, there are $r1$, $r2$, and r . For the applied voltage, there is v ; for the two current components, there are $i1$ and $i2$. The formulas in lines 4 to 6 apply Ohm's law and sum the partial currents to receive at the total current.

The resulting formula for the equivalent resistance r comes to light when the script is called using the `pprint()` (pretty print) function (Figure 3). SymPy has obviously realized that the voltage v can be canceled out of the fraction, and the result no longer depends on it.

If you use Python 3, you can install packages like `sympy` and `matplotlib`, which I'll be using in a moment to illustrate results, easily with the command:

```
pip3 install <package>
```

In order to avoid pulling the carpet out from under the feet of other Python scripts on the same host, I prefer doing that in Python's virtual environment [2].

From Circuit to Bathtub

This formula also works for the bathtub problem mentioned earlier: In the numerator, the filling times per tap get multiplied ($10 * 15$), and the denominator holds their sum ($10 + 15$). According to Adam Riese, this comes out to $w150/25$ (i.e., six minutes), just like in the example solution in the Sunday paper.

By the way, you can also think of the solution as follows: After one minute, the first tap has filled the tub to one tenth and the second to one fifteenth. Both together fill in one minute $1/10 + 1/15 = 3/30 + 2/30 = 5/30 = 1/6$ of the tub, so you can get into the bath after six minutes.

Borderline

What happens if both resistors in the parallel circuit drop down to 0 ohms – that is, they do not slow down the current flow at all, or, in the case of a bathtub with Niagara Falls-style faucets, fill the tub in next to no time? Intuitively, it is clear that this kind of mega-faucet would fill the tub practically as fast when connected in parallel as when operated individually. But if you set the values for $(R1 * R2) / (R1 + R2)$ in the formula $(R1 * R2) / (R2) / (R1 + R2)$ to zero and pass the construct to a Python program for computing, you will experience a fatal error: Computers steadfastly refuse to perform divisions by zero, because it is mathematically undefined.

In the present case, however, both the numerator and the denominator contain a value tending towards zero, which sometimes yields interesting (because finite) results. Mind you, a true zero as denominator defies mathematical definition, but the limit for $R1$ and $R2$ approaching zero can definitely be computed.

SymPy provides the `limit()` function, which takes a symbolic formula,

Listing 1: parallel.py

```
01 #!/usr/bin/env python3
02 from sympy import simplify, symbols, pprint
03 r, r1, r2, v, i1, i2 = symbols("r r1 r2 v i1 i2")
04 i1 = v / r1
05 i2 = v / r2
06 r = v / (i1 + i2)
07 pprint(simplify(r))
```

Listing 2: limit.py

```
01 #!/usr/bin/env python3
02 from sympy import limit, symbols
03 r1, r2, r = symbols("r1 r2 r")
04
05 r = (r1 * r2) / (r1 + r2)
06 r1 = r2
07
08 # r1/2->0
09 print(limit(r, r1, 0))
10
11 # 1/x with x->0
12 x = symbols("x")
13 print(limit(1/x, x, 0))
```

Listing 3: Limit Value

```
$ ./limit.py
0
oo
```

```
$ ./parallel.py
r1 * r2
-----
r1 + r2
$
```

Figure 3: Output via `pprint()`.

a symbol (such as $r1$) and a limit (here θ). Listing 2 defines $r1 = r2$ beforehand, meaning that both variables in the `limit()` function tend towards zero.

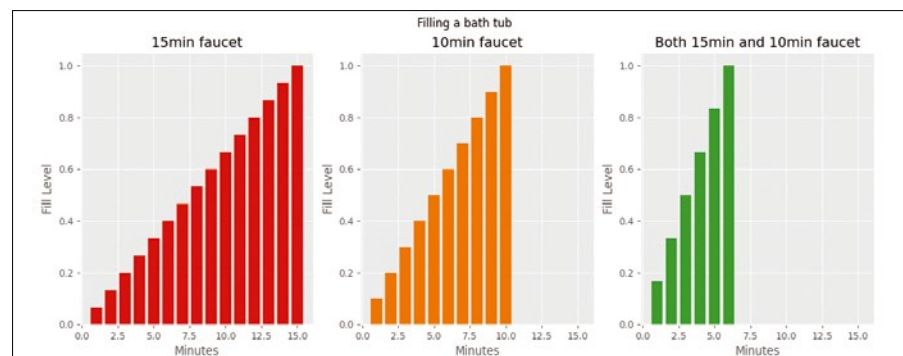


Figure 4: Bar charts generated by `pyplot` for bathtub filling: On the left, the 15-minute tap; in the middle, the 10-minute tap; and on the right, both taps.

The result for the equivalent resistor, as the values for `r1` and `r2` approach zero, is that the script returns, as expected, a value of zero (Listing 3).

The second test case, starting from line 12 in Listing 2, illustrates with an example what happens to another formula, $1/x$, when x tends towards zero. The output in Listing 3 shows here that the result of the formula tends to infinity in this case, which is what the ASCII output `oo` tries to illustrate.

Beautiful and Colorful

To demonstrate how the two faucets interact while filling the tub, the three bar graphs in Figure 4 represent the respective fill levels by minute. In the graph on the left, the bath user only turns on the slow faucet, which fills the tub in 15 minutes; in the middle, the 10-minute fast faucet is active; and, on the right, both faucets are active together.

In Listing 4, starting in line 5, the `fill_tub()` function takes the tub fill rate per minute (`per_min`) as well as the

number of values on the X axis displayed (i.e., the minute ticker). Although the scenarios presented each cover different sets of minute values, the `matplotlib` functions insist that X and Y values be arrays of equal size, or else they will provoke non-intuitive error messages from the depths of the library. The `fill_tub()` function returns two arrays `lx` and `ly` that contain values for the X and Y axes in the bar graph. The two arrays are the same length, and unoccupied Y values with the tub already full are simply set to zero for initialization purposes in line 7.

Painting three different graphs side by side in an image file is a piece of cake for `matplotlib`: The `subplots()` function simply creates a graph grid with one row as well as three columns in line 20 and returns an array of three graph objects in `ax`. With these settings, `matplotlib` manages to arrange the three graphs beautifully, without any manual intervention. Finally, the last line writes the image data to a newly created image file `bars.png`.

Exact Fraction Calculation

To prevent the program from immediately turning a fraction like $1/15$ into a floating-point number, Listing 4 pulls in the `fractions` package, which handles real fractional arithmetic. Thus, line 23 with `Fraction(1, 15)` can actually pass $1/15$ to `fill_tub()` rather than, say, `0.066666666667`. The latter would cause nasty rounding errors when summing up fractional values later and possibly even cause the tub to overflow. The `fractions` package, on the other hand, overloads operators like `+` or `>`. Therefore, `sum += per_min` in line 10 actually adds $1/15$ to the amount of water in the tub, and `sum` inside the `if` condition in line 11 is later exactly 1 when the tub is full, rather than some close, but not quite right floating-point value. ■■■

Info

- [1] Listings for this article: <ftp://ftp.linux-magazine.com/pub/listings/linux-magazine.com/248/>
- [2] Virtual environments: <https://docs.python.org/3/tutorial/venv.html>

Listing 4: bars.py

```
01 #!/usr/bin/env python3
02 from matplotlib import pyplot as plt
03 from fractions import Fraction
04
05 def fill_tub(per_min,xmax):
06     lx=list(range(1,xmax+1))
07     ly=[0] * xmax
08     sum=0
09     for i in range(xmax):
10         sum+=per_min
11         if sum > 1:
12             break
13         ly[i]=sum
14
15     return lx,ly
16
17 xmax=15
18
19 plt.style.use('ggplot')
20 fig,ax = plt.subplots(nrows=1,ncols=3,figsize=(15,5))
21 fig.suptitle("Filling a bath tub")
22
23 lx,ly = fill_tub(Fraction(1,15),xmax)
24 ax[0].bar(lx,ly,color='tab:red')
25 ax[0].set_xlabel("Minutes")
26 ax[0].set_ylabel("Fill Level")
27 ax[0].set_title("15min faucet")
28
29 lx,ly = fill_tub(Fraction(1,10),xmax)
30 ax[1].bar(lx,ly,color='tab:orange')
31 ax[1].set_xlabel("Minutes")
32 ax[1].set_ylabel("Fill Level")
33 ax[1].set_title("10min faucet")
34
35 lx,ly = fill_tub(Fraction(1,10)+Fraction(1,15),xmax)
36 ax[2].bar(lx, ly, color='tab:green')
37 ax[2].set_xlabel("Minutes")
38 ax[2].set_ylabel("Fill Level")
39 ax[2].set_title("Both 15min and 10min faucet")
40
41 plt.savefig("bars.png")
```

The sys admin's daily grind: Users and groups

Who, and with Whom?

This time Charly investigates the three most frequently asked questions about user groups.

By Charly Kühnast

Most people know that every user on a Linux system is also a member of at least one user group. Today we want to look into the three most frequently asked questions about groups: Which groups exist, how many members does a group have and who are those members, and to what groups does a specific user belong?

First off, let's find out which groups exist on our system. There are several ways to do this. One of them is to use

the `groups` command without further parameters; another one is provided by `compgen -g`. The `getent group` (Listing 1, line 1) and `cat /etc/group` commands also return the same result, with some additional information, including the group password. There is usually an `x` here, which means that `/etc/gshadow` takes care of that. This is followed by the numeric group ID and a comma-separated list of members.

The next thing is to find out which members belong to a group. In principle,

we have already done this, because the `getent group` and `cat /etc/group` commands provide this information as well.

Often, however, you need the information to process it in a program. It would be good if you didn't have to disassemble the strings using `awk` or `cut`. A list with one username per line would be far easier to handle. To generate such a list, I first install the `libuser` package (line 7). Now I have the `libuser-lid` command at my disposal, but I have to call it with `sudo` (line 8). The numerical user ID also appears in the output of the command. If desired, this can be disabled with the `-n` parameter.

Finally, the whole thing in reverse gear: Now I want to know the groups to which a certain user belongs, for example *charly*. This can be done quickly and easily with the `groups charly` command (line 11). If you need more information, the `id charly` command (line 13) will provide it. This output also shows the numerical IDs.

Here, too, a list would be the object of desire, with one group name per line. I don't know of a native command for this, but `awk` helps reliably (line 15) and returns the desired listing. If the `awk` syntax seems too unwieldy, just use `cut` instead (line 25) for identical results. Many roads lead to Rome here. ■■■

Author

Charly Kühnast manages Unix systems in a data center in the Lower Rhine region of Germany. His responsibilities include ensuring the security and availability of firewalls and the DMZ.



Listing 1: Users and Groups

```
01 $ getent group
02 root:x:0:
03 daemon:x:1:
04 sys:x:3:
05 admin:x:4:syslog,charly
06 [...]
07 $ sudo apt install libuser
08 $ sudo libuser-lid -g adm
09   syslog(uid=104)
10   charly(uid=1000)
11 $ groups charly
12 charly : charly adm cdrom sudo dip plugdev lxd lpadmin sambashare
13 $ id charly
14 uid=1000(charly) gid=1000(charly) groups=1000(charly),4(adm),24(cdrom),27(sudo),
15   30(dip),46(plugdev),110(lxd),115(lpadmin),116(sambashare)
16 $ grep charly /etc/group | awk -F: '{ print $1 }'
17 adm
18 cdrom
19 sudo
20 dip
21 plugdev
22 lxd
23 charly
24 lpadmin
25 sambashare
26 $ grep charly /etc/group | cut -f1 -d:
27 [...]
```



MakerSpace

Home Assistant makes the Raspberry Pi fit for the Z-Wave world

Home Automaton

Z-Wave components, a RaZberry module, and the free Home Assistant software make the Raspberry Pi a powerful smart home control center. *By Gerhard Schauer*

What is causing the smart home hype? Do you really need socket outlets that can be switched on by your cell phone, and if so, what is it all good for at the end of the day? Starting a Raspberry Pi project to find the answer to such questions has always proved to be a workable approach, thus far. Even if my answer to the “Do I need it?” question was, “No,” one thing is for sure, open source and Raspberry Pi can do it.

Of course, a targeted solution needs to meet a number of criteria. It should be based on standards for which sufficient numbers of different components exist, while being open source. All components need to harmonize with the Raspberry Pi, work as self-sufficient systems without any cloud contact, and avoid transmitting unnecessary data on the network. Ideally, I would want to control the whole setup in a web interface that adapts to the small format of a smartphone display.

Every home automation installation includes sensors, actuators, and a control center that evaluates everything and makes link-ups possible. Sensors and actuators not only need to be available over the counter, they should communicate with the control center wirelessly,

but securely and with encryption. Rewiring the house is out of the question. Furthermore, the specifications state that the individual components must be economical and run on one battery for a long time. Last but not least, the overall solution needs to remain future-proof and expandable for many years.

Component Selection

From the large selection of manufacturers for smart home components, the Z-Wave standard quickly crystallized as a basis that meets all criteria. Under the umbrella of the Z-Wave Alliance [1], more than 600 manufacturers and service providers offer thousands of different components that comply with the standard and communicate with each other.

The consortium of various companies enables cost-effective and flexible installation and integrates the Raspberry Pi. With the help of the RaZberry module [2], the Raspberry Pi can access the Z-Wave network. Z-Wave uses radio frequencies between 850 and 950MHz, which ensures a high range and penetration depth and avoids interference.

That still leaves the question of suitable software. The included with the RaZberry



Figure 1: The Fibaro adapter plug not only switches on the connected device but measures power and consumption.



Figure 2: The multisensor registers temperature, humidity, and brightness and reacts to motion.

module, based on the proprietary Z-Way [3] solution, is disqualified in the context of the project because of its proximity to the operator's cloud. Nevertheless, the software is suitable for first steps and experiments and can be put into operation without much knowledge about Linux or Raspberry Pi OS [4].

A nicer solution would integrate on a Raspberry Pi that is already installed and active in the house. In this way, the usefulness of a Raspberry Pi – used as a monitoring solution, for example – can be expanded without having to purchase another small-board computer.

For this specific use case, Home Assistant [5] is a good choice of software. You can run it in a virtual Python environment, and the configuration is controlled by simple configuration files in the YAML markup language [6]. The system supports a massive range of components, including the RaZberry module.

Home Assistant

For a test setup, I integrated a Fibaro smart wall plug (on/off; measures power and consumption) (Figure 1), a multisensor (temperature, brightness, motion, humidity) (Figure 2), and a radiator thermostat into a rudimentary smart home system. The RaZberry module and Home Assistant software take care of control.

The Home Assistant needs to install in a Python virtual environment. Anyone who

has ever dealt with chroot environments under Linux will see the similarities. You set up an environment within the running system that includes all the required libraries and binaries. This setup also facilitates backups of the Home Assistant installation: An image is kept up and running immediately after a simple restore.

In terms of additional packages from the Raspberry Pi package management system, you only need Python and the *libudev-dev* development package to build the Z-Wave modules. For Home Assistant, set up a separate user account, (e.g., *homeassist*). Because communication with the RaZberry module is over a serial interface, you also need to add the *homeassist* user to the *dialout*, *gpio*, and *i2c* groups.

You can do all of this with the command

```
$ sudo useradd -rm homeassistant \
-G dialout,gpio,i2c
```

Listing 1: homeassist@homeassist.service

```
[Unit]
Description=Home Assistant
After=network-online.target

[Service]
Type=simple
User=%i
ExecStart=/home/%i/homeassistant/bin/hass -c "/home/%i/.homeassistant"

[Install]
WantedBy=multi-user.target
```

Next, you should open the `/boot/config.txt` configuration file for the Raspberry Pi system in a text editor and add the `enable_uart=1` line, edit the `/boot/cmdline.txt` file, and remove the `console=serial0,115200 console=tty1` section from the configuration.

To enable the settings, restart the system, then install Home Assistant in a virtual Python environment:

```
$ python3 -m venv homeassistant
$ cd homeassistant
$ source bin/activate
$ python3 -m pip install wheel
$ python3 -m pip install homeassistant
$ hass --open-ui
```

The command from the last line handles the basic configuration of Home Assistant and completes the basic installation.

Note that when you activate components in the Home Assistant configuration file, any modules that are still

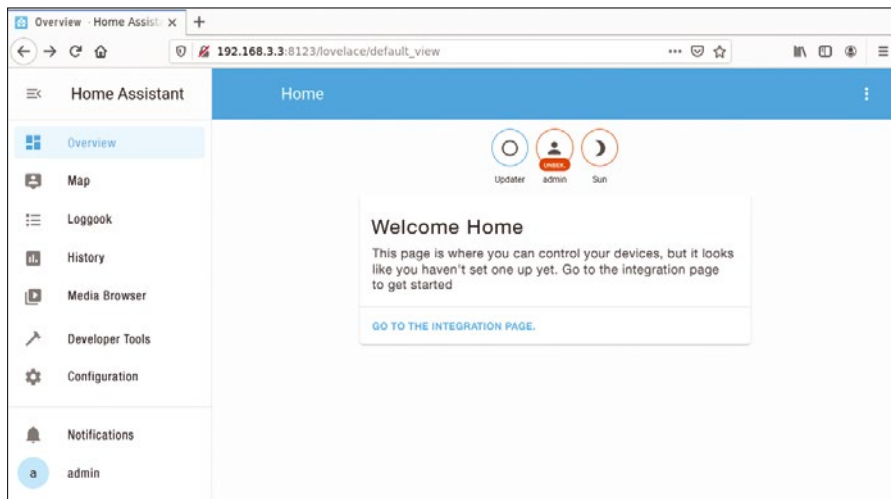


Figure 3: Home Assistant uses a clear-cut web interface.

needed will be reinstalled when you restart, so don't set up the system in a place that doesn't have network access yet – the Raspberry Pi needs an Internet connection until it is fully configured.

So that the system automatically starts the Home Assistant service at boot time, create another service unit for `systemd` by copying the contents of Listing 1 and saving it to `/etc/systemd/system/` in the `homeassist@homeassist.service` file. To enable the service, type:

```
sudo systemctl enable homeassist
```

The service will load automatically at boot time in the future. Finally, start the service manually with

```
sudo systemctl start homeassist
```

just this once.

The first start of Home Assistant takes some time, because the system creates the `configuration.yaml` configuration file and initializes the components. Also, when you add new entries to the configuration that require additional modules later on, Home Assistant needs a somewhat longer pause for thought while it retroactively installs the corresponding modules.

After completing the initialization, the system saves `configuration.yaml` to `.homeassistant/`, the `homeassistant` user's home directory. Its contents looks something like:

```
group: !include groups.yaml
automation: !include automations.yaml
script: !include scripts.yaml
scene: !include scenes.yaml
```

In the same directory, the program also writes a logfile named `home-assistant.log`, which provides information about what is currently happening. If you have a problem somewhere, you will find further details there. On the Raspbian installation for this test setup, for example, I still had to install two packages – `libtiff5` and `libopenjp2-7` – because Home Assistant complained about their absence.

In a web browser, open the administration interface on `http://<RasPi-IP>:8123` (Figure 3). Initially, Home Assistant will prompt you to enter a new username and password, which you will later use to authenticate yourself to the service. Additionally, the page asks for language, location, and time zone.

RazBerry

The communication between the Raspberry Pi and the RazBerry module happens over the GPIO serial interface. For the installation,

you only have to plug the RazBerry onto the GPIO bar of the Raspberry Pi, starting at the upper edge of the board (Figure 4).

For the connection, first create an AES key,

```
$ cat /dev/urandom | \
tr -dc '0-9A-F' | \
fold -w 32 | \
head -n 1 | \
sed -e 's/\(.\)/0x\1, \
/g' -e 's/, $//'
```

and enter it into `configuration.yaml` as shown in Listing 2. The `usb_path` is `/dev/ttyACM0` for USB-connected Z-Wave sticks and `dev/ttyAMA0` for the hardware attached on top (HAT) module.

After a reboot, you can reach the Z-Wave controller from the *Integrations* submenu of the web interface. There, in typical Z-Wave fashion, you need to teach the components their installation instructions (Figure 5). In most cases, you need to switch temporarily the component in question to learning mode by pressing a button. The controller should then find




Figure 4: A first-generation Raspberry Pi is sufficient for Home Assistant. The RazBerry module sits on a section of GPIO pins.

Listing 2: configuration.yaml

```
[...]
zwave:
  usb_path: /dev/tty<AMA0>
  network_key: "<0x41, 0xF3, 0xA6, 0x9A, 0x60, 0x56, 0x43, 0xE1, 0x3D, 0x15, 0xBD, 0xF7, 0x5C, 0x5D, 0x0B, 0x8E>"
[...]
```

the sensor or actuator, parse its data, and assign the controller permanently to its own Z-Wave network. After all, several Z-Wave networks might need to share a frequency. The GUI is very intuitive to use. In case of problems, see the `OZW_Log.txt` file in `/home/home-assistant/.homeassistant/` for details of the actions performed by the Z-Wave module.

Each component registers in the system's Home Assistant universe. In the overview, you can see the states and measured values of all integrated sensors and actuators (Figure 6). The individual components can then be managed and controlled in detail, and the individual nodes can be linked together in the scope of automation setups.

The configuration and automation can be set up manually from the `configuration.yaml` and `automations.yaml` files, respectively. However, the graphical user interface under the *Configuration | Automations* menu offers considerably more convenience. Depending on the intended use and connected components, you can implement anything from a simple to an extremely comprehensive scenario. The test setup is limited to a simple function that switches on the light as soon as a motion detector is tripped (Listing 3).

When you create the automation from the GUI, the corresponding data ends up in the `automations.yaml` file. Listing 3 demonstrates the structure of a simple automation: The trigger is set off when

a certain measured value of a sensor is reached. In the example, I want the motion detector to report an activity value greater than 3 (above: '3'). The assignment is made with the `entity_id` of the sensor component of the multisensor and can also be adapted in the configuration if required.

The trigger now initiates an action in compliance with a condition. In the test setup, I want the Fibaro connector to switch to *On*; the example does not set a condition. For example, it would be conceivable for Home Assistant to switch on the light only after the sun has dropped below the horizon. The condition in this case would query the `entity_id`: `sun.sun` and respond to state: `"below_horizon"` [7]. Incidentally, the entity for the position of



Our ADMIN Online website offers additional news and technical articles you won't see in our print magazines.

Subscribe today to our free ADMIN Update newsletter and receive:

- Helpful updates on our best online features
- Timely discounts and special bonuses available only to newsletter readers
- Deep knowledge of the new IT

Too Swamped to Surf?

OPEN SOURCE 101

A Virtual Event **TUESDAY, MARCH 30**

40 Speakers / 40 Open Source Topics / 1 day

Learn More & Register for **FREE!**

ADMIN

Network & Security

ADMIN Update – Hottest Links

- Building Safe Containers
- Linux Foundation Creates New Code Signing Solution
- The Linux-Next Kernel Finally Gets a Bit "Rusty"
- What's New in PHP 8?
- Monitoring Events with Apache Kafka

Highlights

Building Safe Containers

The basic container images on which you base your work can often be out of date. We show you how to solve this problem and create significantly leaner containers. (more)

Linux Foundation Creates New Code Signing Solution

The Linux Foundation has created a signature to serve as the Let's Encrypt of code signing. (more)

The Linux-Next Kernel Finally Gets a Bit "Rusty"

Rust is finally coming to the Linux kernel, by way of Linux-Next, which has many developers cheering. (more)

What's New in PHP 8?

After about two years of development, version 8 of the popular PHP scripting language has been released. It comes with a number of innovations and ditches some obsolete features. (more)

Monitoring Events with Apache Kafka

Apache Kafka reads and writes events virtually in real time, and you can extend it to take on a wide range of roles in today's world of big data and event streaming. (more)

Most Read Articles

Real-Time Log Inspection with Teler

Teler is an intrusion detection and threat

Further Reading

- Linux Exploit for Spectre Flaw Discovered
- Keycloak Identity and Access Manager
- Yet Another Botnet is Targeting Linux
- Calendar Sync without Google
- Three Major Threats to Linux Discovered

Grow your network and collaborate with others in the Drupal open source community.

REGISTER NOW

Linux Shell Handbook 2021 Edition

bit.ly/HPC-ADMIN-Update

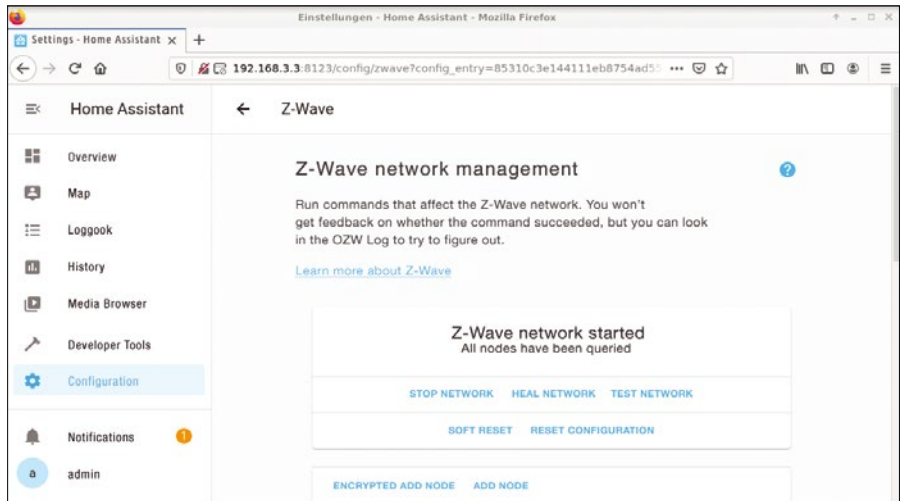


Figure 5: Integrate the Z-Wave devices into the Home Assistant system in learning mode.

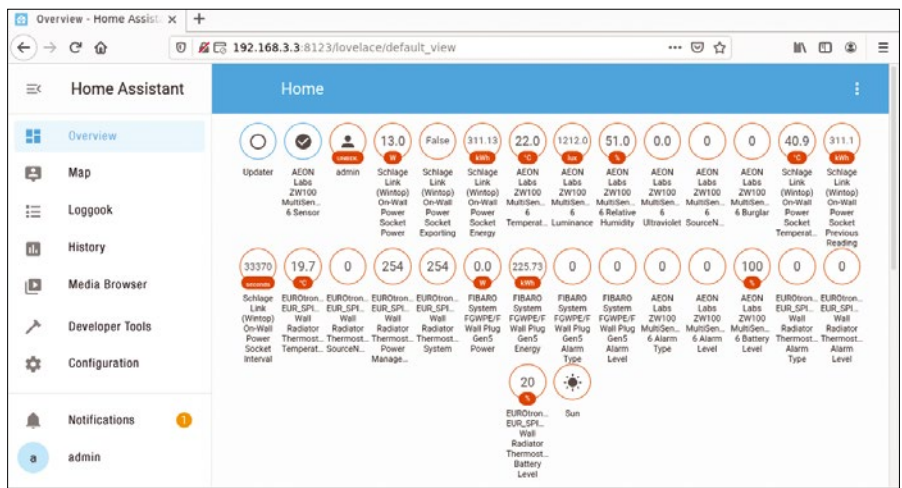


Figure 6: The overview shows the measured values and states of all Z-Wave sensors and actuators.

Listing 3: automations.yaml

```

- id: '1590605895467'
  alias: snap_burg
  description: 'switch light on movement'
  trigger:
  - above: '3'
    entity_id: sensor.aeon_labs_zw100_multisensor_6_burglar
    platform: numeric_state
  condition: []
  action:
  - data: {}
    entity_id: switch.fibaro_system_fgwpe_f_wall_plug_gen5_switch
    service: switch.turn_on
    
```

the sun is one of the automatically added sensors.

The example described here gives you an idea of the wide range of possibilities offered by automation. The multisensor in particular offers numerous options for triggering actions based on light incidence, temperature, or motion. The example still lacks a function to turn off the light again automatically. To do this, you need to add a sensor-controlled or time-controlled shutdown to the routine you have already created or create another automatic routine.

Conclusions and Outlook

The smart home system consisting of Z-Wave components, a Raspberry Pi, the RaZberry module, and Home Assistant implements a newcomer-friendly home automation setup. Only standardized components are used and controlled by open source software that does not require a connection to the cloud.

For advanced users, the solution also opens up the possibility of programming more extensive automations that do not require prefabricated industrial components. This flexibility allows more ambitious hobbyists to further expand the smart home and integrate additional functions at a manageable cost. ■■■

Info

- [1] Z-Wave Alliance: <https://z-wavealliance.org>
- [2] RaZberry: <https://z-wave.me/products/razberry>
- [3] Z-Way: <https://z-wave.me/z-way>
- [4] “Setting up a smart home command center with Z-Wave” by Christoph Langner, *Linux Magazine*, issue 230, January 2020, pg. 32, <https://www.linuxpromagazine.com/Issues/2020/230/Z-Wave>
- [5] Home Assistant: <https://www.home-assistant.io>
- [6] YAML: <https://www.home-assistant.io/docs/configuration/yaml>
- [7] Conditions: <https://www.home-assistant.io/docs/scripts/conditions/>



CLEAR OFF YOUR BOOKSHELF WITH DIGITAL ARCHIVES

Complete your collection of *Linux Magazine* and *ADMIN Network & Security* with our Digital Archive Bundles.

You get a full year of issues in PDF format to access at any time from any device.



Lead Image © enki, 123RF.com

<https://bit.ly/archive-bundle>

2020
Archives
Available
Now!



MakerSpace

Explore gaming on the Lakka console DIY Retro Console

The Lakka Linux distribution comes with everything you need to play retro games and lets you install games directly in the user interface. All you need is a Raspberry Pi and, ideally, a simple gamepad. *By Anzela Minosi*

In the pandemic, its not just Netflix series that are booming, but computer games, too. However, state-of-the-art gaming consoles like Nintendo Switch or PS5 are expensive, which makes Linux distributions like Lakka [1] all the more worthwhile. It even runs on the Raspberry Pi 3, which currently retails for \$35 (EUR35, £34). In other words, you only pay a fraction of what you would have to fork out for a high-end console from Nintendo or Sony for a retro console that you can build yourself.

DIY retro game consoles require very little manual work. You simply plug the Raspberry Pi into a suitable enclosure and connect it to the monitor or TV. With Lakka, you don't even need a keyboard and a mouse; input can be managed with a gamepad like the Logitech F310 [2]. However, you will not want to do without a keyboard for convenience sake. It not only makes it easier to enter text, but you can also assign additional functions to the keys.

Features

The Lakka distribution is based on the LibreELEC [3] media player. The operating system does not need much space on the hard disk, with a pleasingly compact

ISO image of about 400MB. LibreELEC also ensures that all files, including the kernel, are made current during the update, removing the need for you to install individual packages. RetroArch [4] is used as the graphical user interface, which results in Lakka self-configuring thanks to RetroArch's autoconfig feature. Gamepads, for example, are ready for use straightaway.

Lakka supports various computers [5], including the different variants of the Raspberry Pi, from the Pi Zero to the current Raspberry Pi 4. According to Lakka, the Raspberry Pi 3 has good compatibility with Lakka. However, performance-hungry games, such as those for the PlayStation, do not play smoothly.

Installation

A download wizard on the project's website helps you choose the right image [6], which you then unpack for the Raspberry Pi before writing it to a memory card. Working as a Linux user, you then transfer the system to an SD card with the commands

```
$ gunzip Lakka.img.gz
# dd if=Lakka.img of=/dev/<sdcard>
status=progress; sync
```

The correct device identifier of the target drive can be determined with the `lsblk` command. Users of other operating systems turn to external tools such as Etcher [7].

After writing to the microSD card, you will find two partitions. For the monitor and audio settings, you might need to edit the `config.txt` file on the partition labeled `LAKKA`; that said, the display typically self-adjusts to suit the monitor in use. As soon as you boot the Raspberry Pi with the new image, Lakka automatically enlarges the partition with the user data and game ROMs (`LAKKA_DISK`), so you don't have to do this manually.

Configuration

Thanks to Lakka's automatic configuration, changes to the settings are largely unnecessary. However, some settings are worth adjusting. With the arrow keys on the keyboard or the D-pad on the game controller, you move from menu item to menu item. Pressing the Enter key jumps to the submenu or (de)activates settings.

With a gamepad, you input with the *A* switch; the *B* switch goes back one step. When "typing" text (e.g., a username) from the on-screen keyboard, finish typing by pressing Start. Lakka has also published a diagram of a prototype gamepad on its own website [8], which you can use as a template for configuring your own.

For the configuration, first enable the *Settings* | *User Interface* | *Show Advanced*

Settings option. Not everyone likes the Electric Blue color that Lakka uses as its default color scheme (Figure 1). This aspect can be changed in *Settings* | *Menu* | *Menu Color Theme*. With so many settings, it is worth saving the configuration to a file. To do this, switch to *Configuration* | *Save Current Configuration* from the main menu.

Update

At irregular intervals, Lakka comes up with a new release. You can start the update from the graphical user interface. Under the main menu, go to *Online Updater* | *Update Lakka* and select the appropriate image. If the system does not show new versions, although the developers of Lakka have published a new release, then check the links in *Settings* | *Network* | *Online Updates* [9].

The safest approach to a full update is to reinstall. As soon as a new release for the gaming computer appears, download the image and proceed as described in the "Installation" section above. Be sure to back up any existing data, including screenshots, ROMs, and save and play files. If necessary, check for the location of the `LAKKA_DISK` partition and the corresponding directories in *Settings* | *Directories*.

If you prefer to do the update on the Lakka console, you can only access it over SSH. You need to enable the service up front in *Settings* | *Services* | *SSH Enable*. You also need the IP address of

the Lakka game console, which the system outputs in the main menu under *Information* | *Network Information*. Armed with this information, type the lines

```
$ ssh root@<IP address>
$ lakka-update
```

on the console of a Linux machine to start the update. The password for the root user is `root` in the default configuration.

Settings

In network games, it is important that other players identify you. Lakka will use the name you enter in *Settings* | *User* | *Username*. More settings for netplay can be found in the *Settings* | *Network* | *Netplay* menu. Among other things, the IP address of the server and its port can be specified there.

If the Raspberry Pi does not output sound when playing games, make sure onboard sound has been enabled. To do this, remove the memory card from the device and open the `config.txt` file on the `LAKKA` partition on a conventional PC. You need to comment out the `dtoverlay=audio=on` entry with a hash symbol (#). In Lakka itself, in the *Settings* | *Audio* menu, the *Audio Enable* switch must be turned on. You can select a different audio driver in the same submenu under *Audio Device*.

The game controllers can be configured automatically from the Lakka prototype mentioned earlier. To do this, set the switch for *Settings* | *Input* | *Autoconfig Enable* to *On*. If you do not like the assignments of the buttons on the gamepad, you can change them later in *Settings* | *Input* | *Input <user> Binds*.

ROMs

Lakka lets you load games off the web; open the *Online Updater* | *Content Downloader* option in the main menu. When you get there, select an emulator and then the available games. For example, if you choose the *Sheep It Up!* game for the Atari 2600 emulator (Figure 2), Lakka automatically downloads the selected game and saves it to the `LAKKA_DISK` partition in the `roms/downloads/downloads` path. To start the game, go to *Load Content* | *downloads* in the main menu and select the appropriate ROM.



Figure 1: You can navigate Lakka's menus with just a gamepad. With a PC keyboard, use the arrow keys, the Enter key, and the Backspace key.




Figure 2: The Sheep It Up! game can be sourced through Lakka. Don't let the retro style fool you: The game developed for the Game Boy is from 2018.



Figure 3: The character Zelda is constantly under attack in The Adventure of Link.

Alternatively, various websites (e.g., Rom Hustler or Vimm's Lair) offer ROMs of classic game consoles for download. However, when downloading, pay attention to the legal status of the ROMs – only a few ROMs have really been cleared for free distribution. For some, however, the status is unclear because the original developer studio no longer exists (abandonware), and what is on offer is often unlicensed copies.

After downloading, several options are available to load the ROMs in Lakka. Provided you have enabled SSH on the Lakka Raspberry Pi and are using a Linux machine, a file manager is a good choice. Most Linux file managers support opening remote filesystems. On computers with a Gnome desktop, such as Ubuntu, open the address bar in the file manager by pressing Ctrl + L, call the Lakka system with

```
ssh://root@<IP address>
```

and copy the ROM to `/storage/roms/`. Alternatively, transfer the ROM with

```
scp <ROM>.zip root@<IP address>:/storage/roms
```

directly from the terminal.

For Windows users, Samba is an alternative. To begin, turn on the service in *Settings* | *Services* | *Samba Enable*; then, open Windows Explorer and enter the IP address in the address bar in the form `\<IP address>`. It works in a similar way on Linux and macOS systems. The address for accessing the Samba share on the Lakka game console is `smb://<IP address>`. As before in the file

browser, go to the `/storage/roms/` folder and paste the downloaded ROMs. You do not need to unpack the archives beforehand.

Lakka also can load ROMs from any folder of a USB stick. Start by plugging the stick into the gaming console, go to *Load Content* | *Scan Directory* and select the USB stick in the menu. Once you have opened the folder with the ROMs, check the *Scan this Directory* entry. If the ROMs pass Lakka's check, the system inserts the corresponding emulators as menu items. You will then get access to the installed games.

Extras

Games with many levels sometimes seem impossible to complete; your character constantly takes the path to digital nirvana. This problem obviously bothered the makers of Lakka, who therefore implemented some features to keep the game fun. Suppose you have decided to work your way through the levels in *Zelda II – The Adventure of*

Link (Figure 3) but find you are not making much progress. In this case, press F2 to save the current game state. If a mishap occurs in the game, you can press F4 to return to the last saved state without losing a life.

Even more convenience is provided by the Rewind function, enabled in the context menu of the respective game, which removes the need to save the game. You can get there by pressing the RGUI button on the gamepad, usually located between Start and Select and imprinted with the manufacturer's logo. Then go to *Rewind | Rewind Enable* to turn on this feature for the game. If your character meets an untimely end, press R on the keyboard repeatedly until you end up in a safe scene again.

Older games like Formula 1 (Figure 4), Sonic the Hedgehog (Figure 5), or The Legend of Zelda (Figure 6) look quite pixelated, especially on larger monitors; they were originally programmed to be displayed on a TV. This circumstance can be mitigated by switching to *Shader Options | Load*



Figure 4: The Formula 1 arcade game looks pretty pixelated. Test the different shaders to optimize the display on a large monitor.

Shop the Shop → shop.linuxnewmedia.com


Discover the past and invest in a new year of IT solutions at Linux New Media's online store.

Want to subscribe?

Searching for that back issue you really wish you'd picked up at the newsstand?

➤ shop.linuxnewmedia.com

DIGITAL & PRINT SUBSCRIPTIONS



SPECIAL EDITIONS






Figure 5: If you were a teenager in the 1990s, you will know Sonic the Hedgehog very well. Thanks to Lakka, the game can be run on the Raspberry Pi.




Figure 6: The very first version of Zelda, The Legend of Zelda, has quite a low resolution. Here, too, shaders help to improve the rendering.

Shader Preset in the context menu of the corresponding game and choosing a shader. For example, *xsoft/4xsoft* smoothes the pixels, and other shaders make the game look sharper.

Conclusions

Lakka can be operated for the most part from the graphical user interface, so that part, at least, is quite intuitive. It is also ready for use after the first boot, which is not the case for many home-built consoles. Because games can be saved, Lakka is all the more desirable. Additionally, features like shaders keep the game fun longer. Especially with the advanced settings, you can get even more out of Lakka, although it is not mandatory. All told, the distribution is well worth the effort for retro gamers who want to get started right away without tinkering. ■■■

Info

- [1] Lakka: <http://lakka.tv>
- [2] Logitech F310: <https://www.logitech.com/en-us/products/gamepads/f310-gamepad.940-000110.html>
- [3] LibreELEC: <https://libreelec.tv>
- [4] RetroArch: <https://www.retroarch.com>
- [5] Hardware support: <http://www.lakka.tv/doc/Hardware-support/>
- [6] Download Lakka: <http://lakka.tv/get/>
- [7] Etcher: <https://www.balena.io/etcher/>
- [8] Gamepad prototype: <http://lakka.tv/doc/Input-settings/>
- [9] Lakka images: <http://le.builds.lakka.tv/>

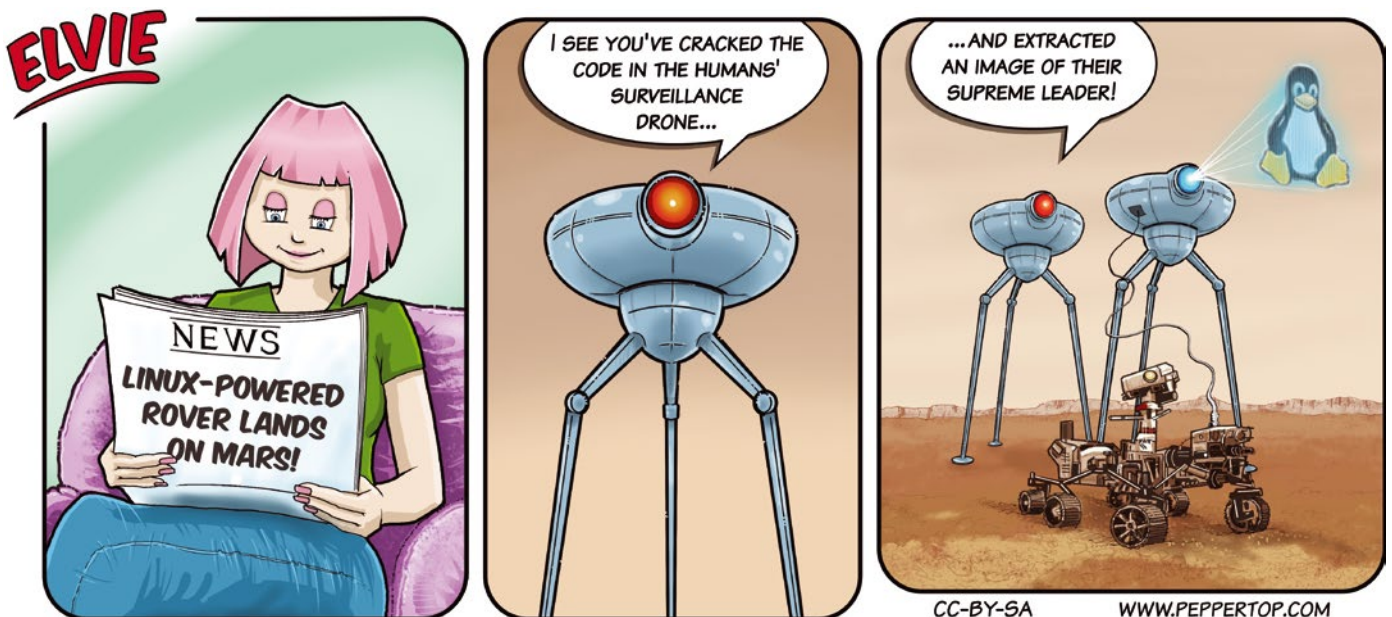
Privacy isn't just a thing you add at the end. In the best case, privacy starts with a set of guiding principles – a philosophy you build into the infrastructure from the ground up. The big social networking companies put all your data on a server, and that server is the site for all the data mining and other monkey business that comes with commercial social media. But, contrary to what you might hear from the giants, social networking doesn't even need a server. Several free tools offer instant messaging and other social services in a peer-to-peer context. This month we look at peer-to-peer messaging with the Tox protocol, and we study some of the leading client apps for configuring and using Tox services. Also in this month's Linux Voice, we show you how to add power and clarity to your written reports with the Zotero research assistant.



Image © Alexandr Moroz, 123RF.com

LINUXVOICE

Doghouse – Distributions	69
<i>Jon "maddog" Hall</i>	
Which distribution is the best choice for you? maddog offers some suggestions.	
Micro	70
<i>Scott Nesbitt</i>	
Micro brings the ease and intuitiveness of a graphical editor to the Linux terminal.	
Tox	74
<i>Christoph Langner</i>	
The Tox protocol uses file-sharing techniques for messaging and audio-video chats.	
FOSSPicks	82
<i>Graham Morrison</i>	
This month Graham looks at Vital, Shortwave, doggo, Bonzomatic, Ultimate Plumber, Open Brush, and more!	
Tutorial – Zotero	88
<i>Anna Simon</i>	
Use this reference management tool to organize citations and create bibliographies.	





Linux Magazine is your guide to the world of Linux. Look inside for advanced technical information you won't find anywhere else!

Expand your Linux skills with:

- In-depth articles on trending topics, including Bitcoin, ransomware, cloud computing, and more!
- How-tos and tutorials on useful tools that will save you time and protect your data
- Troubleshooting and optimization tips
- Insightful news on crucial developments in the world of open source
- Cool projects for Raspberry Pi, Arduino, and other maker-board systems

If you want to go farther and do more with Linux, subscribe today and never miss another issue!

Subscribe now!
shop.linuxnewmedia.com/subs

GET IT NOW!
FAST DELIVERY
WITH OUR PDF
EDITION

MADDOG'S DOGHOUSE



Jon "maddog" Hall is an author, educator, computer scientist, and free software pioneer who has been a passionate advocate for Linux since 1994 when he first met Linux Torvalds and facilitated the port of Linux to a 64-bit system. He serves as president of Linux International®.

Which distribution is the best choice for you? It's not an easy question, but maddog offers suggestions for finding the answer. BY JON "MADDOG" HALL

Choosing the Best Distribution

Almost exactly eight years ago, I wrote an article [1] about why I never answer the question "What distribution do you use?" The reasons are many, but the reason I gave at the time (which is still valid) is that I am a consultant and a writer, and I typically use the distribution that my customers use, not necessarily one of the options that is best tailored for me. Secondly, I recognized that in most cases, the question really being asked was "What distribution should I be using?" Usually I could not answer that either, because I did not know enough about the questioner to give a good answer.

Yesterday in a virtual FLISoL conference, I was asked (again) exactly that same question and gave exactly the same answer. If nothing else, I am consistent.

However I gave very little information in that article about how the questioner should choose their distribution. Also, many things have changed over the past eight years. So in this article I will give some tips for how to choose that first (or 10th or 20th) distribution that you might want to use.

First of all, do you belong to an organization that is already using GNU/Linux on the desktop? Some companies already support GNU/Linux as an operating system through their IT departments. If you are going to start using a distribution, you might as well use the distribution that they support. They can recommend the hardware and configuration you might want to use and perhaps even do the installation for you and set it up. Your laptop or desktop will then fit in well with their environment.

Even if your organization does not officially support GNU/Linux, perhaps some of the administrators use it, and they can give you help with distributions that they use.

If you are attending a university, and particularly if you are studying computer science or computer engineering, you will likely find other students already using GNU/Linux, and they can give you recommendations on distributions and/or help in installing and configuring the system. Modern-day GNU/Linux distributions are typically not that hard to install and configure, but in addition to smoothing the way, finding others who use the same system as you gives you people to discuss features and ways of doing things with, which many users of GNU/Linux like to do.

If you have looked around and found you are a "lone wolf," your next stop might be several of the sites on the web that cater to GNU/Linux beginners. Websites like *linux.com*, *cloud-academy.com*, *itsfoss.com*, or even (cough) *www.linux-magazine.com*, will have articles giving advice on both selecting a distribution and getting started with that distribution. Even a simple

web search of "Linux beginner" turns up many articles about what to select for your first distribution.

But what if you have already been using a "first distribution," perhaps one of the many great "community" distributions, and you want to go deeper?

For that you can go to *distrowatch.org*. This is a website that lists hundreds of "open source" (mostly GNU/Linux and BSD, but occasionally other) systems, tells their features, and provides release dates, popularity ratings, and reviews. Often there are screenshots of the default desktops, but of course most distributions allow you to tailor this, so the defaults often do not mean much.

Of course you will want to have a distribution that matches your architecture, and you need to decide if you want one aimed toward being a desktop, or server, or high performance computer system.

Many of the distributions are what I would call a "spin" from another distribution. The distribution creators started with another distribution and gradually (or not so gradually) morphed it into something that may have a completely different "look and feel" and installation.

As an example, Ubuntu started out using an underlying Debian distribution. They used the Debian package manager and changed the installation, the graphical desktop, and other factors to create their own distribution. However, people using Ubuntu had access to the underlying Debian repositories to bring a wealth of functionality to what was, at that time, a fledgling distribution.

Likewise, later, the Linux Mint distribution was based off Ubuntu, and they utilized many features of the Ubuntu distribution to build their own look and feel and functionality.

In choosing your distributions, you may wish to stick with distributions in a particular "family," using a particular package manager rather than switching between families. I am sure you will find large amounts of functionality in each family of systems.

In choosing any distribution, please look carefully at items like "last release date" and "activity," since many distributions are no longer maintained. DistroWatch does have policies for "dormant" and "discontinued" distributions, but checking dates and downloads is also useful.

Finding the distribution that meets your needs is part of the adventure of open source. ■■■

Info

- [1] "Community Notebook: Distribution of Choice" by Jon "maddog" Hall, *Linux Magazine*, issue 152, July 2013, pg. 94, <https://www.linux-magazine.com/Issues/2013/152/Doghouse-Distro-Decisions>

Editing text at the command line with micro Small and Simple

Editing text at the command line doesn't have to be daunting. Micro brings the ease and intuitiveness of a graphical editor to the Linux terminal.


BY SCOTT NESBITT

Linux users are spoiled for choice when it comes to many types of software. That's especially true for text editors. When I first started writing this article, I tried to count as many editors as I could just off the top of my head. I stopped when I got to around 32.

Quite a few of those editors run at the command line. While I've used several, I've never had great relationships with command-line text editors. I'm not sure why that is. Maybe it's just we don't have enough in common with each other to form a strong bond. Recently, though, I was working with a command-line application that played better with terminal text editors than with graphical ones. So, I duly set my default editor to the venerable GNU nano editor, which was the only terminal editor installed on my computer. While I've used nano in the past, I was quickly reminded that it isn't for me.

Instead of using software I don't particularly like, I searched around for an alternative and came across micro [1]. Billed as "a modern and intuitive terminal-based text editor," micro turned out to be great substitute for nano. In this article, I look at how any Linux user, regardless of their level of comfort with the command line, can start using micro and configure it to meet their needs.

Figure 1: A new and empty file, started in micro.



Setting Up

There are several ways you can install micro. The easiest of those is to install it using your Linux distribution's package manager. If micro isn't available in your package manager, run this command in a terminal window to download a prebuilt binary:

```
curl https://getmic.ro | bash
```

From there, move the file (conveniently named `micro`) to a folder in your path.

Another option is to install micro as a snap package by running this command in the terminal:

```
snap install micro --classic
```

If compiling software is your thing, you can download micro's source code from its repository on GitHub [2]. The editor is written in the Go programming language. You'll need Go v1.11 installed on your computer, and you must enable Go modules as well.

Start a New File

So micro is on your computer. Now what? Get going by opening a terminal window and typing `micro` or `micro <name of file to edit>`. Figure 1 shows micro editing a new, empty file.

From there, start typing – nothing complex or out of the ordinary for any text editor, whether on the desktop or at the command line.

Micro uses many (if not all) of the same keyboard shortcuts that you're familiar with from using desktop applications – although the notation in the help section may differ slightly. For example, if you're used to seeing `Ctrl+X`, note that `Ctrl-x` is just another way to write the same key combination.

Copy text by highlighting it and pressing `Ctrl+C`. Paste that text by pressing `Ctrl+V`. Or press `Ctrl+O` to open a file and `Ctrl+S` to save it.

Press `Ctrl+E` and then type `help defaultkeys` to display a list of micro's keyboard shortcuts.

You can also use your mouse to select and highlight text. There's no right-click support (as

```

mand line application that seemed to play
. So, I duly set my default editor to the
l editor installed on my computer. While
isn't for me.
15
16 Instead of using software I don't particu
tab tabmove tabswitch
> tab

```

Figure 2: Tab completion can help you find a command.

you'd find in a desktop application), but pressing the scroll wheel on your mouse pastes whatever is in your clipboard into your file.

Using the Command Bar

I've already mentioned the Ctrl+E shortcut. When you press that combination of keys, the editor's command bar displays at the bottom of the screen. In the command bar, you can type commands like `replaceAll` to find and replace all instances of a word or phrase, `set softwrap true` to turn on word wrapping, or `help commands` to get a list of all of micro's commands.

If you only remember part of the name of a command, type it and press the Tab key to complete it. If there are two or more possible completions, micro displays them and you can press the Tab key again to choose one (Figure 2).

For example, when you press Ctrl+O to open a file, you type the name of the file and, if it's in another directory, the path to that file on the command bar. Press the Tab key to complete the names of the directory and of the file.

Using Tabs

A nifty feature of micro is the ability to have multiple files open at the same time. If you have one file open in micro, you can open another one in a tab by pressing Ctrl+E and then typing `tab` followed by the path to that file.

The idea is more or less like the tabs you find in a graphical editor.

The tabs that you have open are listed along the top of the editor, as shown in Figure 3.

You can switch between tabs by pressing Ctrl+E and then typing `tabswitch` followed by the number of the tab. The numbering of tabs starts from the left: The tab furthest to the left is 1, and the numbers increment as you move right. Or, you can click on a tab with your mouse. I prefer doing the latter.

```

micro Snesbitt-micro.txt
Snesbitt-micro.txt  micro-query.txt  [103-on-carving_out_time.md]
## Think Small
To do the things that you want to do, you don't need to make radical changes to your life. Take small steps. By that, I mean pacing yourself in everything that you do or attempt. Someone once told me that overnight success takes a long time. You won't reach your destination in a day or a week or a month. Nothing happens all at once. No matter what you're doing &mdash; building a career, creating a habit, learning a language, picking up a new skill &mdash; you need to take *small, steady steps* to reach the point that you want to reach.
Tweak your day. Don't hack it. Don't analyze your day to within a second of its life, don't deconstruct it down to the minute. Instead, get an overview of your typical day. Focus on portions of that day out of which you can carve any amount of time. Chances are, those portions will be the same most days.
Then, try to fit what you want to do into those blocks of time each day. You might not be able to complete certain tasks during a block, but you can probably make solid headway. Let's say you want to read Thomas Piketty's *Capital in the Twenty-First Century*. It's a long book &mdash; the hardcover edition weighs in at almost 700 pages. By devoting 30 minutes a day to reading 103-on-carving_out_time.md (17,1) | ft:markdown | unix | utf-8  Alt-g: bindings, Ctrl-g: help

```

Figure 3: Tabs in micro are similar to what you would find in a graphical editor and allow you to have multiple files open at the same time.

Configuration Options

One of micro's selling points is that you can configure it to make the editor work the way that you want it to. The quickest way to do that is to open the file `settings.json` found in the `.config/micro` folder in your `/home` directory.

That file stores configuration information in a format called JavaScript Object Notation (JSON). If you're not familiar with it, JSON is a data format that both people and computers can easily read. Don't let that intimidate you. JSON uses what's called "attribute/value" pairs to store information. In micro's `settings.json` file, the attribute is a configuration option and the value is often either `true` or `false`. Sometimes, it's a number or something else specific to an option.

Micro has 60 configuration options. Those include options to automatically save files, to automatically match braces when writing in a markup or programming language, to display line numbers or a scrollbar, or to use tabs instead of spaces.

Figure 4 shows what my (rather basic) settings file looks like.

Figure 4: An example of a settings file.

```

micro
"basename": true,
"colorscheme": "material-tc",
"ft:markdown": {
  "aspell.check": "on"
},
"ruler": false,
"scrollbar": true,
"softwrap": true
}
settings.json (1,1) | ft:json | unix | utf-8  Alt-g: bindings, Ctrl-g: help

```

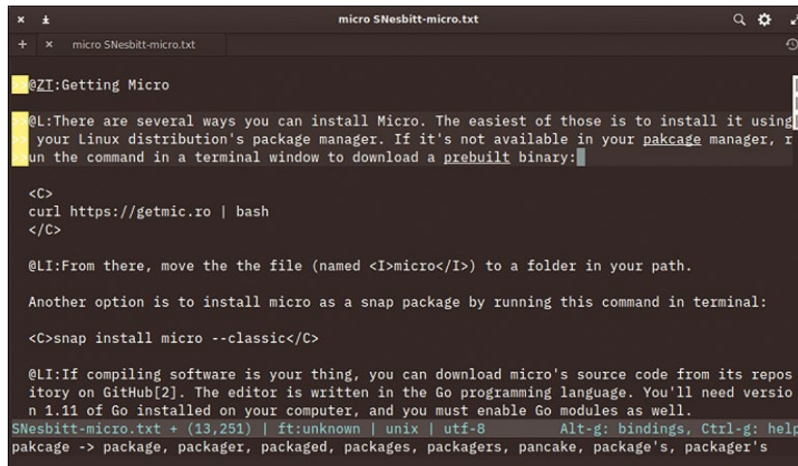


Figure 5: Micro's spellchecking plugin in action.

You can find out more about the options by opening micro's command bar and typing `help options`.

Adding Plugins

Micro lives up to its billing as being extensible thanks to its support for plugins. At the time of writing, there are 27 plugins [3] that you can add to the editor. Those plugins run the gamut of tasks like correcting commonly misspelled words, adding bookmarks to a file, pretty-printing code, and adding support for the Go programming language (which is the language used to develop micro).

Install a plugin by running this command in a terminal window:

```
micro -plugin install <name of plugin>
```

Let's say you want to install the `aspell` plugin, which checks spelling as you type. To do that, run this command:

```
micro -plugin install aspell
```

Figure 5 shows the `aspell` plugin in action.

Some plugins, like `aspell`, run automatically. Others you need to launch yourself. To do that, open micro's command bar, type the name of the plugin to run, and press Enter.

Admittedly, 27 isn't all that many plugins. There's a chance you might not find one that suits your needs. However, if you have a few chops with the Lua programming language, you can write your own plugins. There's information about that in micro's code repository on GitHub [4].

Changing Your Keybindings

Maybe you're moving to micro from another text editor – one that uses a different set of keystrokes to access commands. You can change micro's keybindings by editing a file named `bindings.json` in the folder `conf/ig/micro/` in your `/home` directory.

As with micro's configuration options, you use a set of attribute/value pairs to change the editor's keybindings. In this case, the attribute is the key-stroke combination that you want to change, and the value is the name of the command for that keystroke.

Let's say you're used to pressing `Ctrl+Y` to paste text (remember that out of

the box, micro does that when you press `Ctrl+V`). To change the command, add this to the file `bindings.json`:

```
"Ctrl-y": "Paste"
```

Final Thoughts

Editing text at the Linux command line doesn't need to be intimidating. It doesn't need to involve using arcane keystroke combinations to access the editor's commands. Micro brings the ease of using a desktop text editor to the command line.

Micro not only packs a solid set of features, it wraps those features in a small package. It's one of the fastest and easiest-to-use command-line text editors that I've used over the years.

Micro lives up to its billing as an intuitive terminal-based text editor. It's definitely worth having around for those times when you need to edit text files at the command line. ■■■


Info

- [1] micro: <https://micro-editor.github.io>
- [2] GitHub repository: <https://github.com/zyedidia/micro>
- [3] micro plugins: <https://micro-editor.github.io/plugins.html>
- [4] Developing plugins for micro: <https://github.com/zyedidia/micro/blob/master/runtime/help/plugins.md>

The Author

Scott Nesbitt is a freelance journalist, technical writer, and essayist. He's been using FLOSS since the mid-1990s and writing about it since the early 2000s. Despite having few technical chops, Scott often finds himself working at the command line and loving it.

20 YEARS LINUX MAGAZINE



LINUX MAGAZINE 20 YEAR ARCHIVE DVD

ORDER NOW!
<https://bit.ly/Archive-DVD>



Tox: Text, voice, and video chats without a central server

Peer Group

The Tox protocol uses file-sharing techniques for messaging and audio-video chats, which gives users a greater degree of privacy and freedom.

BY CHRISTOPH LANGNER

The popular messaging application WhatsApp has recently alienated many users with its new terms of service and privacy policy. It's a situation that once again illustrates that when you put yourself in the hands of a proprietary service, you put yourself at the mercy of the operator. If you are dissatisfied with a change, either you grudgingly continue using the service, or you move to a different application with similar features.

In the case of WhatsApp, there are quite a few alternatives, and many were suddenly overrun with WhatsApp users looking for another option. This article looks at the different categories of messaging services, and then delves into

one particular option – using peer-to-peer messaging apps based on the Tox protocol.

Open Source Clients

Signal [1], Threema [2], and Telegram [3] are widely regarded as good alternatives to WhatsApp, partly on the grounds that they operate more freely and openly than the Facebook service. This typically means releasing the app source code under open source licenses and disclosing the underlying protocol so that other app developers can develop alternative clients (see Table 1).

However, discussions of social media alternatives often leave out one very important component: the server. All providers keep a watchful eye on the

server and therefore on the data. Telegram and Threema do not disclose the source code for the server. Signal does provide the source code for the server under a free license, so users could theoretically build their own Signal servers. However, the software is designed in such a way that the different instances cannot talk to each other. If the same principle applied to email, then users of one service (such as Yahoo Mail) would not be able to send messages to those of another service (e.g., Gmail).

Federated Systems

True freedom exists only if the service discloses the

Table 1: Instant Messengers Compared

Service	Server	Client(s)	Protocol	E2EE*
Centralized				
WhatsApp	Proprietary	Proprietary	Proprietary	Yes
Signal	Open source	Open source	Signal	Yes
Threema	Proprietary	Open source	App Remote Protocol	Yes
Telegram	Proprietary	Open source	MTPProto 2.0	Yes
Federated				
Dino	Open source	Open source	XMPP	Yes
Quicksy	Open source	Open source	XMPP	Yes
Element	Open source	Open source	Matrix	Yes
Peer-to-Peer				
Briar	None	Open source	Bramble	Yes
Jami	None	Open source	OpenDHT	Yes
Tox	None	Open source	Tox	Yes

*E2EE = end-to-end encryption

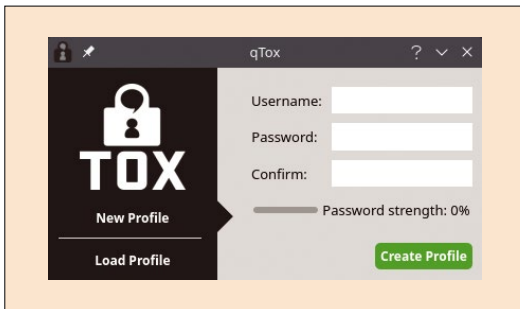


Figure 1: Most Tox clients allow alternating different Tox IDs (quasi user accounts).

source code of the server and client(s) and the underlying protocols. In addition, the network systems must operate in a federated manner, that is, in a distributed manner (federation). However, this kind of freedom goes against the business interests of commercial providers. Therefore, none of the major providers offer this level of freedom. What remains are services and programs such as Dino [4] and Quicksy [5], which are based on the classic XMPP (formerly Jabber) [6], and applications such as Element, which is based on the Matrix [7] open communications protocol.

Peer-to-Peer Networks

Instant messengers can do without any kind of a central authority. They work in a similar way to file-sharing programs like eDonkey or BitTorrent. Messages or voice and video chats travel directly from computer to computer (peer-to-peer) over the network. The contacts are organized with the help of a distributed hash table (DHT). End-to-end encryption ensures that privacy is maintained and that only the sender and the recipient can read the messages.

In this category, the selection of programs is thinned out even further. The last tools standing are Briar [8], Jami [9], and Tox [10]. Briar is available exclusively for Android-based smartphones and tablets. I have already covered Jami in detail in a previous issue [11], so now it is time for a closer look at Tox.

Tox in Practice: qTox

Compared to the other two peer-to-peer messengers, the main advantage of Tox is that there are a number of different client programs, including at least two specifically for smartphones. These various options will be discussed later in the article, but to evaluate the practicality of Tox, we chose to test the default client qTox [12], which completely covers the functional range of Tox.

The current version, qTox 1.17.3, has not yet made it into the package sources of all common distributions. If you're using another version, make sure it's at least qTox 1.17.0, which saw many new features added and bugs fixed. You

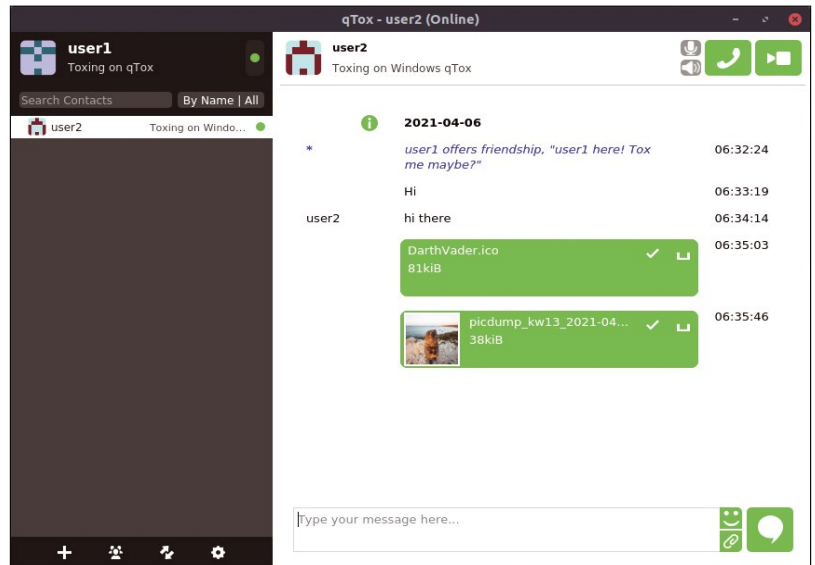


Figure 2: Incoming images, as shown in the example here, must first be accepted by the addressee.

can use the package manager for the installation; the package name is usually *qtox*.

At first startup, create a new Tox profile with the *New Profile* tab. You are free to choose your username and password (Figure 1). A password is not mandatory; however, it ensures that qTox encrypts all your data such as your contact list and chat histories for storage on the hard disk. You can change your username later if necessary, but the actual identification in the Tox network is done by a 76-digit ID. After clicking *Create Profile* the program loads the actual application window (Figure 2).

When tested on Manjaro Linux, qTox had English as the default setting. If you prefer a different language, you can access the settings via the

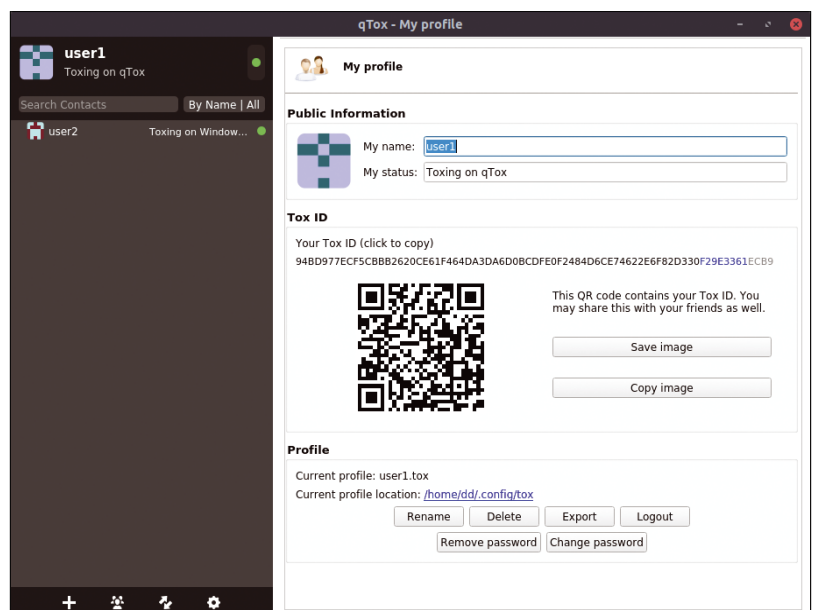


Figure 3: On the Tox network, you are identified by a cryptic Tox ID. For contact requests, you need to know the Tox ID of your future Tox friends.

gear icon below the contact list. To open your profile, tap on the generic user picture or the username. You can customize your profile image (by clicking on it), the username, and your status (Figure 3) here.

Below you will find the Tox ID associated with the profile. To communicate with a partner, you need to send this 76-character mess over a secure channel. For users of one of the smartphone apps, there is an option to save the ID in the form of a QR code. The image can then be scanned directly with a phone using a barcode scanner app, eliminating the need to type in the ID.

Setting Up Contacts

To create a new contact yourself, tap the plus icon below the contact list. Then enter the ID in the *Add a friend* tab, and if necessary add the message to the friend you want to invite. Optionally, in the *Import contacts* tab, you can import a whole bunch of friends in one go. For this qTox reads in a text file in which you list the contacts in the form of their Tox IDs line by line. Last but not least, in the *Friend Requests* tab, you will find all the contact requests that have not been answered yet (see the box, “Spambots,” for dealing with suspect friend requests).

After sending a friend request, the user receives a notification. If the person has accepted your invitation, the contact automatically appears in the contact list with that person’s username and picture. If you click on the entry, a message area typical for instant messengers opens. At the bottom of the text field, enter your message; next to it, you have the option to add emojis or send files.

Tox transmits text messages directly after clicking the speech bubble or pressing the Enter key. The recipient first has to accept the transmission of files such as images. A timestamp shows you when the message was sent. A rotating circle indicates that qTox still has to transmit the message. Tox does not have a comprehensive status display like WhatsApp (message not yet sent, transferred, or read).

Audio and Video Calls

Above the message history, next to the contact’s profile picture, you will find the contact name. If you have trouble remembering which friend goes by which username, click on the entry, and you can give your contact a name you will recognize. Next to it are the switches for starting an audio or video call. You can set the devices used during calls and their properties (for example, the sound quality of the microphone or the resolution of the webcam) in the settings in the *Audio/Video* tab.

In our test with a second PC running qTox as well as a smartphone using a Tox client called Antox, phone calls worked very well and without much delay (Figure 4). Both via a broadband connection and the mobile network (4G), the sound quality was similar to a conventional phone call via landline or a WhatsApp call.

For video chats, qTox can’t quite keep pace with proprietary solutions. The video image is pixelated for a short while at first. After the connection is stabilized, the image fully composes itself. The picture also lost its cohesion for a short time during abrupt changes in the picture, such as fast pans, in our tests. However, these are minor issues; video calls basically work.

Additionally, qTox supports screen sharing, although that will require some setting changes, (see the “Setting Up Screen Sharing” box).

Group Chats

Modern messengers need to support conversations with multiple contacts in addition to communication

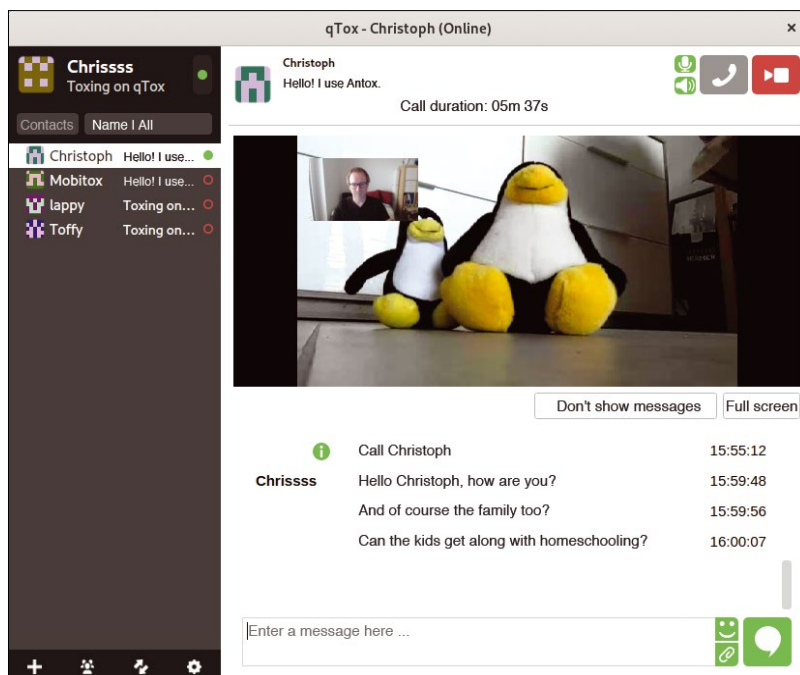


Figure 4: In addition to classic, text-based chats, many Tox clients also support audio and video calls.

Spambots

The open and non-monitored Tox network offers many advantages. However, there are also downsides, such as spambots. Since no one controls Tox, there is no one to throw spammers off the network. If you are inundated with suspicious friend requests, qTox offers the option to switch to a NoSpam-ID in the *Privacy* tab settings. This will allow you to continue communicating with your existing friends; however, your previous Tox ID will become invalid and spam requests will go nowhere. In the same area, you also have the option to blacklist individual IDs.

between two people. qTox offers the possibility to set up group chats for this purpose. The initiator of the chat group must have all the desired members as friends in their contact list. Now, tap the group chat icon below the contact list and then select the option to create a new group. The default name of the chat is *Group name #1*; use the context menu to rename the group appropriately.

In order for your friends to join the group, you have to send them an invitation. To do this, right-click the desired contact and select *Invite to Group* from the context menu. You can either use one of the previously created groups or create a new one directly. Repeat this step with all the contacts that you want to join the group chat. qTox will now notify your friends about the invitation. Once they accept the request, they can chime in. There is no chat log: If one of the contacts switches off their computer and thus the Tox client, they miss all the group messages sent in between.

Besides the classic text-based group chat, qTox also supports telephone conferences. To do this, tap the green phone icon above the chat history. At first glance, nothing changes, but as soon as the contacts in the group chat have enabled the telephony function, you can hear each other. Only people who enable this on their phones can hear the conversation. An option for video calls in a group is not available for qTox so far. The button for the video function is grayed out in the current version.

Advanced Functions

Tox's operating principle currently requires that both contacts on a call be online at the same time. The online status is visualized by a filled green or empty red dot to the right of the username in the contact list. Offline messaging is not yet supported by Tox. One of the reasons for this is that there is simply no central server that temporarily stores messages that cannot be delivered at the

moment and then delivers them as soon as the addressee gets back online.

However, Tox users do not have to do without offline messages entirely. As a workaround, most Tox clients (including qTox) support pseudo-offline messages [13]. Write a message to a contact currently displayed as offline, and then the client keeps the message in memory. As soon as the contact gets back online, the program sends the reset message. Until then, qTox shows a rotating circle next to the text, which signals that delivery is still pending. However, that requires you to stay online as well.

The integration of real offline messages is on the developers' to-do list, as well as the possibility to use a Tox profile on multiple devices at the same time – for example, with qTox on the PC and in parallel with Antox on the smartphone. Currently, a profile can be exported and transferred to another computer, but clients with a copied profile must not be online at the same time. If you try it, Tox gets confused: Contacts come online and go offline again permanently, or messages and data transfers are lost.

Other Tox Clients

The Tox project maintains a list of compatible clients [14] on its website. The open source protocol underlying Tox allows independent developers to program additional Tox clients, and users have the opportunity to select a client that optimally fits their requirements. However, it is important to be careful when testing: Not all of the alternative clients have the same functionality as qTox. Depending on the stage of development, they may only support a feature subset (see Table 2).

Here are a few other things to know about these clients. Both qTox and uTox are reference implementations of the Tox protocol [15], and both applications are available for Linux, FreeBSD, macOS, and Windows. uTox has fewer dependencies than

Setting Up Screen Sharing

One common feature in modern videoconferencing is the ability to route the full screen content or selected application windows to the participants in the call. qTox also supports this, but very unfortunately hides the function in the depths of the settings.

You enable screen sharing like an audio or video call, but you have to change the camera beforehand. To do this, open the settings via the gear icon below the contact list and then switch to the *Audio/Video* tab. There, in the video settings, select the *screen* as the camera. If necessary, additionally specify a region

or – for a multi-monitor setup – a single screen in *Camera resolution*.

Like many other screenshot and screencasting programs, qTox does not yet work with the Wayland display server used in GNOME. When launched in a standard GNOME session, the preview remains blank. You need to select *GNOME on Xorg* when logging into the desktop in the GDM display manager (use the gear menu bottom right), for desktop sharing to work. On KDE Plasma or Xfce, users do not have these problems.

Function	qTox	uTox	Venom	Toxic	aTox	Antox
Contact management	+	+	+	+	+	+
Chats	+	+	+	+	+	+
Group chats	+	+	+	+	-	-
Group chats (audio)	+	+	-	-	-	-
Data transmission	+	+	+	+	+	+
Screen sharing	+	+	-	-	-	-
Audio chats	+	+	-	+	-	-
Video chats	+	+	-	+	-	-

qTox as it does without the Qt framework and is more frugal with resources.

The Venom [16] client is optimized for Gnome. Venom is currently available as version 0.5.5 and integrates cleanly into the Gnome desktop as a GTK3 application, but voice and video chat support is not planned until release 0.6. Overall, the program cuts a good figure despite its young age. It clearly organizes the contact list and chats and already implements group chats (Figure 5). In the default configuration, Venom does not save anything except for the contact list; even the chats disappear into a black hole after the program is closed. To prevent this, go to the settings and then in the *General* tab, move the slider next to the *Save chat history* option to the right.

Toxic [17] is a text-based console client that can be found in the package sources of all major distributions. After installing the package of the same name, call Toxic from a terminal window via `toxic`.

Toxic requires a little more explanation, because it's not immediately obvious how to access the different functions. An initial overview is provided by the help feature, which you call by typing `/help`. For

example, you can invite contacts via `/add <Tox-ID> "Message"`. You can discover your own ID using the `/myid` command. `Ctrl+O` and `Ctrl+P` let you switch between the status window, the contact list, and the individual chats, which are organized in tabs in Toxic.

Since Toxic does not necessarily require a graphical interface, Toxic can run on a server that is accessible via SSH, for example. You will always be online there and continuously reachable, so that you will never miss any group message. On a computer with a graphical interface, you do not have to do without audio and video chats despite working with a command-line program. `/lsdev in` and `/svdev in`, respectively, tell Toxic to list the input devices for image and sound that exist on the system. You can then use `/sdev in <ID>` and `/svdev in <ID>` to select the microphone and webcam.

You can then start a video call (with a picture) from an active chat with `/ca11` or `/vca11`. Toxic plays the corresponding videos in windows that pop up (Figure 6). In the default configuration, the program initially displays only the image of the other person; you need to explicitly enable your own camera with the `/video` command. Toxic uses the minimum resolution of the webcam. With a command like `/res <800> <600>` you can adjust the quality to utilize the available bandwidth and the technical limits of the webcam.

Antox [18] and aTox are clients for Android smartphones. Antox copes well with phone calls and video calls, but does not yet support group chats (Figure 7). If you try to invite a contact to a group via qTox, which is connected to the Tox network via Antox, this fails. The corresponding menu item is grayed out in the dialog of the program. This smartphone app has the advantage of being virtually always online. The capability for pseudo-offline messages is therefore of little consequence here. However, Antox still has a few issues to work out. Sending and receiving messages worked reliably in our tests, but opening a sent picture always resulted in a crash.

The alternative aTox [19] proves to be more stable: Chats and sending pictures and data worked without complications with the Android app. However, it

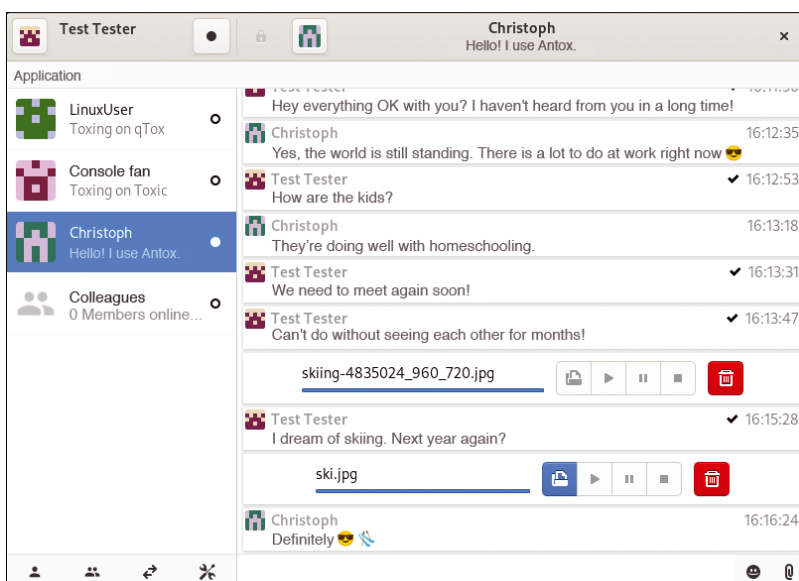


Figure 5: Venom, which is still in development, offers many features, but support for audio and video calls is still missing.

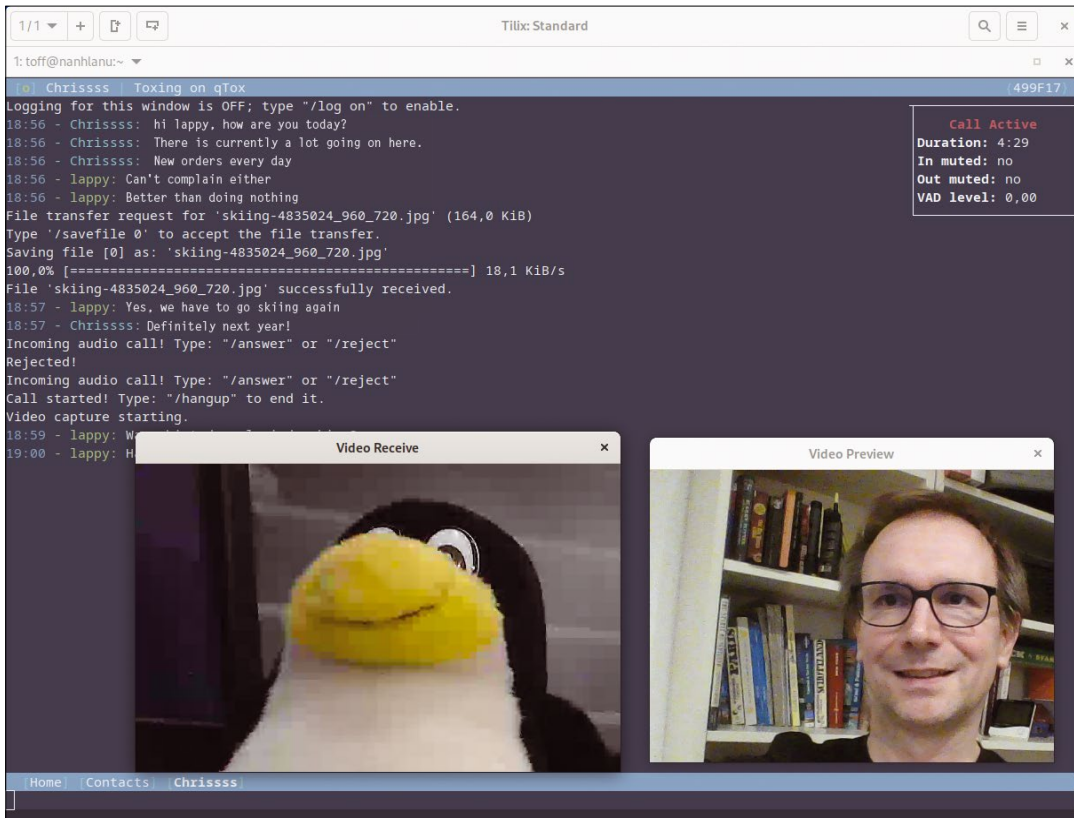
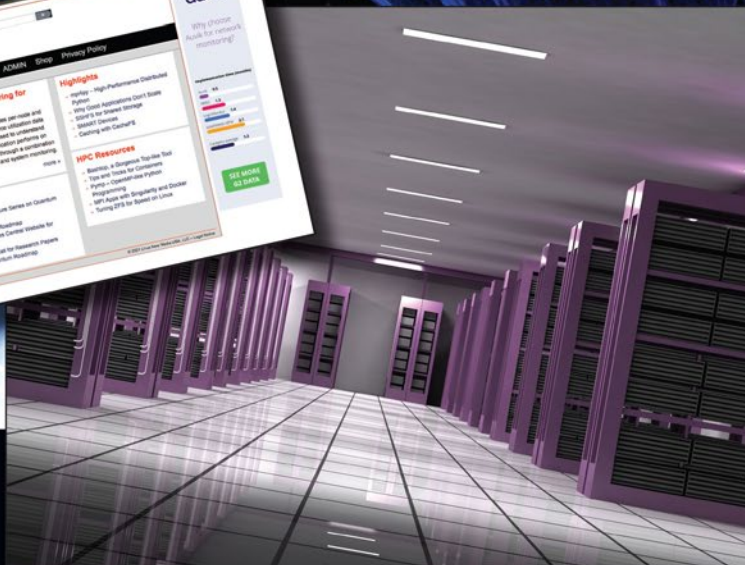


Figure 6: Toxic, intended for use in the terminal, supports pretty much everything. Launched on a system with a graphical desktop, the client even supports video calls.

A Webzine for High-Performance Computing Specialists



ADMIN
Network & Security

If you work with high-performance clusters, or if you're ready to expand your skill set with how-to articles, news, and technical reports on HPC technology.

admin-magazine.com/HPC

cannot handle audio and video chats, nor can it handle groups. aTox does not even tell you that it does not support these functions. In qTox, for example, the invitation function (as opposed to an Antox contact) remains enabled, but the invitation simply fails. Moreover, the program is not yet available from any app store, neither on Google Play nor on the open-source F-Droid marketplace. Installation therefore relies on the APK file provided on the project's GitHub page.

Conclusions

Many users have extremely high expectations of a potential WhatsApp alternative. It has to work as well and be as easy to use as the original, but has to do without companies and servers, as well as offer maximum control and give users all freedoms.

Reconciling all these conditions, however, is a difficult undertaking. It is the lack of a centrally organized user directory and switchboard that makes contact list maintenance more complicated for both developers and users. There are also restrictions when establishing a connection and sending data – as a user, you will notice this when establishing a video call or when you need to use pseudo-offline messages, for example.

Despite all this, Tox works and offers almost all the functions that WhatsApp and similar

applications have in their repertoire. Only a few features are missing, such as viral message forwarding. However, if you consciously choose a P2P messenger like Tox, you will probably not have this on your list of requirements anyway.

One thing has to be said though: Tox does not explicitly guarantee the user's privacy. You are anonymous on the Tox network, and there is no mandatory registration. The messages are transmitted peer-to-peer and directly and encryption ensures that no one can sniff them. Nevertheless, the recipient sees your IP address, and the network operator could also track with whom (more precisely: with which IP) you are communicating on the Tox network. In the FAQ, the Tox developers are open about this issue [20]. They even provide instructions on how to route Tox through the Onion router Tor, if necessary [21]. ■■■

Info

- [1] Signal: <https://signal.org>
- [2] Threema: <https://threema.ch>
- [3] Telegram: <https://telegram.org>
- [4] Dino: <https://dino.im>
- [5] Quicksy: <https://quicksy.im>
- [6] XMPP: <https://xmpp.org>
- [7] Matrix: <https://matrix.org>
- [8] Briar: <https://briarproject.org>
- [9] Jami: <https://jami.net>
- [10] Tox: <https://tox.chat>
- [11] "Private Conversations" by Christoph Langer, *Linux Magazine*, issue 233, April 2020, pg. 73-75, <https://www.linux-magazine.com/Issues/2020/233/Chat-Freely-with-Jami>
- [12] qTox: <https://github.com/qTox/qTox>
- [13] Pseudo offline messaging: https://wiki.tox.chat/users/offline_messaging
- [14] Tox client list: <https://wiki.tox.chat/clients>
- [15] Tox protocol specification: <https://toktok.ltd/spec.html>
- [16] Venom: <https://github.com/naxuroqa/Venom>
- [17] Toxic: <https://github.com/Jfreegman/toxic>
- [28] Antox: <https://github.com/Antox/Antox>
- [19] aTox: <https://github.com/evilcorp/ld/aTox>
- [20] Tox and IP addresses: <https://tox.chat/faq.html#tox-leak-ip>
- [21] Tox via Tor: https://wiki.tox.chat/users/tox_over_tor_tot

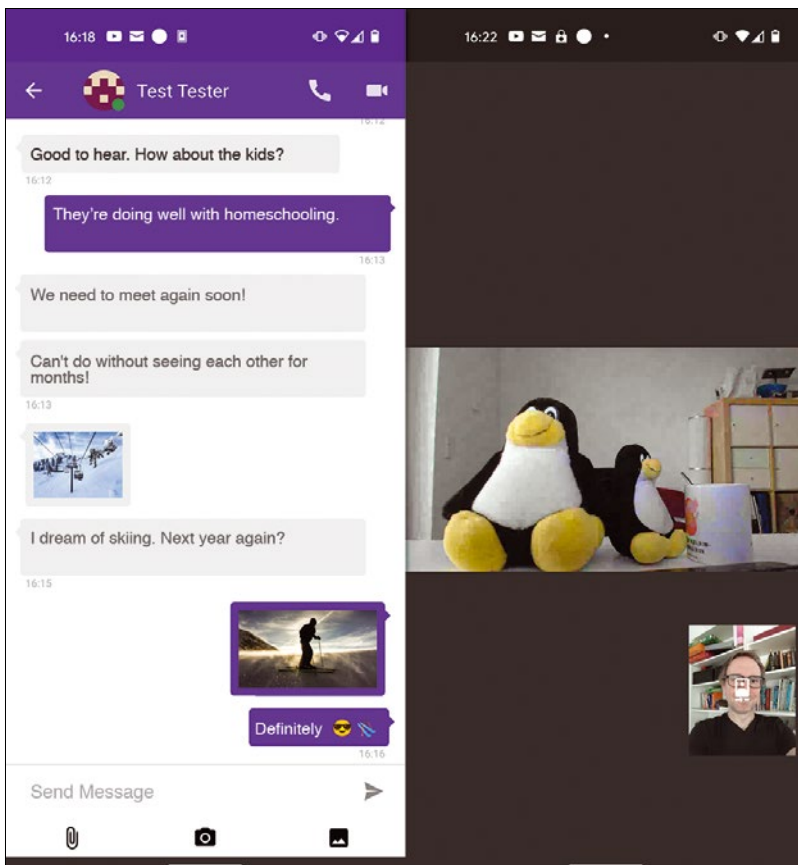



Figure 7: The Antox Android client copes well with audio and video calls but does not yet have group chats.

Shop the Shop
shop.linuxnewmedia.com

Become a LibreOffice Expert



Explore the **FREE** office suite used by busy professionals around the world!

Create Professional:

- Text Documents
- Spreadsheets
- Presentations
- Databases

Whether you work on a Windows PC, a Mac, or a Linux system, you have all you need to get started with LibreOffice today. This single-volume special edition will serve as your guide!



Order online:
shop.linuxnewmedia.com/specials

For Windows, macOS, and Linux users!

FOSSPicks

Sparkling gems and new releases from the world of Free and Open Source Software



After watching Ubuntu help NASA with its first controlled flight on another planet, Graham spent far too much time this month visiting Mars in Elite Dangerous, via Proton on Linux. **BY GRAHAM MORRISON**

Wavetable synthesizer

Vital

This is an incredible open source software synthesizer, created by Matt Tytel. He's the developer behind Helm, the former best open source synthesizer you could run on Linux, and Vital improves on Helm in every single way. It now competes directly with the hugely popular and influential (and costly proprietary) Serum synth, which has been used on countless tracks to provide everything from phased string stabs to the repeating motifs of

Stockhausen's musique concrète. That Vital has been released under GPLv3 is a testament to Matt's commitment to open source and his faith in committed users purchasing a subscription to fund development. We sincerely hope it works, because this is an unprecedented step and paying for Vital is absolutely worth it.

Vital does everything a traditional synth does and more. What makes Vital unique, though, is its sound engine. And it sounds amazing. In any synth, the origin of all sound is one or more oscillators, either voltage controlled (VCO) in analog synths or digitally controlled (DCO) in digital and software types. The vast majority

of oscillators generate waveforms with either a triangle, sawtooth, or square (aka pulse) shape. But Vital is a wavetable synthesizer, and that means the shape of any of its three oscillators comes from a 2D or 3D matrix of data. A 2D wavetable could have a triangle, sawtooth, or square shape plotted on the X and Y axes, but it could equally be a sample, a drawing, or something entirely algorithmic. A 3D wavetable adds to this a Z axis of data, which can be selected manually or moved through via a modulation source.

The previous paragraph might sound complex, but the most important thing about Vital is that it always seems to produce great results. The user interface does a brilliant job of showing what's happening with its animated and informative style. Wavetables are shown in 2D or 3D, for instance, with the selected Z always highlighted and animated if it changes. The three envelopes and four LFOs show their current positions and the two filters even as they change shape. Best of all, you can drag and drop between control sources and destinations, allowing you to experiment and play, such as changing the filter depth from an LFO. It's even possible to create bipolar modulation sources and negative output to inverse values, which is perfect for changing an envelope's attach value according to velocity, for instance.

Vital is also one of the few synths to support MPE, allowing you to control the sound from a physical device in unprecedented ways. It adds per-note pitch, slide, and pressure control, all of which can be sent to whatever destination you want via the drag and drop modulation routing. With no extra hardware, you can still play notes from your computer keyboard. The sound engine is also well integrated into the Linux desktop, with support for ALSA, PulseAudio, and JACK in the standalone version and for LV2 and VST3 in the plugin variation. The free download includes dozens of sounds and wavetables, and you can easily create or import your own from any source. But whatever sound you play, it always sounds incredible.

Project Website
<https://vital.audio/>



1. Open source: Everything but the presets and wavetables of the commercial version are included. **2. Spectral wavetables:** Generate simple and complex audio from spectral wavetable data. **3. Effects:** Add reverb, chorus, delay, compressor, distortion, eq, filter, flanger, and phaser effects. **4. Oscilloscopes:** View the output waveforms of the sounds you're generating. **5. Envelopes and LFOs:** Control anything with these modulation sources. **6. MIDI input:** See incoming signals and assign them to modulation destinations. **7. CPU management:** Manually adjust the output quality and number of voices to control CPU usage. **8. MPE:** Full per-note pitch bend and modulation control. **9. Rich filters:** Use formant and complex comb filters alongside traditional analog models.

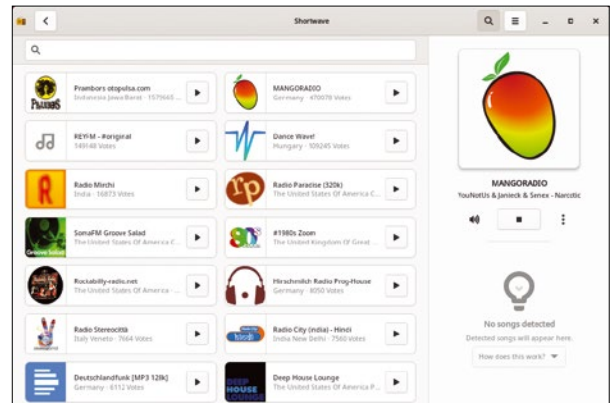
Internet radio

Shortwave

Shortwave is a perfectly named Gnome application. As its FAQ explains, the name comes from shortwave radio signals that, because of their long reflection properties and transmission distances, many of us will know from radio receivers before the digital age. Pre-Internet, a shortwave radio was a gateway to a world of international radio stations, letting teenagers anywhere listen to music from different continents or even from the middle of the ocean when it came to pirate radio in the UK. This Gnome application lets you do the same, albeit with the convenience of digital reception, 25,500 preset stations, and a complete lack of static noise from solar flares. The trade-off of access to such a huge curated list of stations is that

they're grabbed via API from *radio-browser.info* and stored as local metadata. As a result, they're uneditable, and you can no longer contribute to the upstream database after misuse of a previous feature. Hopefully this will change in the future, but ultimately it doesn't matter. Shortwave is still brilliant.

Shortwave is beautifully designed, with a user interface that looks a lot like a modern music playing or streaming service. A panel on the right shows the thumbnail for the station currently playing along with its playlist. On the left, there's an image carousel and a selection of the most popular stations, all of which can be played instantly by pressing their associated play buttons. Most impressive, though, is that this interface is built using the shiny new GTK4, as the



You can't edit the station list in Shortwave, but with more than 25,000 stations you have plenty of choice.

application has been successfully ported from GTK3+. This helps it smoothly scale like an adaptive web page from portrait to landscape and between a smartphone form factor and a regular Linux desktop. There's also a bitmap panel that mimics the Formica laminate and buttons of a classic radio when you don't want the full window. All of this helps Shortwave fit into whatever amount of space you can afford while providing seamless access to music from around the entire globe, which is exactly what you want from a music player.

Project Website

<https://gitlab.gnome.org/World/Shortwave>

DNS Lookup

doggo

Dig is one of those perennial Linux commands that many of us have used at some point to help fix a network issue. It's a relatively simple command that can query a DNS server in the same way a web browser does, asking the server for the IP address associated with a specific domain name. It's usually the first thing to try when your browser can't load a few unassociated web pages but you can still ping IP addresses. Typing `dig` followed by the domain name, for example, will query your configured DNS and return whether one or more answers were received and what addresses they return, alongside their response times. More than one address could indicate a load balancing system or addresses

that differ according to your geographical location and network performance. Or they could indicate a misconfigured local DNS.

Doggo attempts to modernize dig's functionality in both its breadth and its output. It can query across IPv4 and IPv6 networks and request specific record types such as MX for mail exchanges. It supports DNS over TLS and DNS over HTTPS and can format its output as JSON. The latter is great for automated testing, integrating with dashboards, and for more predictable processing with the `jq` command. But the best thing about doggo is the output. Dig's output can be confusing whereas doggo adds color and better organization while cramming in more details. Without

```

$ ./doggo amazon.co.uk @https://1.1.1.1/dns-query --time --debug
DEBUG[2021-04-20T01:11:06+01:00] Initiating DoH resolver
DEBUG[2021-04-20T01:11:06+01:00] Starting doggo
DEBUG[2021-04-20T01:11:06+01:00] Attempting to resolve
  nameserver="https://1.1.1.1/dns-query" ndots=0
NAME      TYPE      CLASS    TTL     ADDRESS      NAMESERVER      TIME TAKEN
amazon.co.uk.  A        IN       52s     54.239.33.58  https://1.1.1.1/dns-query  135ms
amazon.co.uk.  A        IN       52s     178.238.7.229 https://1.1.1.1/dns-query  135ms
amazon.co.uk.  A        IN       52s     54.239.34.171 https://1.1.1.1/dns-query  135ms
$ ./doggo linuxvoice.com
NAME      TYPE      CLASS    TTL     ADDRESS      NAMESERVER
linuxvoice.com. A        IN       86400s  213.138.101.172 127.0.0.53:53
$ ./doggo bbc.co.uk CNAME
NAME      TYPE      CLASS    TTL     ADDRESS      NAMESERVER
./doggo bbc.co.uk
NAME      TYPE      CLASS    TTL     ADDRESS      NAMESERVER
bbc.co.uk.  A        IN       108s    151.101.0.81   127.0.0.53:53
bbc.co.uk.  A        IN       108s    151.101.192.81 127.0.0.53:53
bbc.co.uk.  A        IN       108s    151.101.128.81 127.0.0.53:53
bbc.co.uk.  A        IN       108s    151.101.64.81  127.0.0.53:53

```

Doggo is easier on the eyes and can replace dig for general DNS testing.

further arguments, it will list the requested name alongside each address and name server location. You can easily use another name server with the `@` symbol, add round trip timing information, and send encrypted queries. Thanks to being written in Go, it feels fast and modern, especially when used with `jq` to parse the JSON output in your own scripts, and the output is much clearer and easier to understand.

Project Website

<https://github.com/mr-karan/doggo>

Shader IDE

Bonzomatic

Pixel shaders are amazing. They're perhaps the closest modern technology we have to the custom graphics processors and coprocessors found in many of those 1980s home computers. Shaders came to life as self-contained functions that helped 3D surfaces look better, such as the simplistic Phong or Gouraud shading you might be familiar with from the original Blender. Pixel shaders are also distinct from the ray tracing or light casting that goes into building a scene, because, as their name describes, they operate on the final pixel, dynamically changing their color according to their position within the environment. This is why most are considered post-processing effects that can be applied to any scene or game

after the main artwork has been added.

Rather than shaders being calculated by your CPU, they have become discrete units in a GPU where they are programmed directly from OpenGL using a special shader interpreter language called GLSL. But because they're self-contained, they're relatively easy to hack on – some applications and games will even import raw GLSL from a configuration file. But what you really need is a system that lets you write GLSL while seeing the results instantaneously, and that's what Bonzomatic does. Bonzomatic is a simple command-line text editor that runs in full screen mode and gives you almost immediate feedback on your GLSL



After pressing F5 to recompile a shader (instantly), its effects are rendered across the background of the editing environment.

tinkering. It's ideal for people who know nothing about shaders and want to experiment, as well as for developers who want to tweak the best possible output and performance from every byte of code. Bonzomatic lets you quickly experiment and compile your shaders, with the effects shown in the background and across the editing environment unless you disable them with the F11 key. It's even been used as the platform for coding challenges, and it's a genuinely fun way to play with code you may have found online or when trying to learn how shaders work.

Project Website

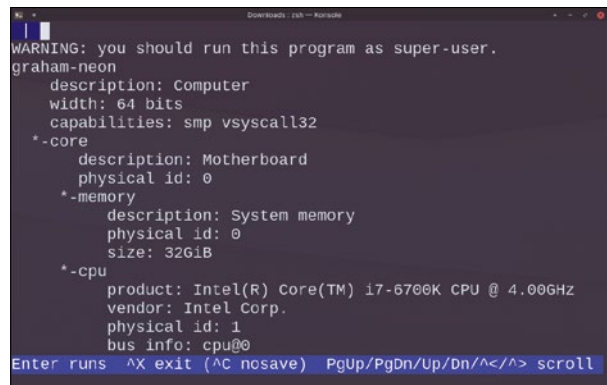
<https://github.com/Gargaj/Bonzomatic>

Pipe command helper

Ultimate Plumber

We can all agree that the Linux command line is wonderful. It's often easier to use than the desktop, because it lets you do all kinds of things without the mental burden of coordinating pointer movement with multiple clicks across different panes and windows. Instead, you can construct simple commands to do simple jobs, one at a time. Creating folders, moving files, searching for things, and editing configuration files can all be done easily. But when you want to do more than one thing at once, it can then sometimes become easier to use the desktop. One reason for this is that the command-line tools that can help, such as pipes with the "|" symbol, are a little difficult to use, and

you can't always be confident about what will happen when you press Enter. And of course, the command line has no undo feature so the jeopardy is real. This is where the Ultimate Plumber (aka up) can help. It's a command-line tool that makes pipes dramatically easier to use. It does this by making them interactive in the way you construct a command and in the way you view the output. You start by piping your initial output into the up command. From there you can use Page Up and Page Down keys to browse through the output. But if you look carefully, you'll see there's also an input box at the top of the output that will take further commands. You can enter any Bash commands into this and press Enter



Use command-line pipes interactively with the Ultimate Plumber tool – also known as up.

again to run them. The output will immediately update to reflect the output from your new command, but unlike on the command line, you can continue to edit the command and add further pipes while continuously updating the output. This lets you confidently construct complex commands, which can then be saved when you quit the up environment with Ctrl+X.

Project Website

<https://github.com/akavel/up>

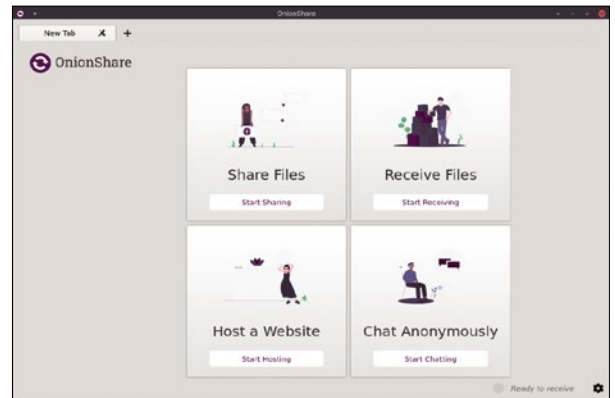
Secure communications

OnionShare

We often recommend the tool Wormhole for informal file sharing between machines. It uses end-to-end encryption and can be easily installed by both parties. It's as easy to use as sharing a temporary passphrase through a secure channel. But it's also very basic in the features it provides, and it's relatively unproven. We recently looked at a far more ambitious solution, Retrosahre, which adds all kinds of secure communication channels, such as chat and forums, alongside its file-sharing functionality. It also supports Tor for maintaining online anonymity. But it is probably over-engineered as a replacement for Wormhole, which is where OnionShare might help. OnionShare is an application that helps you spin up one or more

secure communication channels over Tor and promises to keep your data 100 percent safe and private. Unlike many applications that make similar claims, though, OnionShare has the best possible kind of provenance. It was created by Micah Lee, the (then) staff technologist at Glenn Greenwald's investigative journalism portal, The Intercept, in response to Greenwald's widely reported difficulty creating his own setup to communicate securely with Edward Snowden.

With just a couple of clicks, OnionShare can set up local portals to help with sending a file, receiving a file, chatting anonymously, and even setting up a website from a client's fully anonymized Tor-compatible browser. Clicking on *Share Files*, for example, allows you to drag and drop anything into



For many, including journalists working in difficult parts of the world, secure and foolproof communication can make the difference between life and death.

the main window and instantiate an aggregated share, which generates an anonymized Tor address that can be accessed from any Tor-compatible browser. *Host a Website* does the same, only with HTML files that the client browser will interpret as a website, and both the *Chat Anonymously* and *Receive File* functions operate in the same way. It's an easy and foolproof way to access these secure facilities and highly recommended if you ever need something like this.

Project Website

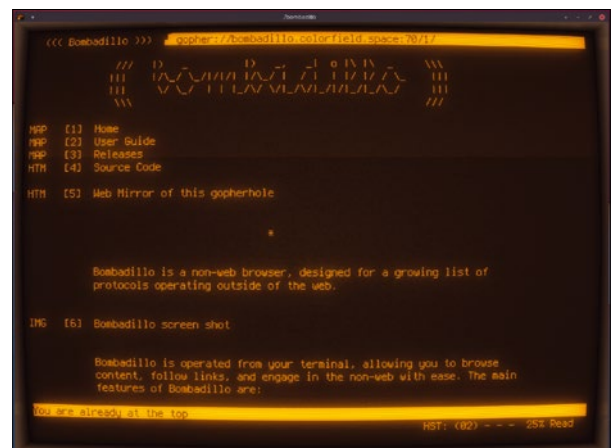
<https://onionshare.org/>

Gopher client

Bombadillo

The Internet has grown to become a huge distraction. The content generation and dynamic linking algorithms on domains like Facebook, reddit.com, YouTube, and Google are purposefully designed to keep you engaged and looking at the screen, tracking your Internet use while surreptitiously inserting discrete advertising. But it wasn't always like this. The early World Wide Web was designed to be rich in accessible and portable content, augmented at the time by a host of other protocols that provided alternative ways of retrieving information. One of these protocols was Gopher, a prototype web that was contemporary with Tim Berners-Lee's HTTP, but easier to deploy and access. It provided a simple and navigable

way to present a collection of documents and even jump between them. It was a perfect way of accessing text on low bandwidth connections with underpowered hardware from a global network. It was important enough that web browsers even supported the Gopher protocol and could display their pages inside the main window. Remarkably, there is still Gopher content available, and while modern browsers can no longer access it, there does exist a modern Gopher client, which is what the Go-coded Bombadillo is. Bombadillo runs in a terminal, takes a `gopher://` protocol address as an argument, and uses many keyboard controls borrowed from Vim. You can press `:` to open the command input, for instance (or



There are still some great sites available via Gopher, such as Project Gutenberg, unencumbered by web bloat, images, and advertising.

press space), and use *jk/h* for movement, with *b* and *f* for going backwards and forwards through your page history. This being Gopher, links actually appear as numbers, much like footnotes in a book. You press the number to follow the link, which takes you to another page of text. And that's what remains refreshing about Gopher. It's an Internet without distraction, where you can focus purely on the details that interest you.

Project Website

<https://bombadillo.colorfield.space/>

VR drawing

Open Brush

This is a wonderful piece of software that has a couple of specific hardware requirements. You will need Valve's SteamVR with a supported virtual reality headset, and perhaps most challenging at the moment, a reasonable graphics card. But Open Brush is easily worth the effort. It's best described as an immersive sketching and drawing tool where you, the artist, draw in the three-dimensional space around you – just as you might by drawing an imaginary line in front of, above, and around you. For many years, it was a commercial Windows-only product called "Tilt Brush," purchased from the original engineers by Google. Its development corresponded with Google's big push into VR, first with Cardboard and then with Daydream, between 2016 and the latter's cancellation at the end of 2020. But Tilt Brush remained an independent product and was initially released with the HTC Vive before going through years of incremental updates. However, there was never

a Linux version, and it was never open source.

But unlike many of Google's other abandoned projects, Tilt Brush wasn't destined to become yet another entry on *killed-by-google.com*. Instead, Google released the entire project under the Apache 2.0 license and gave it to the community. The only restriction was the Tilt Brush trademark, which meant new builds of the project needed a new name. And this is what Open Brush is – a community built progression of Tilt Brush, completely open source and more importantly, now available for Linux, albeit in a now experimental and unstable form. Open Brush does for VR what Deluxe Paint did for the Amiga. It's the perfect ambassador for the underlying product, because it's easy to use yet shows the potential of the platform. With two controllers in the 3D space in front of you and a palette of colors, brushes, and effects attached to your left hand, you simply start drawing. You can use your controller to change the brush size and color, and there are many



There are many hidden features in Open Brush, including a charades game, 3D model import and export, and audio animations.

brush types, including animated brushes, types that respond to sound input, and tools with various particle effects.

You typically start with Open Brush by mimicking the 2D pictures you'd draw on a piece of paper, and this is apparently what the original version of the application was built to do. But you soon learn that real inventiveness comes when you pull backwards or forwards into the space around a single plane. Try drawing a 3D cube, for example, or a three-dimensional tree, and before long you'll be drawing a house around you, or a camp, or a garden, or a spaceship. The whole experience is very therapeutic and accessible. But it's also capable of some very serious results. There are examples where 3D designers have used the Open Brush environment for real world architecture, visualization, and model design. You could do all of this with the original Tilt Brush, but Open Brush is also adding experimental features of its own, including multiple mirrors, new brush types, and of course, this Linux version. It's fantastic to see such a professional, important, and broad application like this get both an open source release from Google and such strong community support.

Project Website

<https://github.com/icsa-gallery/open-brush>



If you can imagine drawing with your fingers in the space around you, you can use Open Brush in virtual reality.

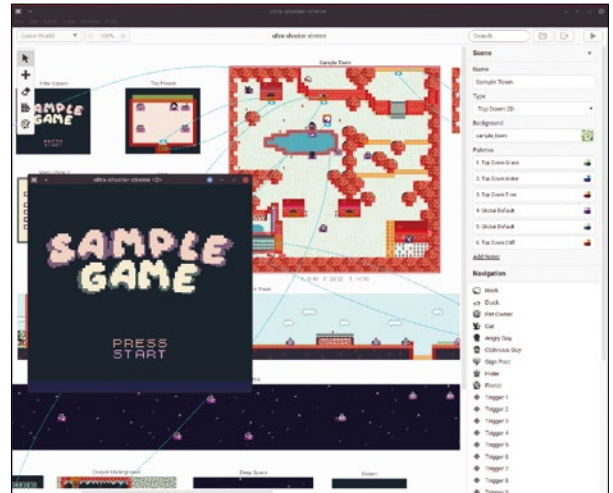
Retro game IDE

GB Studio

There are quite a few retro-orientated games development frameworks, such as the proprietary Pico-8. These typically exist to artificially impose some of the restrictions of old hardware, such as a low resolution and color palette, to help foster creativity. They've produced some amazing results. GB Studio is similar but differs in one important way: The restrictions it imposes aren't artificial but required, because the resultant binary from your game design is intended to run on a bona fide 1980s-era Nintendo Game Boy 8-bit console. The main application even includes a play button to preview your project on an integrated Game Boy emulator so you can be sure it will work. This may seem like a crazy idea, but

there's a booming cottage industry in building games for old hardware, both for the pleasure of creation and because of these enforced limitations.

GB Studio is a far cry from what Game Boy development must have been like in the late 1980s. It's a beautifully engineered Electron application that aims to help the game creator at every step, from providing templates for common game types and an active community to making it possible to almost drag and drop a complete game. You start by creating assets externally or using the brilliant built-in template. You then add these to a new scene, which is equivalent to a screen in the game, alongside actors and event triggers. Events can be one of many types, such as game text,



GB Studio includes three layers of background parallax, a tracker for music, and the ability to sync events to musical queues.

scene modifiers, variables, and various control flow elements such as joystick control, and it's here that you build your game logic. There is a lot to learn, but the process is always visual and very intuitive, and you don't need to be a programmer, making it ideal for anyone who wants to try their hand at retro game development for real hardware.

Project Website

<https://www.gbstudio.dev/>


Voxel game engine

Minetest

At first glance, Minetest looks like a clone of the now venerable Minecraft, and to some extent it is. It features voxel graphics, there are both creative and survival game modes, and there's a huge modding community of people that spend their time designing and building unique environments within the game. All of this will be familiar to Minecraft players. But Minetest offers much more than this. It's open source (unlike Minecraft, which despite Notch's early promises is not), and this means that you can invest your time in the code and the game engine sure in the knowledge that it won't become redundant when development drops or moves to a different platform. The Minetest community has

fully embraced this, creating many of their own games, mods, and textures, many of which can be downloaded for free from the main website.

Minetest improves on Minecraft in other important ways too. There's a Lua-based API that can be used to augment your creations in only the way a programming language can. Maps can be huge, incorporating 62,000 cubic voxels, and there's a gallery of user-created games and maps that can often be freely built upon or incorporated into your own worlds. While it may not have the depth of crafting mechanics of Minecraft, Minetest conveys a greater sense of jeopardy because the worlds are larger and feel more unknown, especially in survival



Abandon Java (and Microsoft!) by switching to Minetest and its open source gaming community.

mode where you have to explore and craft your way through the night and from one virtual day to the next. It's also a great way to play with new game ideas separate from the mining, crafting, and trading of the traditional mechanics. The environment is perfect for building scripted adventures, or an RPG, or even simple collections games, all of which can be found in the database of games created by users.

Project Website

<https://www.minetest.net/>

Literature management with Zotero

Perfect References

Reference management software like Zotero helps you easily record and organize citations and create bibliographies.

BY ANNA SIMON

Programs for managing technical literature, often known as reference managers, are important tools for students and scientists. For larger projects such as a master's thesis, you would find it hard to do without this kind of software.

Launched back in 2006, Zotero [1] is a powerful open source literature management program that harmonizes perfectly with LibreOffice and Microsoft Word. It is especially suitable for fields that do not usually use LaTeX for writing (i.e., the humanities, economics, health, and some natural sciences). A special version of Zotero, called Jurism [2], is adapted specifically for the field of law.

Zotero not only runs on Linux but also on macOS and Windows. It also supports mobile operating systems. Several (unofficial) apps for Android and iOS let you import and edit title data, including scanning ISBNs.

Quoting with Style

Zotero was originally created as an extension for the Firefox web browser. However since 2017, the software has been available as a standalone program. It can be used with browsers such as Firefox, Chrome, and Opera, as well as with office suites like LibreOffice and Microsoft Office. You do not need an account to download and use Zotero, but you will need one if you want to access cloud storage. Creating an account is especially worthwhile if you want to access your literature collection from several devices or work on bibliographies in teams [3].

The free basic account provides 300MB of storage in the Zotero Storage cloud. While this may not sound like much, it is sufficient even

for extensive projects such as dissertations, as long as you only store bibliographic info. If you want to back up sources to the cloud in the form of PDFs, images, or other larger files, the free cloud storage is probably not enough. Upgrading to 2GB costs \$20 a year, while you pay \$60 for 6GB, and unlimited capacity is available for \$120 [4].

Setting up Zotero

While you are unlikely to find Zotero in the official package sources, you should be able to install it without problems on most distributions. The easiest place to set it up is on Arch Linux and its offshoots like Manjaro, where you install it from the Arch User Repository (AUR).

Users can download Zotero directly from the project's homepage in the form of a tarball. To install, open a terminal window, change to the directory containing the downloaded file, and run the commands from Listing 1. The commands unzip the program to `/opt/zotero/` and then link it to `~/local/share/applications/` so that you can access Zotero from the application menu.

When first launched, the program wants to set up connections for the browser and LibreOffice. It's a good idea to install at least the Zotero Connector [5] browser extension right away. The program automatically opens the relevant web page in the system's default browser.

To use data synchronization via Zotero Storage, enter the info for your Zotero account. Go to the Preferences window, and choose the Sync tab. You'll need to enter your username or email address and password to set up syncing. If you don't have an account with Zotero yet, you can find the links for creating one here. You can also choose whether or not to upload file attachments to the Zotero cloud.

Importing Data

Using Zotero is intuitive. The first step is to search for literature on the Internet. Switch to a web browser, but leave the Zotero window (Figure 1)

Listing 1: Installation

```
$ sudo tar -xjf Zotero*.tar.bz2 -C /opt
$ sudo mv /opt/Zotero_linux-x86_64 /opt/zotero
$ sudo /opt/zotero/set_launcher_icon
$ ln -s /opt/zotero/zotero.desktop ~/.local/share/applications
```

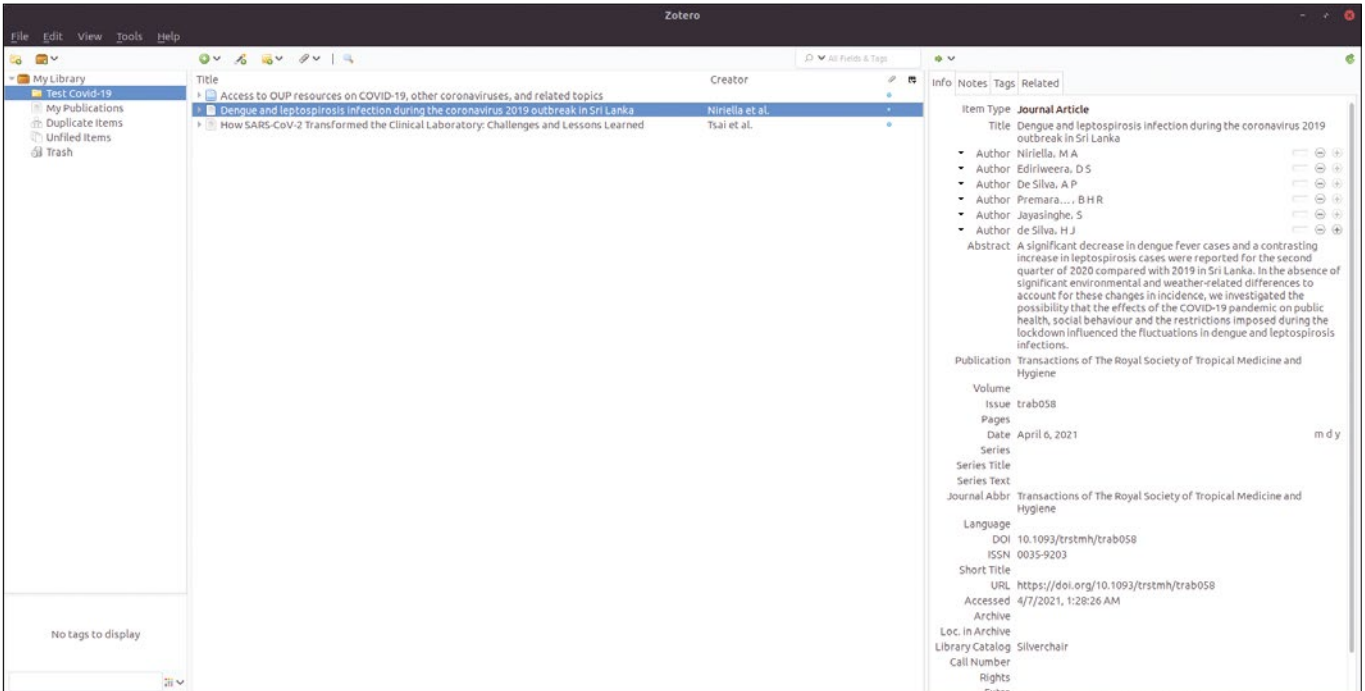



Figure 1: The main Zotero window is divided into three columns.

open. If the Zotero extension is installed, you will see the Zotero icon in the top right corner of the Firefox toolbar. In Chrome, you can unhide the icon by clicking on the extensions icon in the upper right corner.

Use the Zotero icon to import bibliographic data. The appearance of the icon changes depending on what information the browser is currently displaying. An icon like a small sheet of paper indicates a web page without bibliographic information.

Zotero is even more helpful when you search the online catalog of a library. Use the catalog's search box to search for one or more terms, and as soon as the search results load, the Zotero icon in the browser changes so that it now looks like a small, yellow folder. This means that the software has identified several importable elements. As soon as you click on the icon, a small window will appear, listing the search results separately once again. Now check the titles you want to import and press *OK*.

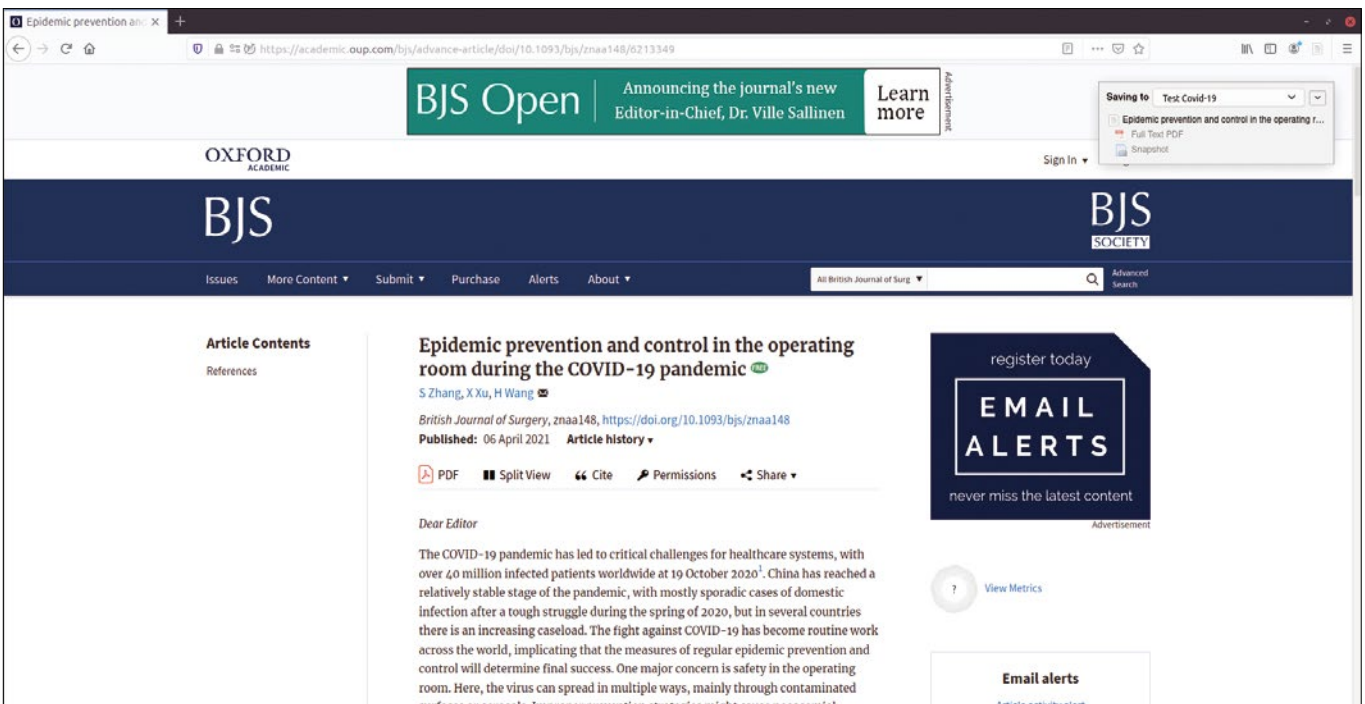


Figure 2: When importing literature, you can immediately classify it into collections.

Alternatively, import the publications one by one after you have looked at them more closely. Click on one of the search results in the library catalog hit list. Zotero then loads more detailed information about the publication. The Zotero icon now typically looks like a small blue book or a written piece of paper. This indicates that Zotero can import a single item, such as a book or magazine article.

To import the publication into your Zotero library, click the icon. A small dialog box appears in the top right corner (Figure 2), where you can instruct Zotero to place the publication in a specific collection or to tag it with keywords. If you do nothing, the dialog will disappear after a few seconds and the publication will be placed in your library, where you can add these options later. If the main program is not currently running in the background, Zotero offers to use the web service when importing. If you chose this option, a dialog will then prompt you to log in to Zotero.

Editing Entries

After importing new literature, switch back to the main Zotero program and click *My Library* on the far left. You will now see all previously added elements in the middle column. Check to be sure that the import worked correctly (Figure 3). Occasionally, individual pieces of data are transferred incorrectly or are even missing completely. To check, click on one of the imported entries in the middle column. On the right, under the *Info* tab, you will then see all the details for the entry, and you can correct any

mistakes by clicking and editing the individual database fields.

Importing bibliographic data works almost perfectly for many catalogs and databases, but it is less successful for others. Databases where the import works very reliably include the JSTOR [6] journal database and Amazon, the online retailer. With the catalogs of large university and national libraries, the imported information is usually almost error-free.

Importing information about portions of books, such as a single essay from a longer book, often proves to be more complicated because Zotero may treat them as independent publications. As a result, information about the book in which the essay is located may be missing. To address this problem, import both the essay and the publication it is in, and then duplicate the entry representing the parent publication. Edit the duplicated entry as needed, adding any needed information about the essay and changing the entry type in the top database row to *Book Part*. Finally, delete the first entry, which contains only the title and author of the essay. However, keep in mind that if there are separate entries in the bibliography for the essay and the book containing it, two corresponding entries must also exist in Zotero.

Sometimes you will have to create entries manually, for example, if you find an unpublished document in an archive. Click on the icon with the small green circle and the white plus sign in the Zotero toolbar, and you will then see a list from which you can select the type of entry. Next, a new, empty entry appears in the

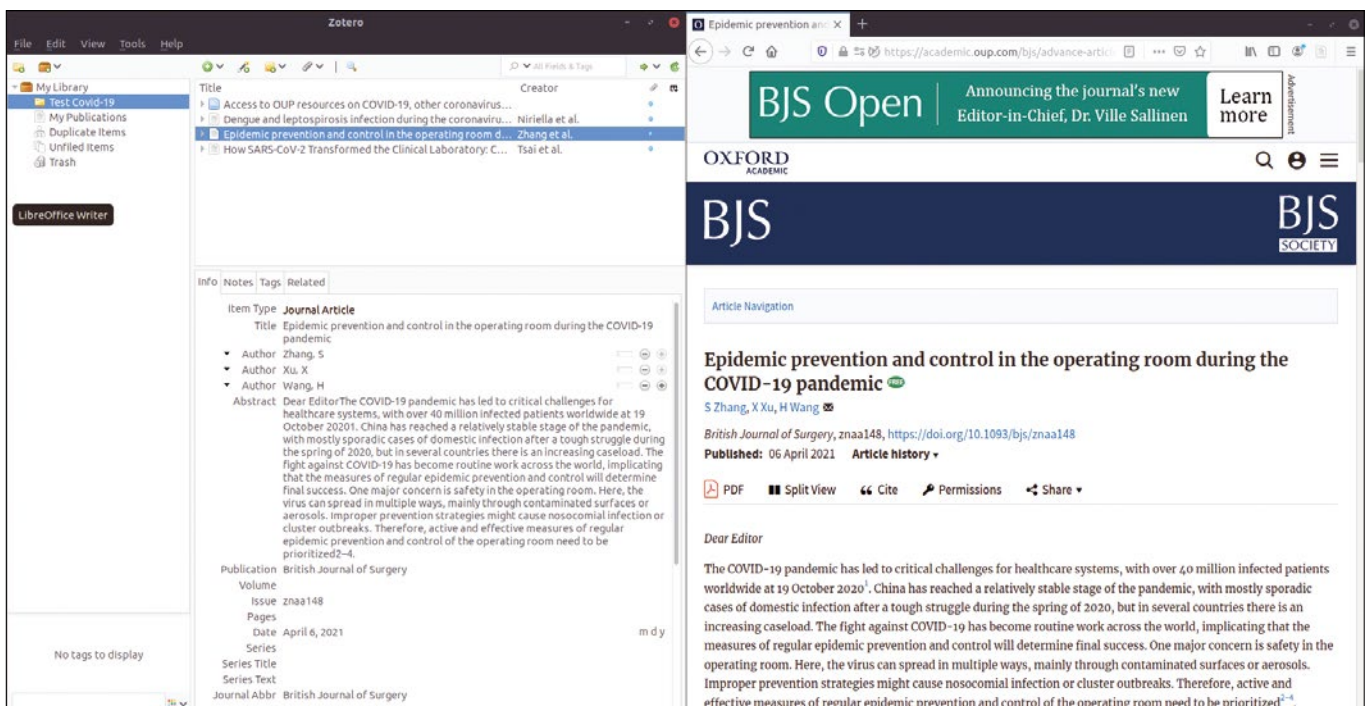


Figure 3: When reviewing imports, it's best to arrange the browser window and Zotero side by side.

middle column of the Zotero window, and you can enter the details of the title, author, and so on into the database fields you see in the right column.

Attaching files to entries is simple. Select an entry and click on the icon with the paper clip in the toolbar. Zotero copies the files to the `~/Zotero/storage/` directory. If you have not disabled synchronization for attachments, the program will also upload the attachment to the cloud storage.

Collections and Bibliographies

When working on a specific project, consider creating a separate collection for it to make it easier to generate a bibliography. Right-click *My Library* in the left column of the Zotero window and click *New Collection* in the menu that appears. In the dialog box that then appears, type a name for the collection.

The new collection now appears under *My Library*. To assign titles to this collection, select *My Library* to view the importable entries in your library in the middle column. Narrow down these entries by typing a search term in the box to the right above the large center column. Click on an entry and drag it to the left to the appropriate collection.

To create a bibliography for a project, right-click the name of the collection in the left column, and then click *Create Bibliography from Collection*. A dialog box appears in which you will first choose a citation style. This determines how Zotero orders

or formats the individual entries. You do not need to change the output mode – Zotero sets *Bibliography* by default. Finally, you have the choice of saving the bibliography as an RTF or HTML file, copying it to the clipboard, or printing it. Look carefully at the bibliography you have created and correct any errors in Zotero, if possible, rather than in the word processor.

Zotero and LibreOffice

It's a good idea to use the Zotero extension for LibreOffice both for inserting references and citation sources into the text and for creating bibliographies. This will help you avoid missing entries in the bibliography that you refer to in the text – a common mistake even among experienced researchers.

Before installing the LibreOffice extension, check if the use of a Java runtime environment is enabled in the office suite. If there is no Java Development Kit on your system, install the appropriate package from the package sources of the distribution you are using – the package usually has a *jdk* component in its name.

Then tell LibreOffice the directory in which the Java runtime environment resides. Call up the LibreOffice settings below *Tools | Options...* You will find the Java settings under *LibreOffice | Advanced*. Under *Java Options*, check *Use a Java Runtime Environment*, click *Add...* below it, and then select the directory. Close the *Options* dialog via the *OK* button and exit LibreOffice.

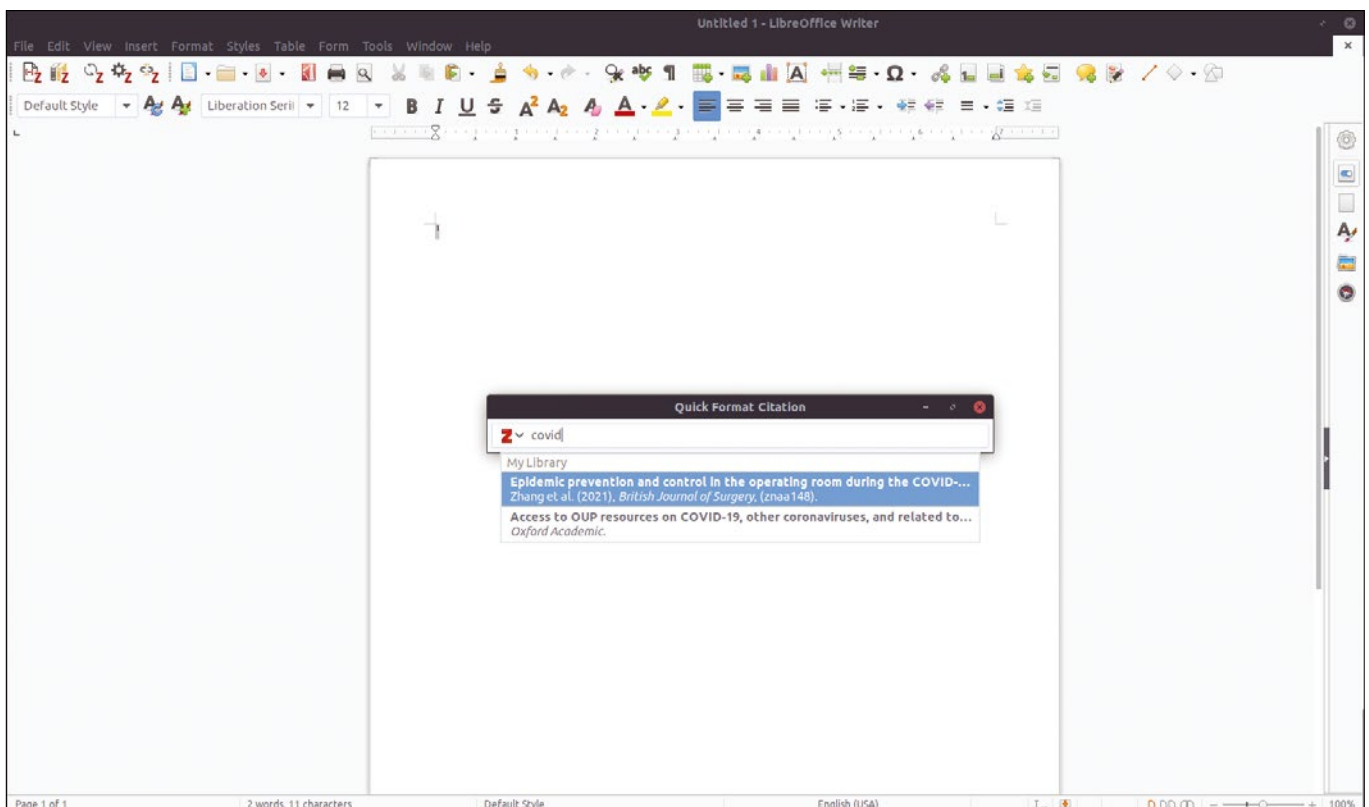


Figure 4: Use the Zotero extension for LibreOffice to quickly insert references into the text.

Go to the Zotero settings to start the installation of the LibreOffice extension. Switch to the *Cite* tab, click on the tab for word processors, and click the button to install the LibreOffice plug-in. The dialog windows of the installation wizard should show the directory where LibreOffice is located, but if needed, you can specify the directory manually. After launching LibreOffice, you will see – usually on the far left – a small additional toolbar where each of the five buttons contains a small red Z. Click on the button which includes a small gear, and in the following dialog, select the citation style you want to use.

To insert a literature reference, click the button with the document icon next to the Z. The *Quick Format Citation* dialog appears in the center of the window. Type a keyword, and select a publication from the search results. Zotero automatically transfers the reference to the desired location in the text (Figure 4).

Once you have added all the references to your work, you can create an up-to-date bibliography with one click. Place the cursor where you want the bibliography to be and click on the second icon in the toolbar, which has a Z next to two books. Zotero then automatically generates the bibliography from the references in the text.

Installing Citation Styles

The citation style determines the appearance of the source references in the text and the bibliography.

This includes whether literature references are in brackets or appear as footnotes, whether the author's name is formatted in small caps or the title is in italics, whether commas or periods separate individual elements, and more.

In the default installation, Zotero brings several citation styles, including *Chicago Manual of Style 17th Edition (author-date)* and *Chicago Manual of Style 17th Edition (note)*. These two multilingual citation styles are suitable for many subjects. However, you can install additional styles from Zotero's citation style database, which now includes many thousands of styles, or create your own.

To install a new citation style, go to Zotero's preferences, switch to the *Cite* tab, and select the *Styles* option. Below the style manager, click *Get additional styles...* . This opens the *Zotero Style Repository*, which allows you to search for new styles (Figure 5).

Type a search term in the box at top left, or click on one or more of the keywords you see on the right. Below that, Zotero lists the styles it found. Mouse over one of the styles, and you will see a preview of the citation style. Click on the desired entry to apply the style.

You can preview the installed citation styles directly in Zotero. Select multiple entries that represent different publication types, for example, a journal article and a book. Then, in Zotero

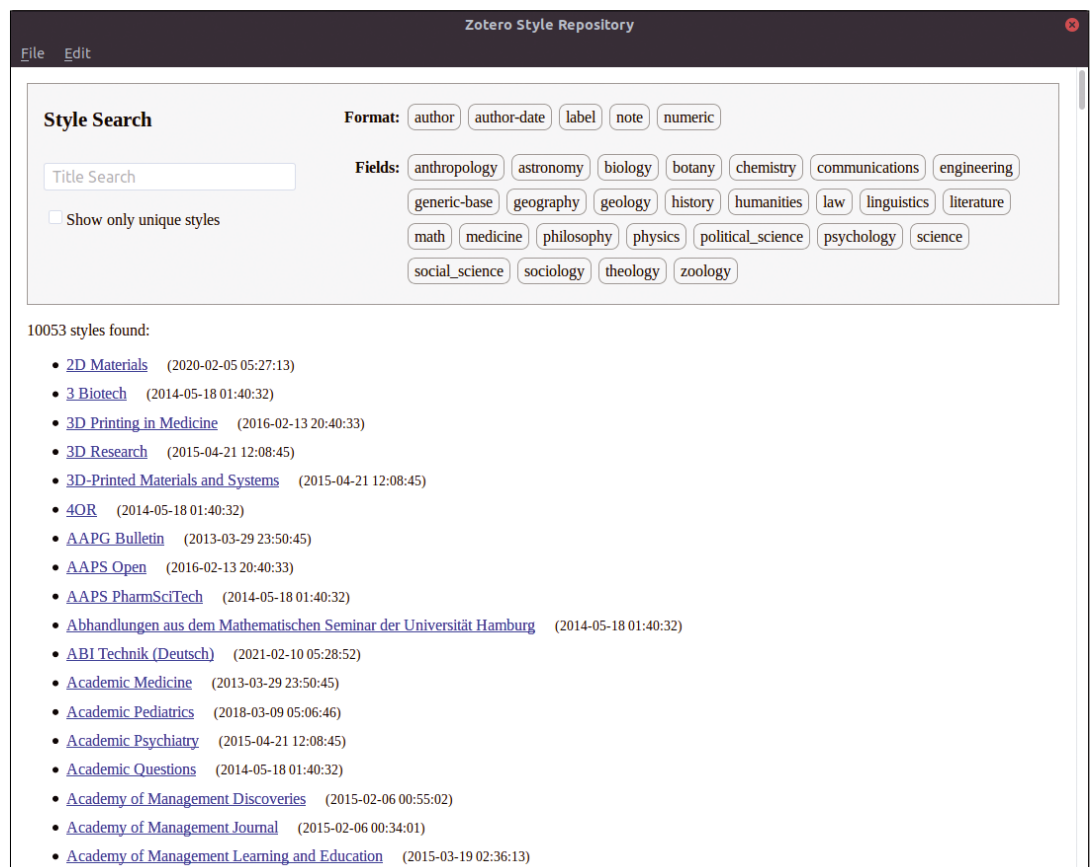


Figure 5: The Zotero Style Repository already contains almost 10,000 citation styles.

Settings, switch to the *Cite* tab, and click *Style Preview*. In the preview window (Figure 6), click *Refresh* in the upper left corner to create the preview.

Customizing Citation Styles

If you can't find a citation style that exactly suits your needs, find a similar style and edit it, for example, by pulling in elements of another style. Zotero uses the Citation Style Language (CSL) to describe citation styles; this is a variant of the XML markup language that is also used by other literature management programs. For experienced programmers, CSL is a largely self-explanatory programming language. However, with a little application, you can edit citation styles, and even create new ones, without any knowledge of programming or XML. The CSL documentation and command reference can be found in the specification [7].

Conclusions

The advantages of Zotero are obvious: The free tool makes it easier to keep track of your re-

search and eliminates a lot of typing and formatting work. Support for some catalogs is not yet perfect, and creating a new citation style can be challenging. Nevertheless, Zotero comes across as a professional, mature, and easy-to-use literature management program that stands up well against its commercial competitors. ■■■

Info

- [1] Zotero: <https://www.zotero.org>
- [2] Jurism: <https://juris-m.github.io>
- [3] Register account: <https://www.zotero.org/user/register>
- [4] Zotero Storage: <https://www.zotero.org/storage>
- [5] Zotero Connector: <https://www.zotero.org/download/connectors>
- [6] JSTOR: <https://jstor.org>
- [7] CSL 1.0.1 Specification: <https://docs.citationstyles.org/en/stable/specification.html>

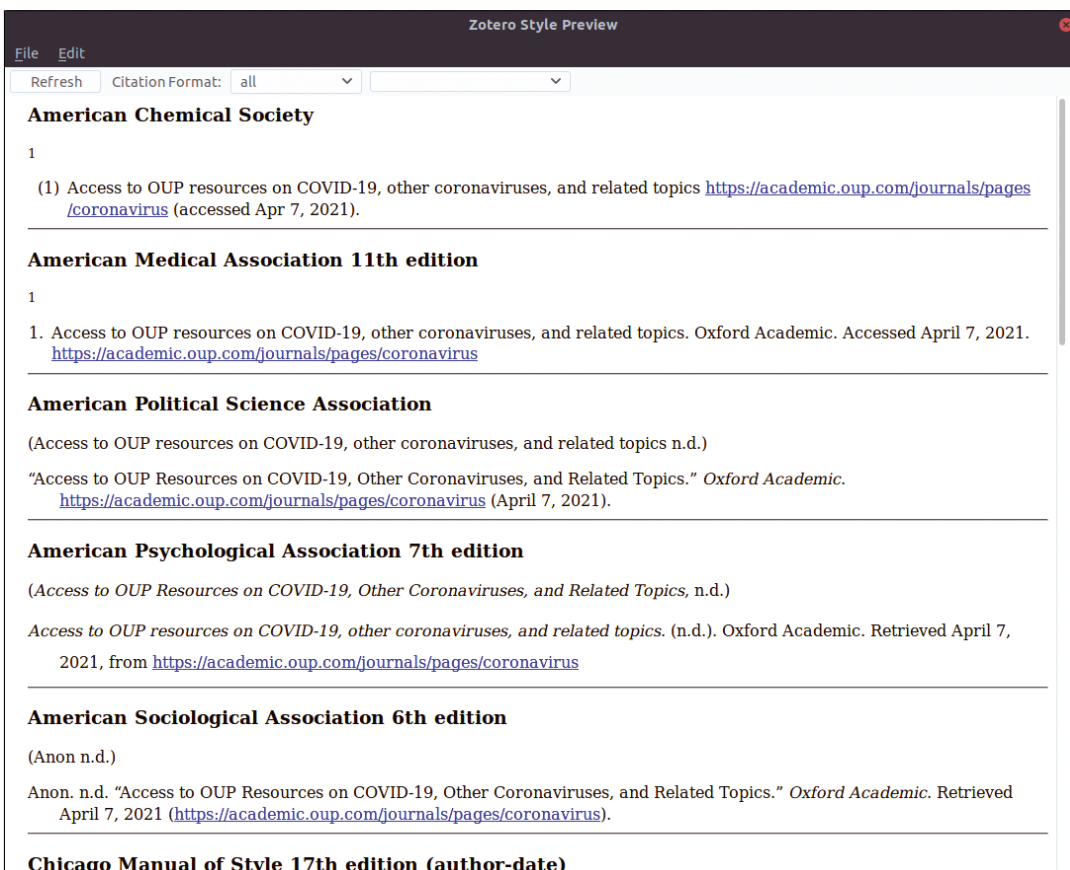


Figure 6: The Zotero Style Preview shows a preview of all installed citation styles.



Hone your skills with special editions!




Get to know Shell, LibreOffice, Linux, and more from our Special Edition library.

The *Linux Magazine* team has created a series of single volumes that give you a deep-dive into the topics you want.

Available in print or digital format

Check out the full library!

shop.linuxnewmedia.com



LINUX NEWSSTAND

Order online:
<https://bit.ly/Linux-Newsstand>

Linux Magazine is your guide to the world of Linux. Monthly issues are packed with advanced technical articles and tutorials you won't find anywhere else. Explore our full catalog of back issues for specific topics or to complete your collection.



#247/June 2021

Post-Quantum Encryption

Quantum computers are still at the experimental stage, but mathematicians have already discovered some quantum-based algorithms that will demolish the best of our current encryption methods. What better time to look for quantum encryption alternatives?

On the DVD: Knoppix 9.1 and ZORIN OS 15.3 Core



#246/May 2021

Faster Startup

Weary of waiting for a login window? Your driver-drenched Linux distro was configured for all systems, not for your system. This month we show you how to optimize your system for faster startup.

On the DVD: Manjaro KDE Plasma 20.2.1 and Clonezilla Live 2.7.1



#245/April 2021

Choose a Shell

You're never stuck with the same old command shell – unless you want to be. This month we review some of the leading alternatives.

On the DVD: Arch Linux 2021.02.01 and MX Linux mx-19.03



#244/March 2021

Stream Processing

The explosion of real-time data from sensors and monitoring devices is fueling new interest in alternative programming techniques. This month we waded into stream processing.

On the DVD: FreeBSD 12.2 and GhostBSD 20.11.28



#243/February 2021

iNet

With Linux, more innovation is always on the way. This month we take a look at the iNet wireless daemon, a new wireless client that is poised to replace the venerable WPA Supplicant.

On the DVD: Linux Mint 20 and Kali Linux 2020.4



#242/January 2021

3D Printing

The weird, wonderful, futuristic world of 3D printing is waiting for you right now if you're willing to invest a little time and energy. This month we help you get started with practical 3D printing in Linux.

On the DVD: Ubuntu 20.10 "Groovy Gorilla" and Fedora 33 Workstation

FEATURED EVENTS



Users, developers, and vendors meet at Linux events around the world. We at *Linux Magazine* are proud to sponsor the Featured Events shown here. For other events near you, check our extensive events calendar online at <https://www.linux-magazine.com/events>.

If you know of another Linux event you would like us to add to our calendar please send a message with all the details to events@linux-magazine.com.

NOTICE

Be sure to check the event website before booking any travel, as many events are being canceled or converted to virtual events due to the effects of COVID-19.

USENIX ATC '21

Date: June 18-20, 2021

Location: Virtual Event

Website:

<https://www.usenix.org/conference/atc21>

USENIX ATC brings together leading systems researchers for the presentation of cutting-edge systems research and the opportunity to gain insight into a wealth of must-know topics, including virtualization, system and network management, troubleshooting, cloud and edge computing, security, privacy and trust, mobile and wireless, and much more.

GUADEC 2021

Date: July 21-25, 2021

Location: Virtual Event

Website:

<https://events.gnome.org/event/9/>

GUADEC is the GNOME community's largest conference, bringing together hundreds of users, contributors, community members, and enthusiastic supporters for a week of talks and workshops. View the GUADEC schedule and start planning your GUADEC experience.

Events

stackconf online 2021	June 15-16	Virtual Event	https://stackconf.eu/
openSUSE Virtual Conference 2021	June 18-19	Virtual Event	https://events.opensuse.org/
Akademy 2021	June 18-25	Virtual Event	https://akademy.kde.org/2021
ISC High Performance 2021 Digital	June 24-July 2	Virtual Event	https://www.isc-hpc.com/
USENIX ATC '21	July 14-16	Santa Clara, California	https://www.usenix.org/conference/atc21
GUADEC 2021	July 21-25	Virtual Event	https://events.gnome.org/event/9/
Embedded Linux Conference North America	August 3-6	Vancouver, British Columbia	https://events.linuxfoundation.org/
Open Source Summit North America	August 3-6	Vancouver, British Columbia	https://events.linuxfoundation.org/
USENIX Security '21	August 11-13	Vancouver, British Columbia	https://www.usenix.org/conferences
Kubernetes Community Days	September 9-10	Amsterdam, Netherlands	https://sessionize.com/kcdams2021/
DeveloperWeek Global: Cloud	September 14-15	Virtual Event	https://www.developerweek.com/global/
KVM Forum	September 27-29	Dublin, Ireland	https://events.linuxfoundation.org/
Embedded Linux Conference Europe	Sept 28-Oct 1	Dublin, Ireland	https://events.linuxfoundation.org/
Open Source Summit Europe	Sept 28-Oct 1	Dublin, Ireland	https://events.linuxfoundation.org/
DrupalCon Europe 2021	October 10-7	Virtual Event	https://events.drupal.org/europe2021
IEEE Quantum Week 2021	October 18-22	Virtual Event	https://qce.quantum.ieee.org/

CALL FOR PAPERS

We are always looking for good articles on Linux and the tools of the Linux environment. Although we will consider any topic, the following themes are of special interest:

- System administration
- Useful tips and tools
- Security, both news and techniques
- Product reviews, especially from real-world experience
- Community news and projects

If you have an idea, send a proposal with an outline, an estimate of the length, a description of your background, and contact information to edit@linux-magazine.com.



The technical level of the article should be consistent with what you normally read in *Linux Magazine*. Remember that *Linux Magazine* is read in many countries, and your article may be translated into one of our sister publications. Therefore, it is best to avoid using slang and idioms that might not be understood by all readers.

Be careful when referring to dates or events in the future. Many weeks could pass between your manuscript submission and the final copy reaching the reader's hands. When submitting proposals or manuscripts, please use a subject line in your email message that helps us identify your message as an article proposal. Screenshots and other supporting materials are always welcome.

Additional information is available at:

http://www.linux-magazine.com/contact/write_for_us.

Authors

Mats Tage Axelsson	46
Erik Bärwaldt	14, 20, 24
Zack Brown	11
Bruce Byfield	6, 32, 42
Joe Casad	3
Mark Crutch	67
Jon "maddog" Hall	69
Frank Hofmann	28
Charly Kühnast	55
Christoph Langner	74
Vincent Mealing	67
Anzela Minosi	62
Goran Mladenovic	36
Graham Morrison	82
Scott Nesbitt	70
Gerhard Schauer	56
Mike Schilli	52
Anna Simon	88
Jack Wallen	8

Contact Info

Editor in Chief

Joe Casad, jcasad@linux-magazine.com

Copy Editors

Amy Pettle, Megan Phelps

News Editor

Jack Wallen

Editor Emerita Nomadica

Rita L Sooby

Managing Editor

Lori White

Localization & Translation

Ian Travis

Layout

Dena Friesen, Lori White

Cover Design

Lori White

Cover Image

© o du van, 123RF.com

Advertising

Brian Osborn, bosborn@linuxnewmedia.com
phone +49 8093 7679420

Marketing Communications

Gwen Clark, gclark@linuxnewmedia.com
Linux New Media USA, LLC
4840 Bob Billings Parkway, Ste 104
Lawrence, KS 66049 USA

Publisher

Brian Osborn

Customer Service / Subscription

For USA and Canada:
Email: cs@linuxpromagazine.com
Phone: 1-866-247-2802
(Toll Free from the US and Canada)

For all other countries:
Email: subs@linux-magazine.com

www.linuxpromagazine.com – North America

www.linux-magazine.com – Worldwide

While every care has been taken in the content of the magazine, the publishers cannot be held responsible for the accuracy of the information contained within it or any consequences arising from the use of it. The use of the disc provided with the magazine or any material provided on it is at your own risk.

Copyright and Trademarks © 2021 Linux New Media USA, LLC.

No material may be reproduced in any form whatsoever in whole or in part without the written permission of the publishers. It is assumed that all correspondence sent, for example, letters, email, faxes, photographs, articles, drawings, are supplied for publication or license to third parties on a non-exclusive worldwide basis by Linux New Media USA, LLC, unless otherwise stated in writing.

Linux is a trademark of Linus Torvalds.

All brand or product names are trademarks of their respective owners. Contact us if we haven't credited your copyright; we will always correct any oversight.

Printed in Nuremberg, Germany by hofmann infocom GmbH.

Distributed by Seymour Distribution Ltd, United Kingdom

LINUX PRO MAGAZINE (ISSN 1752-9050) is published monthly by Linux New Media USA, LLC, 4840 Bob Billings Parkway, Ste 104, Lawrence, KS 66049, USA. Periodicals Postage paid at Lawrence, KS and additional mailing offices. Ride-Along Enclosed. POSTMASTER: Please send address changes to Linux Pro Magazine, 4840 Bob Billings Parkway, Ste 104, Lawrence, KS 66049, USA.

Published monthly in Europe as Linux Magazine (ISSN 1471-5678) by: Sparkhaus Media GmbH, Bialasstr. 1a, 85625 Glonn, Germany.

Approximate

UK / Europe July 03

USA / Canada July 30

Australia Aug 30

On Sale Date

Please note: On sale dates are approximate and may be delayed because of logistical issues.

Issue 249 / August 2021

Your Smartphone as a Linux PC

Our smartphones are getting faster and more powerful every year, so it is fair to ask, is it possible to boot a full version of Linux on an Android phone? The answer is yes, and next month we'll show you how.



Preview Newsletter

The Linux Magazine Preview is a monthly email newsletter that gives you a sneak peek at the next issue, including links to articles posted online.

Sign up at: <https://bit.ly/Linux-Update>

Lead Image © Daniil Peshkov and georgejmlittle, 123RF.com



Stealth Gaming

TUXEDO Book XP15 & XP17



Intel Core i7-10870H
8 Cores | 16 Threads



GeForce RTX 3080 Max-Q
16 GB GDDR6



73 Wh
Li-polymer battery



G-SYNC, 300 Hz oder 4K
Wide display choice



100%
Linux

5

Year
Warranty



Lifetime
Support



Built in
Germany



German
Privacy



Local
Support


TUXEDO

COMPUTERS

[tuxedocomputers.com](https://www.tuxedocomputers.com)

HETZNER

DEDICATED ROOT SERVER DESIGNED FOR PROFESSIONALS



Dedicated Root Server AX41-NVMe

- ✓ AMD Ryzen™ 5 3600
Simultaneous Multithreading
- ✓ 64 GB DDR4 RAM
- ✓ 2 x 512 GB NVMe SSD
- ✓ 100 GB Backup Space
- ✓ Traffic unlimited
- ✓ Location Finland and Germany
- ✓ No minimum contract
- ✓ Setup Fee £34



monthly from **£30**

Dedicated Root Server AX51-NVMe

- ✓ AMD Ryzen™ 7 3700X
Simultaneous Multithreading
- ✓ 64 GB DDR4 ECC RAM
- ✓ 2 x 1 TB NVMe SSD
- ✓ 100 GB Backup Space
- ✓ Traffic unlimited
- ✓ Location Finland and Germany
- ✓ No minimum contract
- ✓ Setup Fee £51



monthly from **£47**

All prices exclude VAT and are subject to the terms and conditions of Hetzner Online GmbH. Prices are subject to change. All rights reserved by the respective manufacturers.

www.hetzner.com