

FREE DVD

LINUX MAGAZINE PRO



GPX VIEWER

Map your tracks and view GPS data in a web browser

BACKUP

Practical tools for protecting your data browser

LINUX MAGAZINE PRO

OCTOBER 2019

BACKUP


Practical tools for protecting your data

Unbound

Set up a local server for faster and more secure DNS

Password Generators

An easy way to better security



LibreELEC 9.0

Turn your Rasp Pi into a media center

Doomsday Rule

A Go program for finding days of the week

Jade Desktop

Will this innovative desktop environment change the game?

LINUXVOICE

- Linux Photo Tools
- maddog on licensing and copyright law
- TaskBook: Task manager for the command line



FOSSPicks

- Cartes du Ciel
- Newsboat
- PyRadio

Tutorial

- Decentralized Social Media with Mastodon

Issue 227
Oct 2019
US\$ 15.99
CAN\$ 18.50



AN EXTREME REFRESH
WITH OCTA-CORE POWER!

DEDICATED ROOT SERVER EX62

High speed performance with the new
Intel Core i9-9900K octa-core processor

ENTERPRISE
HDDs or
NVME SSDs



Dedicated Root Server EX62

- ✓ Intel® Core™ i9-9900K Octa-Core
incl. Hyper-Threading-Technology
- ✓ 64 GB DDR4 RAM
- ✓ 2 x 8 TB SATA Enterprise Hard Drive 7200 rpm
- ✓ 100 GB Backup Space
- ✓ Traffic unlimited
- ✓ Location Germany or Finland
- ✓ No minimum contract
- ✓ Setup Fee \$78

monthly from \$ **72.50**

Dedicated Root Server EX62-NVMe

- ✓ Intel® Core™ i9-9900 Octa-Core
incl. Hyper-Threading-Technology
- ✓ 64 GB DDR4 RAM
- ✓ 2 x 1 TB NVMe SSD
- ✓ 100 GB Backup Space
- ✓ Traffic unlimited
- ✓ Location Germany or Finland
- ✓ No minimum contract
- ✓ Setup Fee \$78

monthly from \$ **72.50**

WE ARE HERE

Dear Reader,

As one who has been watching Linux for a long time, I'm saddened to mark the end of *Linux Journal*, which was just announced this week. I know what you're thinking – aren't they a competitor? But it isn't like that. The community spirit of the open source world is an umbrella for everyone in it – including publishers. If anything, there has always been a kind of camaraderie among the Linux mags.

I knew about *Linux Journal* long before I had this job. In fact, my associations with *Linux Journal* date to an earlier era in IT publishing – back to when I was working for other magazines, like *SysAdmin* and *C Users Journal*. *Linux Journal* has been around for so long that I was especially surprised to read the comments on their site that said, "What? You're gone already? I just found out about you!"

Huh, you use Linux, and you've never heard of *Linux Journal*?

That made me consider how important it is to spread the word about worthy IT publications. In the old days, publishers spent a lot of money on branding and direct marketing to be sure everyone knew they were out there, but all that ended a couple market crashes back. Nowadays, most of our readers hear about us from other readers, unless they happen to live near a full-service IT newsstand, which are also getting rarer in this world.

Magazines play a special role in the open source community. Microsoft and Apple have lots of budget for sponsoring tutorials, white papers, help lines, and websites devoted to promoting their systems and the accompanying applications. The Linux world has no real equivalent to that kind of deep-pocket investment. At the same time, the sprawling and free-wheeling Linux environment prioritizes freedom over uniformity, which means there is no standard answer for how to do anything, and there are always lots of different tools and configuration details to talk about.

Rather than following some canonical rule book, Linux users find their answers through a variety of informal sources. If you know exactly what problem you are trying to solve, you call up the man page, read a wiki, or post the question to a forum. But what if you don't know what you don't know? Magazines (in print and online) play an important role in calling attention to tools you might not know about and techniques that raise your general understanding of Linux – they help you be ready with the solution before the problem arises.

Linux Journal was nearly as old as Linux itself, and it was instrumental in building mindshare for Linux since the very beginning. As for us, we're still hanging in there after 19 years, thanks to our readers, our authors, and our amazing and extremely experienced staff. We are going to keep doing what we do for as long as readers like you want to keep reading it, but it always helps to have more readers like you. So pass the word: We are here!



Joe Casad,
Editor in Chief





WHAT'S INSIDE

The bad news is that disasters happen. The good news is that Linux has a variety of powerful backup tools that will protect you when disaster strikes. This month, we cover some backup apps for the Linux environment, including Rclone, restic, and rsnapshot. Elsewhere inside:

- **Local DNS with Unbound** – improve performance and enhance security with a local DNS server (page 50).
- **Jade Desktop** – the creators claim this new desktop environment will “change the way you interact with your desktop” (page 56).

See our MakerSpace section for a look at the LibreELEC media center distro, and check out LinuxVoice for a study of some useful and little-known photo tools.

SERVICE

- 3 Comment
- 6 DVD
- 96 Featured Events
- 97 Call for Papers
- 98 Preview

NEWS

08 News

- Linux Mint 19.2 “Tina” Released
- Gnome and KDE Coming Together
- Fedora CoreOS Preview Released
- SUSE Appoints New CEO
- GitHub Blocks Access to Private Repositories in Certain Countries
- A New Ransomware Targeting Linux-Based NAS Devices

12 Kernel News

- Communicating with Linux During Bootup
- Pruning SuperH
- Bug Hunting for Steam

16 Interview – prpl Foundation

Mirko Lindner of the prpl Foundation discusses how vendors and chipset makers are coming together to make IoT more interoperable and secure.



COVER STORIES

20 Rsync for the Cloud

Backups for critical data have to withstand major disasters. Rclone, a descendant of Rsync, enables easy backups to the cloud and supports many services.

26 restic

Many users still find it difficult to regularly back up their data. Thanks to restic, and its graphical front end, Restatic, a backup is quite easy to manage – even for a beginner.

30 rsnapshot

A recent backup is more reliable than any kind of data rescue. But for many users, a backup won't happen unless the process is easy to manage.

IN-DEPTH

34 GPX Viewer

With GPX Viewer, you can map GPX tracks and view GPS data in a web browser. It's a simple way to revisit a recent vacation, organize your photos, or map your favorite bike routes.

40 Command Line – Password Generators

CLI tools for generating passwords have many options that can help you strike a balance between ease of use and security.

IN-DEPTH

- 44 Programming Snapshot – Go**
Math wizards amaze audiences by naming the day of the week for any date called out by the audience. Mike Schilli writes a training program in Go to turn amateurs into headline performers.



- 50 Local DNS with Unbound**
You don't have to be satisfied with your ISP's slow and cumbersome DNS server. Your own Unbound server could improve performance as well as security.
- 56 Jade**
This new desktop environment strives to offer something different, and while the design isn't for everyone, it does introduce a couple of interesting features.



- 58 firewalld**
Managing a firewall can be a hassle, but it's worse to manage a breach because you didn't have one.
- 60 Charly's Column – f.lux**
You don't have to be a vampire to be sensitive to bright light at night. Sys admin columnist Charly, who – as regular readers know – is a practicing light conservationist, now makes it clear to his desktop PC that it's not good to be too dazzling at night.


MAKERSPACE

- 62 Open Hardware – Crowdfunding**
Open hardware and crowdfunding are a natural fit. While their relationship has evolved over time, crowdfunding continues to play a big role in open hardware development.
- 66 Kodi 18.0 and LibreELEC 9.0**
Thanks to Kodi 18.0, LibreELEC 9.0 now supports the DRM encryption used by many streaming services. However, integrating Netflix, Amazon, and other streaming services does involve some manual work.



LINUXVOICE

- 73 Welcome**
This month in LinuxVoice.
- 75 Doghouse – Copyright and Licensing**
Despite the common assumption, everything online is not in the public domain.
- 76 Photo Tools**
DIY tools and shell scripts can make for a smooth photographic workflow on Linux.
- 80 TaskBook**
Boost your productivity from the Linux command line with the TaskBook task manager and note board.
- 84 FOSSPicks**
This month Graham looks at Cartes du Ciel, Foliote, Newsboat, xinput-gui, DevolutionX, Performous, and much more!



- 90 Tutorial – Fediverse**
If you're looking for social media options where the user has more control, you'll find a range of options to explore in the Fediverse, including the popular Mastodon.

On the DVD

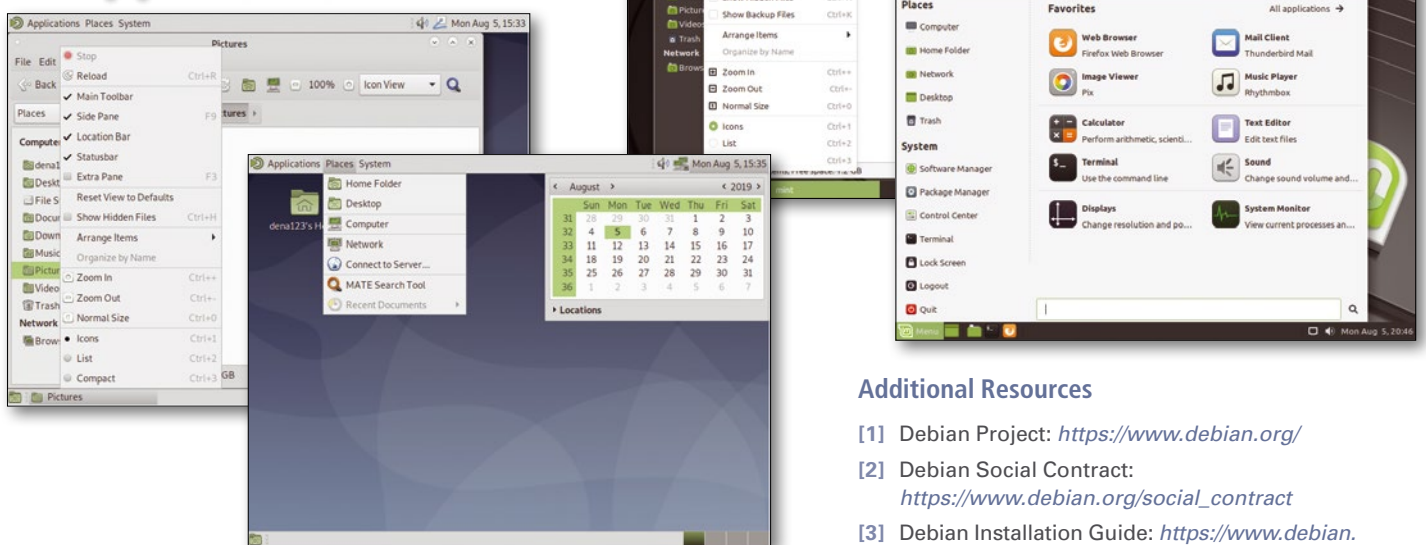


Debian 10 "Buster"

The legendary Debian is an independent, community-based distro with an impressive worldwide network of volunteers. Careful testing and an extra-long development cycle give Debian a reputation for stability. Debian, which contains no proprietary components, serves as the basis for several popular Linux systems, including Ubuntu, Knoppix, and more. The latest release comes with both Wayland and Xorg display servers. AppArmor and the nftables firewall are included by default, along with simplified UEFI support and updates to thousands of packages in the vast Debian repositories.

Linux Mint 19.2 "Tina" MATE Edition

Linux Mint is a user-friendly desktop Linux based on Ubuntu. This month's DVD comes with the MATE edition, featuring the Gnome 2-based MATE desktop. The last version includes improvements to the update manager and several configuration tools. A rebuilt System Reports tool offers more convenient access to troubleshooting information.



Additional Resources

- [1] Debian Project: <https://www.debian.org/>
- [2] Debian Social Contract: https://www.debian.org/social_contract
- [3] Debian Installation Guide: <https://www.debian.org/releases/stable/amd64/index.en.html>
- [4] Debian Wiki: <https://wiki.debian.org/>
- [5] Linux Mint: <https://linuxmint.com/>
- [6] Linux Mint Installation Guide: <https://linuxmint-installation-guide.readthedocs.io/en/latest/>

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.



Join us at DrupalCon Amsterdam!

REGISTER NOW ON
[EVENTS.DRUPAL.ORG/AMSTERDAM2019](https://events.drupal.org/amsterdam2019)

DrupalCon Amsterdam 2019 will take place **from 28 – 31 October at the RAI Amsterdam**, with high quality Keynote Speakers like **Dries Buytaert, Dr. Sue Black and Boris Veldhuijzen van Zanten**



NEWS

Updates on technologies, trends, and tools

THIS MONTH'S NEWS

- 08 • Linux Mint 19.2 "Tina" Released
- Gnome and KDE Coming Together
- 09 • Fedora CoreOS Preview Released
- SUSE Appoints New CEO
- More Online
- 10 • GitHub Blocks Access to Private Repositories in Certain Countries
- A New Ransomware Targeting Linux-Based NAS Devices

Linux Mint 19.2 "Tina" Released

The Linux Mint project has announced the release of Linux Mint 19.2, codenamed Tina. It's an LTS release, which will be supported until 2023. The release comes with updated software and refinements. It also introduces some new features to desktop users.

Some under-the-hood improvements bring good news to those who play at the kernel level. The Update Manager continues to get better; it now shows the kernel life cycle. Also, if there are old kernels that you want to remove or update, you won't have to install or remove kernels one by one. Users can queue installations and removals and install and remove multiple kernels in one go. It also adds support for custom kernels. "If you've got a non-generic kernel installed, the Update Manager will show a combo box, so you can switch between flavors," said the project page.

In terms of performance, the new version of Cinnamon improves RAM consumption. "On a test VirtualBox virtual machine, Cinnamon 4.2 uses approximately 67MB RAM (compared to 95MB RAM for Cinnamon 4.0)," said the blog. Optimizations made to the Muffin window manager reduce input lag and make windows feel smoother and lighter.

Linux Mint 19.2 is available for free download. If you are running a previous version of Linux Mint, there is a relatively easy path to upgrade from one version to the next.



Gnome and KDE Coming Together

Linux dominates the world, except for the desktop. One of the problems associated with the Linux desktop is fragmentation.

In an exclusive interview, Linus Torvalds told me that, "fragmentation of the different vendors have held the desktop back."

Primarily this fragmentation comes from competing and often conflicting desktop projects. The good news is that two major Linux desktop communities are working on joining hands to eliminate this fragmentation. The Gnome Foundation



and KDE e.V. have announced Linux App Summit (LAS) 2019, which will be held in Barcelona from November 12-15, 2019.

LAS is the first collaborative event cohosted by the two organizations since the Desktop Summit in 2009. "With the joint influence of the two desktop projects, LAS will shepherd the growth of the FOSS desktop by encouraging the creation of quality applications, seeking opportunities for compensation for FOSS developers, and fostering a vibrant market for the Linux operating system," said the press release.

Fedora CoreOS Preview Released

The Fedora community has released a preview of Fedora CoreOS, a Linux-based operating system designed to run containers.

Red Hat acquired CoreOS, an open source company, last year and has been integrating CoreOS products and services with its own. CoreOS used to have a distribution with the same name, which was later renamed to Container Linux.



Benjamin Gilbert of Red Hat wrote on the mailing list that Fedora CoreOS is built to be a secure and reliable host for compute clusters. "It's designed specifically for running containerized workloads

without regular maintenance, automatically updating itself with the latest OS improvements, bug fixes, and security updates," he said.

The initial preview release of Fedora CoreOS runs on bare metal, Qemu, VMware, and AWS, on x86_64 only. It supports provisioning via Ignition spec 3.0.0 and the Fedora CoreOS Config Transpiler, automatic updates with Zincli and rpm-ostree, and running containers with Podman and Moby.

The community will be adding more features and supported platforms to Fedora CoreOS. If you want to test Fedora CoreOS, you can download it from the official site, but keep in mind that this is a preview version that isn't approved for production.

SUSE Appoints New CEO

SUSE has appointed Melissa Di Donato as the next CEO of the company to succeed Nils Brauckmann, who will be retiring and leaving SUSE.

"I am incredibly proud of SUSE's progress and growth over the last eight years, which has culminated in it securing independent status," Brauckmann said.

Prior to SUSE, Di Donato was chief operating officer and chief revenue officer at SAP, where she was responsible for the worldwide revenue, profit, and customer satisfaction of the company's digital core solutions.



"SUSE is at the cusp of a historic shift as open source software is now a critical part of any thriving enterprise's core business strategy," said Di Donato in a press release. "We are well-positioned to emerge as the clear leader of this shift, with our ability to power digital transformation for our customers at their own pace and with agile, enterprise-grade open source solutions, edge to core to cloud."

Di Donato also held senior executive positions at Salesforce and was recognized for her contributions to growing global organizations by winning the 2018 Digital Masters Award for Excellence in Commercial Management.

MORE ONLINE

Linux Magazine

www.linux-magazine.com

Linux Administration Focus

<http://www.linux-magazine.com/tags/view/administration>

Getting to Know Firewalld • Ken Hess

Managing a firewall can be a hassle, but it's worse to manage a breach because you didn't have one.

Manage Logs with logrotate

Mayank Sharma

Take charge of your installation's logfiles with logrotate.

ADMIN HPC

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

High-Performance Python – Compiled Code and C Interface • Jeff Layton

Python is one of the fastest growing languages for computing, the number one language for deep learning, and in the top three for machine learning.

ADMIN Online

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

Five Kubernetes Alternatives

Martin Loschwitz

Many admins consider Kubernetes the obvious choice for managing containers; however, don't ignore the highly efficient alternatives just because they are less prominent.

New Features in PHP 7.3 • Tim Schürmann

The new PHP 7.3 simplifies string handling, supports PCRE v2, adds LDAP controls, improves logging, and deprecates some features, functions, and syntax elements.

Software Inventory with PowerShell

Thomas Wiefel

PowerShell easily combines different interfaces for querying or reporting data.

GitHub Blocks Access to Private Repositories in Certain Countries

GitHub announced that it has restricted access to some of its commercial services in countries that are on the US sanctioned list. The countries include Iran, North Korea, Syria, Crimea, etc.

Nat Friedman, CEO of GitHub, wrote on Twitter "It is painful for me to hear how trade restrictions have hurt people. We have gone to great lengths to do no more than what is required by the law, but of course, people are still affected. GitHub is subject to US trade law, just like any company that does business in the US."

Many developers complained that they lost access to their repositories when they tried to access them from these countries. GitHub has blocked access

based on the IP address of the user. It's certainly going to flag many accounts in error. GitHub has created a form that users can complete to review their restrictions. GitHub also confirmed that due to US trade laws, only commercial services are blocked. Users from these countries can freely access pages and public repositories. Users can also access GitHub's self-hosted platform to run their private repositories locally.



Image © Igor Zhuravlov, 123RF

A New Ransomware Targeting Linux-Based NAS Devices

Linux-based network-attached storage (NAS) devices made by QNAP Systems are under a new ransomware attack, allowing bad actors to hold users' data hostage for ransom.

According to *The Hacker News*, independently discovered by researchers at two separate security firms, Intezer and Anomali, the new ransomware family targets poorly protected or vulnerable QNAP NAS servers either by brute-forcing weak SSH credentials or exploiting known vulnerabilities.

The ransomware implementations are called QNAPCrypt by Intezer and eCh0raix by Anomali. Written in the Go programming language, the ransomware encrypts files with targeted extensions using AES encryption and appends an .encrypt extension to each.


For some unknown reason, the ransomware is being merciful to NAS devices located in Belarus, Ukraine, and Russia. "The ransomware terminates the file encryption process and exits without doing any harm to the files," reported *The Hacker News*.



Image © Oleksandr Omelchenko, 123RF.com

Get hyped.

Register for SC19 today and discover why HPC is now.



Calling all HPC researchers, engineers, and students – don't miss the broadest program in HPC, the massive exhibition, networking opportunities, and more. Big HPC adventures await in Denver this fall!

November 17–22, 2019

The International Conference for High Performance Computing, Networking, Storage, and Analysis



SC19

Denver, CO | hpc is now.

Sponsored by:

Zack's Kernel News



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

By Zack Brown

Communicating with Linux During Bootup

The open source development model is geared to maximize everyone's contribution with the two assumptions: If you need something, it's probably a good idea to implement it, and if you implement it, there are probably other developers interested in helping you get your patch into an acceptable state. So the amount of wasted work is probably very low, though I'm not sure how to measure that precisely.

It's definitely the case that sometimes you can put in a lot of work, only to discover that other developers are simply opposed to whatever it is that you need.

Masami Hiramatsu recently tried to extend DeviceTree to pass debugging configuration information into the kernel. This could also be done on the kernel command line, but as Masami himself pointed out, "the kernel command line is a command line." It's not the ideal format for passing structured data.

His excuse was that the command line actually didn't support passing all the information he wanted to pass. So, why not look for something else? And since DeviceTree already existed for the express purpose of passing structured data into the kernel, making a small extension for this use case seemed like a perfect fit.

He posted his patch in good faith and immediately ran into a fairly common syndrome in Linux kernel development: People need the better thing that's also harder, so no one can do anything in that region of code except implement the good thing.

No easy hacks allowed! Frank Rowand replied to Masami's initial patch submission, pointing out some of the problems. DeviceTree was not something to be extended on a whim. It was a community standard with many corporate stakeholders. And when it comes to corporate stakeholders, you are likely to see each of them salivating to add their own favorite extensions in order to hit some short-term milestone or other. Even though the official standard was only

version 0.2, it was still something that needed careful thought and policing.

Frank explained that DeviceTree existed to describe hardware to the kernel that could not otherwise be discovered. It was not a configuration tool. Its purpose was specifically to free the kernel from having to be hard-coded with tons of hardware descriptions and instead to pass that data in during bootup.

The possibility of extending DeviceTree in the way Masami wanted, had actually been discussed and rejected at the Linux Plumbers Conference in 2015, Frank said.

And this was the problem – lots of people wanted to pass configuration and other data into the kernel at boot time, using something better than the command line. But, Frank said, DeviceTree was not the solution. Yes, it was a mechanism for passing data into the kernel; but it had a specific and inviolable purpose. If Masami and others wanted to pass configuration data and other forms of data into the kernel at boot time, they would have to develop a full-featured and generic solution of their own. They couldn't extend DeviceTree.

So there you have it: A policy statement – something that might be overridden eventually. But for the moment at least, that was Frank's position, and the position of the main DeviceTree stakeholders.

But no one ever takes it lying down, the quiet wheel never gets the grease, and you don't find out if a standard is truly locked down if you don't do your best to bust it open.

To some extent, Masami simply pleaded. It seemed obvious to him that the command line was a terrible data conduit, and that standards existed to be extended. He knew already that DeviceTree was not currently used for configuration data, but he said, "Can't we talk about some future things?" And when he offered up as an unrealistic suggestion the possibility of encoding large blobs of data in some sort of ASCII-safe format and including it on the command line, he was surprised when Frank said this

would be acceptable – and certainly preferable to extending DeviceTree.

But Frank really wanted “people to come up with an additional boot-time communication channel or data object to support this use case. So far, no serious proposal that I am aware of.”

It was unpleasant. Masami’s initial patch had been 500 lines of code; while implementing a general-purpose boot-time communication channel would be some kind of ungodly amount of code, with many subtleties, security considerations, bizarre-yet-crucial use cases, and so on.

But there were other reasons not to use DeviceTree. Rob Herring said that not all bootloaders supported it; for Masami’s purpose, it would definitely be best to use a mechanism from which everyone else could benefit. Rob said he would actually be fine with extending DeviceTree, if it could also be extended to work with all bootloaders.

Eventually Masami did abandon his patch in favor of trying to develop a new communication channel, as Frank had originally suggested.

Not all developers are as willing to change direction. Richard Gooch maintained DevFS for a very long time as a separate patch from the kernel, but was never able to win over the other kernel maintainers, particularly Greg Kroah-Hartman who supported his own udev project. I don’t think we can generalize from Masami’s experience to say that nowadays disputes over technical direction are handled more agreeably, but it’s nice to see it when it happens.

Pruning SuperH

Where oh where has SuperH gone – the RISC chip for embedded systems that came out in the 1990s? Recently Christoph Hellwig pointed out that the maintainers of the Linux port were not so responsive, and there hadn’t been any pull requests coming from that project in quite awhile. He suggested that the architecture was dead and could just be removed wholesale from the kernel.

This kind of bell-ringing is always only partly a genuine suggestion to ditch a kernel feature – it’s also intended to alert any users of that feature that they should speak up quickly. If a feature has users, it won’t be removed no matter how old it is. But if no one speaks up for it, then out the window it goes.

In fact, it’s true that even very old features will be kept if they still have users. But it’s also true that even a very popular feature will be removed if it constitutes a security hole. Security is always the top priority, even to the point of accepting massive slowdowns and other inconveniences. That’s one of the differences between Microsoft’s development philosophy and that of the open source community. Microsoft will allow long-term security flaws in their operating system and software, just because users like those features. Linux won’t.

But security issues were not one of the things pushing SuperH out the door. It just seemed to Christoph that the code was rotting, and there were no users.

John Paul Adrian Glaubitz was the one who came to SuperH’s rescue, saying that the architecture – at least version 4 of the chip – was indeed still available as a Debian package and was most definitely still in use. Maybe the code was getting old, but the compiled kernels would still run.

There was some skepticism. Adam Borowski suggested that merely being a Debian package didn’t mean anyone used it. He pointed to the Debian Popularity Contest package, which showed a bottom-of-the-barrel score for SuperH; and he said he hadn’t heard anyone talking about SuperH on any of the Debian mailing lists.

But John Paul proved his point, giving a link to a recent bug report from an industry user at Dell [1].

So that settled it. No security issues plus an active user meant that some support for SuperH would remain in the official Linux kernel. But that didn’t mean absolutely all of SuperH had to remain. Arnd Bergmann in particular suggested that the SuperH maintainers should go through the code and identify which parts were truly necessary and which could go. Obviously SuperH v4 had users, so it would stay. But maybe other versions (he mentioned version 5) could be removed, since those parts “don’t build, or are incomplete and not worked on for a long time.”

Rich Felker and Yoshinori Sato were the official SuperH maintainers. Rich came into the discussion now, saying that version 5 could definitely be removed. It had already lost support in the C compiler, and he seemed to recall that the hardware chip itself had never really been available. Rich could also see a

bunch of other infrastructure-related code that was probably broken – if it had ever run at all.

Rich also said that he had originally maintained the architecture as part of his job, but as his situation had changed, he was no longer being paid to work on it. So his motivation had declined in favor of other projects. But he said he still wanted to be involved in any remaining team of SuperH maintainers. In fact, he had some patches queued up and ready that he just hadn't gotten around to, because they'd require fixing other parts of the code that had also rotted.

Yoshinori, meanwhile, said he still had some actual SuperH hardware, specifically versions 2, 2A, and 3, and he'd be happy to update the code for those versions.

So SuperH will get a little love, even though only a handful of people still use it. Dead code will be removed. Living code will be fixed and stabilized, at least to the point of being easily maintainable over the long term.

I almost said that eventually the last users would move on, and the architecture would be removed at last. But this is not certain at all. You can never tell what ancient hardware will develop a hobbyist following and live on through the years, even coming to exist only as a software simulation with a modern kernel running both below (on the host system) and above (on the simulated hardware). Maybe SuperH will become one of those some day and remain in the official kernel tree for decades to come.

Bug Hunting for Steam

Linus Torvalds sometimes involves himself in tracking down a specific bug, either because it's code he particularly cares about, or for some other reason. Recently he helped track down a bug in the official 5.1.11 kernel release, because it affected a large group of users – all those running Valve's Steam platform for gaming.

It's unusual to have such a significant breakage slip through the release candidate process, but it does happen. This time, Pierre-Loup A. Griffais reported a lot of discussion traffic on GitHub and reddit, from users who couldn't run Steam on Linux 5.1.11.

The entire bug hunt, including pushing the fix out to both the development and stable Linux kernels, took less than two days. Linus shepherded the process along

himself and just by his presence was able to ensure that every step went as quickly as possible.

At first, Greg Kroah-Hartman, the stable kernel maintainer, asked for more information, specifically whether Pierre-Loup had been able to bisect the kernel git repository to identify the exact patch that caused the problem.

But this wasn't as easy to do as it might seem. Linus pointed out that not everyone was experiencing the same problem or else experienced it only intermittently. This suggested it was a timing issue. But if it was intermittent, it might make it harder to find a consistent way to reproduce it during the bisection process.

Since the advent of git, bisecting the tree has been one of the primary ways kernel developers identify the origins of bugs in the code. It's an incredibly convenient technique and, in many cases, can even be automated.


In this case, the problem turned out to be a security fix from Eric Dumazet. This complicated things slightly, because you can't simply abandon a security fix. Eric's patch prevented a denial-of-service attack. Without the patch, a hostile user could trigger a kernel panic by overloading the system with networking requests. But the patch also made it so Steam would no longer work.

In fact, it turned out that Eric's patch could be replaced by something much simpler – and that wouldn't break Steam – if the patch simply did nothing but tackle the security flaw and left out some of its other subtle details. And so the new fix went into the tree – both Linus's tree and Greg's and was made available to users right away.

It's interesting, because the pace of the bug hunt was so rapid. During normal development, with release candidates coming out once in awhile, no one worries about getting fixes into the tree as soon as possible, because only kernel developers use release candidates. The fact that this bug affected an official kernel release, meant that real users would be experiencing real problems. In that situation, Linus will always try to move as fast as possible, consistent with actually getting the fix right. ■■■

Info

[1] SuperH bug report: <https://marc.info/?l=linux-sh&m=155170489401832>



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



The prpl Foundation

Making Connections

Mirko Lindner of the prpl Foundation discusses how vendors and chipset makers are coming together to make IoT more interoperable and secure. *By Mayank Sharma*

The prpl Foundation [1] is a non-profit consortium of hardware vendors, software vendors, service providers and developers. It describes itself as an open source, community-driven organization whose members [2] collaborate to develop standards and software to enable security and interoperability of embedded devices. We talk with the group's recently appointed Program Director, Mirko Lindner, to better understand the objectives of the foundation and how they'll impact end users like us.

Mayank Sharma: Let's start with your role in the project and how you got involved with prpl. You were once the Project Lead at OpenMoko?

Mirko Lindner: That was my first job in an open source company, basically running the development of the open source UI for the phone. After that I went to work with Vodafone in different capacities – first built some mobile applications, then joined their innovations team and ultimately ended up in the fixed-net department, which is the team that builds all of the routers and gateways. There we realized that if you look at how ISPs and service providers in general build devices, especially in the router space, that they work with a large number of different suppliers and there is no standardization.

Basically think back 15-20 years where every mobile phone manufac-

turer had their own operating system. Applications like WhatsApp weren't even on the horizon and the best online messenger was the Blackberry Messenger. The situation on embedded devices today is very similar. We did a huge analysis and ultimately we came back with the need for standardization and open source development. Out of that we started talking to the prpl Foundation, and we worked out how that effort could manifest in a project that would involve not only Vodafone but many other service providers as well as software vendors, hardware vendors, and chipset makers.

I worked for a while with the prpl Foundation from Vodafone, and then I left Vodafone because I wanted to explore Asia and work a little closer with hardware manufacturers. Then about 15 months ago, the prpl Foundation was looking for someone to take up more of an operational position, and I joined prpl part time as Program Director to bring the different working groups and projects into shape. This July I took the next step, switching perspective yet again, switching to the software vendor side. I joined one of prpl's members, SmartRG, who have made quite a substantial contribution to the foundation both in resources and product development. My role is to serve as a bridge between a vendor with a very strong background in open source and a foundation on the other side. So now I do both – I take some

daily work items here at SmartRG, but also run the operational side of things in prpl for many of the working groups. I try to think of myself more as a product manager than a program manager.

MS: What prompted the formation of the foundation and how has it evolved?

ML: When the foundation was created about six years back, it was created with a strong focus on silicon and software that powers systems on a chip. But over the last two years, the prpl Foundation has become more of an Open Source CPE (customer-premises equipment) software foundation. If you look at our current membership, you'll find lots of hardware manufacturers from Asia and all around the globe that build devices for ISPs, you'll find the ISPs themselves and of course software vendors alongside the chipset vendors. So it's more about CPE, which is basically an evolution of the foundation.

The problem was that if you want to move a whole industry, one service provider, one chipset manufacturer, or one software vendor is never big enough to have the scale of impact you need. So what you need is basically an alliance, a place where people and companies can come together and solve common problems in an open and standardized way. That's why the prpl Foundation is so powerful, because it is a place where technical peo-

ple from different companies who usually wouldn't meet each other can [discuss] what their problems actually are. That's the whole point of the prpl Foundation.

MS: Why would competing companies want to get together and collaborate?

ML: Very often those companies are not in direct competition but rather have the same problem that they all need to solve. Like WiFi for example. So they all need to make sure the WiFi connection is stable and that you can connect 20 devices to your router in the home and still do your WhatsApp'ing and chatting while watching Netflix or IPTV streams. Those are common problems that every company's R&D department has to solve before they can productize. And the reason for them to come together is then they can solve this problem once and everyone can benefit from the common solution. Each company can then focus on building features that only they have, what we call the "differentiating features," where you can see that one company

focussed on routers for IPTVs, while others are focussed on the gamer market and things like that.

MS: How do your efforts translate to bettering the experience for the end user?

ML: For the end users two things are going to happen. One is that service quality is going to improve. For example, let me stick to WiFi as that is a problem that is relevant to all users. Right now you have a maximum of 100-200 engineers per company that are trying to solve all of these WiFi problems that we are talking about – common things like speed of client's connection. But we are also talking about less common use cases – what happens if I have three different smartphones and my IPTV setup box connected and someone begins a download? In this case, whose service gets downgraded in the service performance?

With prpl, we might have thousands of engineers working on the same set of problems, that will definitely translate

into an increase in quality and security. It might take two to three years before it reaches a large number of customers, but by then you'll feel the difference.

The second thing that comes right after is more innovation. You'll see more new services being rolled out, more integration between different companies. Basically something similar to the smartphone ecosystem where you can think of third-parties being able to roll their services on top of devices that you already have in your home. If you have security services in the home, or parental control services, those usually come with their own set of devices. That means if you add even one more device to your home that increases your electricity bill and takes space; it's just clunky and inefficient. All of that is going to go away and we're going to see a service ecosystem in the home that can run on already existing devices. So I think there are lots of advantages to what we do, and I do hope we'll be able to communicate to consumers how that is better for them, because that's why we're doing this.

HAPPY SYSADMIN DAY!

ADMIN
Network & Security

SPONSORED BY
Linux Professional Institute

10 Terrific Tools
FOR THE BUSY ADMIN
2019 EDITION

Yet one more killer collection of time-saving tools!

Find a free tool to help you

- Create an IP blacklist
- Analyze logfiles from the command line
- Edit file names from a directory-style listing
- Discover the best path for a DNS query

Bonus articles

- Sudo Voodoo
- Deleting Old Kernels

WWW.ADMIN-MAGAZINE.COM

US\$ 7.95

10 Terrific Tools
for the
Busy Admin

FREE Download!

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

2019 Edition
Sponsored by LPI

MS: One of your missions is to “enable the interoperability of embedded devices.” Could you explain what these words mean to prpl?

ML: Probably the most obvious example is something we call prplMesh [3]. You see over recent years how different companies came up with a combination of several devices to give you perfect WiFi coverage in your home. We used to have a single router creating the WiFi for the whole home, now you have a main access point and a couple of satellite devices that extend that main WiFi. Until recently, you could only buy the whole pack from one company.

What the Wi-Fi Alliance has done is come up with the standards, working with all of these companies, that allow you to mix and match solutions and make the different WiFi satellites talk to each other. So you can get a device from your ISP that supports mesh, and then you can decide if you buy the mesh extender from Company A or Company B. prplMesh is the open source reference implementation of exactly that. So we are taking this one step further and saying that instead of only having the standards, let’s then implement those common features. Let’s make them available in open source to every vendor and anyone who wants to build anything on top and thereby increase the interoperability.

MS: What are the major action areas that prpl’s working on to get us to that ideal world?

ML: Prpl’s efforts largely fall into one of two areas. One is API definition and API standardization work. Here our low-level API efforts, we will simplify and dramatically shorten the time it takes to integrate different chipsets and new hardware features while our high-level API layer will enable software modules to be developed more quickly and exchanged more easily between vendors and solutions. The other work area is the open source implementation of those APIs and specifications.

Those are the two main groups and these unfold into different flavors. One of these, on the implementation side, is prplMesh. Then we have a project

where we are taking OpenWRT [4] and basically adding features that we then upstream as much as possible. This allows vendors of commercial software stacks for CPE devices to work directly with the open source project without having to build their own set of tools around OpenWRT. This enables more software and more patches and fixes and security fixes from company development into the open source project that they all utilize for their products.

MS: Your objectives are translated into action by the Working Groups. Could you give an overview of what these are and why they are important?

ML: We have five or six active working groups depending on how you count. We have one for each of the APIs that I’ve just mentioned. Then there’s prplMesh, which is one of the largest working groups. Then we have a working group for our OpenWRT related efforts and CPE stack. And a Security working group [5] that’s specifically looking at how we can make the home more secure and how we can use the devices like Internet access gateways and WiFi routers to make the whole network more secure.

MS: The Security Working Group gets the most attention from what I could tell.

ML: It depends on who you talk to. There’s a lot of talk about the Security Working Group because it’s a very important topic for all parties involved, but prplMesh and prplWRT [6] are very prominent projects. Maybe less so publicly, but certainly when you look at the size of these working groups and the amount of work that is being done inside them. For security to receive the attention it needs and deserves, it takes a lot of education and conversations. We have to make sure everyone puts security as their number one requirement and that means we have to talk to everyone. In contrast, the other working groups involve more negotiations between the companies, so it’s a bit more collaboration and writing software. But all of our working groups are well-staffed and quite active.

MS: Consumer use of smart devices, particularly with IoT, has already come a long way in the last couple of years. What are the major challenges that you think still need to be addressed to usher in the next generation of IoT devices?

ML: I can only speak for the router and the CPE side of things and not so much about the individual end devices. The biggest challenge is the integration of different services and how you can ensure, a lightbulb, for example, can be used by any software services that you choose as a consumer instead of you having to buy into a whole ecosystem. We need those cross-company ecosystems that you find in most other industries. I think that will unlock IoT going forward because the innovation cycle of a software company is completely different from that of a hardware company, and right now those two are too tightly coupled. That’s what we are working on, and once we unlock that you’ll see a whole new breed of hardware.

MS: What can we expect from prpl in the near future?

ML: In the near future, I do hope that we get to harmonization on the CPEs. So from a consumer perspective, you’ll find more third-party services on the devices at home, dramatically increased security, and an improved overall customer experience. ■■■

Info

- [1] prpl Foundation <https://prplfoundation.org/>
- [2] Foundation members <https://prplfoundation.org/current-members/>
- [3] prplMesh <https://prplfoundation.org/prplmesh/>
- [4] OpenWRT <https://openwrt.org/>
- [5] prplSecurity <https://prplfoundation.org/working-groups/prplsecurity/>
- [6] prplWRT <https://prplfoundation.org/working-groups/prplwrt-carrier-feed/>



ORDER NOW

Get 7 years of *Ubuntu User*
ON ONE DVD!



THE COMPLETE
UBUNTU
user
ARCHIVE

Over
3,000
PAGES!
7 GREAT YEARS
OF UBUNTU
USER

Searchable DVD!
All Content Available in Both HTML and PDF Formats

A cartoon mole wearing a magnifying glass, looking at a DVD disc on the ground.A cartoon squirrel wearing a red shirt, holding a yellow measuring tape.

shop.linuxnewmedia.com



Rclone as a helper for external backups

Rsync for the Cloud

Backups for critical data have to withstand major disasters. Rclone, a descendant of Rsync, enables easy backups to the cloud and supports many services. *By Ferdinand Thommes*

We all live in the hope of never getting into a situation where our belongings are destroyed by theft, fire, water, or natural disaster, but there are no guarantees. If you are serious about preserving your important digital data, you should secure it outside your home, preferably in several places at the same time. The Rsync-based Rclone [1] tool helps with off-site backup scenarios.


Rclone is a command-line program that performs one-way synchronization between locally defined datasets and a cloud. If you run it a second time, it searches the local filesystem for changes and then incrementally uploads only the changes to the cloud of your choice. Rclone implements most of Rsync's options and syntax. In addition, it offers additional commands to optimally support the individual services. Rclone offers a graphical front end for people who prefer not to work at the command line.

Rsync's two-way synchronization is not part of Rclone's feature set. This means that it deletes any locally-deleted files from the cloud, but not vice versa. In addition, the program does not save older versions when changes are made to files.

It overwrites changed files, unless the selected service has its own versioning system, such as

Google Drive or Dropbox. One way to avoid this is to use the `backup-dir` function, which moves changed or deleted files to a separate directory [2].

The tool, written in Go, which can be found in the archives of most distributions, supports numerous cloud service providers and protocols, including the open source clouds Nextcloud, ownCloud, and OpenStack Swift, as well as the commercial offerings Amazon S3, Box, Dropbox, Google Drive, pCloud, and Microsoft OneDrive [3].





```
n) New remote
d) Delete remote
r) Rename remote
c) Copy remote
s) Set configuration password
q) Quit config
e/n/d/r/c/s/q> n
name> gdrv.
Can't use "gdrv." as it has invalid characters in it.
name> gdrv2
Type of storage to configure.
Enter a string value. Press Enter for the default ("").
Choose a number from below, or type in your own value
1 / A stackable unification remote, which can appear to merge the contents of several remotes
  \ "union"
2 / Alias for an existing remote
  \ "alias"
3 / Amazon Drive
  \ "amazon cloud drive"
4 / Amazon S3 Compliant Storage Providers (AWS, Ceph, Dreamhost, IBM COS, Minio)
  \ "s3"
5 / Backblaze B2
  \ "b2"
6 / Box
  \ "box"
7 / Cache a remote
  \ "cache"
8 / Dropbox
  \ "dropbox"
9 / Encrypt/Decrypt a remote
  \ "crypt"
10 / FTP Connection
  \ "ftp"
11 / Google Cloud Storage (this is not Google Drive)
  \ "google cloud storage"
12 / Google Drive
  \ "drive"
```

Figure 1: In the interactive configuration dialog, create a name for the remote access and select the desired service or protocol from the list.

Rclone also supports the HTTP, FTP, SFTP, and WebDAV protocols, so web and file servers, as well as NAS devices, can be integrated as back ends using standard protocols. The program is available in variants for x86, ARM, and MIPS architectures. Not only for Linux, Rclone also is available for macOS, Windows, BSD, Plan 9, and Solaris, as well as NAS servers from QNAP and Synology.

Rclone v1.48 is the current version. Your distribution may use an earlier version (e.g., Debian “Stable” uses version 1.35). However, you can install the latest version of Rclone directly with:

```
curl https://rclone.org/install.sh | sudo bash
```

You may have to install cURL first by typing:

```
sudo apt install curl
```

Table 1: Important Rclone Commands

Command	Function
rclone config	Opens the configuration
rclone copy	Copies data from the local system to cloud storage
rclone sync	Synchronizes source and target
rclone mount	Mount the drive locally into the filesystem
rclone move	Moves data from source to destination
rclone delete	Removes the contents of a path
rclone purge	Removes the path and its contents
rclone ls	Lists all objects in the path with a size specification
rclone md5sum	Creates md5sum checksums for all objects in the path
rclone size	Number and total size of all objects in <Remote>:<Path>
rclone cryptcheck	Checks the integrity of an encrypted remote

These and other commands show their options after entering -h.

Creating Remotes

You control Rclone with the command of the same name (rclone) and an attached command line including various options (see Table 1). To get started, call the configuration with rclone config. The setup wizard offers a number of options: Use S to set a password and Q to quit the program.

Then use N to set up the first remote connection (“remote” here refers to the connection to a specific folder on a specific cloud service) – select this option first. Since Rsync lets you configure multiple remotes, make sure you use a meaningful name in the first step. Then select the desired service or protocol using a number (Figure 1).

Google Drive Tested

In my hands-on test, I set up Rclone with Google Drive as a proprietary provider on a laptop. Of course, I do not want to store private data on Google’s servers without encryption. Therefore,

after an initial open test, I encrypted the data to be backed up before uploading with Rclone.

In the test, the remote repository is named gdrv. You then need to select the corresponding number for the desired service (*Google Drive* here) from the list by entering 12. For each provider and each protocol, the documentation offers a detailed manual [3], including Google Drive. The following two queries after client_id and client_secret can be skipped by pressing Enter.

Then decide whether Rclone should have full access to all files of the online repository or only access its own uploaded data. In addition, access can be restricted so that Rclone can only read. I selected full access for the

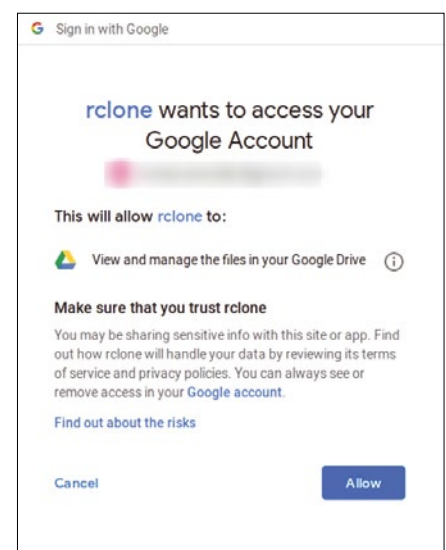


Figure 2: The dialog gives Rclone access to your Google Drive instance. If the connection fails to open, entering the link manually can help.


```
ft@dev:~$ rclone copy meintext.txt gdrv:meintext.txt
ft@dev:~$
ft@dev:~$ rclone copy -P meintext2.txt gdrv:/texte/meintext2.txt
Transferred:      2.803k / 2.803 kBytes, 100%, 58 Bytes/s, ETA 0s
Errors:           0
Checks:           0 / 0, -
Transferred:      1 / 1, 100%
Elapsed time:     49.4s
ft@dev:~$
```

Figure 3: Rclone works very slowly without optimization. Google will slow you down here unless you specify additional parameters.

Accelerating Rclone

You can specify the number of parallel file transfers via `--transfers`. For example, to back up a folder with 50 files, specify `--transfers=25`. As a rule, a value that is too high does not do any damage.

The `--checkers` parameter defines the number of test processes running in parallel. The processes identify all the files to be uploaded and compare them with what is already in the cloud store. In this way, the system excludes unchanged data from the transfer. The value is based on the `--transfers` parameter.

Using `--drive--chunk-size`, you can specify how much memory Rclone uses when uploading. The developers recommend 16MB, which you specify as `--drive--chunk-size=16384k`. However, experimenting with higher values will not do any harm. If you also specify the `--progress (-P)` parameters, you can see the exact values for the upload's duration. Rclone provides even more information with the `-v` or `-vv` options.

test with the `l` option. Then skip the question about the `root_folder_id`, as well as the question about the `service_account_file` – Rclone can determine the root itself.

Can It Do That?

Finally, it doesn't hurt to look at the advanced configuration for Google Drive to learn about the additional options detailed in the documentation [4].

You are then taken to the point where Rclone can access your Google Drive account. To do this, say `Y` to the `auto config` prompt, and open your Google login in a browser window. After logging in, confirm that Rclone is allowed to access your drive (Figure 2). In the terminal, a *Got Code* message will appear; you can now close the browser again.

Having said this, the automatic method did not work for me. It was possible to register with Google, but the Rclone connection did not happen. A second attempt where I manually entered the URL, `http://127.0.0.1:53682/auth`, was then successful and let me complete this remote's configuration. To run Rclone in headless mode (i.e., without the option of opening a browser), create an API key in JSON format via a Google Service Account [5] and enter it in the advanced configuration.

Listing 1: Testing Acceleration Options

```
$ rclone copy -P --transfers=10 --checkers=10 --drive-chunk-size=16384 beautiful_images gdrv:/images/beautiful_images
```

This completes the setup wizard. You now have the choice of configuring additional remotes or protecting the configuration with a password. To do this, restart the configuration and select 5 for “set configuration password”. Instead of a password, you can also make the configuration file readable for only your user with:

```
chmod 600 ~/.config/rclone/rclone.conf
```

First Test

An initial test will show you if everything worked out so far. Select a small local text file and copy it to the cloud repository with:

```
rclone copy mytext.txt gdrv:mytext.txt
```

This may take a while, because Rclone works very slowly without options. My text only reached a transmission speed of 58bps without parameters. I learned this by specifying the parameter `-P` (Figure 3). The “Accelerating Rclone” box explains parameters that you need to pass in to accelerate the Rclone commands.

Step on the Gas

In order to investigate the effect of the acceleration options, I had Rclone back up a folder with 10 images with a size of 40MB in a folder to be created on the target (Listing 1). The average data rate during the backup fluctuated in the test at around 500KBps; doubling of the chunk size returned about 620KBps.

Then I added three files with a size of 12MB to the locally stored image folder and restarted the backup with the same command. Since Rclone now only saved the newly added data, the process was completed in about 25 seconds. When restoring Google Drive to the local instance, no parameters are required. The command is otherwise the same as the copy command; you only have to swap source and destination.

Copy and Sync

In addition to copying, Rclone also lets you synchronize files between a client and a server. The difference between copy and sync is that the latter compares the target with the source file by file and deletes everything on the target that is not in the source directory from which you issue the command.

If you call a sync command from the wrong directory, you risk losing important data. The sync option should therefore be tested with the `--dry-run` parameter. copy does not delete any data locally; however, the process overwrites the backup in the cloud. If you want to keep a backup, you can bypass this process by using the backup directory mentioned above.

Listing 2: Mounting with FUSE

```
$ rclone mount gdrv:Pictures ~/cloud
```

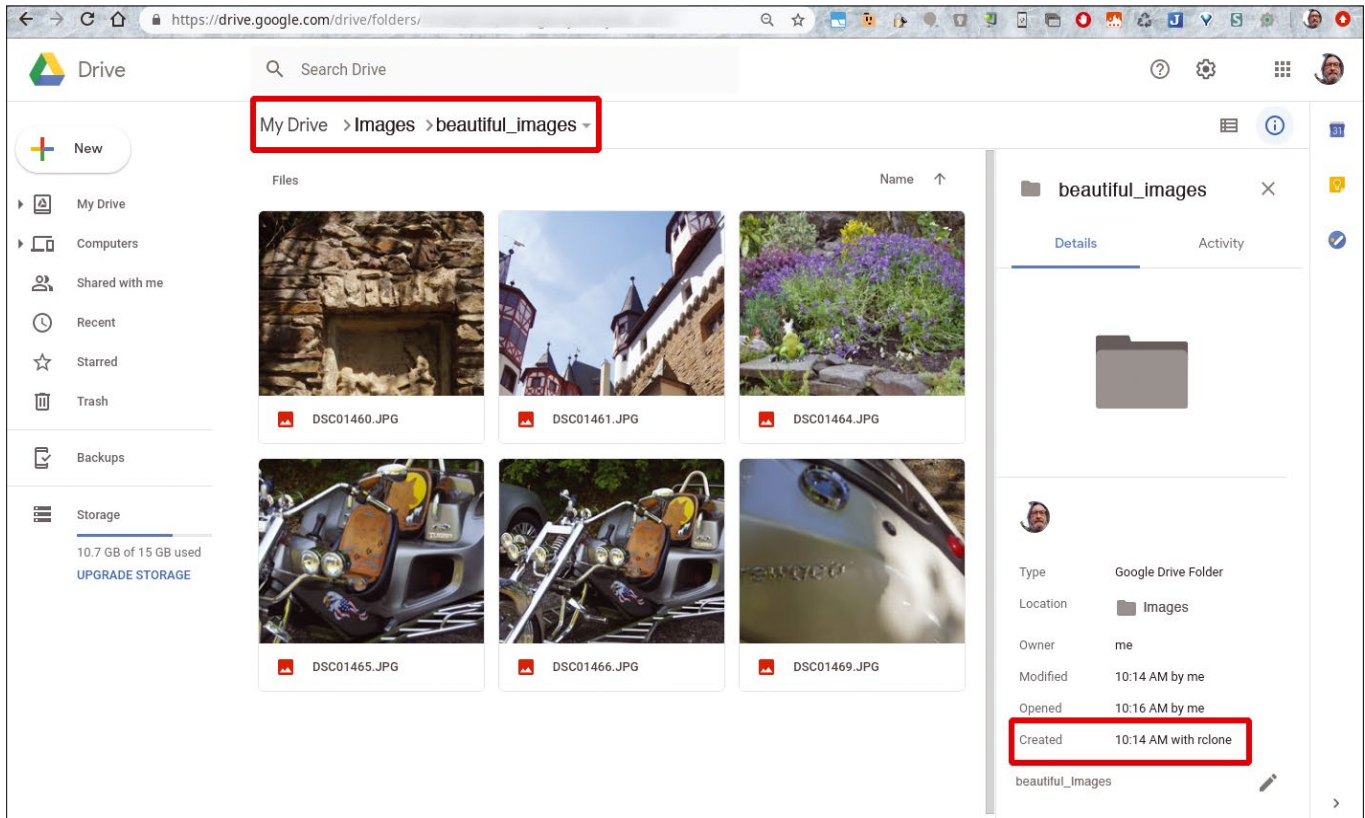


Figure 4: After the upload, you can check in the browser whether Rclone has uploaded the data completely.

Another way to connect to Rclone is to mount Google Drive as a filesystem in your home directory using Filesystem in Userspace (FUSE) [6]. In my example, I created the `cloud/` folder there and then mounted the previously configured Google Drive account (Listing 2).

Browser or Terminal

After copy or sync actions, you can see in the browser whether the transfer was successful. Simply call the corresponding ser-

vice's web front end (Figure 4). For a faster response, use the following command at the command line:

```
rclone lsd gdrv:/Pictures
```

This call tells Rclone to display the last action in relation to the specified folder.

For more information, use

```
rclone ls gdrv:/Pictures
```

which lists each transferred file and its size (Figure 5). If you omit the folder for both commands (assuming that you selected the `1` option for full access to cloud storage during configuration), Rclone will list all of the drive instance's folders or files.

Encryption Must-Have

As previously mentioned, any data you store on Google Drive should be encrypted to protect your privacy. Now that you know that Rclone is working as expected, it's time for automatic encryption during the upload [7].

Call `rclone config` again, and select `9` to create a new remote of the type *Encrypt/Decrypt a remote*. The configuration is similar to before. In addition to

```
ft@dev:~$ rclone lsd gdrv:Images
-1 2019-03-24 10:32:28      -1 beautiful_images
ft@dev:~$
ft@dev:~$ rclone ls gdrv:images
4161536 beautiful_images/DSC01423.JPG
3964928 beautiful_images/DSC01424.JPG
3735552 beautiful_images/DSC01425.JPG
4259840 beautiful_images/DSC01460.JPG
3833856 beautiful_images/DSC01461.JPG
4751360 beautiful_images/DSC01462.JPG
6193152 beautiful_images/DSC01463.JPG
5537792 beautiful_images/DSC01464.JPG
3604480 beautiful_images/DSC01465.JPG
4030464 beautiful_images/DSC01466.JPG
4325376 beautiful_images/DSC01467.JPG
4358144 beautiful_images/DSC01468.JPG
2883584 beautiful_images/DSC01469.JPG
```

Figure 5: If you are dealing with machine-readable-only data that you can't check visually, the terminal offers an easier option than the browser.

passwords, you are also prompted to decide whether the system should encrypt the file and folder names (Figure 6). During the test, I named this remote `secret`.

As the target, I specified `gdrv:Images` and told Rclone to encrypt file names but not folder names. Then I used the command from the first line of Listing 3 to save the test data that I had previously deleted from the online storage for more clarity. After that, the encrypted file names were visible in the terminal (line 2), as well as in the browser (Figure 7).

Working Graphically

If you want to use Rclone with a graphical user interface (GUI), the Rclone Browser [8] front end lets you do so; it covers most of the Rclone functions (Figure 8). The program lets you upload and download data and mount Google Drive in your home directory, supporting many parameters. If several remotes are used, you manage them together within the GUI.

Rclone Browser can be found in the package sources of Debian and Ubuntu. However, if you have built Rclone yourself, you will also want to compile the latest version of Rclone Browser. The package manager will again try to set up the Rclone version from the repository together with the GUI, which will tend to lag behind the current state of development.

The GUI turned out to be bottleneck in the test. In tests over several days, Rclone Browser only achieved a little more than half the throughput that I achieved at the command line. I

couldn't find an explanation for this. You should therefore always start a big initial upload via the terminal.

Conclusions

In my lab, Rclone proved to be a reliable backup client for the cloud. When uploading data to Google Drive, however, Google slows down the connection despite sufficiently fast data lines. Apart from this, Rclone offers an easy way to store

```
[crypt]
type = crypt
remote = gdrv: Bilder
filename_encryption = standard
directory_name_encryption = false
password2 = *** ENCRYPTED ***
-----
y) Yes this is OK
e) Edit this remote
d) Delete this remote
y/e/d> [ ]
```

Figure 6: Rclone uses its own remote to encrypt your data during the upload to provide protection against unauthorized access even on proprietary clouds.

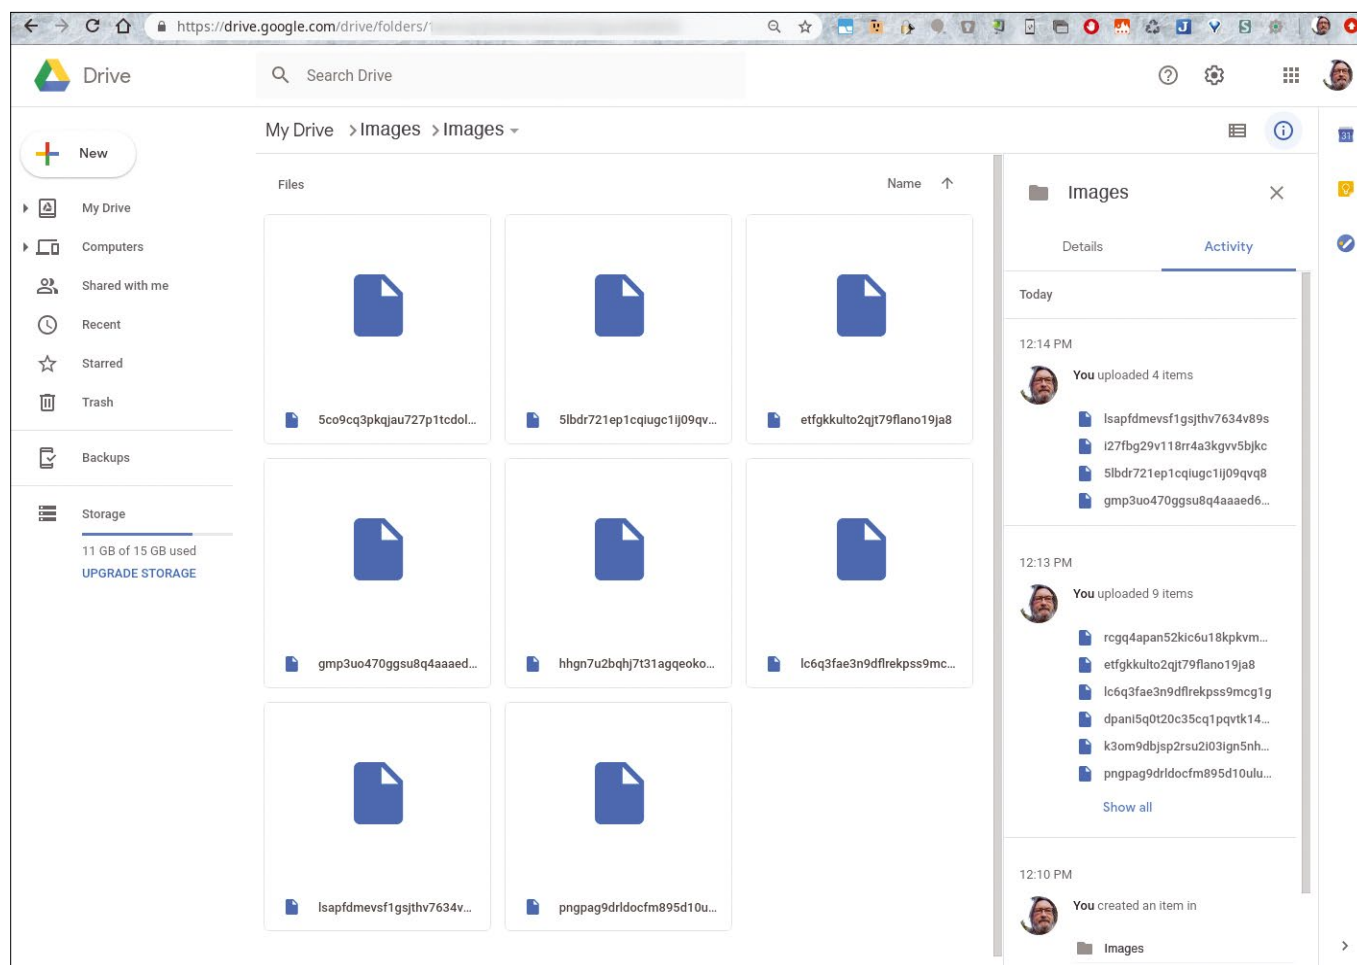


Figure 7: After uploading with encryption enabled, the browser and thus the cloud storage provider only see cryptic file names.

important data encrypted with various protocols with various providers.

Rclone is already integrated in many other projects [9]. It can be called by Cron and systemd timers or scripted (jobber.sh [10] provides examples of this). In addition, Rclone can synchronize data between two servers without detouring via a local instance. I was only able to scratch the surface of this powerful application's options here; the detailed documentation delves deeper into the matter. And if you are looking for alternatives to Rclone, check out odrive, GoodSync or RaiDrive. ■■■

Listing 3: Making Encrypted File Names Visible

```
$ rclone copy_beautiful_images/ secret:Images
$ rclone ls gdrv:Images
3834832 Images/5co9cq3pkqjau727p1tcdolha8
4359248 Images/51bdr721ep1cqiugcl1j09qvq8
3965936 Images/bpbpke61suvv7p5tdcq5ac080c
[...]
```

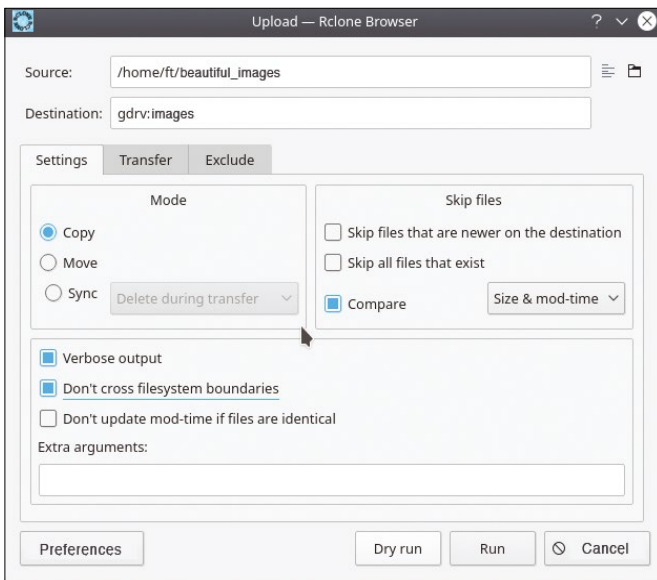


Figure 8: When uploading, Rclone Browser supports the most important options of the command-line application.

Info

- [1] Rclone: <https://rclone.org>
- [2] backup-dir: <https://rclone.org/docs/#backup-dir-dir>
- [3] Services: <https://rclone.org/docs>
- [4] Google Drive: <https://rclone.org/drive>
- [5] Google service account: <https://console.developers.google.com/projectselector2/iam-admin/serviceaccounts?pli=1>
- [6] FUSE lets users with default privileges mount filesystems in the system's folder structure. The technology is generally used for mobile mass data storage devices or network drives.
- [7] Rclone crypt: <https://rclone.org/crypt>
- [8] Rclone Browser: <https://github.com/mmozeiko/RcloneBrowser>
- [9] Rclone projects: <https://github.com/ncw/rclone/wiki/Third-Party-Integrations-with-rclone>
- [10] jobber.sh: https://github.com/wolfv6/rclone_jobber

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?

DIGITAL & PRINT SUBSCRIPTIONS



SPECIAL EDITIONS





Backup with restic

Data Safe

Many users still find it difficult to regularly back up their data. Thanks to restic, and its graphical front end, Restatic, a backup is quite easy to manage – even for a beginner. *By Erik Bärwaldt*

Creating data backups for emergencies has always been one of the most unpopular tasks in the home office. The available applications mostly target enterprise use and spoil the fun with feature overkill, or they are simply too complicated to use for occasional backups. A small command-line program by the name of restic [1] proves that there is another way, and it is suitable for many different scenarios.

Restatic is a graphical user interface (GUI) for restic designed for home users who don't want to mess with a command-line interface. The Restatic front end is currently under development, and you'll need to install it separately.

Restic can store data on the local PC, as well as on a computer on the home network or in the cloud. It encrypts all the data with AES256 regardless of where they are stored, ruling out viewing by unauthorized third parties. In addition, the software is extremely fast as it eliminates any graphical overhead, and this makes it interesting for occasional users, too.

You can pick up restic either from the project's website or from the software archives of almost all the major distributions. Since the application is written in Go, you need a compiler to install it; on Debian/Ubuntu and its derivatives you will find this in the *golang-go* package. After the install, you call the software with the `restic` command at the prompt to see an overview of the available parameters.

Backup

A local backup with restic requires only two steps. First, create a repository for the datasets to be archived (Listing 1, line 1), and then save the data in the newly created archive (line 2).

When creating the archive, you can specify a password for the repository, which you must confirm by retyping. During the backup, restic may request the password for the newly created archive and will then start to store the specified data in the repository in an encrypted format. While doing so, the software displays the progress of the storage process as a percentage and in absolute figures (Figure 1).

To store the backup on a remote server, use the SFTP protocol. Instead of the local backup directory, enter the server name, including the target path to be created on it, at the command line. The further operating steps, including password entry, are like those on a local system.

For each new backup, restic creates a new backup file within the archive path below `snapshots/`. To display its content, use the command from line 3 of


```
$ restic -r /home/erik/restic init
enter password for new repository:
enter password again:
created restic repository bc0025030a at /home/erik/restic

Please note that knowledge of your password is required to access
the repository. Losing your password means that your data is
irrecoverably lost.
$ restic -r /home/erik/restic/ backup /home/erik/Downloads/
enter password for repository:
password is correct
scan [/home/erik/Downloads]
scanned 3112 directories, 7920 files in 0:01

[2:29] 100.00% 15.258 GiB / 15.258 GiB 11032 / 11032 items 0 errors ETA 0:00

duration: 2:29
snapshot 97c706c7 saved
```

Figure 1: Restic is very informative during backup.



Listing 1. The corresponding table shows the ID number of the backup, as well as the host. The tool lists the backup's original path in full, so that you can quickly see which backup it is. If an archive with a directory tree of the same name contains data from different host computers, restic lists it in order, but does not display individual files in any case.

Granular Details

To view individual files and directory structures in a repository, you first need to mount the repository. This feature lets you copy individual files from a backup instead of restoring all the data. To mount a backup, it is best to use a temporary directory with an arbitrary name as the target.

If the directory does not exist, restic asks if it should create the directory when it calls the `mount` command (Listing 1, line 4) and mounts the target directory after you confirm. However, it mounts the temporary directories as read-only, so your options for working with the files in the directories are limited. In addition, the command for mounting the backup remains active in the terminal, so all activities performed with the files have to be carried out in a separate tab or window. After you finish this work, you need to unmount the temporary directory; restic will not quit until you do this.

Additional parameters let users perform various additional tasks with the backed up datasets. For example,

you can remove individual backups from a repository using the `forget` option. This helps to keep the archive clean by removing obsolete individual backups, especially from scheduled backup runs.

Since restic always creates full backups, the differences between individual backups cannot be seen at first glance. The `diff` option helps users here; it shows differences in the data between two backups. `check` lets you verify the data integrity.

The exact syntax for the use of individual subcommands is revealed by entering the following command at the prompt:

```
restic <Command> --help
```

The software then displays detailed descriptions for using the option in question.

Restore

To restore a backup on demand, you will need the ID number – assuming you have several backups. To do this, first display the existing backups (Listing 1, line 3). From the full file paths that are displayed, you can see which backup has which ID number. Then specify the command from Listing 2 to restore a backup. Restic now returns the specified backup to the specified target path, automatically creating the corresponding subdirectory if it does not exist (Figure 2).

Scripts

Because restic can also handle variables, you can use the tool to regularly back up data using scripts. The scripts can then be executed as cron jobs on a time-controlled basis. Ideally, you should set an environment variable for the password entry and for the storage path.

The fact that the software can basically access all your data media and filesystems opens up the possibility of transferring data to a NAS system or to a server via SFTP, for example. If

```
$ restic -r /home/erik/restic/ snapshots
enter password for repository:
password is correct
ID          Date           Host                               Tags      Directory
-----
97c706c7   2019-04-19 17:44:19 erik-HP-Z600-Workstation          /home/erik/Downloads
7eaa77db   2019-04-19 18:09:43 erik-HP-Z600-Workstation          /home/erik/Bilder
-----
2 snapshots
$ restic -r /home/erik/restic/ restore -t /home/erik/sicherung1 7eaa77db
enter password for repository:
password is correct
restoring <Snapshot 7eaa77db of [/home/erik/Bilder] at 2019-04-19 18:09:43.285706675
+0200 CEST by erik@erik-HP-Z600-Workstation> to /home/erik/sicherung1
```

Figure 2: Saving the data with restic just requires a little typing.

Listing 1: Creating a Repository

```
01 $ restic -r /Path/Archive init
02 $ restic -r /Path/Archive backup /Path/Files
03 $ restic -r /Path/Archive snapshots
04 $ restic -r /Path/Archive mount /Temporary directory
```

Listing 2: Restoring a Backup

```
$ restic -r /Path/Archive restore -t /Path/To/Restore ID number
```

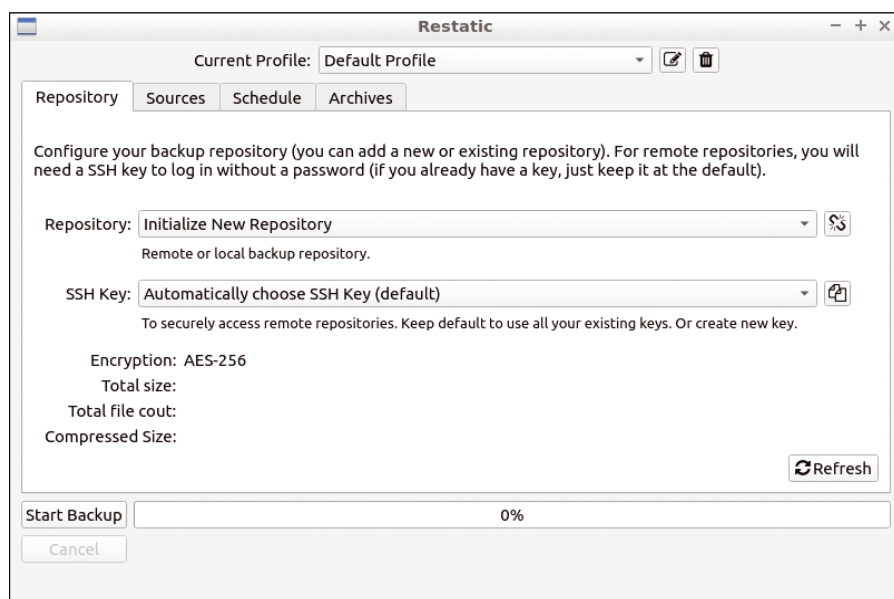



Figure 3: Restatic covers all the restic functions in a few tabs.

you have slow storage or narrow bandwidths, however, the backup speed is naturally significantly lower than backing up the local filesystem.

GUI

Restatic [2], a GUI for restic, has been under development for some time now. You will find the application programmed in Python on GitHub. As its underpinnings, the front end needs Python v3.6 or higher. Many distributions with long-term support come with even older Python versions, which you have to update beforehand in order to be able to install Restatic.

After setup, call the software at the prompt by entering `restatic`. In the program window that then opens, you can access the individual functions using a clear-cut tab structure (Figure 3). You always need to follow the same order as for the command-line program. In addition, you can define

lower window segments, you can then define exclusion criteria for folders and files that you are not going to back up.

After pressing *Start Backup* (bottom left), Restatic runs the backup and displays a progress bar. You then need to click on the *Refresh* button in the *Archives* tab to see the backups in the repository. The *Extract*, *Check*, *Mount*, and *Prune* buttons let you work with the backups. In this dialog, you can specify how many old archives you want to keep for multiple backup files.

To run the backup automatically in the future, use the *Schedule* tab. When you get there, you can define, at the push of a button, how often Restatic should save data from the previously specified source folders to the active repository. This avoids the need to write a shell script or create a cron job for automatic backups. After pressing *Apply*, Restatic applies the data and activates your automatic backup actions (Figure 4).

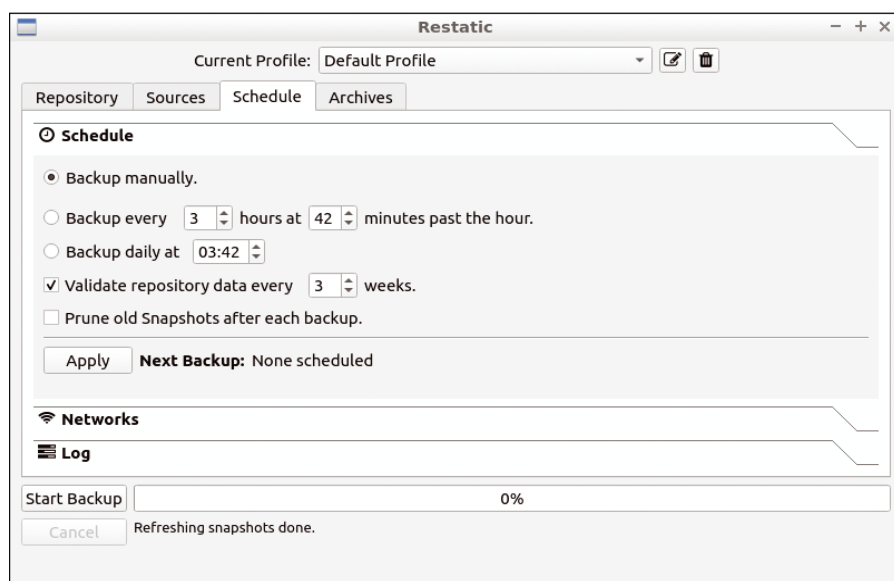


Figure 4: Restatic can automatically back up data in a time-controlled manner without any problems.

profiles at the beginning using the *Current Profile*: field; this is particularly useful for larger infrastructures with heterogeneous databases.

Then define a repository for the backup. To do this, click the *Repository* field in the *Repository* tab and select one of the options. To let you integrate existing repositories, a separate dialog box appears; you can use this to connect to the archive in question. Restatic supports connections to remote servers.

In the same way, you can use the *SSH Key*: field to specify in a context-dependent manner whether to use an existing key or generate a new key. Then switch to the *Sources* tab where you specify the directories to be backed up. To do this, use the *Add Folder* button to add individual folders to a list view. In the two

Conclusions

Restic takes away users' reservations about the command line by consistently adhering to a simple syntax for each function. The backup tool is very appealing due to its fast action, and above all the default encryption, which means that there is nothing standing in the way of backups to the cloud. With automatic, time-controlled backups and the option of using local SFTP servers as storage media, the software is also suitable for larger infrastructures. The Restatic GUI also lets users create backups easily, in this case, at the push of a button. ■■■

Info

[1] restic: <https://restic.net>

[2] Restatic: <https://github.com/Mebus/restatic>

REAL SOLUTIONS *for* REAL NETWORKS

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

Learn the latest techniques for better:

- network security
- system management
- troubleshooting
- performance tuning
- virtualization
- cloud computing

on Windows, Linux, and popular varieties of Unix.

**GET IT
FAST**

with a digital subscription!

6 issues per year!

ORDER NOW

shop.linuxnewmedia.com



Automatic backups to external media

Safe House

A recent backup is more reliable than any kind of data rescue. But for many users, a backup won't happen unless the process is easy to manage. *By Bernhard Bablok*

The simplest form of backup is to copy the data to a mobile device, such as a USB stick or an external USB hard drive. You can perform this kind of backup either with the file manager or with a copy command like:

```
cp -a source_directory target_directory
```

This command uses the archive option `-a` to recursively copy the files, and it also includes the file permissions. If the target directory does not exist, the command creates it as a copy of the source and its contents. If, on the other hand, it already exists, it will contain a subdirectory named `source_directory` after the copy operation.

The `cp` command with a single option is easy to remember, but it has a number of drawbacks. Exotic files and attributes can be too much for `cp`. On top of that, the process is inefficient, because it stubbornly copies all the source files. Furthermore, there is no verification of the results.

That's why professionals use the more complicated command shown in Listing 1, with `rsync`, which is installed by default on almost every distribution. Here, too, there is the archive option `-a`. The other options relate to extended attributes (`-X`), special attributes (such as access control lists, `-A`), hardlinks (`-H`), and sparse files (`-S`). In addition, the command does not copy across partition boundaries (`-x`). `Rsync` creates the target directory if it does not already exist. A slash after the name of the source directory ensures that `rsync` does not copy the directory itself (line 1), but only its contents including subdirectories (line 2).


If you run the command more than once, `rsync` will copy only modified files. It also checks whether the copy was actually successful. If you add the `--delete` option, the command will even delete files on the target that no longer exist in the source – so you get a perfect one-to-one copy.

Listing 1: Copying with rsync

```
01 $ rsync -aXHSx <source> <target>
02 $ rsync -aXHSx <source>/ <target>
```

Copies Are Not Enough

The problem with this type of backup method is that a simple replication of





the data also copies defective files, or files encrypted by Trojans. With Linux, the danger is less from viruses and Trojans than simply the user sitting in front of the keyboard. If you quickly delete or edit a file by mistake, then only notice this after the backup copy, the old version is long gone.

A genuine backup therefore also includes archiving old versions. At this point, things usually get complicated, and many users give up. But, fortunately, there are a number of projects dedicated to this task. In most cases, the solutions cause some overhead during the initial setup, but the daily backup then becomes all the easier.

You have to overcome two hurdles during rsnapshot's initial setup: creating an appropriate backup medium and installing the software and its configuration. All in all, however, this takes about fifteen minutes of prep work.

Backup Medium

For the backup medium, use either a USB stick with enough capacity, or an external USB hard disk. For mobile use, I use an 8GB USB stick, which is fine for my home directory with my DIY programs and articles. For a large media collection, a 4TB disk is recommended. Make sure that the backup medium does not contain any important data. (Note: Moving forward, when I refer to a backup stick, I also mean the alternative hard disk option as well.)

USB sticks and hard disks usually come preformatted with the FAT32 or NTFS filesystem, both are predominant in the Windows world. Neither is suitable for backing up Linux systems. In the case of sticks, you also need to consider whether there is a partition, depending on the capacity. If in doubt, repartition the stick so that it contains exactly one partition. The desktop environment typically comes with a tool for this; alternatively, you will find an application like GParted [1] in the package sources.

After plugging in the stick, first check in the file manager whether the Linux system has automatically mounted the partition on the medium. This may not be the case after repartitioning, but it can happen. In this case, unmount the partition and format it with ext4.

This can be done in a GUI with GParted (Figure 1) or at the command line with the command from Listing 2, line 2. Replace `/dev/sd<XY>` with the device name of the medium (e.g., `/dev/sdb1`). If in doubt, it is best to check the device name before formatting with `lsblk` (line 1). The command outputs the partitions, the sizes, and the corresponding mount points, if any (Figure 2). In Figure 2, the stick can be easily identified by its size of almost 8GB.

Listing 2 also assigns a label for the partition. Instead of `BACKUP` you can specify something else here. You will need the identifier again later, because you only want the backup

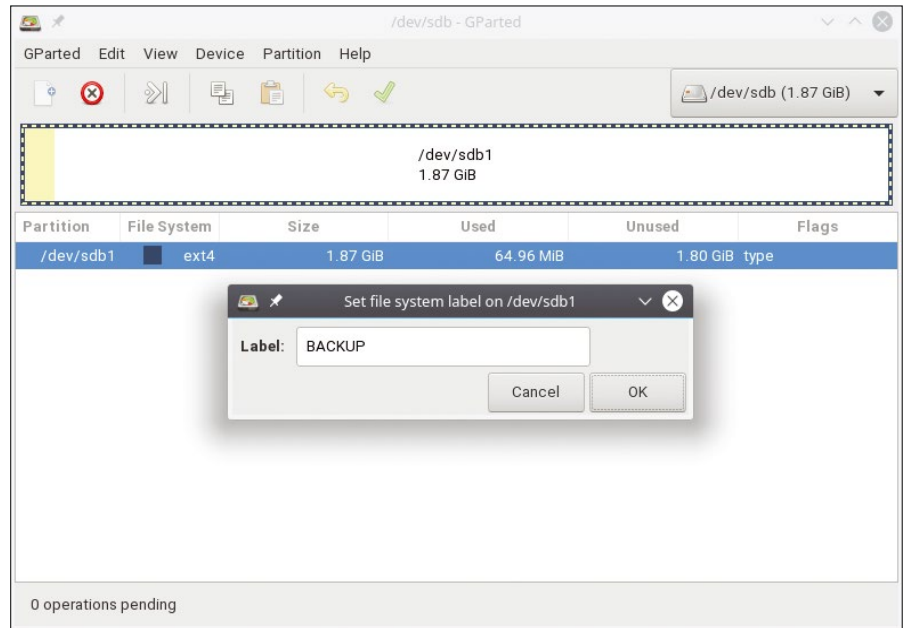


Figure 1: Using GParted, you can quickly partition and format disks and assign a label to the new partition.

to start if you plug the correct medium into the system. With GParted, you will find the function *Label File System* within the *Partition* menu.

System Adaptations

After setting up the backup media, some packages need to be installed and configured. First you import the *rsync*, *git*, and *rsnapshot* packages in the package manager. Line 1 of Listing 3 shows the command for Debian, Ubuntu, and their derivatives. If *rsnapshot* is missing from the repositories of your distribution, you can download the script from the project page [2].

You also need my *autobackup* script and configuration files [3], which you download and install with the commands from lines 2 to 4 of Listing 3. If you are using a Debian-based system, such as Ubuntu or Linux Mint, the command from line 4

Listing 2: Formatting the Stick with ext4

```
01 $ lsblk /dev/sd?
02 $ sudo mkfs.ext4 -L "BACKUP" /dev/sd<XY>
```

```
[bablokb@toliman:~] > lsblk /dev/sd?
NAME                                MAJ:MIN RM  SIZE RO  TYPE MOUNTPOINT
sda                                  8:0    0 238,5G  0  disk
├─sda1                               8:1    0    2G  0  part /boot
├─sda2                               8:2    0    8G  0  part [SWAP]
└─sda3                               8:3    0 228,5G  0  part
   ├─vg1-root2                       254:0    0   10G  0  lvm /
   ├─vg1-root1                       254:1    0   10G  0  lvm
   ├─vg1-root3                       254:2    0   10G  0  lvm
   ├─vg1-usr_local                   254:3    0   20G  0  lvm /usr/local
   ├─vg1-home                       254:4    0    5G  0  lvm /home
   └─vg1-data                       254:5    0   15G  0  lvm /data
sdb                                  8:16    1   7,6G  0  disk
└─sdb1                              8:17    1   7,6G  0  part
```

Figure 2: The output from `lsblk`: The backup USB stick can be quickly identified based on partition size.



Listing 3: Installing rsync, git, and rsnapshot

```
01 $ sudo apt install rsync git rsnapshot
02 $ git clone https://github.com/bablokb/autobackup-service.git
03 $ cd autobackup-service
04 $ sudo tools/install
05 $ sudo cp /etc/rsnapshot.conf.autobackup-service /etc/rsnapshot.conf
```

will also install the *rsync* and *rsnapshot* packages. If you do not want to do this, comment out the command `install_packages` at the end of the `tools/install` file.

When you plug in the backup stick later on, the kernel detects this and creates a udev event (Figure 3). The constantly running udev daemon then starts the `/usr/local/sbin/autobackup` script. It checks whether the backup stick is actually plugged in, and whether a corresponding backup already exists for the current date, current week, or current month. If this is

not the case, the script starts the *rsnapshot* script with the correct parameters. To make all of this work, you need to configure the *autobackup* script as shown in Listing 4. The configuration file can be found in `/etc/autobackup.conf`. Here you normally only have to enter the LABEL, which in our example is `BACKUP`. You can

set the value of `SYSLOG` to 1. The script then writes detailed messages to the system log, so that you can view more information if problems occur. Normally the script will only make one backup per day; `force_daily=1` changes this. When making any changes, be careful not to violate the shell syntax; especially do not put any spaces before or after the equals signs. The second step is to create the `/etc/rsnapshot.conf` file. On installing *rsnapshot*, you will usually have an example, as with installing the *autobackup* script. Copy the `/etc/rsnapshot.conf.autobackup-service` file to `/etc/rsnapshot.conf` (Listing 3, line 5). You only have to make adjustments in the middle and at the very end of the file (Listing 5). In the middle, you need to define how many generations you want to store for each interval. The identifiers (day, week, month, and year, with the latter commented out) must match the values in the last four lines of Listing 4.

At the end of the file, you specify which directories the script should save. *Rsnapshot* offers several possibilities to exclude files – but this is rarely worth the effort.

There is a pitfall in `/etc/rsnapshot.conf`, however: You have to separate the single words with tabs, not spaces. You will want to call the `sudo rsnapshot check` command after editing. The command checks the syntax and determines whether an editor has tacitly replaced tabs with spaces.

Listing 4: Configuring autobackup

```
01 # /etc/autobackup.conf
02 LABEL=BACKUP
03 SYSLOG=0
04 # wait for device x second(s)
05 WAIT_FOR_DEVICE=2
06 # If 1: Daily backup whenever plugged in
07 force_daily=0
08 # Backup name (according to /etc/rsnapshot.conf)
09 daily="day"
10 weekly="week"
11 monthly="month"
12 yearly=""
```

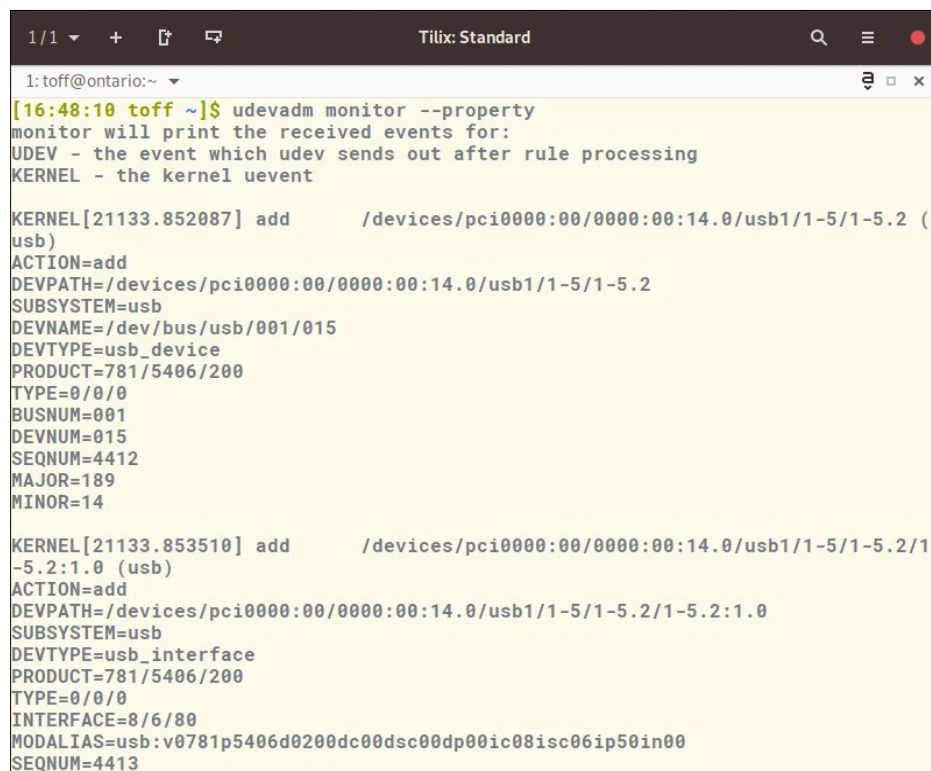


Figure 3: The command `udevadm monitor --property` lets you observe the events under the hood of udev in real time.

All Systems Go

After these preparations, restart the system. The installation has installed a new rule for the udev daemon; the restart ensures that this rule is automatically applied. If you now plug the backup stick into a free USB port, the backup should start automatically. When the stick's LED stops flashing, the backup is finished, and the stick is unmounted. If you leave it plugged in and mount it again manually, you will see the first backup in the directory `day.0` in the file manager (Figure 4).

The script creates additional backups on each subsequent day. On the second day, the content moves from `day.0` to `day.1` and `day.0` contains the most recent backup. This continues for a full seven days – even if you skip individual days. Then the script creates a weekly backup: `day.6` becomes `week.0`. This takes place just once a calendar week; after four weeks, the scheme continues with monthly backups.

When using *rsnapshot* for a server that runs every day, entries in the



crontab ensure an orderly workflow – such as a weekly backup every Sunday evening and a monthly backup on the first of every month. With desktops and laptops that are not consistently used every day, this is unlikely to work. For this reason, it makes sense to start the backup precisely when the backup medium is plugged into the computer.

If you are wondering why backups are so fast from the second day onward and why the disk space used by the backups only increases very slowly even though the files are stored in so many daily and weekly directories: This is the “magic” behind rsync. The tool copies only the changed files. For files that do not change, rsync only creates a link, or more specifically a hard link. Then what looks like a file `day.1/home/bablokb/old-text.odt` is in reality only a pointer to the location of the file `day.0/home/bablokb/old-text.odt` and `day.1/home/bablokb/old-text.odt` on the stick.

The procedure offers a number of advantages. Instead of copying large files via USB, rsync just creates the entry – it’s fast and saves a large amount of space. Nevertheless, each daily backup shows you exactly the status of the directories on the backup date. If you want to see what has changed between today and the day before yesterday, just compare `day.0/[...]/<file>.txt` with `day.2/[...]/<file>.txt`.

Conclusions

With a little preparation, regular backups are performed automatically. This procedure is particularly recommended for people who are already in the habit of copying their data to an external medium from time to time. As always in the open source world, you can adapt the process to your specific needs.

For example, you could replace the trigger for the backup (inserting the stick) with something else – for a laptop, for example, logging into your home WLAN. The backup should then also logically be created on a network drive.

No matter which method you choose, it is crucial that you make regular backups. Don’t forget to check from time to time that the backups do what they promise. In my personal experience, there have been cases where the data were backed up daily to the same magnetic tapes for years. Unfortunately, the tapes were only designed for a low, three-digit number of write cycles, which would have made successful data recovery impossible in an emergency. With modern flash media this scenario is not a worry – but, as the saying goes: It’s better to be safe than sorry. ■■■

Info

[1] GParted: <https://gparted.org>

[2] rsnapshot: <http://www.rsnapshot.org>

[3] autobackup: <https://github.com/bablokb/autobackup-service>

Listing 5: Adjusting /etc/rsnapshot.conf

```
# /etc/rsnapshot.conf
[...]
#####
# BACKUP LEVELS / INTERVALS
# Must be unique and in ascending order
# e.g. alpha, beta, gamma, etc.
#####
retain day 7
retain week 4
retain month 3
#retain year 3

#####
# BACKUP POINTS / SCRIPTS
#####
# LOCALHOST
backup /etc/ ./
backup /home ./
```




Figure 4: The autobackup script creates a new directory for each backup. You always have different versions of your data at hand.

LINUX UPDATE

Need more Linux?

Our free Linux Update newsletter delivers insightful articles and tech tips to your mailbox twice a month.

- Original articles on real-world Linux
- Tips on Bash scripting and other advanced techniques
- Discounts and special offers available only to newsletter subscribers
- Linux news



<https://bit.ly/Linux-Update>



GPX Viewer

On Track

With GPX Viewer you can map GPX tracks and view GPS data in a web browser. It's a simple way to revisit a recent vacation, organize your photos, or map your favorite bike routes. *By Dmitri Popov*

Tracking and storing your movements using a dedicated GPS logger or an Android application can come in handy in many situations. As a photographer, you can use the saved tracks in the GPX format to geo-correlate your photos. If you are into biking or hiking, you can analyze the route. And when you travel, you can have a record of the places you've visited. Besides a device or an app that lets you record tracks in the GPX format, you also need a tool to view and analyze the recorded tracks. There are several applications that can do the job. But if you want to publish the mapped tracks on the web, or you just want to use the browser as your preferred tool, then GPX Viewer [1] is just the ticket. This JavaScript library makes it possible to create a read-to-publish map and plot GPX tracks on it. Better still, the library allows you to extract and display the data stored in the GPX tracks as well as map geotagged photos. More importantly, GPX Viewer greatly simplifies the process of creating simple and advanced maps, and you only need a text editor and a working knowledge of HTML to create sophisticated maps with GPX tracks.

GPX Viewer Basics

Being a regular JavaScript library, GPX Viewer does not require installation or

configuration. Grab the latest version from the project's website, unpack the downloaded ZIP archive, and you can start working on your first GPX Viewer project. To keep things tidy, create a separate directory on your web server for your first project, and move the *GM_Utils* folder into it. This folder contains the JavaScript files that power GPX Viewer. Place the desired GPX track in the project folder, create an HTML file, open it for editing, and paste the code in Listing 1 into it.

The code is rather straightforward, and there are only a couple of points that require a brief explanation. The `<html lang="en">` line is important, because GPX Viewer uses it to determine which language to use. If not specified, GPX Viewer defaults to German. In addition to English, GPX Viewer also supports French. The *class* and the *style* attributes in

```
<div class="gpxview:tracks.gpx" style="width:800px;height:600px">
```

Listing 1: Simple HTML Page

```
01 <!DOCTYPE html>
02 <html lang="en">
03   <head>
04     <meta charset="utf-8">
05     <link rel="stylesheet" href="https://unpkg.com/sakura.css/css/sakura.
      css" type="text/css">
06     <meta name="viewport" content="width=device-width, initial-scale=1">
07     <title>GPXViewer</title>
08     <script src="GM_Utils/GPX2GM.js"></script>
09   </head>
10   <body>
11     <h1>Example</h1>
12     <p>Simple map</p>
13     <div class="gpxview:tracks.gpx"
      style="width:800px;height:600px">
      <noscript><p>Enable JavaScript to
      view the map</p></noscript></div>
14   </body>
15 </html>
```

Map with tracks

This is a simple example of using the GPX Viewer JavaScript library to generate maps with GPX tracks.

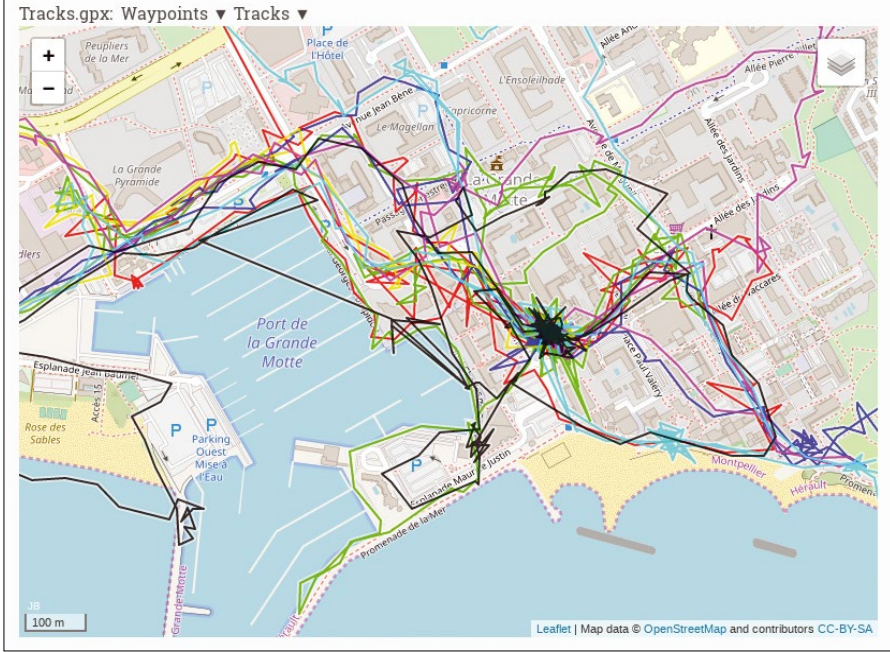


Figure 1: A simple map with GPX tracks generated by the GPX Viewer JavaScript library.

specify the GPX track file and dimensions of the map. By default, GPX Viewer uses OpenStreetMap's default map layer, but you can change that by specifying the desired mode by appending it to the name of the GPX track as follows:

```
class="gpxview:tracks.gpx:OPENTOPO"
```

This instructs GPX Viewer to use the Open Topo layer, while `gpxview:tracks.gpx:Satellit` opens the map in the satellite view (note the German spelling *Satellit*).

Now open the created HTML file in the browser, and you should see your first map (Figure 1). Although the map you've created is rather basic, it does have a few clever features. First of all, you can switch to a different layer using the icon in the upper-right corner of the map. You can use the mouse to navigate around the map, as well as zoom in and out. The *Tracks* drop-down list at the top of the map displays a list of all GPX tracks, and you can use it to show and hide individual tracks (Figure 2). When you mouse over a track, it's automatically highlighted, which can come in useful when you have multiple intersecting

tracks on the map. Click anywhere on a track, and a pop-up window shows

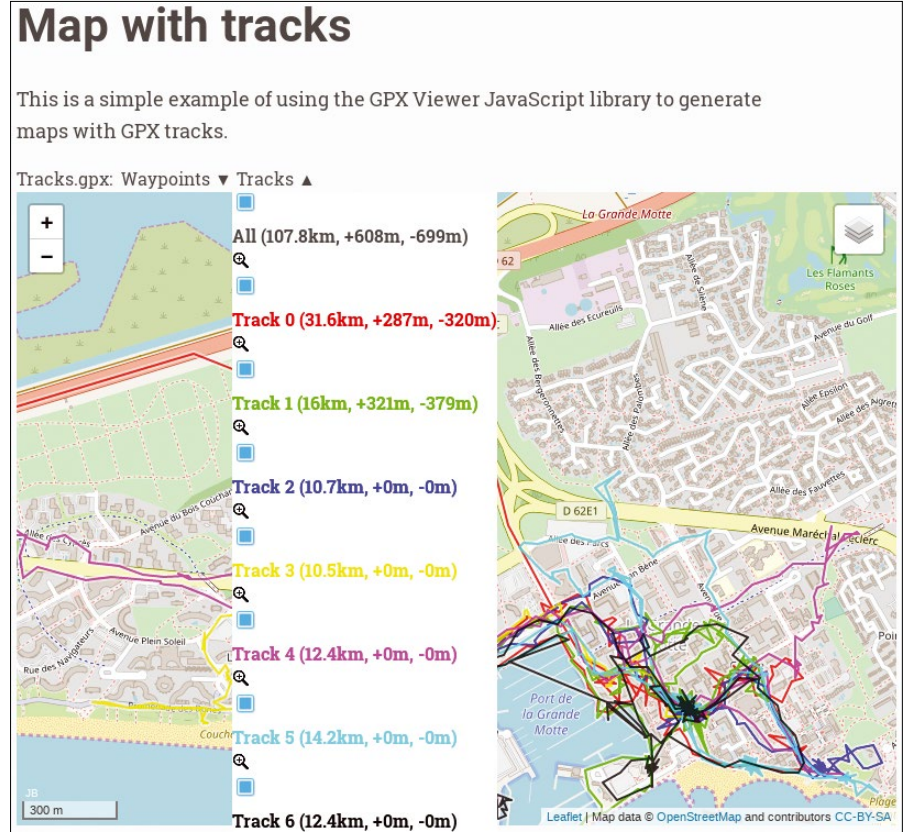


Figure 2: You can show and hide individual GPX tracks on the map.

you key info for the specific point, including distance from the starting point, elevation, and speed.

Some GPS logging devices allow you to add waypoints to the currently recorded track, and GPX Viewer automatically detects and displays waypoints as markers on the map. You can use the *Waypoints* drop-down list to show and hide individual waypoints.

In our simple example, the name of the GPX file is hardwired into the HTML file. If you want to show multiple tracks on a single map, you have to specify them manually as follows:

```
<div class="gpxview:track1.gpx, track2.gpx, track3.gpx" style="width: 800px; height: 600px">
```

This approach works, but it's not very practical. First, you have to make sure that the names of the GPX files match the names in the HTML file. Secondly, if you upload a new GPX file, you have to manually add it to the HTML page, too. You can merge multiple GPX files in one as explained in the Merge GPX Tracks boxout, but a dash of PHP provides a more elegant solution to the problem. Change the file extension of the HTML page to `.php`, and

then insert the code below into the page where you want the map to appear:

```
<?php
$files = glob("*.gpx");
$list=implode(", ", $files);
echo "<div id='map' class='gpxview:2
$list' style='width:800px;height:2
600px'></div>";
?>
```

From now on, when you open the page, the PHP code finds all .gpx files and adds them to the map.

Buttons, Graphs, Oh My!

With the basic map in place, you can add some useful features. For example, you can easily add buttons that reset the map and load additional GPX tracks. To add a button that resets the map, insert the following line into the HTML file where you want the button to appear:

```
<button type="button" class="2
" gpxview:map:skaliere">2
Reset position and zoom</button>
```

Adding a button that loads another map with a different GPX track is equally easy. The code below loads a map with the track-2.gpx GPX file:

```
<button type="button" class="2
" gpxview:map:lade:track-2.gpx">2
Load Track 2</button>
```

Alternatively, you can load maps with specified GPX tracks via a drop-down list.

```
<select class="gpxview">
<option value="map:track1.gpx" 2
selected="selected">Track 1</option>
<option value="map:track2.gpx">2
Track 2</option>
<option value="map:track3.gpx">2
```

Merge GPX Tracks

Instead of dealing with multiple GPX tracks, you can merge them into one and then use the resulting file with GPX Viewer. The GPSBabel tool, which is available in the official repositories of most mainstream Linux distributions, can help you with that. Use the code in Listing 2 to create a script that merges all .gpx files in the current directory and saves the result as a tracks.gpx file.

```
Track 3</option>
</select>
```

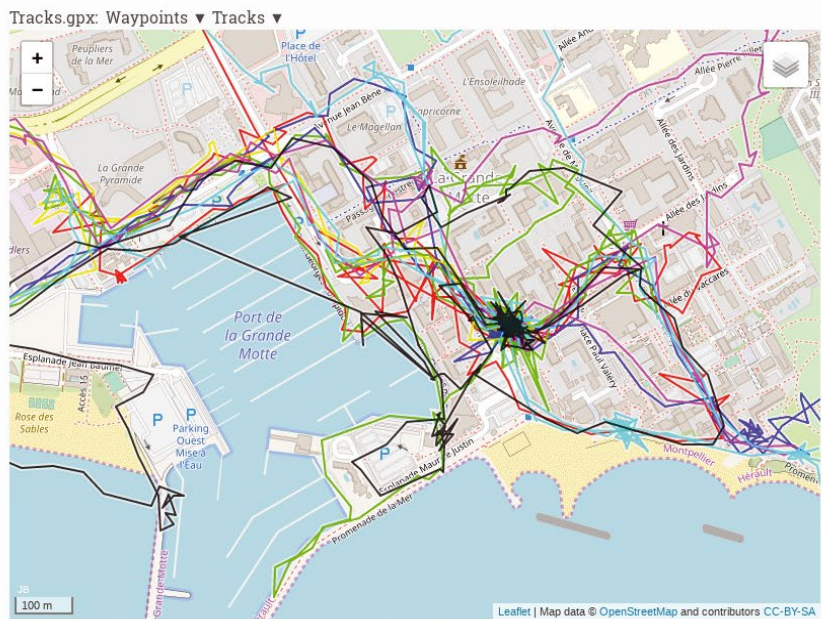
This is a much more practical solution when you want to include multiple GPX tracks, as you don't have to define buttons for each map. Here too, you can

add some PHP magic to generate a list of all GPX tracks in the project directory automatically:

```
<select class="gpxview">
<?php
$files = glob("*.gpx");
```

Map with tracks

This is a simple example of using the GPX Viewer JavaScript library to generate maps with GPX tracks.



Speed/Distance




Figure 3: GPX Viewer makes it supremely easy to add buttons and generate graphs.


```
foreach($files as $track) {
echo "<option value='map:$track'>?
  Track: $track</option>";
}
?>
</select>
```

A regular GPX track usually contains additional data, such as speed and elevation. Mousing over the track displays this information for the given location in a pop-up window, but GPX Viewer also makes it possible to generate graphs and embed them into the page using a single line of code. All you have to do is to specify one of the supported classes for the `div` element and provide the graph's dimensions:

```
<div id="map_vp" style="width:800px;?
height:600px">
```

In this case, the `map_vp` identifier generates a speed/distance graph (Figure 3). GPX Viewer supports a number of other identifiers, such as `maps_vpt` (speed/distance), `map_hp` (elevation/distance), and `map_hpt` (elevation/time).

By default, the generated graph displays data from all included tracks using different colors. If you disable all but a single track, you can mouse over the graph and view info for the current point directly on the map.

Listing 2: Shell Script for Merging GPX Tracks

```
01 #!/usr/bin/env bash
02 ff=""
03 for f in *.gpx
04 do
05     ff="$ff -f $f"
06 done
07 gpsbabel -i gpx $ff -o gpx -F "tracks.gpx"
08 gpx=$(pwd)/tracks.gpx"
```

Listing 3: Shell Script for Processing Photos

```
01 #!/usr/bin/env bash
02 rm photos.txt
03 for f in *.jpeg
04 do
05     lat=`exiftool -n -p '$GPSlatitude' $f`
06     lon=`exiftool -n -p '$GPSlongitude' $f`
07     desc=$(exiftool -Comment "$f" | cut -d":" -f2 | sed -e 's/^[[:space:]]*/')
08     echo "<a href=\"$f\" data-geo=\"$lat:$lat,lon:$lon\">$desc</a>" >> photos.txt
09 done
```

Map with tracks

This is a simple example of using the GPX Viewer JavaScript library to generate maps with GPX tracks.



Figure 4: GPX Viewer makes it possible to map geotagged photos.

If there is a named waypoint on the map and an image with the matching name residing in the project directory, GPX Viewer automatically displays the image on the map using an appropriate waypoint icon. You can click on the icon to view a thumbnail version of the image in a pop-up window. But what if you only have geotagged photos without matching waypoints? GPX Viewer provides a solution for that, too (Figure 4). You can add photos manually as follows

(the `xxx.xxxxx` placeholders refer to actual latitude and longitude values):

```
<a href="foo.jpg" data-geo=?
"lat:xxx.xxxxx, lon:xxx.xxxxx">?
Description</a>
```

This solution requires some manual work. First, you need to extract the latitude and longitude values from each photo. Then, you need to manually add entries to the HTML file. However you can easily automate this manual process using the shell script in Listing 3. The script goes through all the files with the specified file extension (in this case, it's `.jpeg`) and extracts geographical coordinates and descriptions from each photo. The script outputs the results to the `photos.txt` file, which you can then copy and paste into your HTML file.

Keep in mind that the script relies on ExifTool, so make sure that it's installed on your system. There is one more thing you need to take care of when adding photos to the HTML file. By default, GPX Viewer displays the links to the added photos right below the map, which doesn't serve any prac-

tical purpose. Use a simple CSS trick to hide them. In the <head> section of the HTML file, specify the following style:

```
<style>
  #map_img { display:none }
</style>
```

Then wrap all the image items into the <div> element:

```
<div #map_img>
  <a href="foo.jpg" data-geo="
    "lat:xxx.xxxxx,lon:xxx.xxxxx">
    Description</a>
</div>
```

Naturally, you can also add PHP code that processes the photos and generates

the required code, but this requires more than just basic PHP coding skills. Moreover, the PHP library for working with Exif metadata can be a bit finicky. Since you already have a shell script that does the job of processing photos and generating a separate file, you can run it either manually or via a cron job and use a single line of PHP code to include the resulting file into the HTML page:

```
<div id="map_img">
  <div><p><?php include(
    ('photos.txt'); ?></p></div>
</div>
```

This not particularly elegant hack assumes that you can run shell scripts on your web server. With all the bits and

code snippets added to the page, the final result looks like Listing 4.

Final Word

Whether you are looking for an easy way to publish GPX tracks on the web, map geotagged photos, or analyze GPS data, GPX Viewer is just the tool for the job. This JavaScript library works with any modern browser, and it allows you to generate maps and graphs with a minimum of effort. The library's documentation is available in German only, but the *Translate* button on the project's website should give you a usable translation. ■■■

Info

[1] GPX Viewer: <https://www.j-berkemeier.de/GPXViewer/>


Listing 4: Final Page with a Mix of HTML and PHP

```
01 <!DOCTYPE html>
02 <html lang="en">
03   <head>
04     <meta charset="utf-8">
05     <link rel="stylesheet" href="https://unpkg.com/sakura.css/css/sakura.css" type="text/css">
06     <meta name="viewport" content="width=device-width, initial-scale=1">
07     <title>GPXViewer example map</title>
08     <script src="GM_Utils/GPX2GM.js"></script>
09     <style>
10       #map_img { display:none }
11     </style>
12   </head>
13   <body>
14     <?php
15       $title="Map with tracks";
16       $blurb="This is a simple example of using the GPX Viewer JavaScript library to generate maps with GPX tracks.";
17       $files = glob("*.gpx");
18       $list=implode(", ", $files);
19       echo "<h1>$title</h1>";
20       echo "<p>$blurb</p>";
21       echo "<div id='map' class='gpxview:$list' style='width:800px;height:600px'></div>";
22       echo "<noscript><p>Enable JavaScript to view the map.</p></noscript>";
23       ??
24     <p></p>
25     <button type="button" class="gpxview:map:skaliere">Reset position and zoom</button>
26     </noscript>
27   </div>
28   <div id="map_img">
29     <div><p><?php include('photos.txt'); ?></p></div>
30   </div>
31   <h3>Speed/Distance</h3>
32   <div id="map_vp" style="width:800px;height:600px">
33     <noscript>
34       <p>Enable JavaScript to view the graph.</p>
35     </noscript>
36   </body>
37 </html>
```


Shop the Shop

shop.linuxnewmedia.com

DISCOVER LibreOffice




Explore the **FREE** office suite used by busy professionals around the world!

Create your own:

- Word processing docs
- 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!



Memorable but
secure passwords

Balancing Act

CLI tools for generating passwords have many options that can help you strike a balance between ease of use and security. *By Bruce Byfield*

How should you generate passwords? You probably know the standard advice of using a variety of characters, but the resulting passwords are hard to remember, especially since the recommended length keeps getting longer as cracking methods become more sophisticated. It's enough to make users choose convenience over security and use the same password everywhere – which is just about the worst thing you can do. Yet even if you choose unique passwords in every circumstance, how random can passwords invented by one person truly be? And what strategies can you use to make random passwords

memorable? A password manager can ease the daily burden, but what happens if your password manager becomes corrupted? What were simple issues a couple of decades ago have become complicated today, with no easy answers.

Ask security experts, and the answers to these questions can be conflicting. However, that doesn't stop the developers of password generators illustrating possible solutions. Many implement traditional answers, but others are starting to implement ways of making passwords more memorable, either by tweaking tradition in various ways or taking a new approach altogether.

PWGen

PWGen is one of the oldest password generators for Linux [1]. It is based on the traditional password advice, providing eight character strings for passwords that include lower- and upper-case letters, numerals, and special characters. This pattern can be modified by the options in Table 1 to produce an easier to remember password, but remember that excluding types of characters will probably result in a weaker password. For maximum security, specify `--secure (-s)`, which will maximize the mixture of characters (Figure 1).

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/>.

Table 1: PWGen Customization Options

<code>--no-numerals (-0)</code>	Exclude numbers
<code>--no-capitalise (-A)</code>	Exclude uppercase letters
<code>--ambiguous (-B)</code>	Exclude confusing characters, like upper case <i>O</i> and zero (<i>0</i>).
<code>--capitalise (-c)</code>	Include at least one capital
<code>--numerals (-n)</code>	Include at least one number
<code>--no-vowels (-v)</code>	Exclude vowels and numbers that could be mistaken for vowels like <i>1</i> or <i>0</i>
<code>--symbols (-y)</code>	Include at least one special character


```
bb@nanday:~$ pwgen
Kae8yi9i Roph2gei zuo0Aeno goC2niqu oozah3Qu ros3Xeeh Io6rah5i Nai5eef8
qui7Baey oquooX8u ei5i0egh ahy2neeR ij1rox80 Aip9aife aiz6Zo8A Shiehah4
shuh6Voh Sohhd1jie yahd8aiB ahteir9a ahzah5Qu Zie7Aech aMa90hca aXo0feeg
Ipu7ei9 oz3ooTha Uukixai3 sah1looV AhChaex0 cae4Uaye ierah3Sh Tae1chib
uKailai5 rae3Cho1 aingia3X rai40too dew1eyiV dasoh6So Ahs7uif8 kah1Piu7
bei2ooNg wail4ied joo7deiV pee30hsh eehi3Eel Eishah1u Ie4teeSh ahthee80
```

Figure 1: PWGen produces a list of moderately secure passwords from which to choose.

```
bb@nanday:~$ pwmake 128
UkUtabUsnUk$Ez7oHyn0tqyLAD
```

Figure 2: Pwmake is a simple but effective password generator.

```
bb@nanday:~$ gpw 4 12
forelidgesad
klialtshounr
counseggings
entalabuslat
```

Figure 3: Gpw produces pronounceable passwords that should be easy to recall.

Results are displayed in columns by default. However, you can use `--lines (-l)` to display results one per line. You can also control the number of results with `--num-passwords=NUMBER (-N)`.

PWGen provides a wealth of options, but notice that its passwords are eight characters long. That used to be an adequate length and still provides moderate security. However, using traditional password standards, an increasing number of experts suggest a minimum of 12 characters. If the password you are generating is important to you, you might want to use another password generator.

pwmake

`pwmake` is a simple utility available in the `libpwquality-tools` package [2]. It includes only one option: the amount of entropy bits used to generate the password (Figure 2). The minimum number of bits is 56. The man page suggests 64 bits when attackers do not have direct access to the password hash file, or 80-128 bits when attackers might have access or the password is used as an encryption key. However the number of bits can go even higher,

if desired. The entropy is generated by `/dev/urandom`. After generating the password, you can use `pwscore` or `pam_pwquality`, two other utilities in `libpwquality-tools`, to check the strength of the generated password. Since the passwords are random and long, you will probably only want to use `pwmake` with a password manager.

gpw

`gpw` [3] is another simple password generator. It uses two options: the number of passwords generated, and the length of each password. Unlike `pwmake`, `gpw` generates passwords that are easy to pronounce for English speakers and should therefore be easy to remember. The price for easier to remember passwords, of course, is less security (Figure 3).

otp

Readers of thrillers may remember that one time pads are used for coded messages intended to be used only once. The sender uses the top password or encryption, then discards it, and the receiver discards their copy after receiving a message. `otp` [4] has no direct connection to actual one time pads, except that the name adds drama to security. Contrary to what the name seems to imply, there is no limit to how often you can use the passwords produced by `otp`. Nor should this little script from the Debian repositories be confused with the similarly named Red Hat tool.

What `otp` does offer is a number of simple controls for generating passwords. Its options consist of a format, followed by the number of characters in the generated password. The default is all uppercase passwords, but more options are easily added to modify results.

```
bb@nanday:~$ otp -c12 -s4 -n5
1) WJJH-XJKV-KYTV 2) FRAY-EEFY-RLFY 3) JOBA-ZEZZ-PSXA
4) YJHG-UMBV-ACCU 5) SOFP-MGNP-DGUY
```

Figure 4: Despite the name, `otp`-generated passwords can be used more than once.

For example, `-c14` produces a password consisting only of letters that is 14 characters long. Similarly, users can opt for a password consisting of numbers (`dCHARS`) or letter groups that are easy to pronounce (`eCHARS`). For ease of use, `-sCHARS` can also be used to specify the spacing of hyphens throughout the password. If no options are specified, the default is passwords of eight characters with a hyphen every four characters.

Also, `otp` includes an option to specify the number of keys generated (`-nNUMBER`). In addition, it can also create an output file that can be used to verify incoming passwords (Figure 4).

Diceware

Diceware [5] gets its name from a method of generating results by rolling dice. The numbers on the dice are assembled as a number that is used to look up a word in a dictionary or word list that corresponds to that number. A number of words – five by default – are run together to produce the password. By default, each word begins with a capital letter unless the `--no-caps` option is used. The number of words that comprise the password can be set with `--num 'NUMBER'`, and special characters added with `--specials 'NUMBER'`. A delimiter between words can be set with `-d 'CHARACTER'`. The Diceware application is unique in that its option `--dice-side NUMBER` can be used so that results are not necessarily based on six-sided dice. As well, `--randomsource SOURCE` can be set, so that the randomness is generated by your operating system (Figure 5).

Diceware's original dictionaries have inspired a number of refinements (see `xkcdpass`). Diceware itself includes `en` (English), `en_eff` (based on Electronic Frontier Foundation modifications), `en_orig` (the original Diceware dictionary), and `en_securedrop` (English designed for security), which is the default. Each dictionary lists one word per line, prefaced with a sequential number, making the creation of a custom list an easy task.

xkcdpass

`xkcdpass` [6] is a Python script inspired by a comic strip from the geekily popular `xkcd` comic (Figure 6). Instead of the usual mixture of characters, the strip advocates strings of words, maintaining that these strings are just as secure as a


```
bb@nanday:~$ diceware -d'&'
Bye&Basis&Cosmos&Beget&Wince&Optimal
```

Figure 5: Diceware includes an option to add delimiters between words.

traditional password, and much easier to remember. `xkcdpass` is designed to generate these strings [7].

`xkcdpass` works by default with a word list called `eff-long` [8], which was released by the Electronic Frontier Foundation under a Creative Commons Attribution license for the specific purpose of generating passwords. `eff-long`, in turn, was originally a modification of Alan Beale’s `12Dicts` package for `Aspell` [9], which itself was based on the standard word list for Diceware. `12Dicts` consists of common English words of varying lengths originally derived from 12 different dictionaries, with outdated works, jargon, and scientific terms excluded. `eff-long` consists of 7,776 words, listed one per line, with the first line numbered `1111` and the rest continuing in sequence. Generally, `eff-long` is all that anyone needs, but other dic-

tionaries are also installed: `eff-special`, which contains 1,296 memorable words that are easier to remember but provide less security, and `eff-short`, in which each word begins with a unique three-letter prefix that could be used one day for autocompletion. Dictionaries for Finnish, French, Italian, German, Norwegian, Portuguese, and Spanish are also available. Those who want greater security can also produce longer, more specialized lists if desired. All word lists are stored in `/usr/lib/python3/dist-packages/xkcdpass/static/`.

The number of words in a password is five by default. However, `--numwords=NUMBER` can be used to change the default, and `--min=NUMBER` or `--max=NUMBER` can be specified to control the length of each word. Still another way to customize the resulting password is to specify a regular expression with `--var-char=REGEX`. For ease of memory, `--acrostic=WORD` can be set, so that the

first letter of each word spells out another word. For example, if the word supplied is “chaos,” `xkcdpass` might supply the password Church Hermann Auvergne Orthodox Sculptor (Figure 7).

Those who are security-conscious can include `--verbose` to read the level of security supplied by a specific password. Yet another convenience is `--interactive`, which continues to generate passwords until you accept one.

Choosing a Password Generator

Deciding which password generator to use is no easy task for most of us. Few of us are experts, and sorting through the arguments long enough to develop an informed opinion takes time. I suspect `xkcdpass` will be tempting, both because of its association with the comic and its deliberate efforts to make the passwords it generates easier to remember. However, cautious users may want to stay with a more traditional approach to passwords, just in case.

In the end, only one thing is certain: The question of how to produce secure passwords is a problem that is not going away in a hurry. ■■■

Info

- [1] PWGen: <https://sourceforge.net/projects/pwgen-win/>
- [2] libpwquality-tools: https://debian.pkgs.org/10/debian-main-amd64/libpwquality-tools_1.4.0-3_amd64.deb.html
- [3] gpw: <https://packages.debian.org/search?keywords=gpw&searchon=names&suite=stable§ion=all>
- [4] otp: <https://packages.debian.org/search?keywords=otp&searchon=names&suite=stable§ion=all>
- [5] Diceware: <https://packages.debian.org/search?keywords=diceware&searchon=names&suite=stable§ion=all>
- [6] `xkcdpass`: <https://pypi.org/project/xkcdpass/>
- [7] Explanatory comic: <https://www.xkcd.com/936/>
- [8] `eff-long`: https://www.eff.org/files/2016/07/18/eff_large_wordlist.txt
- [9] `12Dicts` package: <http://wordlist.aspell.net/12dicts/>

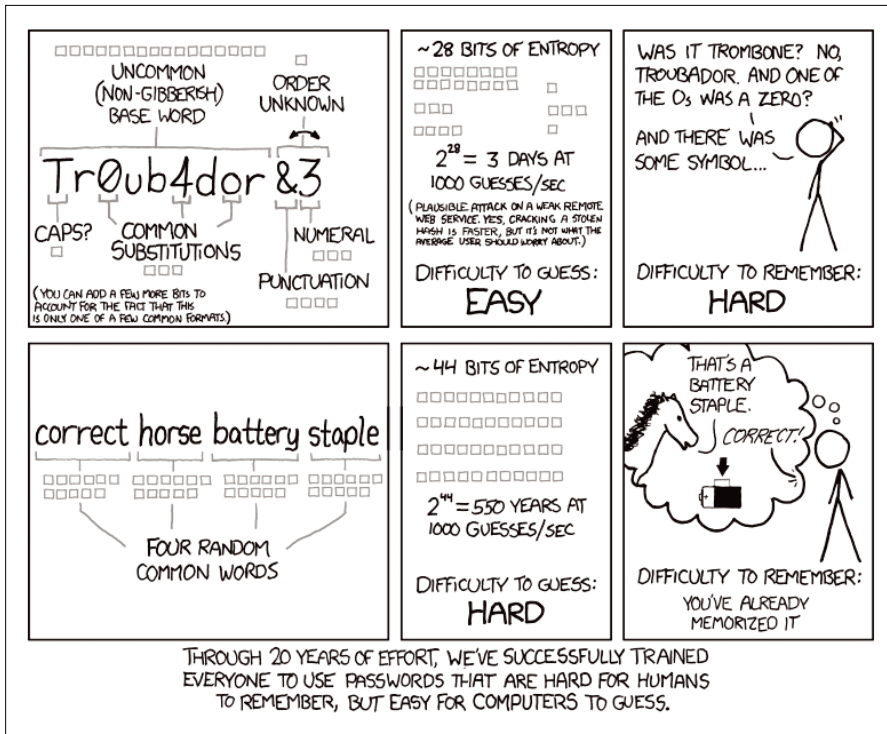


Figure 6: The comic strip that inspired `xkcdpass`.

```
bb@nanday:~$ xkcdpass --acrostic=CHAOS
Church Hermann Auvergne Orthodox Sculptor
```

Figure 7: Based on a popular comic, `xkcdpass` can use acrostics to make passwords easier to remember.

Complete Your Open Source Library with Archive DVDs!

Each fully-searchable archive DVD includes past articles so you can find the content you need quickly.

Save hundreds off the print and digital copy rate with a convenient archive DVD!



Order Your DVD Now!
shop.linuxnewmedia.com



Calculating weekdays and dates with Go

Magic Day

Math wizards amaze audiences by naming the day of the week for any date called out by the audience. Mike Schilli writes a training program in Go to turn amateurs into headline performers. *By Mike Schilli*

Math geniuses can do this: Someone from the audience calls out “December 12, 2019,” and the numbers wizard announces “Thursday!” after just a few seconds. How did he do that? Does the entertainer have a photographic memory, or some kind of calendar function built into his head? The solution is surprisingly simple: He only has to go through a few rules that are easy to remember and, with a little bit of practice, can come up with the day of the week for any given date.

Years ago in this column, I introduced a similar mental arithmetic method for calculating weekdays, but with more elaborate steps [1]. A reader then replied that the method was unnecessarily com-

plex and referred me to the simpler Doomsday rule [2], which I will use here to create a Go training program for weekday prediction.

Last Day

According to the Doomsday rule, the following days of a year always fall on the same weekday: 5/9, 9/5, 11/7, and 7/11 (using the Month/Day format). You can easily memorize this with the formula “9-5 at 7/11” (i.e., the typical nine-to-five workday at 7-Eleven, the US convenience store chain).

Many other Doomsday days fall on day-month duplicates: 4/4, 6/6, 8/8, 10/10, and 12/12. January, February, and March are the only exceptions; in non-leap years, the day Doomsday falls on is January 3, February 7, and March 7

(Figure 1). In leap years, Doomsday changes to January 4 and February 8, while March 7 stays the same.

The Doomsday for 2019 is Thursday according to a separate procedure (2018 was a Wednesday; 2020 will be a Saturday; see Figure 2). So, if someone asks you for April 4, 2019, the answer is obvious: Thursday, because April 4 is the Doomsday.

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.



Month	Doomsday	Leap Year	Memory Aid
January	1/3	1/4	three years 3, in the fourth 4
February	2/7	2/8	
March	3/7		divisible by 7
April	4/4		
May	5/9		nine to five
June	6/6		
July	11/7		seven-eleven
August	8/8		
September	9/5		nine to five
October	10/10		
November	11/7		seven-eleven
December	12/12		

Figure 1: If you have these easy-to-remember rules in your head, you can calculate the day of the week for any date.

What about April 25, 2019? Which weekday was this date? Again, it's a Thursday, of course, because the 25th is exactly 21 days after the 4th (i.e., exactly three weeks later to the day). How about November 12, 2019? Because of the "9-5 at 7/11" rule, November 7 is a Thursday, so November 12 is five days (or one week minus two days) later, and therefore a Tuesday.

Or, you can either count the weekdays in your head, on your fingers, or by numbering the weekdays Sunday to Saturday from zero to six and then knocking off the remainder after dividing by seven to reach a result. Thursday is day four in this scheme; five added to it gives you nine, and after dividing by seven, you are left with two: So, it's a Tuesday.

Doomsdays for the Gregorian calendar													
Mon.	Tue.	Wed.	Thu.	Fri.	Sat.	Sun.	Mon.	Tue.	Wed.	Thu.	Fri.	Sat.	Sun.
1898	1899	1900	1901	1902	1903	→	1904	1905	1906	1907	→	1908	1909
1910	1911	→	1912	1913	1914	1915	→	1916	1917	1918	1919	→	1920
1921	1922	1923	→	1924	1925	1926	1927	→	1928	1929	1930	1931	→
1932	1933	1934	1935	→	1936	1937	1938	1939	→	1940	1941	1942	1943
→	1944	1945	1946	1947	→	1948	1949	1950	1951	→	1952	1953	1954
1955	→	1956	1957	1958	1959	→	1960	1961	1962	1963	→	1964	1965
1966	1967	→	1968	1969	1970	1971	→	1972	1973	1974	1975	→	1976
1977	1978	1979	→	1980	1981	1982	1983	→	1984	1985	1986	1987	→
1988	1989	1990	1991	→	1992	1993	1994	1995	→	1996	1997	1998	1999
→	2000	2001	2002	2003	→	2004	2005	2006	2007	→	2008	2009	2010
2011	→	2012	2013	2014	2015	→	2016	2017	2018	2019	→	2020	2021
2022	2023	→	2024	2025	2026	2027	→	2028	2029	2030	2031	→	2032
2033	2034	2035	→	2036	2037	2038	2039	→	2040	2041	2042	2043	→
2044	2045	2046	2047	→	2048	2049	2050	2051	→	2052	2053	2054	2055
→	2056	2057	2058	2059	→	2060	2061	2062	2063	→	2064	2065	2066
2067	→	2068	2069	2070	2071	→	2072	2073	2074	2075	→	2076	2077
2078	2079	→	2080	2081	2082	2083	→	2084	2085	2086	2087	→	2088
2089	2090	2091	→	2092	2093	2094	2095	→	2096	2097	2098	2099	2100

Figure 2: The table reveals that the Doomsday in 2019 is Thursday.

Easy Learning Method

What about January 1, 2020? Next year, the Doomsday is a Saturday, according to Figure 2, so January 4 (watch out – it's a leap year!) is a Saturday and New Year's Day thus a Wednesday. Slowly the mists clear, and the truth comes to light: There is no magic involved, just simple mnemonic rules that anyone can easily practice before going on stage.

To train would-be number wizards, the Go program presented here selects a random date in the current year and lets the user choose between seven weekdays. If you click on the right day after applying the formula in your head, you win a point, and the counter in the display's upper-right corner is incremented by one (Figure 3). The display changes to a new date, and the game resumes.

If, on the other hand, the player miscalculates and bets on the wrong day of the week, a penalty follows: All the points you scored so far expire, and the counter drops back to zero (Figure 4). Afterwards, you can try again and hopefully choose the right day of the week; again you score a point and can slowly climb to a new high score.

The game runs in a terminal user interface (UI) after you launch it at the

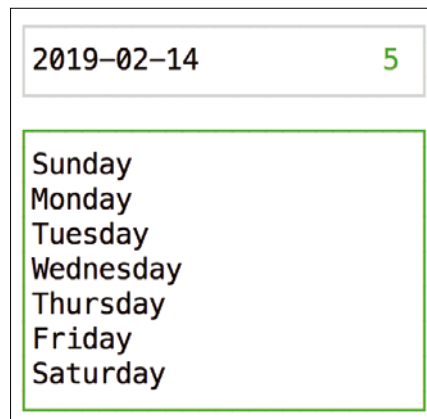


Figure 3: The player has determined the day of the week correctly and earns a point.




Figure 4: Oh no, wrong guess! The counter is reset to zero.

Listing 1: dateday.go

```

001 package main
002
003 import (
004     "errors"
005     "fmt"
006     ui "github.com/gizak/termui/v3"
007     "github.com/gizak/termui/v3/widgets"
008     "math/rand"
009     "strings"
010     "time"
011 )
012
013 var wdays = []string{"Sunday", "Monday",
014     "Tuesday", "Wednesday", "Thursday",
015     "Friday", "Saturday"}
016
017 func main() {
018     year := time.Now().Year()
019     wins := 0
020
021     if err := ui.Init(); err != nil {
022         panic(err)
023     }
024     defer ui.Close()
025
026     task := randDate(year)
027
028     p := widgets.NewParagraph()
029     p.SetRect(0, 0, 25, 3)
030     displayTask(task, wins, p)
031
032     days := widgets.NewParagraph()
033     days.Text = fmt.Sprintf(
034         "[%s](fg:black)",
035         strings.Join(wdays, "\n"))
036     days.SetRect(0, 3, 25, 12)

```

Listing 1: dateday.go (continued)

```

037 ui.Render(p, days)
038
039 uiEvents := ui.PollEvents()
040 for {
041     e := <-uiEvents
042     switch e.ID {
043     case "q", "<C-c>":
044         return
045     case "<MouseLeft>":
046         wdayGuess, err := wdayPick(
047             e.Payload.(ui.Mouse).Y)
048         if err != nil { // invalid click?
049             continue
050         }
051         wdayName := wdays[task.Weekday()]
052
053         if wdayGuess == wdayName {
054             days.BorderStyle.Fg =
055                 ui.ColorGreen
056             task = randDate(year)
057             wins++
058         } else {
059             days.BorderStyle.Fg = ui.ColorRed
060             wins = 0
061         }
062
063         displayTask(task, wins, p)
064         ui.Render(p, days)
065         go func() {
066             <-time.After(
067                 200 * time.Millisecond)
068             days.BorderStyle.Fg =
069                 ui.ColorWhite
070             ui.Render(days)
071             }()
072     }
073 }
074 }
075
076 func displayTask(task time.Time,
077     wins int, widget *widgets.Paragraph) {
078
079     widget.Text = fmt.Sprintf(
080         "[%d-%02d-%02d] (fg:black)" +
081         "%s[%3d] (fg:green)",
082         task.Year(), task.Month(), task.Day(),
083         " ", wins)
084 }
085
086 func wdayPick(y int) (string, error) {
087     if y > 10 || y < 4 {
088         return "", errors.New("Invalid pick")
089     }
090     return wdays[y-4], nil
091 }
092
093 func randDate(year int) time.Time {
094     start := time.Date(year, time.Month(1),
095         1, 0, 0, 0, time.Local)
096     end := start.AddDate(1, 0, 0)
097
098     s1 := rand.NewSource( // random seed
099         time.Now().UnixNano())
100     r1 := rand.New(s1)
101
102     epoch := start.Unix() + int64(r1.Intn(
103         int(end.Unix()-start.Unix())))
104     return time.Unix(epoch, 0)
105 }

```

command line. Even exhausted datacenter system administrators therefore can take a little break to relax. Go and the `termui` library, introduced in a previous column [3], run on all conceivable platforms including Linux, but also on other Unix derivatives and macOS – it even runs on Windows.

To create an executable binary from the Go code for Listing 1 [4], first create a new Go module (using `go-1.12` or later) and then start the compilation process with `build`; this automatically retrieves all libraries identified as dependencies off the web and compiles them, too:

```

go mod init dateday
go build dateday.go

```

As the installation process in Figure 5 shows, `go build` takes a whole bunch of libraries as source code from their GitHub repositories and bundles them all in one binary, which is not overly large at 2.8MB.

Opening and Closing

Line 6 of Listing 1 adds the code for the terminal UI library under `ui`. Its `Init()` function switches the terminal window into graphics mode in line 21 and delays a clean teardown until the end of the main program with `defer` in line 24. The UI in Figures 3 and 4 consists of two stacked `Paragraph` widgets from the `termui` widgets library.

The upper widget shows the date to be guessed; on the right-hand side, you

can see the number of successful guesses in the `wins` variable. The lower section shows a static string that displays the days of the week from Sunday through Saturday separated by newline characters. The `SetRect()` method sets the size of the widgets in rows and columns that each can hold precisely one character.

In order for the UI framework to render the widgets on the terminal interface, it notifies the rendering engine via `ui.Render()` in line 37. That's all there is to drawing the GUI. Line 39 then opens a channel with a call to `ui.PollEvents()`; it reports UI events like key presses, mouse clicks, or window resize actions. Line 41 blocks until an event occurs, while the subsequent switch statement checks whether the user has pressed `Ctrl + C` or `Q` (i.e., to end the program).

Terminal with Mouse Input

The `MouseLeft` event comes up if the user clicks and releases the left mouse button. The payload enclosed with the event indicates

the coordinates at which the mouse pointer was located at the time of the click. Line 47 needs to dynamically convert the event received as a generic `interface{}` type to a `ui.Mouse` type; it then uses `Y` to access the (vertical) `y`-value of the click position.

The `wdayPick()` function in line 86 determines the clicked weekday. It can do that because the program knows from the UI layout that Sunday is at the click position with a `y`-value of 4, Monday at 5, and so on through Saturday at position 10. The function discards any other position, returning an error, so that the main program simply ignores the click.

If the user has done their mental arithmetic correctly and selected the right weekday, the condition in line 53 is true,


```

$ go mod init dateday
go: creating new go.mod: module dateday

$ go build dateday.go
go: finding github.com/gizak/termui/v3/widgets latest
go: finding github.com/gizak/termui/v3 v3.0.0
go: downloading github.com/gizak/termui/v3 v3.0.0
go: extracting github.com/gizak/termui/v3 v3.0.0
go: finding github.com/cjbassi/drawille-go v0.0.0-20190126131713-27dc511fe6fd
go: finding github.com/mattn/go-runewidth v0.0.2
go: finding github.com/nsf/termbox-go v0.0.0-20190121233118-02980233997d
go: finding github.com/mitchellh/go-wordwrap v0.0.0-20150314170334-ad45545899c7
go: finding golang.org/x/arch v0.0.0-20181203225421-5a4828bb7045
go: downloading github.com/cjbassi/drawille-go v0.0.0-20190126131713-27dc511fe6fd
go: downloading github.com/mattn/go-runewidth v0.0.2
go: downloading github.com/nsf/termbox-go v0.0.0-20190121233118-02980233997d
go: downloading github.com/mitchellh/go-wordwrap v0.0.0-20150314170334-ad45545899c7
go: extracting github.com/mitchellh/go-wordwrap v0.0.0-20150314170334-ad45545899c7
go: extracting github.com/cjbassi/drawille-go v0.0.0-20190126131713-27dc511fe6fd
go: extracting github.com/mattn/go-runewidth v0.0.2
go: extracting github.com/nsf/termbox-go v0.0.0-20190121233118-02980233997d

$ ls -lh dateday
-rwxrwxr-x 1 mschilli mschilli 2.8M Jun  8 11:22 dateday

```

Figure 5: A newly created Go module automatically retrieves the source code for required libraries from GitHub during the build process and compiles it to create a worry-free binary.

and the `BorderStyle.Fg` attribute sets the weekday widget's frame to green. The following call to `randDate()` in line 56 fetches a new date task, and the success counter `wins` is incremented by one. However, if the user guesses wrong, line 59 colors the widget's frame red and sets `wins` back to zero. Life can be hard!

To refresh the success counter graphically, the call to `displayTask()` in line 63 changes the widget's contents, and the call `ui.Render()` with both widgets as parameters updates the current status in the terminal. To make sure that the red or green border that provides user feedback only appears briefly and then disappears again, line 65 starts a Go routine that runs in parallel; it first sleeps for 200ms thanks to `time.After()`, then resets the window's frame to the original white (which actually looks gray), and displays the whole thing with `ui.Render()`. That's how dynamics come into play, thanks simply to the standard parallelism features built into Go.

Text in Color

The text displayed in the widgets can also be colored using `termui` if so desired – not with an attribute like the widget borders, but using special tags in the text string on display. The date to be calculated appears in black thanks to this, and the number of successfully completed tasks in green; lines 80 to 81 include corresponding color commands in the text to be displayed with `(fg:black)` and `(fg:green)`.

Tricky Date Arithmetic

Selecting a random day of the year is more difficult than you might think. Sure, most years have 365 days, but in a leap year with 366 days, even February 29 should occur every now and then. Thanks to the `Duration` type from the `time` package, Go offers a method to calculate the time span between January 1 of the year investigated and the same date of the following year. Unfortunately, it refuses to express this in days, but uses hours instead.

The reason for this is the mess that occurs when a leap second [5] or the summer/winter time changeover occurs between two dates. Does that count as a fraction of a day or not? Go forces the programmer to multiply the hour difference by 24 to get the day number – as an indication that this may not be entirely true.

It is easier to use the Unix epoch time on Linux; this gives you the number of seconds that elapsed since January 1, 1970. It does not include leap seconds that have occurred, but gives the same timestamp to the times before and during a leap second. The `randDate()` function from line 93 in Listing 1 determines the Unix time for January 1st of the investigated year as well as the time stamp for January 1 of the following year and determines the time difference in seconds.

The `rand.Intn()` function then selects a random number between 0 (inclusive) and the timestamp of January 1 of the following year (exclusive) from the

`math/rand` package and adds the value to the start date. The result is a second timestamp sometime in the current year, which the `time.Unix()` function converts into a `time.Time` object from the Go standard library, whose month and day are output by `Month()` and `Day()`. Since Unix seconds are available as `int64` and the random functions from `math/rand` expect and return normal ints, line 102 needs to mediate between the two and convert the types accordingly.

The random number generator provided by Go still has the unpleasant feature of supplying the same random values whenever you call the program; this is unlikely to offer users any lasting training benefits in the long run. For this reason, line 98 taps into a new entropy source, initializes it with the current time in nanoseconds as the seed, and thus creates the new random number generator in the `r1` variable. Its `Intn()` method therefore comes up with new random sequences each time the `dateday` binary is called.

Stage Fright

The preparations are complete, and your training program can begin! Advanced students are also welcome to extend the program by training for several changing years with different Doomsdays. If you want to simulate the time pressure that a magician experiences in front of an impatient audience, you can add a progress bar [3] to the game as an additional stress factor; this leaves the candidate only a short time to choose the day of the week. If the time expires, the score drops to zero points, which you could even accompany by a game show klaxon sound by playing an appropriate sound file. Enjoy! ■■■

Info

- [1] "Perl Script Reveals Math Trick" by Mike Schilli, *Linux Magazine*, issue 87, February 2008, p. 80
- [2] Doomsday rule: https://en.wikipedia.org/wiki/Doomsday_rule
- [3] "Develop a DIY Progress Bar" by Mike Schilli, *Linux Magazine*, issue 220, March 2019, pp. 46-49
- [4] Listings for this article: <ftp://ftp.linux-magazine.com/pub/listings/linux-magazine.com/227/>
- [5] Leap second: https://en.wikipedia.org/wiki/Leap_second

PIRATEBOX
Handy anonymous file server for parties and meetups

MYCROFT
An open source personal assistant

Linux
MAGAZINE
JUNE 2018

MYCROFT
Can an open source personal assistant compete with Alexa?
Scanners and Linux
Peppermint OS
Peppy distro built for the cloud
Meltdown & Spectre
Kernel developers take swift action to protect Linux
Knight's Tour
Classic chess puzzle gets a Python twist

Node-RED
Control a Rasp Pi using text messages

Linux
MAGAZINE
JULY 2018

INTERNET OF THINGS
Tools for home automation

Internet of Things
Tools for home automation

Bluetooth Tricks
Track movements within your house

Linux on Old Hardware

Tablets on Linux
Configure graphics settings with xsetwacom

Raspberry Pi 3 B+
Does the world need a new Rasp Pi?

Python Gambling
Simulating games of chance

Logism
Design and digital circuits

Linux
MAGAZINE
AUGUST 2018

MAPPING TOOLS
Create and edit digital maps with QGIS and QMapShack

4 Open Source Microblogging Tools

9 Raspberry Pi Projects

QGIS AND QMAPSHACK
Tools for creating and customizing maps

9 MakerSpace Handpicked Pi projects

Kubuntu Up Close
Explore the new 18.04 release

Caddy
HTTP server out of the box

ODROID-C2
This single-board system is twice as fast as the competition

Ogg Vorbis
Free format for audio files

A Python Script
That solves the Chinese ring puzzle

LinuxVoice
FOSSPicks
Tutorials

soundConverter and freac
Full-featured audio converters

systemd
TRICKS WITH THE NEW LINUX INIT SYSTEM

WILL SYSTEMD REPLACE CRON?

Linux
MAGAZINE
OCTOBER 2018

PRIVACY
STOP SNOOPERS AND PROTECT YOUR IDENTITY

PRIVACY
snoopers and protect identity

your music player with free firmware

organize all your news sources

spaces for your Rasp Pi

Ready to Go?
Smart queries in the powerful Go language

HUGE SAVINGS!
\$29.99 VALUE!

ON A SEARCHABLE DISC!

ALL 214 ISSUES

Linux
MAGAZINE
NOVEMBER 2018

.NET CORE
The benefits of dotnet on Linux

.NET CORE
Microsoft's flagship work on Linux

little-known OS to compete with BSD?

and-mapping Tools
suaize decisions and organize your best ideas

Counting Pennies
Tabulating audio feedback with a Raspberry Pi

Electron Framework
Cross-platform apps with JavaScript

ExaGear
Run x86 apps on Rasp Pi and other ARM systems

Linux
MAGAZINE
SEPTEMBER 2018

Systemd
Getting more from the elusive new Linux init system

10 Top PDF Readers

Arduino Programming for Open Hardware Projects

Homegrown Facial Recognition
Build facial recognition into your own Python scripts

Maker Tricks
Monitor a beehive with a Raspberry Pi

Cover Your Ass(ets)
Back up your system files with CIA

FOSSPicks
Cutter Disassembler
Circi MIDI Editor

Tutorials
Using science with BIONIC

Linux
MAGAZINE
OCTOBER 2018

PRIVACY
STOP SNOOPERS AND PROTECT YOUR IDENTITY

PRIVACY
snoopers and protect identity

your music player with free firmware

organize all your news sources

spaces for your Rasp Pi

Ready to Go?
Smart queries in the powerful Go language

FOSSPicks
cheat.sh
HyperRoom

Linux
MAGAZINE
DECEMBER 2018

INNOVATIVE DISTROS
Ultra-safe Linux with runtime protection

these hidden gems:
A simpler enlightenment
Anonymous surfing even onware
UI designers get creative with KDE
OS - Lockdown security
e focus on isolation

Filesystem
w into a g system

cheat.sh
Syntax tips at your fingertips

Git Workshop
Working with remote repositories

Purism LibreL
Linux laptop special security features

Upribox 2
Use a Rasp Pi to filter ads and trackers

Linux
MAGAZINE
JANUARY 2019

NAVIGATION
Free services that extend OpenStreetMap

Kernel Disposal
Are old kernels sucking up valuable space on your hard drive?

Build a Robot Car

Zorin OS
Will this well-appointed Linux succeed in attracting Windows users?

WebAuthn
Authenticate without a password

OOM Killer
What happens when kernel runs out

Put It On!
Build your own wearable with Android and MIT's App Inventor

FOSSPicks
Whalebird
musicube
Imaginary Teleport

Tutorial
Getting started KDE's Plasma

Linux
MAGAZINE
FEBRUARY 2019

CUSTOMIZE THE BOOT MENU
Clean up your startup for clarity and fewer mistakes

Wipe SSDs
Make the storage fit your hardware

IoT Tricks
Store data in memory with Redis

Exploring Ubuntu 18.10

Pulse Sensor
Measure your heartbeat with a Raspberry Pi

EncryptPad
Text editor with easy encryption

Linux
MAGAZINE
MARCH 2019

VirtualBox Hacks
Save time and extend the power of your virtual machines

Scapy
Automate packet analysis with Python

Put a Recording Studio on a Raspberry Pi

Elementary OS
Elegant Linux with an ambitious vision

Create a Cartoon
With open source tools

Tools for Writers
Find words and organize your thoughts

PDF Tricks
Clean up blotches and fading text

FOSSPicks
x86-UI
Aurpo
Calculator

Tutorials
Bash Arrays
Nattion Video Tricks

Go For It! Easy Tool for building ToDo lists

maddog, Diligence and other FOSS resolutions for 2019

WWW.LINUX-MAGAZINE.COM

Linux
MAGAZINE
APRIL 2019

INNOVATIVE DISTRO
These hidden gems:
A simpler enlightenment
Anonymous surfing even onware
UI designers get creative with KDE
OS - Lockdown security
e focus on isolation

Filesystem
w into a g system

cheat.sh
Syntax tips at your fingertips

Git Workshop
Working with remote repositories

Purism LibreL
Linux laptop special security features

Upribox 2
Use a Rasp Pi to filter ads and trackers

Linux
MAGAZINE
MAY 2019

NAVIGATION
Free services that extend OpenStreetMap

Kernel Disposal
Are old kernels sucking up valuable space on your hard drive?

Build a Robot Car

Zorin OS
Will this well-appointed Linux succeed in attracting Windows users?

WebAuthn
Authenticate without a password

OOM Killer
What happens when kernel runs out

Put It On!
Build your own wearable with Android and MIT's App Inventor

FOSSPicks
Whalebird
musicube
Imaginary Teleport

Tutorial
Getting started KDE's Plasma

Linux
MAGAZINE
JUNE 2019

CUSTOMIZE THE BOOT MENU
Clean up your startup for clarity and fewer mistakes

Wipe SSDs
Make the storage fit your hardware

IoT Tricks
Store data in memory with Redis

Exploring Ubuntu 18.10

Pulse Sensor
Measure your heartbeat with a Raspberry Pi

EncryptPad
Text editor with easy encryption

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



Setting up a local DNS server with Unbound

Name Caddy



You don't have to be satisfied with your ISP's slow and cumbersome DNS server. Your own Unbound server could improve performance as well as security. *By Rubén Llorente*

When you turn your home computer on, launch a web browser, and instruct it to visit the `linux-magazine.com` website, your computer sends a DNS query, asking for the IP address associated with the name `linux-magazine.com`. For many users, this query is sent to a DNS server provided by the user's Internet Service Provider (ISP).

Using your ISP's DNS server is an easy and low-stress option, but in many situations, it also has some disadvantages. The most popular reason why some users prefer a non-default DNS server is performance. Simply put: some servers have lower latency and faster query times than others. If your ISP's DNS servers are slow, switching to faster servers will lead to noticeable improvement in your web browsing experience.

Another reason for switching to a different server is to avoid (or enforce) soft censorship. For instance, a school administrator might wish to prevent students from accessing social networking sites such as `facebook.com` during the school day. The easiest way to prevent a user from reaching a website is to instruct the DNS server to return a bogus address or to return an `NXDOMAIN` message, which means the server doesn't think the domain exists. Another option is for the server to return the address of a webpage that displays a message such as "No Social Networking Allowed Here."

Many DNS providers offer anti-advertisement, anti-malware, and anti-phishing protection in such a way that, if your browser tries to resolve the address of some service known to serve advertisements or harmful code, it will be redirected to a bogus address or a site with a warning. Parental controls that filter sites deemed unsafe for kids are also offered by some DNS providers. See Table 1 for a list of some third-party DNS providers.

A third-party provider could help you with performance and parental control, but if you want to customize the DNS environment, you will need to set up your own server.

Running your own DNS server in your own premises gives you a lot of flexibility. You could also get a performance boost. Imagine that there are four family

members in a home LAN, and two are browsing the Internet at the same time. One person visits `linux-magazine.com`. Then a different person, sitting at a different computer, visits the same site. If this LAN has a local DNS server, the server that resolves the address for the first user could cache the address. When the second user visits the same site, the server could provide the address from the cache without having to pull the information from outside the network.

Running your own DNS servers allows you to have custom DNS entries. You can create your own blacklists (that block advertisements, for example). You can also assign names to local resources, such as your printer, your NAS, or your IP cameras.

Another benefit of a local DNS server is that it lets you take advantage of

Table 1: Publicly Available DNS Servers

Address	Operator	Stated Purpose
37.235.1.174	FreeDNS	Offering DNS services that bypass ISP tampering.
37.235.1.177	FreeDNS	Offering DNS services that bypass ISP tampering.
8.8.4.4	Google	Improving DNS query times.
8.8.8.8	Google	Improving DNS query times.
208.67.222.222	OpenDNS	Improving DNS query times.
208.67.220.220	OpenDNS	Improving DNS query times.
208.67.222.123	OpenDNS	Blacklisting adult content.
208.67.220.123	OpenDNS	Blacklisting adult content.
8.26.56.26	Comodo	Blacklisting malware and harmful domains.
8.20.247.20	Comodo	Blacklisting malware and harmful domains.
84.200.69.80	DNS.WATCH	Offering uncensored DNS services.
84.200.70.40	DNS.WATCH	Offering uncensored DNS services.

Lead Image © Natalia Natykyach, 123RF.com

DNSSEC [1], which still has not been implemented by many ISPs. DNSSEC is a security overlay that protects users from having DNS traffic altered by mali-

cious actors. Its actual usefulness is disputed, but some users prefer the protection of DNSSEC. This article describes how to set up your own DNSSEC-aware

DNS configuration using the Unbound DNS server.

Enter Unbound

BIND is the undisputed king of free and open source DNS servers. The BIND DNS server is feature rich, well documented, and available on most distributions. On the other hand, BIND is too heavy for most small LANs, and it has some significant security concerns.

Unbound is a non-authoritative, recursive DNS server, with support for DNSSEC validation (see the box entitled “Authoritative Servers and Recursive Servers.”) It is included in the default installation of the OpenBSD operating system and is available on the repositories of most serious Linux distributions. It is easy to configure and quick to set up.

Usually, local system administrators install their DNS servers on machines that run 24 hours a day. In a home environment, any machine that is continuously running torrent software, or a web proxy, or any infrastructure service, will be a perfect candidate. A small Single Board Computer (SBC) such as a Raspberry Pi will also do the trick, if you don't want to assign the task to a full featured computer.

In Devuan, you can install Unbound using the apt-get utility:

```
# apt-get install unbound
```

Services in Devuan are started upon installation by default. You might want to stop it right away and begin with the configuration:

```
# /etc/init.d/unbound stop
```

One of your first steps should be to install a root hints file. This file contains the location of the root DNS servers of the Internet. Unbound itself comes with a hardcoded set of root servers, but it is always a good idea to provide it with an up-to-date set, just in case. The hardcoded set is prone to become outdated.

You can obtain an updated list as follows:

```
# cd /etc/unbound
# curl -Lo /etc/unbound/root.hints https://www.internic.net/domain/named.cache
```

Authoritative Servers and Recursive Servers

Authoritative DNS servers [2] are servers designed to answer DNS queries pertaining to a specific DNS zone. A DNS zone is a distinct portion of a DNS hierarchy (usually the global Internet one) that has been placed under the administration of a particular entity. This is to say, it is the sort of server Verisign uses to administrate the .com TDL and ensure that sites that use the .com zone can be found on the Internet. Definitely not the sort of DNS server you use in your home LAN.

There are 13 very special authoritative DNS servers, known as the root DNS servers. A root DNS server is a name server located in the root DNS zone, which is the highest zone in the DNS hierarchy (Figure 1). The root DNS servers are authoritative for the whole DNS space.

In theory, the whole Internet could be managed by authoritative servers alone. A DNS resolver that has to depend exclusively on authoritative servers would first contact one of the root DNS servers. If the resolver were trying to solve, as an example, www.example.org, the root DNS server contacted would answer “I don't know the address of www.example.org, but this other authoritative server does” and guide it to some server responsible of the .org TDL. When the resolver asked this new server for an answer, it would either get a result or be told to look for an answer at some other authoritative server placed below in the DNS hierarchy (Figure 2).

If every resolver and DNS server worked like that, the Internet would collapse under the traffic load, since a very small number of root servers and authoritative servers for TDL zones would have to answer so many queries.

The sort of DNS server ISPs provide to their customers is usually a recursive caching server. Recursive servers are designed to answer DNS queries without overloading the root servers. When a DNS resolver asks a recursive server “where is www.example.org?”, the server will ask other DNS servers on behalf of the resolver, and return an already valid and final answer. Since most recursive servers also have big caches, they may be able to extract an answer for the DNS resolver from their caches without having to send additional queries to other DNS servers.

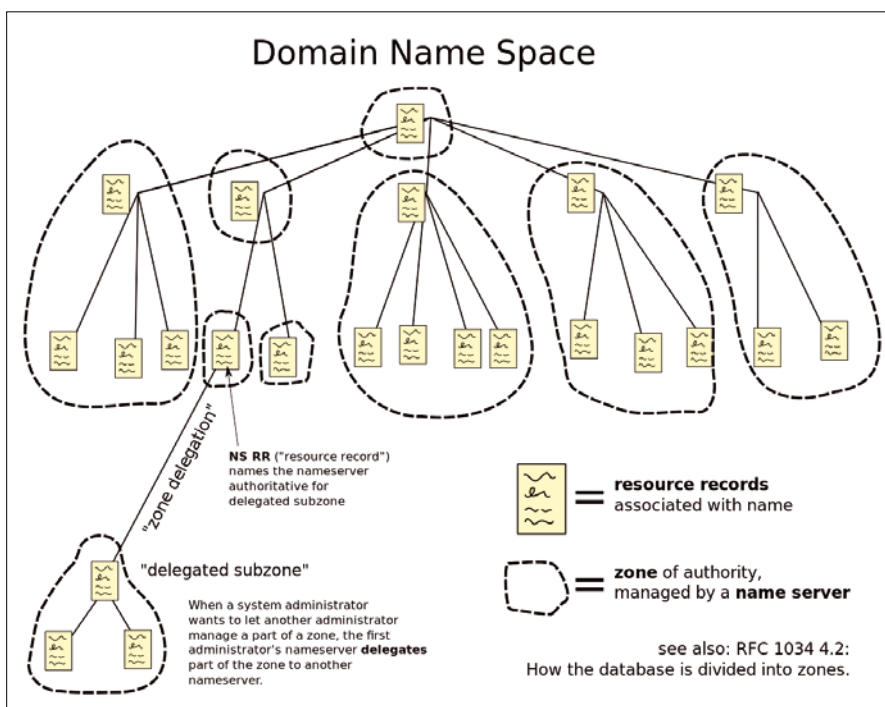


Figure 1: DNS is a hierarchical system. Zone administrators delegate parts of their zones to other administrators below them in the hierarchy.

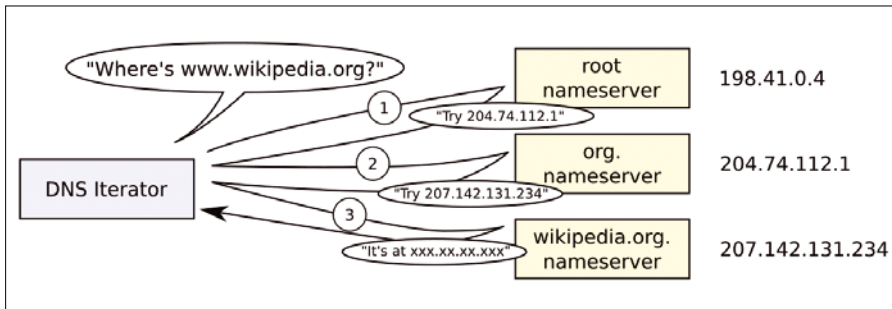


Figure 2: An iterative DNS query against authoritative DNS servers.

This file should be updated periodically. You can execute the update through a cron job or manually. A 6-month interval is adequate for home purposes.

Be careful which root list you use! If your list contains malicious servers, you will be vulnerable to certain attacks, like having DNS queries answered with malicious results.

Listing 1 shows a possible configuration file for Unbound. The `/etc/unbound/`

Listing 1: unbound.conf

```
# /etc/unbound/unbound.conf
# The include directive loads
# configuration parameters
# from the indicated files.
server:

include: /etc/unbound/unbound.conf.d/access_options.conf
include: /etc/unbound/unbound.conf.d/name_solving.conf
include: /etc/unbound/unbound.conf.d/privacy_options.conf
include: /etc/unbound/unbound.conf.d/cache_options.conf
include: /etc/unbound/unbound.conf.d/dnssec_options.conf
include: /etc/unbound/unbound.conf.d/blacklist.conf
include: /etc/unbound/unbound.conf.d/local_names.conf
include: /etc/unbound/unbound.conf.d/opennic_names.conf
include: /etc/unbound/unbound.conf.d/forwarders.conf

remote-control:
    #Disable remote control which we don't intend to use.
    control-enable: no
```

Listing 2: name_solving.conf

```
# /etc/unbound/unbound.conf.d/name_solving.conf

root-hints: /etc/unbound/root.hints
```

Listing 3: dnssec_options.conf

```
# /etc/unbound/unbound.conf.d/dnssec_options.conf

auto-trust-anchor-file: "/var/lib/unbound/root.key"
domain-insecure: "dyn."
domain-insecure: "geek."
```

`unbound.conf` file loads different directives from files located in `/etc/unbound/unbound.conf.d`.

Listing 2 shows an example of how to configure Unbound to use the root hints file. The `root-hints` parameter tells the server the name of the root hints file.

Enabling DNSSEC

DNSSEC is a protocol designed to prevent man-in-the-middle attacks against DNS queries that could result in the DNS resolver getting a maliciously crafted DNS response. See the box "How Does DNSSEC Work?" for more information.

The DNSSEC trust anchor can be retrieved and initialized by issuing the following command:

```
# unbound-anchor 🔗
-a "/var/lib/unbound/root.key"
# chown unbound:unbound 🔗
"/var/lib/unbound/root.key"
# chmod 500 🔗
"/var/lib/unbound/root.key"
```

Listing 3 displays an example configuration that enables DNSSEC validation. The `auto-trust-anchor-file` parameter points at the root key of the DNSSEC trust anchor. This key is stored in a file. As long as the user running the Unbound server has read and write permissions over the trust anchor file and the folder that contains it, Unbound itself will be able to keep the trust anchor updated automatically. The `domain-insecure` parameters in Listing 3 disables DNSSEC verification for the unofficial DNS zones `.dyn` and `.geek`.

Using Multiple Forwarders

By default, your Unbound server will try to resolve DNS entries by asking the root DNS servers for an answer. This is OK, but sometime you want to adopt a different approach.

Listing 4: opennic_names.conf

```
# /etc/unbound/unbound.conf.d/opennic_names.conf

forward-zone:
    name: "dyn."
    forward-addr: 193.183.98.66 #opennic
    forward-addr: 87.98.175.85 #opennic
    forward-addr: 5.135.183.146 #opennic
    forward-addr: 51.254.25.115 #opennic

forward-zone:
    name: "geek."
    forward-addr: 193.183.98.66 #opennic
    forward-addr: 87.98.175.85 #opennic
    forward-addr: 5.135.183.146 #opennic
    forward-addr: 51.254.25.115 #opennic
```

Listing 5: forwarders.conf

```
# /etc/unbound/unbound.conf.d/forwarders.conf

forward-zone:
    name: "."
    forward-addr: 84.200.69.80 #DNS.WATCH
    forward-addr: 84.200.70.40 #DNS.WATCH
    forward-addr: 51.254.25.115 #opennic
```


How Does DNSSEC Work?

The whole DNS system was designed in such a way that every connection was sent unencrypted over the network. Worse yet, there was no way to ensure that the traffic between DNS servers and DNS resolvers was not tampered with. For example, if you want to visit `mybankwebsite.com` and your web browser tries to resolve that domain, it is theoretically possible for some evil entity to intercept that DNS request and give you a fake response.

DNSSEC is an extension for the DNS protocol that lets DNS entries be cryptographically signed in such a way that the validity of a DNS response can be verified.

DNSSEC uses PKI (Public Key Infrastructure) in order to accomplish its goal. Each zone (including the Root Zone) has its public/private keypair. Keypairs are certified in a hierarchical way. The keypair of the `.com` zone has been signed with a key belonging to the administration of the root zone. The children zones directly below the `.com` zone are signed by the administrator of the `.com` zone, and so on. This builds what is called a chain of trust, and the public key that sits at the top of the hierarchy is known as the trust-anchor – it has no superior entity certifying its validity.

Most DNSSEC-aware resolvers use the root DNS public keying as their trust-an-

chor. When such a resolver needs to validate a DNS entry, it asks for the public keys of the administrator of its zone to the adequate authoritative server. If this public key needs validation, the resolver asks for the public key of its parent zone to an authoritative server of the parent zone, and this process continues until the trust-anchor is reached. Once the whole chain is fetched, it can be validated. After the validity of the chain is accepted, the public key at the end of the line may be used to check the authenticity of the cryptographic signature attached to the DNS entry.

DNSSEC has many drawbacks [3][4]. For one, it is argued that private keys related to popular TDLs (like `.com` or `.net`) are indirectly under the control of governments and are therefore compromised. Also, having to retrieve the public keys necessary for validating the chain increases the time the DNS query takes. Finally, many DNS zones have not enabled DNSSEC at all, which means that the domains under their management cannot have their validity verified.

This has led to the creation of many alternatives, with different degrees of success and acceptance. DNS over TLS [5] and DNSCrypt [6] are two of the best known.

You may want to have a list of recursive DNS servers to send DNS queries to, like, for example, the ones provided by your ISP. You can set forward zones, which designate DNS zones that will have queries related to them forwarded to servers you designate.

Suppose that, for any reason, you want all your queries for domains in the `.geek` zone sent to an OpenNIC server instead of to the root DNS servers. The example configuration at Listing 4 ensures that queries for domains under the `.geek` and `.dyn` extraofficial TDLs are forwarded to a DNS server of the OpenNIC Project.

Why is this feature useful? For two main reasons. The first one is that you might want to connect to an alternative, unofficial DNS infrastructure, such as the OpenNIC Project [7], which carries unofficial TDLs that are not recognized by the IANA.

The second reason is that you may not want to use the root servers at all. In fact, end users are not expected to use them and depend on recursive servers instead. The example at Listing

5 instructs Unbound to use two recursive DNS servers from DNS.WATCH [8] when attempting to resolve any domain name located under the root zone, which, in practice, means every domain.

In any case, remember that, as long as DNSSEC is enabled, DNSSEC verification will be attempted in any DNS query Unbound initiates, including queries that use forward servers. If you select forward servers that don't support DNSSEC, the queries will fail because verification is not possible, and your whole name resolution system will become inoperative. As demonstrated in Listing 3, you can disable DNSSEC validation for domains and zones that are problematic.

Blacklists

One big reason for running your own DNS server is to be able to blacklist sites you don't want the users of your LAN to visit. In a home environment, that's advertisers. In a small office, that might include time-wasting sites, such as social networks or digital sport journals. Your

easiest option is to return bogus addresses or NXDOMAIN messages when asked about domains you don't want users to visit.

Something important to take into account is that DNS blacklisting is easy to set up but also very easy to bypass. Users in your LAN may try to configure their computers to use a different DNS server, use Tor, set up a VPN, or use a web proxy. A user can also bypass DNS if they already know the target IP address. A DNS blacklist thus works best when it is combined with other measures.

A simple way of blacklisting a domain is to add an entry like the following to your Unbound configuration:

```
local-zone: "example.org" ↗
always_nxdomain
```

When a client asks the Unbound server where `example.org` is, it will get an NXDOMAIN response.

Adding hosts manually to the configuration files can be tiresome. If you want to have good malware, phishing, and advertisement protection, getting an existing list of bad domains and adapting the list to Unbound is a good start. Many good lists of bad domains exist on the Internet. The StevenBlack blacklist [9] is very complete, so I will use it as a demonstration. The following commands will download the list and convert it to Unbound format:

```
$ curl -o hosts ↗
https://raw.githubusercontent.com/↗
StevenBlack/hosts/master/hosts
$ su
[password]
# grep '^0\.\0\.\0\.\0' hosts |
awk '{print "local-zone: ↗
\""$2"\"always_nxdomain"}' > ↗
/etc/unbound/unbound.conf.d/↗
blacklist.conf
```

The Steven Black site has some tools for customizing the list, which are totally worth the time it takes to check them out.

Configuring Local Zones

Suppose you have a printer in your LAN. You can connect to that printer by using its known IP address, like, for example, `192.168.1.2`. However, wouldn't you

Listing 6: local_names.conf

```
# /etc/unbound/unbound.conf.d/local_names.conf

private-address: 192.168.1.0/24

local-zone: "mylan.dyn." static

    local-data: "gateway.mylan.dyn. IN A 192.168.1.1"
    local-data: "printer.mylan.dyn. IN A 192.168.1.2"
    local-data: "computer.mylan.dyn. IN A 192.168.1.3"
    local-data: "server.mylan.dyn. IN A 192.168.1.100"

    local-data-ptr: "192.168.1.1 gateway.mylan.dyn"
    local-data-ptr: "192.168.1.2 printer.mylan.dyn"
    local-data-ptr: "192.168.1.3 computer.mylan.dyn"
    local-data-ptr: "192.168.1.100 server.mylan.dyn"
```

rather give a human readable name to that printer?

Unbound is not an authoritative server, so it cannot manage a full zone with all its bells and whistles directly. However, it has horsepower enough for managing a small home LAN. Listing 6 shows an example configuration for a home LAN zone. It assumes that the LAN is using 192.168.1.0/24 as the network.

The `private-address` directive prevents addresses in your LAN from being returned for public Internet names. This step prevents DNS rebinding attacks [10].

The `local-zone` directive defines all domains under `mylan.dyn` as local. The `static` word means that the static entries defined in the configuration file are used as DNS entries. Each of the `local-data` entries assigns a name to an address. For example, 192.168.1.2 would be assigned the name `printer.mylan.dyn`. If you queried the Unbound server for a name in the `mylan.dyn` zone that did not exist, it would be answered with a `NXDOMAIN` message. Alternatively, `transparent` could be used in-

Listing 8: cache_options.conf

```
# /etc/unbound/unbound.conf.d/cache_options.conf

prefetch: yes # Fetch things before they expire from cache.
prefetch-key: yes # Fetch DNSSEC keys early in the validation process.
cache-min-ttl: 1200 # Seconds it takes for items in cache to die at minimum.
```

Listing 9: privacy_options.conf

```
# /etc/unbound/unbound.conf.d/privacy_options.conf

hide-identity: yes # If enabled id.server and hostname.bind queries are refused.
hide-version: yes # If enabled version.server and version.bind queries are refused.
qname-minimisation: yes # Send minimum amount of information to upstream servers to enhance privacy.
```

Listing 7: access_options.conf

```
# /etc/unbound/unbound.conf.d/access_options.conf

access-control: "0.0.0.0/0" allow
access-control: "127.0.0.0/8" allow
access-control: "192.168.1.0/24" allow
```

stead of `static`. A transparent local zone is one in which the server tries to resolve the name of a host by other means if it has no static entry for it in its configuration.

The `local-data-ptr` entries are optional and define reverse DNS information. Reverse DNS is, as the expression implies, the opposite of DNS. A reverse DNS query asks “What is the name of the host with the address 192.168.1.2?”

Configuring Access

Listing 7 shows how to grant access to the Unbound server to hosts on your LAN and to the machine running the server. This example assumes that the LAN sits at 192.168.1.0/24.

There are many good reasons for restricting access to your DNS server. The first one is that a DNS server may be used as part of a denial of service attack. A common technique is to send queries with spoofed IP addresses to exposed recursive DNS servers, which will send their responses to what they think is the computer that made the query in the first place. In practice, it means that an attacker can ask the recursive server for a DNS record using a fake IP, and the owner of the IP address that was faked will get the response. This means that an

evil entity can force a recursive server to flood a victim with DNS responses and therefore use the server as a proxy for a denial of service attack. Another reason is that a local DNS server might contain sensitive DNS entries that are not intended to be known by outsiders. If you are using a local zone for naming local resources, such as printers, cameras, and NAS servers, it is better to have that information protected from outsiders.

In addition to the Unbound configuration presented here, it is a good idea to block access to your DNS server by using appropriate firewall rules. DNS servers listen for queries at port 53 and may support both UDP and TCP.

The `access-control` directives are self-explanatory.

Cache Configuration

For most home users, the best reason for using a local DNS server is caching DNS entries and speeding up web browsing. Listing 8 shows an example configuration for a simple, yet powerful, DNS cache.

Domain names have an official *time to live* assigned by the manager of that domain. This time to live is the time that recursive DNS servers are supposed to keep the DNS entries in their caches before deleting them. Imagine that your Unbound server resolved `richard-falken.com`, which at the time of this writing has a TTL of 86400 seconds. It would remember the DNS entry for `richard-falken.com` for 24 hours.

Some domains have very short TTLs. The `cache-min-ttl` directive in the example defines the minimum time a cached DNS entry will be allowed to live in the cache. If Unbound comes across a domain with a TTL shorter than 1200 seconds, the official TTL will be ignored, and 1200 seconds will be used instead. Beware that DNS entries that are conserved in the cache for too long may become stalled and outdated, which

could be counterproductive and lead to problems. Use this directive wisely.

`prefetch-key` instructs Unbound to fetch DNSSEC keys earlier than usual in the DNSSEC validation process. It saves time at the expense of CPU load.

`prefetch` instructs the server to try to resolve cached entries that are about to expire from the cache in order to keep the cache fresh. This option might increase bandwidth consumption by about 10 percent, but response times will be better.

Privacy

Listing 9 has some privacy options. The most important is `qname-minimisation`. Enabling this option makes the queries sent by Unbound to other DNS servers more compact and less prone to leak private information.

`hide-identity` and `hide-version` are less relevant in LAN scenarios, since they prevent the DNS server from replying to internal special queries that attempt to obtain information, such as the hostname of the server or the software version.

Query Times

You may want to check the query times for a server provided by your ISP, another open DNS server on the Internet, or any server you configure on your own LAN.

Proper DNS benchmarking is difficult. A quick and dirty way to check the connection latency for a DNS server is using `icmp echo` requests with the ping utility:

```
$ ping -c 4 $address_of_server
```

```
bash-4.3$ dig @8.8.4.4 operationalsecurity.es
; <<>> DiG 9.10.8-P1 <<>> @8.8.4.4 operationalsecurity.es
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 18586
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 0

;; QUESTION SECTION:
;operationalsecurity.es.                IN      A

;; ANSWER SECTION:
operationalsecurity.es. 3600    IN      A      188.76.45.99

;; AUTHORITY SECTION:
operationalsecurity.es. 86400  IN      NS     ns15.piensasolutions.com.
operationalsecurity.es. 86400  IN      NS     ns16.piensasolutions.com.

;; Query time: 308 msec
;; SERVER: 8.8.4.4#53(8.8.4.4)
;; WHEN: vie mar 29 12:12:24 CET 2019
;; MSG SIZE rcvd: 113
```

Figure 3: This command uses `dig` to test how quickly server 8.8.4.4 can give an answer.

The time stat for each reply indicates the time it took to get a “pong” response from the server. The bigger the time, the longer it takes for the server to reply to you when you ask it to.

The `dig` utility (Figure 3) performs name resolution, and it is very useful for retrieving DNS records and checking query times.

```
$ dig @$address_of_dns_server 2
example.org
```

If you are serious about benchmarking a DNS server, you will need heavier tools for the job. Google’s `namebench` [11] is a popular option.

Finishing Up

When your configuration is ready, just boot your server up with the following command:

```
# /etc/init.d/unbound start
```

This command will bring up a caching DNS server with anti-advertisement, anti-malware, and anti-phishing capabilities, as well as a limited capability for validating the authenticity of the DNS responses it takes. Not bad at all.

The only task left to do is to configure the devices in your LAN to use this server. There are three main ways to do this. You can manually configure each one to use your DNS server, which is usually impractical. The second option is to configure your network router to assign your server via DHCP to each de-

vice that connects to the network. This is the easiest way, and it will work most of the time, as long as your router supports assigning custom DNS servers to the devices in your LAN.

There is a third, evil option, that I like to use when I have to deal with devices that will ignore DNS servers provided by DHCP. This method is to instruct the router to redirect queries to unauthorized DNS servers to your local DNS instance, using `ds-nat`. In practice, you are performing a man-in-the-middle attack by letting your local DNS server pretend to be the DNS server the rogue device is trying to connect to. This requires a router capable of advanced firewall configuration. But such mean deeds are the subject for another article. ■■■

Info

- [1] What is DNSSEC? <https://www.icann.org/resources/pages/dnssec-what-is-it-why-important-2019-03-05-en>
- [2] What Is the Difference between Authoritative and Recursive DNS Nameservers? <https://umbrella.cisco.com/blog/2014/07/16/difference-authoritative-recursive-dns-nameservers/>
- [3] Against DNSSEC: <https://sockpuppet.org/blog/2015/01/15/against-dnssec/>
- [4] Questions and Answers from “Against DNSSEC”: <https://sockpuppet.org/stuff/dnssec-qa.html>
- [5] DNS over TLS Wikipedia article: https://en.wikipedia.org/wiki/DNS_over_TLS
- [6] DNSCrypt official site: <https://dnscrypt.info/>
- [7] OpenNIC Project: <https://www.opennic.org/>
- [8] DNS.WATCH official site: <https://dns.watch/>
- [9] StevenBlack DNS Blacklist: <https://github.com/StevenBlack/hosts>
- [10] Practical Attacks with DNS Rebinding: <https://www.tripwire.com/state-of-security/vert/practical-attacks-dns-rebinding/>
- [11] Namebench: <https://code.google.com/archive/p/namebench/>

Author

Rubén Llorente is a mechanical engineer whose job is to ensure that the security measures for the IT infrastructure of a small clinic are law compliant and safe. He is also an OpenBSD enthusiast and a weapon collector.

Try out the Jade desktop environment

Jade for Jaded Users



This new desktop environment strives to offer something different, and while the design isn't for everyone, it does introduce a couple of interesting features. *By Bruce Byfield*

2019 is not a year for innovation in the desktop environment. Ever since the desktop revolts over KDE 4, Gnome 3, and Unity in 2008-2012, developers have been cautious about innovation and alienating users with too much change. Any innovations have been incremental or minor. Under these circumstances, Jade [1] (which stands for Just Another Desktop Environment) is a welcome development. Although it suffers from a lack of layout knowledge, at least Jade tries and manages one or two promising features along the way.

Jade was started by Vitor Lopes and developed within the Manjaro distribution. To date, Jade is unavailable in any

other distribution, but, since it was developed using standard web-technologies like HTML5, CSS, JavaScript and Python, porting it should be a trivial task, and only a matter of time. On the Manjaro forum [2], it was announced in 2017 as “a completely different DE [desktop environment] concept, that changes the way you interact with your desktop, is made to be easy to use, independently of your computer skill.” According to Lopes, he began the product to “learn Python” and to “keep my coding skills sharp.”

Currently, Jade is available in the Manjaro Webdad Preview [3] version 17.1.11-stable. The preview runs in VirtualBox as an Other Linux – and not as a version of Arch Linux, as you might

expect if you know the origins of Manjaro. In fact, if you try to install it as an Arch variant, the installation may stall, even if you use the fallback theme as suggested. The installer appears to be a modified version of Ubuntu's installer, with the addition of usefully verbose on-line instructions. The only oddity is that, when allotted 15GB for the installation, the automatic partitioning creates a SAP file of 15GB, which seems excessive.

The DVD image is live. However, it is so slow that, even on a recent machine, it took 13 seconds to reach the login screen and another 15 to reach the desktop. Even then, I found it so unresponsive that several times I thought the desktop had frozen. The speed improved somewhat after installation to a drive, although even then I found it so unresponsive that several times I thought the desktop had frozen when I clicked to start applications. Jade is still in development, and the point of installing it is not to use it so much as to see how it is designed (Figure 1).

Touring the Desktop

Jade installs with mostly Gnome technologies plus the Chromium browser, a few Manjaro utilities, and a link to Steam for games. Overall, it is a standard collection, with little that is unusual. The interest lies in the desktop design, rather than the applications. Rather than spending time




Figure 1: Jade's desktop sports a minimalist look, with many features hidden at first glance.




Figure 2: Jade's login screen contains the usual features, but with sometimes cryptic icons.

on the applications, I suggest that curious users log in with the password *manjaro*, bypass the welcome screen, and turn their attention immediately to the desktop (Figure 2). It is not hard to navigate, but superficial appearances are deceiving.

The first impression is of a minimalist desktop. Many of the standard desktop elements appear to be missing. There are no panels as such, although various links are arrayed across the top of the screen, and a Cairo dock is on the bottom. These links are the usual icons for a calendar, sound, Ethernet connection, and updates. In addition, on the top right, there is the useful addition of the available drive space displays.

At first, the menu appears to be missing as well. Instead, it appears to be dispersed to a list of favorites in the dock and a similar one in the center of the screen. Only when you click on the *Applications* link in the top middle or input a search does the entire menu appear, arranged horizontally rather than vertically to take advantage of the standard modern monitor (Figure 3). Click a top-level menu item, and the individual items in the menu appear across the screen. Admittedly, the arrangement does

not reduce the number of clicks needed to navigate the menu, but at least items are not crammed into a vertical space somewhere on the left of the screen. Each top-level menu has its own wallpaper for easy identification, and all its menu items open on the same screen, revealing yet another hidden feature – virtual desktops automatically arranged by general category.

Two other desktop features are also worth noting: First, notifications are primarily audio, which might be a shock the first time you log in. Notifications also linger in text windows on the screen's upper right, just below the listing of hard drive space. Second, if you search the dock, you will find a window that lists recent events and activities, a kind of desktop Undo function. Being text-oriented, I am unsure of the usefulness of the audio notifications, which would only be a distraction for me, but I may be a minority. As for the Recent Events window, it is such an obvious and powerful feature that I am surprised that no one thought of it before. I predict and hope that other DEs will copy the feature.

All of these features seem like ones that I could learn to live with. Unfortu-

nately, though, the appeal of Jade is lessened because of the way features are distributed. To say the least, a note-taking app in the center of the desktop seems an acquired preference. More importantly, the division of the menu on each side of the note taker, with Help, Setting, and System on the left, and Manjaro and Exit on the right seems completely arbitrary. The two menus contain the same kind of housekeeping information, so why not put them all in a single menu? Or better yet, why not add them as top-level items to the Applications menu and rename the menu? Similarly, why repeat icons above the left-hand menu and add the same links in the dock? A quick refresher course in design would remind the Jade designers that two basic principles of design are that related items should be placed together, and separated from unrelated ones. With these two basic concepts applied to the layout, Jade would be far closer to the intuitive DE that the designers intend to build.

Still in the Rough

I came away from exploring Jade with an appreciation for the minimalist appearance and a sense of the potential for some of its design features. However, the layout could do with a general cleaning. Moreover, while much of the appearance can be adjusted within the current theme, Jade's success will depend on how much customization is eventually added. I, for one, would want the option to change the application poised in the middle of the screen or perhaps to eliminate it altogether.

Still, full credit to the developers for making the effort to innovate. If not all their experiments are successful, at least they are trying to adjust the 30-year-old metaphor of the desktop for modern computing. Regardless of whether Jade proves popular, perhaps it will inspire the programmers of other DEs to be less cautious and try something different. If so, it will be a change long overdue. ■■■

Info

- [1] Jade: <https://sourceforge.net/projects/manjaro-webdad/>
- [2] Manjaro forum: <https://forum.manjaro.org/t/new-desktop-jade-in-the-works/17228>
- [3] Manjaro Webdad Preview <https://manjaro.org/download/webdad/>

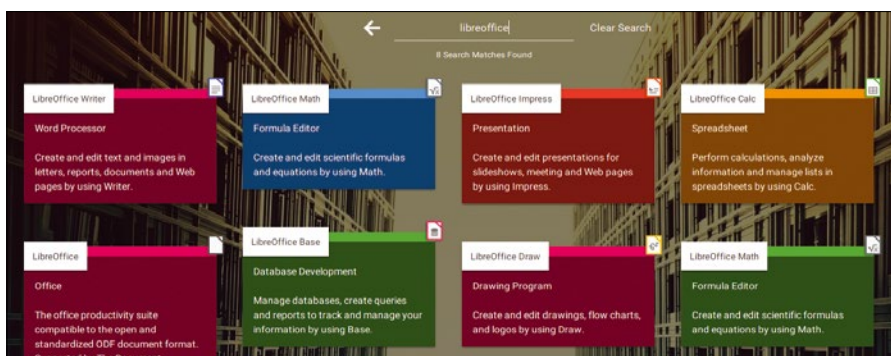


Figure 3: Jade's menu is arranged horizontally, for a widescreen monitor, giving plenty of room for details about items.

Firewall management

Getting to Know firewalld

Managing a firewall can be a hassle, but it's worse to manage a breach because you didn't have one. *By Ken Hess*

A firewall is an important part of a security strategy. However, it is only one component and not a security panacea for reasons that will become clear later in this article. A host-based firewall protects the local system just as a network firewall protects an entire network or part of a network, such as a DMZ.

On CentOS 7 and newer, Red Hat Enterprise Linux 7 and newer, and Fedora 18 and newer, the default firewall is firewalld (see the “Features” box for more

information.) If you use a Red Hat-based distribution, then you probably have it already. If you use other distributions, firewalld is available via git and as a tarball [1]. Firewalld uses zones to define trust levels of network connections or interfaces. (Zones are an advanced topic not covered in this article; look for a future article covering firewalld zones).

Troubleshooting Firewalls

Sys admins of all skill levels have wasted countless hours troubleshooting a problem that ended up pointing to a firewall that has prevented remote access to a service. The term “remote” is important. Firewalls don't prevent access to local services; firewalls prevent access from remote systems across the network but not access from the local system itself. The point of a firewall is to deny everything from the outside except what you specifically allow in. Unfortunately, frustration with firewall rules often ends in the firewall being disabled by an otherwise well-meaning sys admin.

Troubleshooting firewall access is easy. In the following example, the Apache web server, NGINX, or some other web server of your choice has just been installed onto a server system. During testing, you get a

“This site can't be reached” message in the browser. To resolve the problem:

1. Check the system's process list to be sure the service is running.
2. Test the service from the local system. Open a web browser (if you have a GUI, Firefox or Chrome; if not, Lynx, a text-based web browser).
3. Check to see if a firewall is running.
4. Add a firewall rule to allow HTTP (TCP port 80) and HTTPS (TCP port 443) or whatever ports your web service uses. Reload the configuration to enable the rule.
5. Retest from a remote system. This same procedure works for any service accessed over the network.

HTTP and HTTPS Access

To illustrate this troubleshooting procedure, I will demonstrate setting up a web server and then accessing it remotely. First, install a web server on your system. For my Red Hat-based system, the process is simple to install the Apache web server:

```
$ sudo yum -y install httpd
```

Enable the web server to run at startup:

```
$ sudo systemctl enable httpd.service
$ sudo systemctl list-unit-files | \
grep httpd
httpd.service      enabled
```

Listing 1: Verifying the httpd Service

```
$ ps -ef | grep httpd
root      1194      1  0 20:05 ?        00:00:00 /usr/sbin/httpd -DFOREGROUND
apache    1199    1194  0 20:05 ?        00:00:00 /usr/sbin/httpd -DFOREGROUND
apache    1200    1194  0 20:05 ?        00:00:00 /usr/sbin/httpd -DFOREGROUND
apache    1201    1194  0 20:05 ?        00:00:00 /usr/sbin/httpd -DFOREGROUND
apache    1202    1194  0 20:05 ?        00:00:00 /usr/sbin/httpd -DFOREGROUND
apache    1203    1194  0 20:05 ?        00:00:00 /usr/sbin/httpd -DFOREGROUND
```

Features

- IPv4 and IPv6 support
- Ethernet bridging
- IP sets
- Separate run time and permanent configuration options
- No service or system restart required for configuration changes
- Complete D-Bus API
- Predefined zone list
- Simple configuration options
- Flexible enough for complex zone rules
- Direct interface
- Simple log of denied packets
- Application whitelisting
- Automatic kernel module loading
- Puppet integration
- CLI and graphical configuration

Note: A firewall is a set of allow and deny rules that control packet flow to and from networks. A firewalld service is a combination of ports, protocols, modules, and destination addresses.

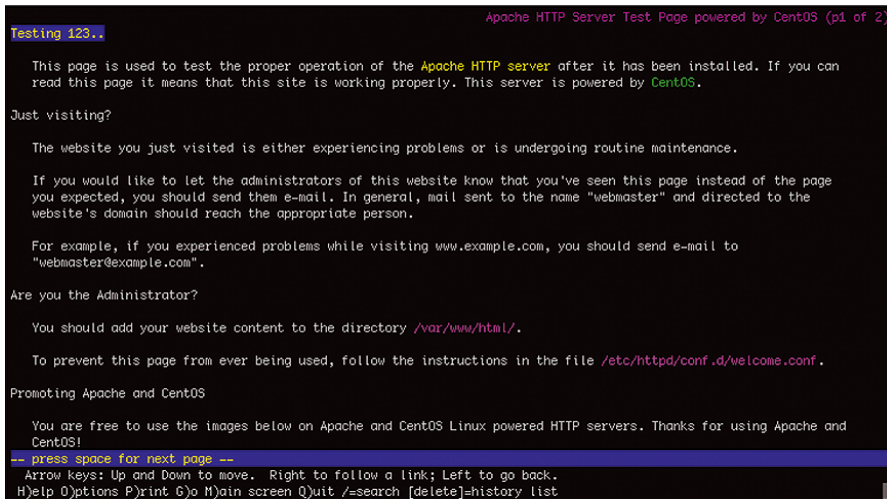


Figure 1: Verifying Apache is running and available on the localhost using Lynx.

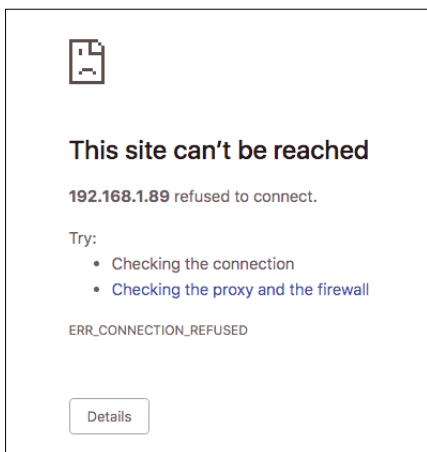


Figure 2: Checking access from a remote computer.

Start `httpd.service`:

```
$ systemctl start httpd.service
```

Check that the `httpd` service is running, as shown in Listing 1. Then, test the web server locally using Lynx or a graphical browser (Figure 1):

```
$ lynx http://localhost
```

Check that you can access the web server from a remote computer; it fails the test, because the firewall is currently blocking all ports (Figure 2).

Check to see if a firewall is running:

```
$ sudo firewall-cmd --state
running
```

To allow access from remote systems, you must enable the ports configured for your web server. In this example, it is

port 80. The `--permanent` switch adds the allowed port to the firewall’s permanent configuration:

```
$ sudo firewall-cmd --permanent \
--add-port=80/tcp
success
```

Reload the firewall configuration:

```
$ sudo firewall-cmd --reload
success
```

Check access again from your remote computer. You should see a web page appear in the browser (Figure 3).

Use this same procedure to configure other ports for your services. You can also add multiple ports before reloading the firewall’s configuration.

Focusing on Security

To reiterate, firewalls are a single part of an overall security strategy, not a panacea. In configuring a firewall exception to allow remote access to port 80 for the Apache web server, you have created a vulnerability on the host system. For the moment, assume that you only allow ports 22 (SSH) and 80 (HTTP) on that server system. That’s two vulnerabilities or what security people call “acceptable risk.” You have to accept some risk when you allow network access to a system’s services.

The reason allowing access creates vulnerabilities is that you’re allowing computers on a network, and possibly the entire Internet, to access this system via port 80. What if the version of Apache you installed has an unpatched security flaw? Your system is exposed and vulnerable to that flaw until it’s patched. The firewall won’t protect the system, because you have allowed access to that port. The door is open.

Is this a real problem? Yes, and no. It is a problem but the alternative is to have no services running on computers, which means you have no customers or employees connecting to those services. That’s not acceptable. There is some degree of risk that you have to accept to run a service. You have to practice due diligence and protect the system in other ways (encryption, application firewall, backups, and monitoring) and routinely patch the system.

Summary

You should have a host-based firewall running on every system on your network – no exceptions. A firewall, as stated previously, is not the perfect security tool, but it does help protect the system from attacks on other services that are not exposed. The same rules apply to network firewalls. This is why you must employ a multilayered approach to security and not rely on any one technology or solution. `firewalld` is installed and enabled by default on all Red Hat-based systems, which should put system administrators at ease that their systems are protected as soon as they’re placed online. ■■■

Info

[1] `firewalld`: <https://firewalld.org/>

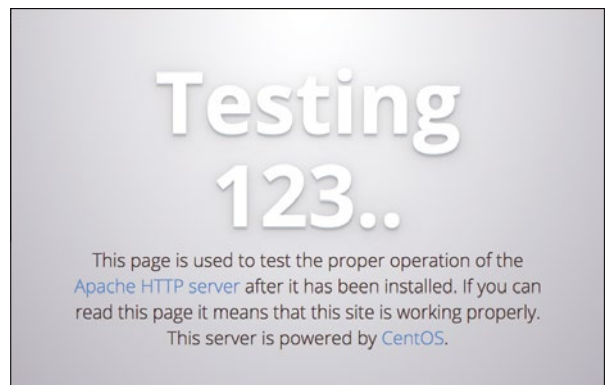


Figure 3: The remote computer’s web service is now accessible.

The sys admin's daily grind: f.lux

Shimmer Dimmer

You don't have to be a vampire to be sensitive to bright light at night. Charly, who – as regular readers know – is a practicing light conservationist, now makes it clear to his desktop PC that it's not good to be too dazzling at night. *By Charly Kühnast*

If you ask me, it has to be possible to adjust lighting to suit your needs in an unobtrusive and fully automated way. I enabled night mode on my Android smartphone. In the evening, the display shows an ever-decreasing amount of blue light as time progresses. At first, this seems a bit strange and takes getting used to, but it is very friendly on the eyes. Without night mode, I get dazzled when I unlock the phone. This is because its display – like most monitors – is adjusted to a color temperature of 6500 Kelvin.

Cirque du Soleil

Light with a temperature of 6500 Kelvin has a higher blue component than sunlight, which does not exceed 5800 Kelvin even on a clear day. For work during the day in a bright room, 6500 Kelvin is completely OK. In the evening with dim light, I feel as if my PC displays are so bright that work becomes tiring. Of course, I could manually adjust the brightness and color temperature using the buttons on my displays – but that is something that I want to be done automatically. My choice of dimmer goes by the name of f.lux.

Installing f.lux

F.lux is simply pronounced “flux”. It is available as a command-line tool [1] or with a GUI [2]. I decided to go for

the graphical variant. F.lux is written in Python 3, which means quickly installing some packages on my Ubuntu test system for it to run:

```
sudo apt install python3-distutils \
  gir1.2-appindicator3-0.1 \
  gir1.2-gtk-3.0
```

Guidance for other distributions can be found at the same source. Then I added the repository and installed the *fluxgui* package:

```
sudo add-apt-repository \
  ppa:nathan-renniewaldock/flux
sudo apt update
sudo apt install fluxgui
```

At startup time, I had to tell f.lux my approximate geographic position (Figure 1). In return, it actually showed me a map, guessing my location based on my IP address. (It was about 100 kilometers away, but that was totally irrelevant from a lighting point of view.)

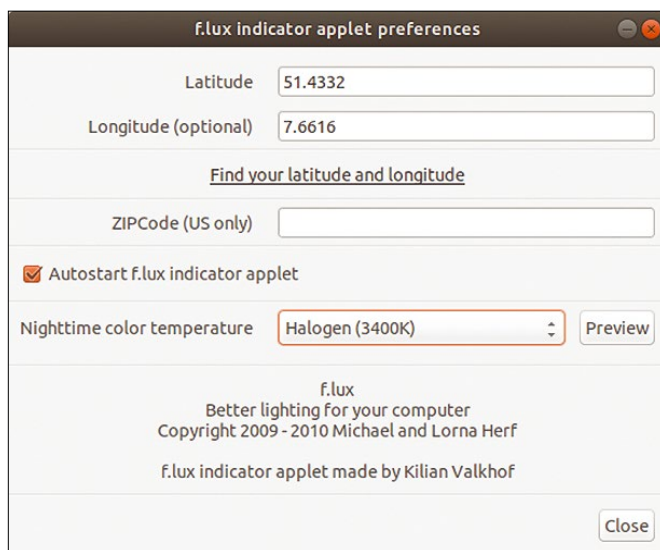


Figure 1: You'll enter the latitude and longitude of your location, as well as the color temperature value the display should switch to at night.

Then I specified the maximum color temperature to which f.lux can reduce the display at night. I initially chose 3500 Kelvin, which seemed to be warm enough in preview mode.

After a few evenings using the computer, I will know whether I need to re-adjust this value. In any case, there are enough options for this (Figure 2). Finally, to be on the safe side, I checked again whether the time zone and time are set correctly; I don't want f.lux to intervene at the wrong time of day.

F.lux is also available for Windows and macOS. Enlightenment thus far: Everything is now to my liking in terms of illumination levels on the desktop – a great step forward for any light-shy Linux user. ■■■

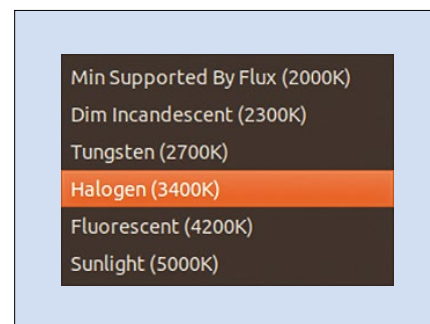


Figure 2: Kelvin home alone – light sources and their color temperatures.

Info

- [1] f.lux: <https://justgetflux.com/linux.html>
 [2] fluxgui: <https://github.com/xflux-gui/fluxgui>

Author

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



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...

shop.linuxnewmedia.com/digisub



MakerSpace

Open hardware and crowdfunding Hand in Hand

Open hardware and crowdfunding are a natural fit. While their relationship has evolved over time, crowdfunding continues to play a big role in open hardware development.

By Bruce Byfield

Open hardware and crowdfunding became popular at about the same time and have been associated ever since. In fact, in the absence of major investment, open hardware would not be where it is today. Crowdfunding has helped to launch several open hardware small businesses like Purism and Keyboardio and has helped existing small businesses to launch new businesses. But how else do the two trends reinforce each other? Today, the answers are different from what they were five years ago.

The close connection between open hardware and crowdfunding should not be surprising. As Josh Lifton (Figure 1),




Figure 1: Josh Lifton, the founder of Crowd Supply.

the founder of Crowd Supply [1] notes, the two “go hand in hand, because they both rely on their community of users. Participating in crowdfunding requires a certain amount of risk. A product being open source mitigates some of that risk and builds trust with the end user.” Consequently, those involved in one of these trends have little trouble accepting the other.

Just as importantly, crowdfunding resolves a long-standing problem. As Eric von Hippel points out [2], innovation often comes from users rather than established companies. As companies get set in their ways, they can become more concerned with growth than with new products, while users want solutions to their own needs. In other words, those who have the money may not be overly concerned with innovation, while those who are interested in innovation generally lack



Figure 2: Like many innovations, Input Club’s Keystone Analog keyboard was developed by a small company rather than a corporation.

Lead image © lightwise, 123RF.com

the money to realize their ideas. Thanks to crowdfunding, those with the ideas have a way to fund themselves that does not depend on corporations or venture capitalism.

Still another consideration is that a large corporation may consider the potential profit from an innovation too small to be worth development. For example, most sellers of keyboards are content with a small range of products: cheap membrane keyboards for the average user, and a few expensive programmable mechanical keyboards for gamers. It takes a small business like Input Club to develop a product like the forthcoming Keystone Analog, whose keys can produce different results depending on how hard they are pressed (Figure 2).

Under these circumstances, it seems no accident that the overwhelming majority of crowdfunding campaigns for open hardware are for only a few hundred thousand dollars, and only a few for over a million. Many are for less than \$50,000. So far, the majority of open hardware products are for niche prod-

ucts that major manufacturers do not consider worth developing for themselves. Yet what is loose change for a corporation can be enough to support a family operation or a business with half a dozen employees – and can be a realistic goal for an average crowdfunding campaign as well.

Analyzing the Odds

Popular coverage of crowdfunding focuses on the successes. However, for open hardware, success is by no means guaranteed. Moreover, the success rate has changed over the years.

Five years ago on Indiegogo [3], the term “open hardware” was not used [4]. However, many of the projects listed under “open source” would be called open hardware today. At the time, there were a few outstanding successes, like the magazine *Linux Voice* (now a part of *Linux Pro Magazine*). However, based on the first 12 pages of results for the term “open source,” only 7.0 percent of campaigns reached their target funding, while 89 percent raised only half their target funding. Searching for

“free software” produced roughly the same results. Neither keywords for searches or gifts for backers had any noticeable effects on success rates, although updates did.

2019 presents a different picture. For one thing, open hardware is now a recognized, if small, category. Crowd Supply has a dedicated page, while a search on Kickstarter [5] divides open hardware into Product Design, DIY Hardware, 3D Printing, and Hardware. Perhaps because of this increased recognition, the success rate is much higher than five years ago. My unscientific analysis shows a success rate of 78 percent on Indiegogo, 72 percent on Kickstarter, and 85 percent on Crowd Supply.

No doubt the increased recognition of open hardware is part of the reason for this difference. However, one noticeable change in the last five years is that open hardware campaigns are far more detailed than in 2014. Certainly, the reason for the high success rate on Crowd Supply is that the site offers detailed advice to fundraisers and guides them through each stage of the campaign. For

FIND THE RIGHT TOOL!

Check out our new Linux Administration corner!

<http://www.linux-magazine.com/administration>

SPONSORED BY



Linux
Professional
Institute

instance, I have yet to see a Crowd Supply without a video, and most of the successful ones have frequent updates, averaging one every two or three days of the campaign. Even backer gifts are encouraged, although from what I can see, they are not a major factor in a campaign's success.

By contrast, campaigns that fail tend to communicate much less, or simply communicate poorly. For instance, I believe that one campaign that I backed failed because it consistently defined the hardware as useful only for a small audience, despite the fact it had a broader appeal. Moreover, on Indiegogo and Kickstarter, many failed campaigns didn't include a video.

One difference between the three crowdfunding sites analyzed was that successful campaigns on Crowd Supply




Figure 3: The RISC-V chip is becoming one of the entry points into open hardware for established companies.

and Kickstarter tended to exceed their goals by under 100 percent, or by a few 100 percent at the most. In comparison, 33 percent of successful Indiegogo campaigns exceeded their goals by over 1,000 percent, with several exceeding by over 2,000%. If this difference is not a coincidence, I suspect that, in Crowd Supply's case, the reason may be that the site helps campaigns set realistic goals. As for Kickstarter, the reason may be that its campaigns are often for smaller goals.

One thing that has not changed is that open hardware on all three sites remains heavily oriented to hobbyists, with monitors, robotics, breadboards, testing hardware, and similar devices. Another general category is various computer parts, such as WiFi. Raspberry Pi accessories seem to have declined since 2014, while 3D printers and utilities have surged, probably because of the decline in prices.

The biggest takeaway is that open hardware is better known and has over 10 times the success rate it had five years ago. Probably, the increase reflects an increased sophistication in using crowdfunding. However, whatever the reason, the tie between open hardware and crowdfunding is closer than ever before.

Next in Open Hardware Evolution

How long open hardware will continue to depend on crowdfunding is uncer-

tain. In the past year or two, established companies have started taking a cautious interest in open hardware. For instance, the RISC-V Foundation [6], which oversees the development of the first open hardware chip, currently has Google and Western Digital on its board of directors (Figure 3). Josh Lifton also observes that a lot of the interest in open hardware "is from companies not traditionally in the hardware space – they seek better security and control over the hardware they are using." Despite many investors' innate conservatism, it is probably only a matter of months before the first major venture capital goes to an open hardware project. Already, Joseph Jacks, the founder of OSS Capital [7], is keeping an eye on open hardware, mostly in manufacturing, but also in consumer products (Figure 4). Almost certainly, he is not the only one. In another few years, open hardware may start trending the same way open source software did around the turn of the millennium.

Yet even if that happens, the association of open hardware with crowdfunding is likely to continue. In just a few years, open hardware has gone from being as hit and miss as any crowdfunding project to having a higher than average success rate. More and more, open hardware developers have learned to game crowdfunding. Others can draw on the experience of experts like those at Crowd Supply. So long as there are enthusiasts who dream of starting a small business, the reliance of open hardware on crowdfunding can only continue to thrive. ■■■

Info

- [1] Crowd Supply: <https://www.crowdsupply.com/>
- [2] Eric von Hippel: https://en.wikipedia.org/wiki/Eric_von_Hippel
- [3] Indiegogo: <https://www.indiegogo.com/>
- [4] Crowdfunding and Open Source: <https://www.datamation.com/applications/crowdfunding-and-open-source.html>
- [5] Kickstarter: <https://www.kickstarter.com/>
- [6] RISC-V Foundation: <https://riscv.org/risc-v-foundation/>
- [7] OSS Capital: <https://oss.capital/>



Figure 4: Joseph Jacks of OSS Capital is one of the venture capitalists starting to watch open hardware.

Too Swamped to Surf?



ADMIN Network & Security

ADMIN offers additional news and technical articles you won't see in our print magazine.

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

© Rachata Tevparait, 123rf.com



<https://bit.ly/HPC-ADMIN-Update>



MakerSpace

Streaming services on LibreELEC 9.0
with Kodi 18.0

Popcorn Cinema

Thanks to Kodi 18.0, LibreELEC 9.0 now supports the DRM encryption used by many streaming services. However, integrating Netflix, Amazon, and other streaming services does involve some manual work. *By Christoph Langner*

One of the most common tasks for the Raspberry Pi is still deployment as a media center. In combination with the Kodi media center software, the Rasp Pi can turn every “dumb” TV into a smart TV. Content from locally connected data carriers, an intranet, or the web can then be played back. These features are now available on even the simplest of modern TVs; however, with appropriate hardware upgrades and advanced configuration, the Rasp Pi media center can also record TV shows or stream them to the

network so you can watch TV on your PC or smartphone.

The LibreELEC Kodi distribution serves as the basis for many Rasp Pi media centers. Under the hood, the operating system available for the Raspberry Pi and other small-board computers (SBCs), as well as for PCs, follows the “just enough operating system” principle – that is, Linux reduced to the bare bones.

Shortly after the release of Kodi 18.0, dubbed “Leia” by its developers, the LibreELEC project also announced the latest edition of its distribution, version 9.0.0. In this article, I take a look at what has happened in both projects and investigate how well new features, such as retro games and streaming services, perform on a Raspberry Pi.

New in Kodi 18.0

The developers of Kodi 18.0 deliver the first major release since February 2017 [1]. The media center was given completely new functions: Classic video games can now be launched and played directly on Kodi. However, Kodi itself does not run the games: The gamer is expected to provide the emulators, games, and ROMs.

Kodi 18 uses the Libretro library [2] from the RetroArch project [3] (Figure 1). Many emulators such as MAME [4] now support the library directly. Additionally, typical gaming input devices such as joy-

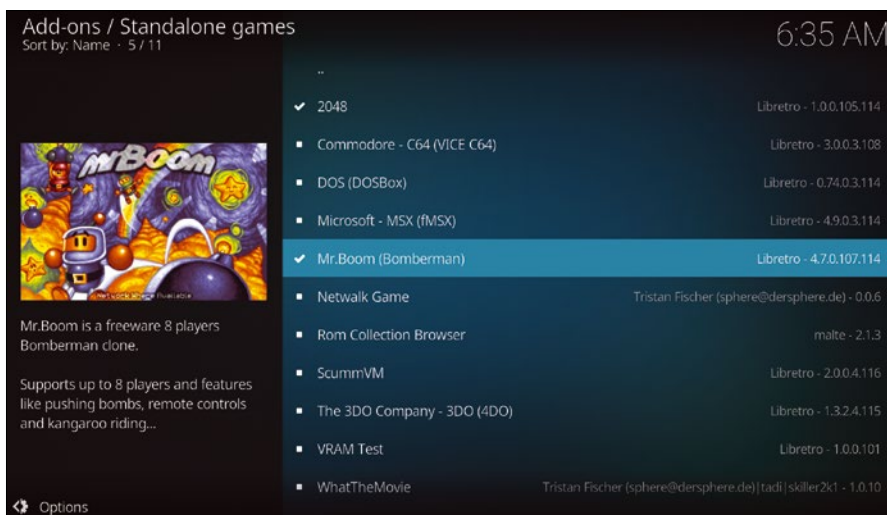


Figure 1: Kodi integrates retro gaming into version 18.0. The system supports various emulators, from the Amiga and C64 to game consoles.

sticks, gamepads, and other controllers can be integrated into the system, which means nothing can get in the way of enjoying a round of a Jump-n-Run classic like Super Mario Bros. or The Great Giana Sisters – if you have the right software offerings.

In addition to your own video collection, streaming services such as Netflix, Amazon Prime Video, and Hulu deliver movies and series in digital form directly to your living room. Zattoo, Pluto TV, and other providers do the same for TV programs; thus, you can receive digital TV without cable access or satellite equipment. However, to protect against unlicensed copying, most services encrypt their content. Kodi 18 now knows how to handle these Digital Rights Management (DRM) routines, although compatibility essentially depends on the hardware and operating system you use.

The developers of Kodi 18 (Figure 2) also improved the music player and music library management. Indexed titles can now be organized more flexibly. For example, you can filter songs by source or the artist's gender. The improved API allows faster access to music collection content, which is particularly beneficial for users of a Kodi app on a smartphone or tablet.

Owners of a full-fledged media center PC will be especially pleased with the improved video player, which can process content faster and better in 4K and 8K resolution, as well as HDR. The media player now has priority over other functions when accessing the CPU

and GPU, so it can play content as smoothly as possible. If the device on which you are using Kodi 18 has radio reception, it now displays station information with the help of the Radio Data System (RDS), which is standard on most car radios and hi-fi tuners today. RDS enables the transmission of additional digital information with analog FM radio.

New in LibreELEC 9.0

LibreELEC builds a fully functional system around Kodi. LibreELEC 9.0 [5], introduced in the wake of the new Kodi version, is based on the Kodi 18.x branch and integrates its new functions, such as DRM support for streaming media, as well as the RetroPlayer framework for running console games (assuming an appropriate add-on). Under the hood, the rapidly updated LibreELEC 9.0.1 MR uses the 4.19.23 kernel. (See the “Installing LibreELEC” box for instructions on getting LibreELEC on your Raspberry Pi.)

As a media center, a LibreELEC installation is in the category “install once, neglect forever”; very few users even think about security. Often, updates are not even installed, let alone default passwords changed. As a result, unprotected LibreELEC systems repeatedly end up directly on the Internet (e.g., when a user digs a tunnel into the network via a VPN).

Therefore, one of the new features of LibreELEC 9.0 that appears in the setup wizard is an interface for changing the SSH password. Alternatively, you can do

this later in the gear menu under *LibreELEC | Services | SSH | SSH password* (Figure 4).

Moreover, a basic configuration of the iptables firewall provides additional security. In home mode, LibreELEC blocks

Installing LibreELEC

To install LibreELEC on an SBC like the Raspberry Pi, I recommend using the LibreELEC USB-SD Creator [6]. The program offered for Linux, macOS, and Windows prepares the required memory card in four simple steps. First, you select the desired version (usually *Raspberry Pi 2 and 3*) and then click *Download* to download the corresponding version from the web. In the third step, you specify the disk on which the wizard will install the image. Proceed with caution here, because all data on the storage medium is lost when you install. Finally, pressing *Write* puts the downloaded image to the SD card (Figure 3).

Now plug your SD card into the Raspberry Pi and boot the system. A setup wizard will guide you through the most important configuration steps. The best way to proceed is to plug a keyboard into the Pi, but you don't necessarily need a mouse. Start by assigning a computer name and setting up the network connection; wireless on a Rasp Pi 3 (RPi3) is supported by default. Next, enable SSH and Samba network services. In the case of SSH, the wizard automatically prompts you for a new password.

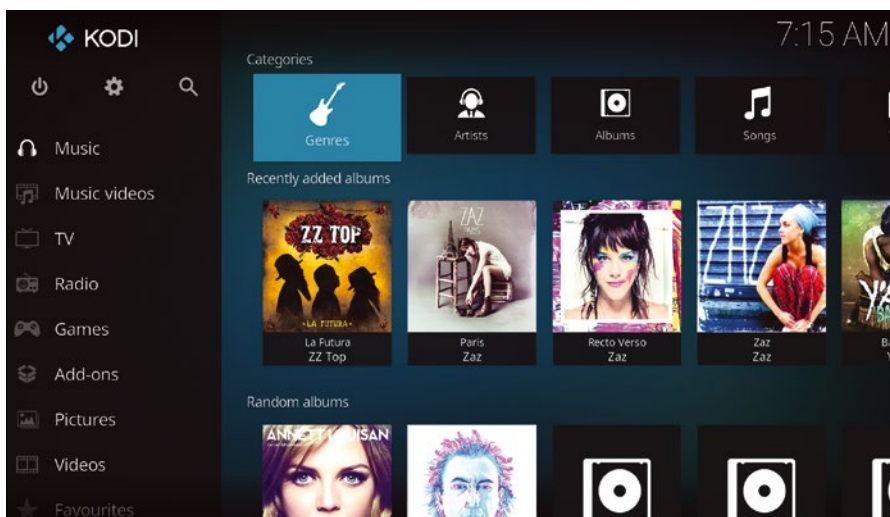


Figure 2: The new Kodi version makes it easier to browse the music database. API access to the music collection has also been optimized.

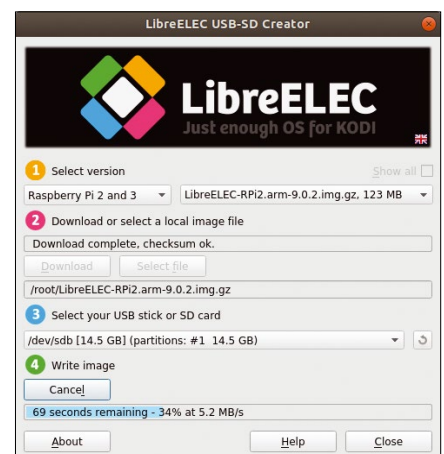


Figure 3: The LibreELEC USB-SD Creator makes installation on an SD card designed for the Raspberry Pi easy. The program is available for Linux, macOS, and Windows.

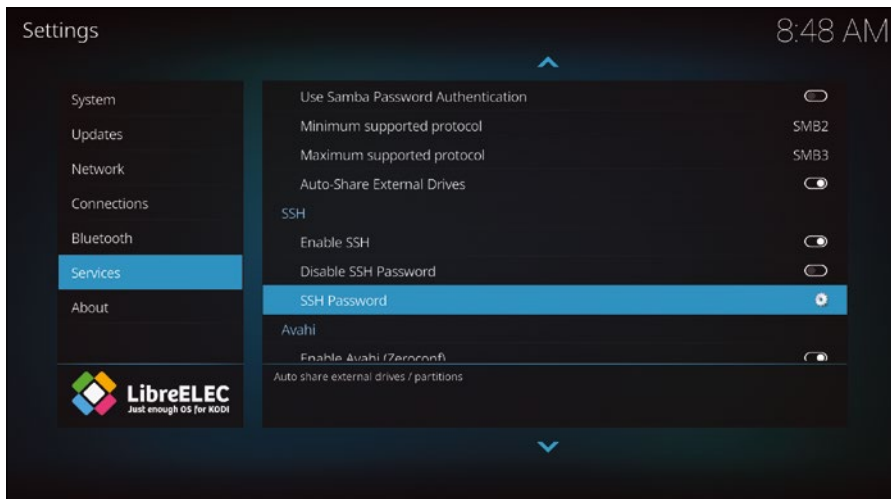


Figure 4: If you enable SSH access from the setup wizard, LibreELEC 9.0 immediately prompts you to change the default password, but you can also change the password from the Settings screen.

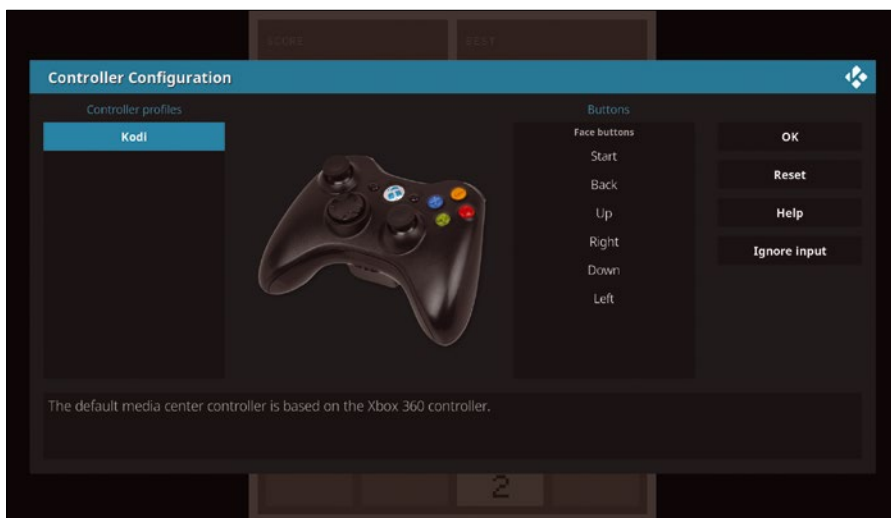


Figure 5: When you connect a gamepad to the system, Kodi automatically opens a wizard to configure the new game input device.




Figure 6: LibreELEC comes with two games out the box: the simple 2048 and the multiplayer Bomberman, which did not work in a test on a RPi3.

incoming connections that do not come from the local network and thus protects active services such as SSH, Samba, and the Internet against unwanted access. If obsolete add-ons repeatedly cause problems when starting the system, a *Safe Mode* now provides a remedy. The system then starts with the basic configuration so that misbehaving add-ons and settings can be disabled.

The Games category will initially be empty. Two games can be installed quite easily for demonstration purposes. Open *Add-on browser* | *Offline Games* and set up *2048* and *Mr.Boom (Bomberman)*, just as you would set up any other add-on. The corresponding game engines, including *DOS (DOS-Box)*, *ScummVM*, and *Commodore-C64 (VICE C64)*, are automatically retrieved by the system.

If you plug a gamepad into the USB port, LibreELEC automatically registers the new device and offers to configure it. Alternatively, you can initialize the input device from the gear menu and the options *System* | *Input* | *Peripherals* | *Configure attached controllers*. Step by step, you can then hardwire the gamepad's buttons to console-specific buttons (e.g., A, B, X, and Y) (Figure 5).

In the RPi3 test, only the very simple 2048 game worked (Figure 6). Although the funnier and more action-packed *Mr. Boom (Bomberman)* is executed by the Rasp Pi, the screen remains blank. After pressing Esc, a menu appears to let you quit the game; strangely, the screen then shows the game itself in the background.

C64 Games on LibreELEC

As a practical test, I installed the *Commodore-C64 (VICE C64)* emulator under *Add-ons* | *Install from repository* | *All repositories* | *Game Add-ons* | *Emulators*. It then appears in the main level below *Games* as an independent extension. However, if you select the add-on directly, you will end up at the light blue prompt typical for the C64 – only insiders will know that you now have to type:

```
LOAD":*", 8, 1
```

It is easier to pick up a game from a collection like the *GameBase64 Collection* [7], transfer the downloaded ZIP file via Samba (or the Windows network environment) to the LibreELEC

Rasp Pi, and unpack it there (e.g., in the Downloads/ folder). If you then select *Games* from the main menu, the library manager appears. Go to *Add games | Browse | Root file system* and select the folder with the uploaded games.

Kodi integrates the folder into the Games menu. When you select it, a file manager appears; you can use the manager to select the game as a disk image (e.g., in D64 format) or a program file (in P00 format) [8]. During the game, Vice forwards all keystrokes to the emulator. Only the Esc key still interacts with Kodi. From the menu, you can now control the game with *Pause/Resume*, *Reset*, *Exit*, or individual *Settings*.

Other innovations in Kodi 18.0 and LibreELEC 9.0 still need a little time to mature, however. The Zattoo add-on, for example, can be activated from the repositories and also accepts your login credentials. The TV section, where the Zattoo add-on adds itself as the PVR service (personal video recorder, which has a mass storage device, instead of a tape drive, and intelligent additional functions), remains empty, however. Even for Netflix, Kodi/LibreELEC does not yet offer a solution that works out of the box.

Netflix

For Netflix on LibreELEC you have to enable the repository first. Download the repository.netflix-1.0.1.zip file (the version available at the time of the test) from the kodinerds repository [9] and copy it over the network to your Kodi Rasp Pi. Select the gear menu and *Add-ons | Install from zip file* to enable the package source in the home folder. You might have to allow the installation of add-ons from unknown sources first; a prompt to this effect will appear automatically.

You can then install the Netflix add-on in the gear menu with *Add-ons | Install from repository | Netflix add-on repository | Video add-ons*. In the settings, you must enter both the login and password for your Netflix account. If you launch Netflix from the Add-ons menu, you can browse through the provider's digital video library (Figure 7). However, as soon as you play a video, error messages appear: The system requires

the *InputStream Adaptive* add-on, which can be found in the add-on manager under *Videoplayer InputStream Add-ons*.

The next time you try to play a Netflix video, Kodi complains about the lack of the Widevine CDM library it needs to decrypt the DRM-protected video streams [10]. For legal reasons, LibreELEC is not allowed to integrate the library directly into the system or keep it in its own package sources. Instead, a wizard jumps in to download a suitable Google

Chrome OS image [11] and extract the ARM version of the library from the image. The process on an RPi3B+ takes about a quarter of an hour and requires 2GB of free disk space.

Finally, LibreELEC plays Netflix videos. However, the computing capacity of the fastest current RPi3B+ model is not sufficient to play the videos in full HD (1080p): Playback is interrupted frequently or stops completely at intermittent intervals. However, the playback resolution cannot be changed

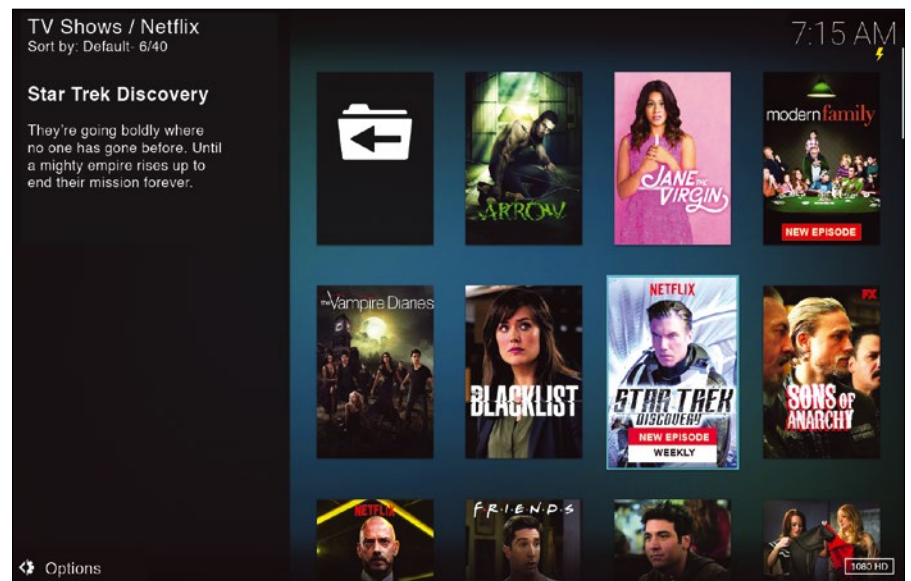


Figure 7: In the Kodi 18.0 announcement, the developers claimed that the media center supported streaming services such as Netflix. However, integrating them still requires a great deal of manual work. (Note: Some shows shown here might not be available in your country of origin.)

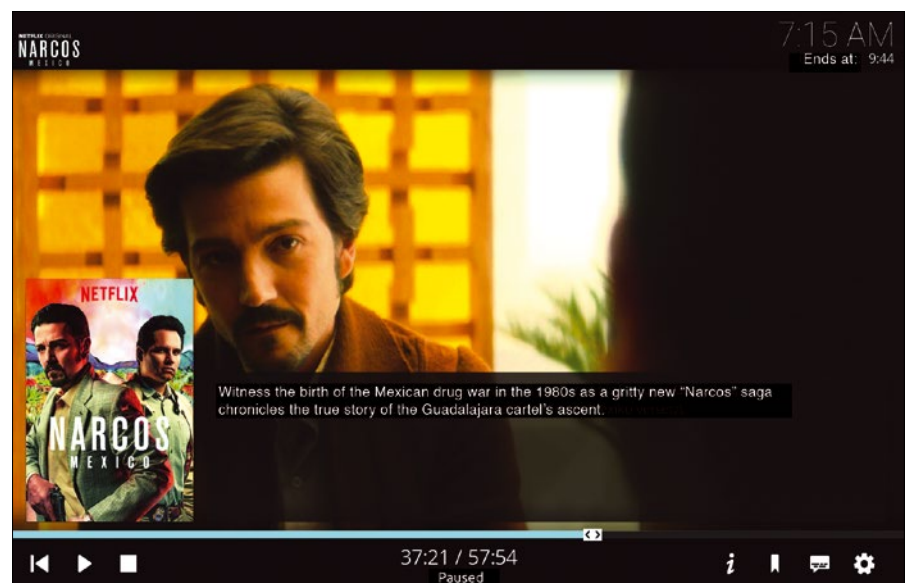


Figure 8: After installing the Netflix add-on, the InputStream extension, and the Widevine library for decrypting streams, the system plays Netflix movies reliably.

Two Amazon Plugins

The *Sandmann79s* repository contains two add-ons for Amazon's streaming service: *Amazon* and *Amazon VOD*. The difference between the two variants is that the *Amazon* add-on (without VOD) maintains a local database, whereas *Amazon VOD* always loads all titles and menus fresh from Amazon. The database-supported version offers the advantage of faster browsing through Amazon videos after the initial fill. However, you do need to refresh the database regularly, which happens automatically with the VOD version [13].

from within the player and is only possible by configuring the InputStream Adaptive add-on (*Add-ons* | *My add-ons* | *VideoPlayer InputStream* | *InputStream Adaptive*). In the *Max. Resolution general decoder* and *Max. Resolution secure decoder* fields, enter the value *720p*. The Netflix videos then play on the RPi3 with virtually no jerkiness (Figure 8).

Amazon Prime Video

Amazon Prime is now also slightly better integrated into LibreELEC – even if you still need to indulge in a great deal of manual work. Just like the Netflix add-on, the Amazon counterpart also needs the InputStream extension and the Widevine library. Here, I assume that you have already experimented with Netflix and have already installed both extensions.

Setting up the Amazon add-on again requires a number of intermediate steps: First, select *Add-ons* | *Install from repository* | *All repositories* | *Add-on repository* and then *Kodinerds add-ons* [12] from the gear menu.

From this package source you can now enable *Sandmann79s Repository* from *Add-ons* | *Install from repository* | *Kodinerds add-ons* | *Add-on repository*. Sandmann79 is the pseudonym for a developer on whose work the integration of video services such as Netflix and Amazon is based. The additional repository finally lets you install the desired *Amazon* add-on under *Installing from repository* | *Sandmann79s Repository* | *Video add-ons* (see also the “Two Amazon Plugins” box).

With the Amazon add-on, you also need to enter your account login and password. The system does not save the data, but creates an individual token used for authentication against Amazon in the future. Like the Netflix add-on, Kodi lists the Amazon extension below *Add-ons*. After launch, you can browse the Amazon movie collection. In contrast to Netflix, Amazon playback is always smooth, but only at 540p and without the possibility of increasing the resolution (Figure 9).

Conclusions

At first glance, Kodi 18.0 appears to be a cautious update without many changes compared with its predecessor. Only if you take a closer look do you notice the many changes under the hood with the InputStream add-on.

DRM-encrypted streams can now be played back, opening the door to streaming services such as Netflix, Amazon, and others. The Widevine decryption routine, however, reminds one a bit of playing video DVDs on Linux with *libdvdcss* [14]. From a technical point of view, it works, but don't expect either official support or ease of installation and configuration.

In terms of integrating games, Kodi continues to work consistently on developing the software into a comprehensive digital entertainment platform. In practice, it is not difficult to get old retro games from C64 times to run. However, documentation of the new functions is still missing. Moreover, help in the form of a games database would be ideal, so that game classics can be installed with just a few clicks. However, with a complex legal situation and unclear copyrights, such convenience remains pretty much a utopian dream. ■■■

Info

- [1] “Kodi 18 is here!”: <https://kodi.tv/article/kodi-180>
- [2] Libreto: <https://www.libretro.com/index.php/api>
- [3] RetroArch: <http://retroarch.com>
- [4] MAME: <https://www.mamedev.org>
- [5] LibreELEC release notes: <https://libreelec.tv/2019/02/libreelec-leia-9-0-0-release>
- [6] LibreELEC USB-SD Creator: https://libreelec.tv/downloads_new
- [7] The GameBase64 Collection: <http://www.gamebase64.com/search.php?h=0>
- [8] C64 disk image formats: https://www.c64-wiki.com/wiki/Disk_Image
- [9] Netflix add-on: <https://github.com/kodinerds/repotree/master/repository.netflix>
- [10] Widevine: <https://www.widevine.com>
- [11] Chrome OS: <https://www.chromium.org/chromium-os>
- [12] Kodinerds: <https://kodi-addons.club/addon/repository.kodinerds/>
- [13] Amazon Prime instant video: <https://www.kodinerds.net/index.php/Thread/44211-Release-Amazon-Prime-Instant-Video>
- [14] libdvdcss: <https://www.videolan.org/developers/libdvdcss.html>





Figure 9: Amazon Prime overhead in LibreELEC is similar to that for Netflix. However, Amazon restricts the quality of the streams to 540p.

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
- Create documents and spreadsheets
- Play games
- Surf the web
- Process photos
- Play music and videos
- and much more!



HELP OTHERS JOIN THE LINUX REVOLUTION!

ORDER ONLINE:
shop.linuxnewmedia.com/specials

THE COMPLETE

LINUXVOICE

ARCHIVE DVD: ALL 32 ISSUES!

LINUXVOICE

ARCHIVE DVD: ALL 32 ISSUES!

LINUX AND OPEN SOURCE LOVE!

3,400
PAGES OF

3,400
PAGES OF
LINUX AND OPEN SOURCE LOVE!

Since April 2014, Linux Voice has showcased the very best that Free Software has to offer. Now you can get it all on one searchable DVD.

Includes all 32 issues in E PUB, PDF, and HTML format!

Order now!

shop.linuxnewmedia.com



Photos are so easy now. Your camera is in your phone, and your phone is in your pocket or purse, following you around through streets and alleys, restaurants, and chance meetings with friends. The photos you take on your daily travels make their way to websites, Facebook pages, email messages, and personal image archives.

All these photos can pass by your eyes without taking much of your attention, but if you're someone who likes to pay attention, the Linux world has a multitude of tools to assist with your photo efforts. Some of the most popular applications are probably already on your system – or at least in easy reach from your distro's repositories. This month, we review some photo tools you might not know about.


Also this month, we explore TaskBook, a handy task manager for the command line, and we introduce you to Mastodon – a free social microblogging platform for FOSS-minded bloggers.



Image © Alexandr Moroz, 123RF.com

LINUXVOICE

Doghouse – Copyright and Licensing	75
<i>Jon "maddog" Hall</i>	
Despite the common assumption, everything online is not in the public domain.	
Photo Tools	76
<i>Dmitri Popov</i>	
DIY tools and shell scripts can make for a smooth photographic workflow on Linux.	
TaskBook	80
<i>Scott Nesbitt</i>	
Boost your productivity from the Linux command line with the TaskBook task manager and note board.	
FOSSPicks	84
<i>Graham Morrison</i>	
This month Graham Morrison looks at Cartes du Ciel, Foliate, Newsboat, xinput-gui, DevilutionX, Performous, and much more!	
Tutorial – Fediverse	90
<i>Paul Brown</i>	
If you're looking for social media options where the user has more control, you'll find a range of options to explore in the Fediverse, including the popular Mastodon.	




CC-BY-SA WWW.PEPPERTOP.COM

GOT CLUSTER?

Tune in to the HPC Update newsletter for news, views, and real-world technical articles on high-performance computing.

<https://bit.ly/HPC-ADMIN-Update>



HPC Up Close

- OpenACC – Parallelizing Loops
- DES Completes Dark Matter Survey
- Open Compute Project Call for Posters
- Dolphin Announces New Switch for Composable Architectures
- Usenet

Highlights

OpenACC – Parallelizing Loops
OpenACC is a great tool for parallelizing applications for a variety of processors. In this article, I look at one of the most powerful...

DES Completes Dark Matter Survey
Scientists will continue to analyze the results until 2021

Open Compute Project Call for Posters
Submissions for the OCP Future Technologies Symposium are due on January 31

Dolphin Announces New Switch for Composable Architectures
The MXS824 will support software-defined configuration and device sharing at low latency

Usenet
Usenet, is a gigantic Internet forum with thousands of subforums. The Usenet system is designed as a federated network, which means you just need...

Most Read

GUI or Text-Based Interface?
Sys admins are like smokejumpers who parachute into fires, fighting them until they are out, or at least under control. When you jump into the...

Exploring the /proc Filesystem with Python
The proc filesystem (procs for short), is a Linux pseudo-filesystem that provides an interface to the operating system's kernel data structures.


Quantum Computing Milestone Achieved
Researchers at ORNL perform independent operations on two qubits encoded on photons of different frequencies.

ISC 2019 Call for Submissions

The ISC High Performance 2019 conference (June 16-20) is currently open to various opportunities for engineers, researchers, and scientists working in high performance computing.
[Learn more here.](#)

Further Reading

- Shared Storage with NFS and SSHFS
- Human Brain Supercomputer Wakes Up
- Resource Management with Slurm
- Intel Unveils New Processor Line
- Discovering ROCm




ADMIN
COMPLETE ARCHIVE DVD
3,000+ articles
Smart and Practical Information for Network Professionals

7 YEARS OF ADMIN On One DVD

Order Now!
Shop.linuxnewmedia.com

2018 Archives Available Now!



Clear off your bookshelf and get every issue from 2018 at 50% off the digisub price!

The 2018 ADMIN Network & Security magazine digital archive bundle is available now (issue #43 through #48). Order now, and get all 2018 issues added to your digital account for access at any time from any device.

Order the archives today!

MADDOG'S DOGHOUSE

Despite the common assumption, everything online is not in the public domain. BY JON "MADDOG" HALL



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 Linus Torvalds and facilitated the port of Linux to a 64-bit system. He serves as president of Linux International®.

Copyright law and licensing on the Internet

"But, but, but ... I found it on the Internet!" Many people today think that just because the image or text is found on the Internet that it is free to be used or copied, as if the copied item were in the "public domain."

Of course, copyright law (and licensing) is determined by the laws of particular countries. What is law in one country may not apply in another country, but in many cases copyright laws are shared between countries via international agreements. As a rule of thumb, I tend to adhere to the strictest set of laws to protect myself, and the strictest set of laws tends to be from the United States.

The United States is so focused on copyright (and patent) law that they built the beginnings of these laws into their Constitution – the same document that supposedly guarantees the freedoms of US citizens also guarantees the existence of copyrights and patents and sets up a copyright and patent office much sooner in the document than the part guaranteeing freedom of speech.

What is not stated in these few words of the United States Constitution are all the terms, conditions, and rules of copyright and patent law [1]. These are regulated by many pages of legal words developed by lawmakers over time. Simplified, these laws state: If you write it, paint it, or take it as a picture, and it is your original "work" (or art), you have a copyright on it.

Period. End of discussion.

Your work is then guaranteed protection for a certain period of time, after which it enters the public domain.

The laws then go on to tell you how you can register your work to better protect it and make sure that people know it is yours. If someone else copies it, you can get even more damages from them.

To better inform people that they need to do something in order to use your work, you can include a copyright statement in or somewhere around your work. This copyright statement might say something along the lines of "This work is copyright" and include a date and perhaps a jurisdiction (i.e., "in the United States") to indicate what set of copyright laws you are following and what length of time the copyright will apply before your work does enter public domain.

Public domain is a real, legal phrase that has true legal meaning. Simplified, the meaning of public domain says the work is available for any type of use, and the original author has absolutely no rights at all to say how it is used, changed, or otherwise distributed.

Which brings us to "licensing," which is how you are going to share this with other people.

Normally a license gives certain rights to another individual or company to do certain things with the work in exchange for some type of compensation. This may be a very long legal contract made between the person wishing to use the work and the owner of the copyright. Many times (but not always) the compensation for using the work is money, but it could also be a "trade" of rights or other type of compensation. Nevertheless, there needs to be a license of some type, or else the user is definitely doing copyright or patent infringement whether or not they know it, and whether or not the original creator is aware of their infringement.

An exception to this (mentioned here to eliminate people screaming at me) is a concept known as "fair use," a fairly complex set of circumstances and rules where a person may use some small part, usually to review or make some point other than blatantly use an artist's whole work. Short quotations from a book or speech, usually with the original artist's name is an example of fair use. These rules may change from country to country.

The path to getting a license from the original author may be difficult. Both for the artist and the people who wish to use the art, "permissive" licenses were developed that give people certain rights to the art.

Examples of permissive licenses are open source licenses, some of which have few rules and regulations about how you use the software, or more demanding licenses, such as the GPL, which have very specific requirements about how you use and distribute the code.

More conducive to art and text are Creative Commons licenses, which can be very permissive (making few, if any demands) or more restrictive, requiring attribution or limiting your rights to non-commercial use.

I am at a conference where this topic came up, and people were genuinely convinced that "if it is on the Internet it is in the public domain." No, to be safe, assume that it is *not* in the public domain, and look to see what rights you have instead of assuming that you have the rights to use it. ■■■

Info

[1] Note: This article is trying to simplify a complex set of laws. It is not a legal discourse on copyright and licensing law. For a complete discussion, please take a course or talk to a lawyer.

DIY photography tools

Follow the Script

DIY tools and shell scripts can make for a smooth photographic workflow on Linux.

BY DMITRI POPOV

Photographic workflow is more than just keeping tabs on your photo library and processing RAW files. There are plenty of other tasks you might want to manage: from backing up RAW files and photos when you are on the move, to publishing your best work on the web.

Developing custom tools that keep a photographic workflow smooth is arguably as much fun as taking photos. So for several years now, I've been working on a number of simple applications and scripts that I use on a daily basis. All the tools described in this article are released under the GPLv3 license.

Little Backup Box

When you are traveling, having a backup of all your photos and RAW files makes a lot of sense. Even if your camera has two card slots, it still pays to have a separate backup on a different storage media. Of course, you can spend serious money on something like WD My Passport Wireless SSD. Aside from the price, this otherwise clever device has one major drawback: You can't tweak, extend, or upgrade it. Theoretically, you can upgrade storage, but doing this requires neurosurgery-level skills and most certainly voids the warranty. And

as soon as the battery inside the device dies, you are left with an expensive paperweight.

Instead of paying through the nose for a commercial product with limited functionality and a nonexistent upgrade path, you can easily transform a Raspberry Pi into a simple yet versatile backup device using Little Backup Box [1].

Opting for the DIY solution offers several advantages. To begin with, Little Backup Box is based on mature and reliable open source tools. Depending on your skills, you can customize, extend, and improve Little Backup Box to meet your specific needs. For example, with a bit of work, you can implement cloud backup functionality, add hardware buttons to trigger a backup action, and much more.

Little Backup Box is designed to run on a Raspberry Pi, which is by far the most popular single-board computer on the market. You can buy it, along with the rest of the required parts, from practically any reputable store. Using the available parts, you can build a tiny Little Backup Box device based on a Raspberry Pi Zero that fits any pocket and can be powered by a small power bank (Figure 1). Or you can opt for a more powerful and versatile device based on Raspberry Pi 3. Better still, thanks to its modular design, you can easily replace the battery, upgrade storage, and swap any faulty component. More importantly, Little Backup Box comes with a script that automates the entire installation and configuration process.

To build your own photo backup device, you need three things: a Raspberry Pi, a power supply or power bank, and a high-capacity USB drive that will act as a backup storage device. Theoretically, it's possible to use Little Backup Box with a conventional hard disk or solid-state drive, but you need to either use a powered USB hub or a hard disk with its own power source.

Transforming a Raspberry Pi into a photo backup device is a three-step procedure. First, create a bootable microSD card with the Raspbian Lite system on it. Second, boot the Raspberry Pi using the created card, run the `raspi-config` tool

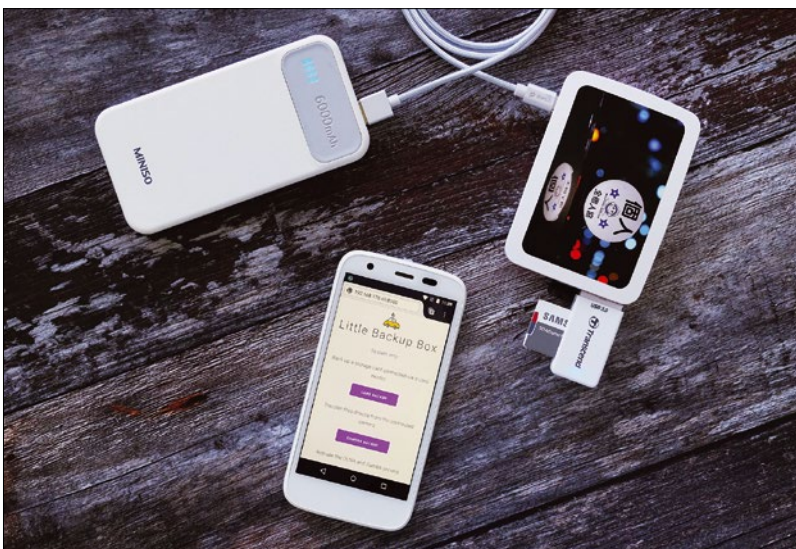


Figure 1: Little Backup Box can be powered by a power bank and controlled via a web interface.

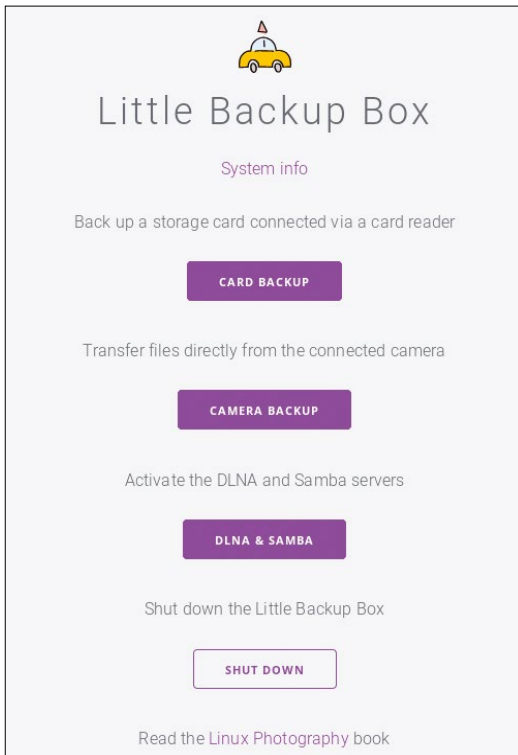


Figure 2: Little Backup Box features a simple web interface.

and configure the system. Also, edit the `/etc/wpa_supplicant/wpa_supplicant.conf` file to connect the Raspberry Pi to the desired wireless network. Third, run the command

```
curl -sSL https://is.gd/littlebackupbox | bash
```

to install Little Backup Box. That's all there is to it.

Little Backup Box supports three backup modes: *Card backup* (automatically backs up the contents of a storage card to an external storage device), *Camera backup* (transfers photos, RAW files, and videos from the camera connected directly to the Raspberry Pi), and *Remote control* (access all features via a web interface). The first two modes make it possible to run backup operations without user interaction, while the *Remote control* mode allows you to activate any of the two modes manually. In addition to that, the *Remote control* mode allows you to enable the DLNA and Samba servers, view system info, and shut down the Raspberry Pi.

During installation, the installer prompts you to choose the default mode. If you want to use Little Backup Box strictly as a backup device, choose either the *Card backup* or *Camera backup* mode. To access all the available features, choose the *Remote control* mode.

Using Little Backup Box in the *Card backup* or *Camera backup* modes couldn't be easier (Figure 2). Boot the Raspberry Pi and plug the backup storage device into it. Depending on which mode Little Backup Box operates in, connect either a

card reader with a card or your camera. The backup operation then starts automatically, and the Raspberry Pi shuts down as soon as the backup is complete.

In the *Remote control* mode, the Raspberry Pi connects to the specified wireless network, and you can access Little Backup Box's web interface using any device on the same network. Point the browser to `http://ipaddress:8000`, where *ipaddress* is the IP address of the Raspberry Pi, and you should see Little Backup Box's web interface.

Otto

When it comes to transferring RAW files and photos from a storage card or directly from your camera, there are plenty of tools at your disposal: from dedicated tools like Rapid Photo Downloader, to the Import module in digiKam. Although these tools excel at what they do, they might not always fit your particular needs. For example, I prefer to store important information about a photo in the Exif metadata *Comments* field. The information includes the camera model, lens, and weather conditions. This way, when I view the photo in digiKam, I can immediately see what camera and lens combination I used to take the photo and what the weather was on that particular day. I also wanted to streamline the import procedure, so it would not only transfer files, but also geotag, rename, and organize them by date – all in one go.

The Otto shell script [2] does all of this with practically no user interaction. During the first run, the script prompts you to specify the required info, such as the destination directory for the transferred RAW files and photos, the desired copyright notice, a key for the Dark Sky weather service API, and a Notify key. After that, using Otto to transfer RAW files and photos from a storage card is as easy as it gets. Plug a card reader with a storage card into your machine, mount the card, and note the path to it. Run the command

```
./otto.sh /path/to/card [OPTION]
```

(where `/path/to/card` is the path to the mounted card). If you run the script without any options, it assumes that the photos it's about to import are already geotagged. If the photos are not geotagged, you have three options. Using the `-g` followed by the name of the city where the photos were taken, you can do coarse geotagging. If you have a GPX file, you can specify the path to it with the `-c` parameter, and the script geo-correlates the photos for you. And in case you have multiple GPX files spanning several days, use the `-m` parameter followed by the path to the folder containing GPX

files. This instructs the script to merge the GPX files before performing geo-correlation. Once RAW files and photos have been transferred and geotagged, the script obtains the required info and writes it, along with the specified copyright notice, into the appropriate fields in Exif metadata. Once this task is completed, the script renames the transferred files and neatly organizes them by date. Finally, if you've provided a Notify key, the script sends a notification to your Android device running the Notify app [3] to let you know that the import process is finished.

To install the script, clone the project's Git repository using the command

```
git clone https://gitlab.com/dmpop/otto.git
```

or download the latest version of the script from the project's page. On openSUSE, move the `otto.sh` script to the `~/bin` directory. On other Linux distributions, you can install the script using the following commands:

```
sudo cp otto.sh /usr/local/bin/otto
sudo chown root:root /usr/local/bin/otto
sudo chmod 755 /usr/local/bin/otto
```

Finally, install the required packages. On openSUSE, this can be done using the command:

```
sudo zypper in getopt bc jg curl ExifTool
```

Konbini

There are plenty of reasons to like KDE. One of them is the ability to add context actions to the Dolphin file browser. Need to quickly resize a

photo directly in the file browser? Create a context action, and you can resize any photo by right-clicking on the file and choosing the action. Need to upload a photo to your blog? Create another action, and you are good to go.

While the process of creating custom context actions is not particularly complicated, it does consist of several steps. If you'd rather not spend time on that, you might want to give Konbini [4] a try. This is a collection of actions that let you perform several useful actions (Figure 3): resizing and recompressing photos, renaming photos using Exif metadata, and converting RAW files into JPEG format. In addition to that, Konbini features tools for adding comments to photos, geotagging photos, as well as viewing the currently selected geotagged photo on OpenStreetMap. If you happen to run openSUSE, Debian, or Ubuntu, you can deploy Konbini using a dedicated installer script. Run the command

```
curl -sSL https://is.gd/konbini | bash
```

in the terminal, and it will install all the required packages and files.

Natsukashii

Love it or hate it, Facebook offers a couple of genuinely nice features, one of them being "photos from the past." Every day, Facebook shows you photos you took on this date in previous years. Google Photos and other services have similar functionality, but what if you don't use any of them and still want to receive small photo greetings from the past? Natsukashii [5] comes to the rescue. This tool consists of two parts: a Bash shell script that finds photos taken on today's date in any previous year, and a simple PHP page that displays the results. The shell script relies on standard tools that are available on most mainstream Linux distributions (ExifTool, `find`, `sed`, `seq`, and ImageMagick). This means that you can deploy Natsukashii on practically any machine running a modern Linux distro. With a little bit of work, it is also possible to run Natsukashii on a NAS appliance.

Installing Natsukashii is easy. Clone the project's repository using the command:

```
git clone https://gitlab.com/dmpop/natsukashii.git
```

Alternatively, download the latest version of the source code as an archive from the project's page. Then install the required packages. To do this on openSUSE, run the command:

```
sudo zypper in exiftool ImageMagick php7
```

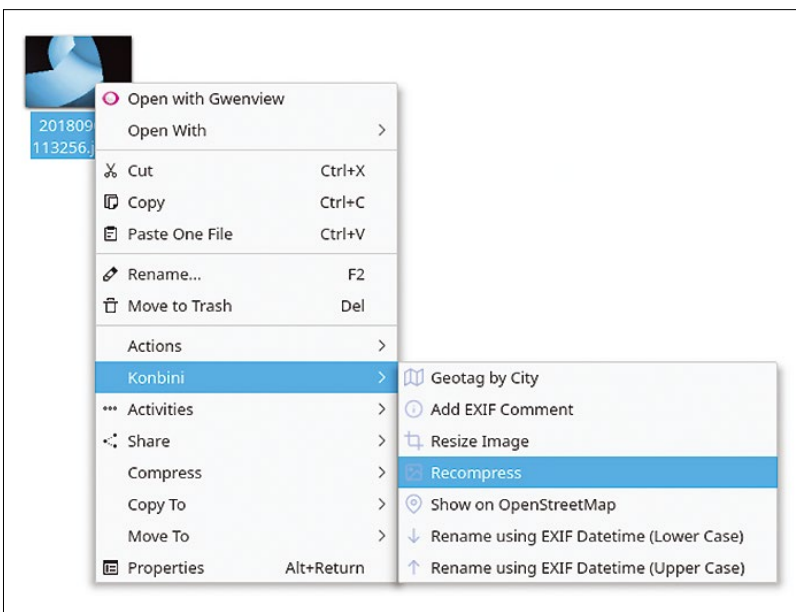


Figure 3: Konbini gives you quick access to common custom actions for working with photos and RAW files.



Figure 4: Natsukashii displays photos from the past using a simple PHP-generated page.

The PHP part of Natsukashii offers password protection, and you should change the default password before you start using the tool. To do this, replace the default password in the `protect.php` and `login.php` files in the `natsukashi i /www` directory. Keep in mind that the underlying password protection mechanism is relatively simple, so you shouldn't rely on it as a sole protection for your private photos.

Similar to Otto, Natsukashii prompts you to provide the required info during first run, and after that it works without any user interaction. While

Natsukashii does the job, it isn't particularly fast. The script goes through the specified directory and all its subdirectories looking for files with the given extension (e.g., JPG). For each photo found, the script extracts its creation date and compares it with the current date. This process takes time and resources. So if you have several thousand photos on your machine, running the script may take awhile. This means that it might be a good idea to set up a cron job that runs the script during the night.

The script not only finds photos from the past (Figure 4), but it also resizes and places them into the `www` directory and serves the result using PHP's built-in server on the specified port. Remember, though, that you should not rely on Natsukashii's unsophisticated password protection mechanism to protect your personal photos.

Conclusion

When it comes to managing your photo collection, it pays to have tools to make the process go as smoothly as possible, and these four custom tools get the job done. ■■■

Info

- [1] Little Backup Box: <https://github.com/dmpop/little-backup-box>
- [2] Otto: <https://gitlab.com/dmpop/otto>
- [3] Notify: <https://github.com/mashlol/notify>
- [4] Konbini: <https://gitlab.com/dmpop/konbini>
- [5] Natsukashii: <https://gitlab.com/dmpop/natsukashii>

Available Now

*** 2018 EDITION ***

Linux Magazine Archive DVD

ORDER NOW!
shop.linuxnewmedia.com



Searchable Linux Archive:

19,000

Pages of Practical Know-How

214 issues of Linux Magazine
on one handy DVD!

- ▶ All the Tricks
- ▶ All the Hacks
- ▶ All the Apps

ISSUE 215 OCT 2018

THE COMPLETE LINUX MAGAZINE ARCHIVE

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 destruction, loss, or damage to data and computer systems related to the use of this disc.

Managing your tasks at the command line with TaskBook

Taskmaster

Boost your productivity from the Linux command line with the TaskBook task manager and note board.

BY SCOTT NESBITT

How do you keep track of what you need to do? Chances are, that involves a calendar, a to-do list, or some arcane and complex mix of software. Although that's one way to do the deed effectively, it might not be the most efficient way to proceed.

Instead, you can manage all your tasks under one roof, or, in this case, from a single terminal window, and you can do it with a touch of visual panache. In this article, I take a look at a simple task management system called TaskBook, which can help anyone get and stay organized quickly and easily.

Why the Command Line?

Working from the command line goes beyond merely grabbing an opportunity to embrace your inner geek. Using the terminal to organize yourself has several advantages over the use of desktop applications.

The command line stays out of your way until you need it. A terminal window is usually small and unobtrusive and doesn't block out windows with your other work. Once you get into the flow of using the command line, you'll find a keyboard is faster than using a mouse.

Working in the terminal doesn't eat up system resource like many graphical applications, which is important for those of us who run Linux on older hardware that doesn't have the

memory and processing power that modern systems pack.

With no notifications and alerts constantly flashing and blaring at you, the terminal represents the ultimate in calm technology, encouraging you to be more mindful and regular in checking what you're supposed to be doing to make sure you're on top of it all.

TaskBook [1] gives you all of that, and a bit more.

Getting Started

To begin, check your Linux distribution's package manager to see whether you can install TaskBook there. If TaskBook is not available, you can use that same package manager to install a piece of software called *npm* that enables you to install software and libraries written with Node.js (an environment for running JavaScript code outside of a web browser). NPM (Node Package Manager) not only puts TaskBook on your computer, it also installs the various bits and pieces that TaskBook needs to run.

With NPM installed, crack open a terminal window and run the command:

```
sudo npm install taskbook
```

The installation process should only take a few seconds. Once that's out of the way, you're ready to go.

Configuring TaskBook

This next step is optional. To configure TaskBook, you can edit a JSON (JavaScript Object Notation) file, which is a data format that both people and computers can easily read. The file in question is called `.taskbook.json`, and you can find it at the top of your `/home` directory to edit in your favorite text editor.

The file has only three options (Figure 1). The first is `taskbookDirectory`, which indicates the path to the folder that contains your task list. The de-

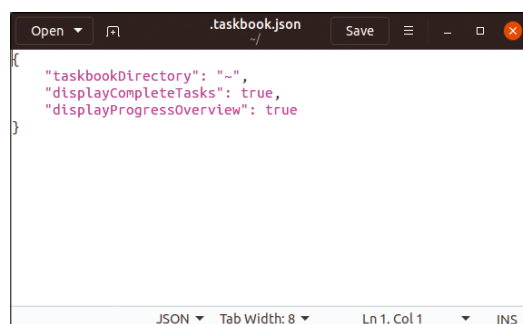


Figure 1: Editing TaskBook's configuration file.

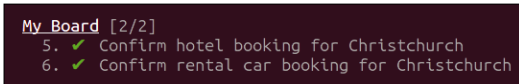


Figure 2: Completed tasks in TaskBook.

fault is the folder `.taskbook` at the top of your `/home` directory. If you prefer, for example, to save your task in a shared folder or sync your tasks with an instance of a file-sharing application like Nextcloud, change the first entry to the path to the file (e.g., `~/Nextcloud/`).

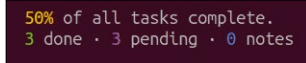


Figure 3: The progress indicator in TaskBook.

The second option is `displayCompleteTasks`, which controls whether TaskBook includes tasks you've finished in your list of tasks (Figure 2). Change this to `false` if you find that completed tasks clutter up the view.

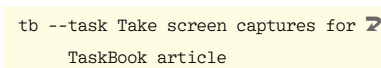
The third option is `displayProgressOverview`, which controls whether TaskBook shows how many tasks are completed or are in progress at the bottom of your list (Figure 3). Again, change this to `false` if you want to hide that information.

Using TaskBook

When you first take TaskBook for a spin, it seems almost too simple – and maybe a bit too limited. Once you become familiar with using TaskBook, you'll appreciate that simplicity.

Interacting with TaskBook is not the same as how you interact with any of the similar programs in the GUI world. Adding your tasks and notes requires just one set of options, and displaying your task uses another command. This simple method doesn't make TaskBook any less useful or effective than other task management tools.

Say, for example, I want to create the tasks I need to accomplish to write this article. One of those tasks is to take some screen captures. To do that, I type:



I continue following this pattern (with a different description for each task, of course) until I have a list of everything I need to do.

TaskBook doesn't include a way to add due dates to tasks, so it lends itself to being more of a daily task list. However, you can indicate due dates by adding `due: <date>` to the description (e.g., `due: 04/24/2019`).

Viewing a Task List

Now that you have a list of tasks, you'll want to see them. To view your

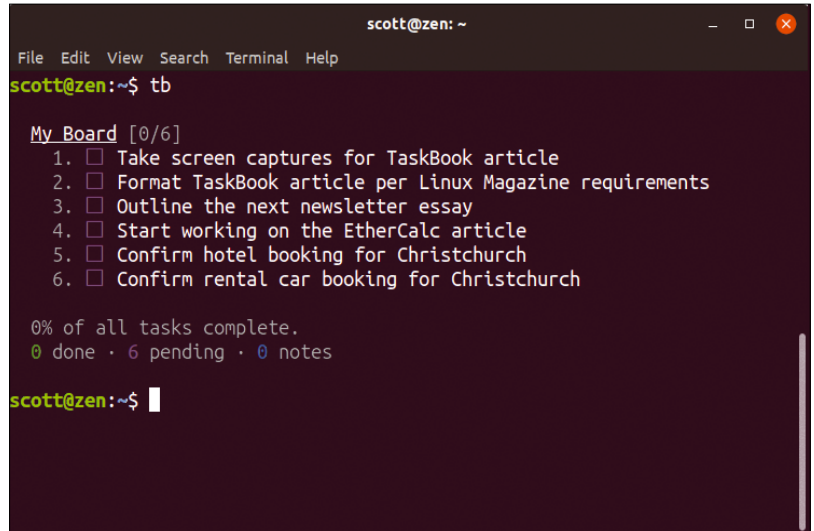


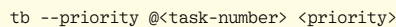
Figure 4: A list of tasks in TaskBook.

task list, type `tb` in a terminal window (Figure 4). If you left the default settings in the configuration file, TaskBook displays a running tally of the tasks you've completed, the tasks you still have to do, and the notes on your board at the bottom. I'll talk more about notes later.

Keeping this window open for reference saves having to reopen a terminal window and typing the command to run TaskBook each time you want to consult your list.

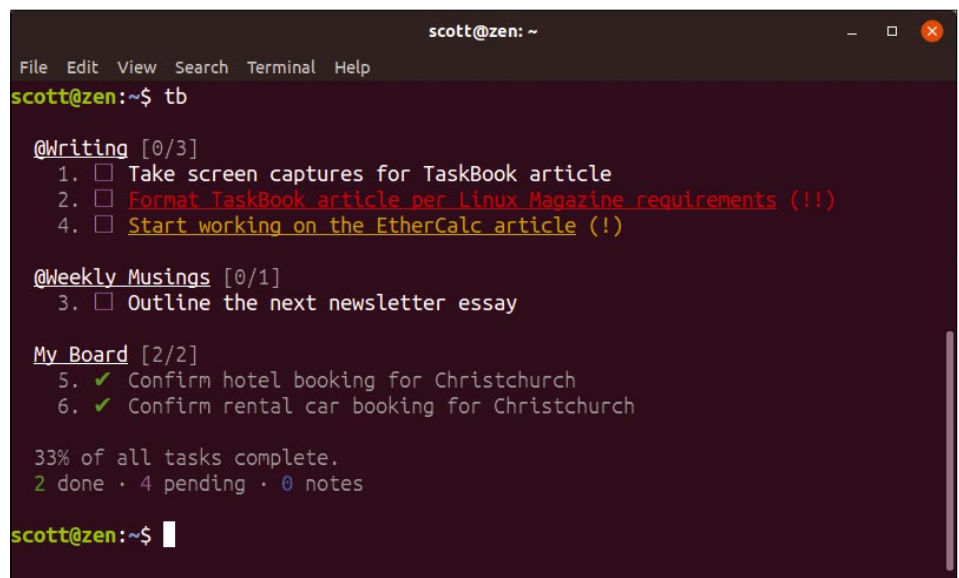
Adding Priorities

Not all tasks are created equal. Some are more important than others. You can indicate which task you need to tackle first by assigning priorities:



Priorities run from 1 (lowest) to 3 (highest). For example, say I want to give the fourth task in my list the highest priority; I type:

Figure 5: A prioritized set of tasks in TaskBook.



```
tb --priority @4 3
```

Figure 5 shows my prioritized task list. Tasks with the highest priority are red, and are followed by a pair of exclamation points. Tasks one level down are yellow, followed by a single exclamation point.

Grouping Tasks

A long list of tasks is not only demoralizing, it can set off a wave of procrastination. To get around that, you can group your tasks so you can focus on specific areas of your work or life and easily find the task you need to perform. For example, to move the first task in my list to a group called *Writing*, I typed

```
tb --move @1 Writing
```

and continued doing that until I had all my tasks in the appropriate groups. Any tasks that aren't in a group are lumped together into a generic group called *My Board*.

Changing and Completing Tasks

Sometimes, tasks change. Their scope might widen or narrow, or you might want to change the wording of a task to make it clearer. The general format is:

```
tb --edit @<task-number> <new description>
```

When I was writing this article, I was also working on a long essay for my email newsletter. The task I had in TaskBook was third in the list, and I needed to change its description, which read *Outline the next newsletter essay*, to change the wording of that task, I could enter:

```
tb --edit @3 Finish the outline for the next newsletter essay
```

When you're heads-down in work, knocking an item off your task list is a great feeling. When I finished the fifth item on my list, I marked it as done by typing:

```
tb --check 5
```

TaskBook added a green checkmark beside the task; the running tally of tasks *pending* at the bottom of the board decremented by one and the number of tasks *done* incremented by one.

Adding Notes to Your Board

TaskBook isn't just about managing your tasks, you can also add short notes to your task list or to your groups. Notes can be reminders of deadlines or to contact someone, and you can use them to remind yourself about why a task list exists. Your notes can be anything you need them to be.

I live on the other side of the dateline from the publications and clients with whom I work. Being a day ahead can sometimes become an excuse to procrastinate. An excuse I embrace all too often. To head that excuse off at the pass, I add a note like this to my *Writing* group in TaskBook:

```
tb --note @Writing Remember that the deadline for the Linux Magazine article is 24 April
```

Figure 6 shows the note in the *Writing* group (Figure 6).

Final Thoughts

TaskBook isn't for you if you need your task management tool to send regular reminders, integrate with apps on your phone, or telnet your Internet-connected coffeemaker to start brewing. If, on the other hand, you need a tool that is quick, simple, and visually appealing, then TaskBook is a great choice to keep you on task and organized. ■■■

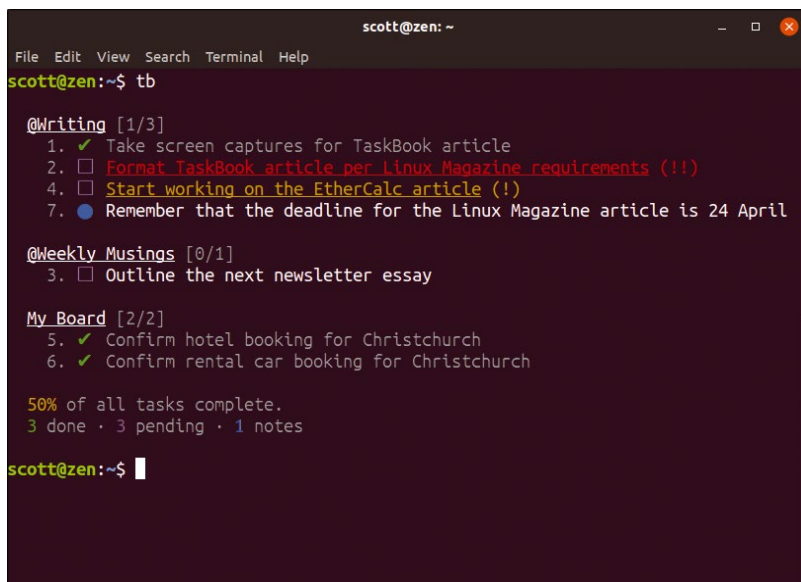
Info

[1] TaskBook GitHub repository: <https://github.com/klaussinani/taskbook>

The Author

Scott Nesbitt is a freelance journalist, technical writer, and essayist. He is also a Community Moderator for Opensource.com and has been using FLOSS since the mid-1990s and writing about it since the early 2000s.

Figure 6: A list of tasks with a note.



9 Years of ADMIN on One DVD

ADMIN

Network & Security

COMPLETE ARCHIVE DVD
ISSUES 00-49

ISSUE 50/2019

While this ADMIN magazine is not responsible for any damage to your system, we cannot be held responsible for any data loss, or damage to data and computer systems related to the application.

DVD

OVER
3,700
PAGES OF UP-CLOSE
SYSTEM
ADMINISTRATION

Smart and Practical Information
for Network Professionals



This searchable DVD gives you 50 issues of ADMIN, the #1 source for:

- network security
- system management
- troubleshooting
- performance tuning
- virtualization
- cloud computing

Clear off your bookshelf and complete your ADMIN library with this powerful DVD!

ORDER NOW!

shop.linuxnewmedia.com

FOSSPicks

Sparkling gems and new releases from the world of Free and Open Source Software



Despite both the Bug Report and Linux Voice podcasts being on hiatus, you can still hear Graham every two weeks on the Late Night Linux podcast. **BY GRAHAM MORRISON**

Astronomy app

Cartes du Ciel

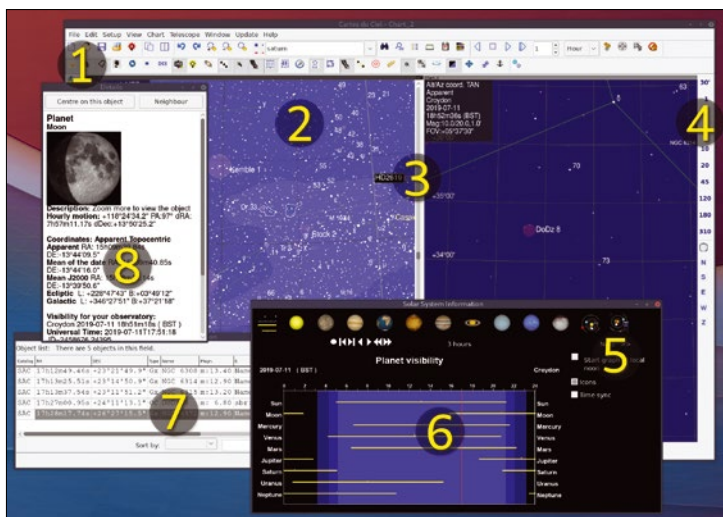
The English translation for the name of this application is “Sky map” or “Sky chart,” but this app’s French name, Cartes du Ciel, sounds so much more romantic that we’ll stick with it. And the name tells you all you need to know about this great application, an application that’s been in development since the early 2000s. Like Stellarium and KStars, Cartes du Ciel is a tool that maps a star’s position in the night sky. But rather than go for Stellarium’s realistic approach, it attempts to make

their positions clear and discoverable. This is what you need if you’re planning a night of real-life observations with a telescope, rather than a virtual preview of what the sky above you might look like (this is something that KStars also does very well).

It’s this clarity that makes Cartes du Ciel a standout. The charts, for example, resemble the kind you used to see printed in newspapers and still see in astronomy-related publications and websites. It’s diagrammatic more than realistic, and it’s the best ap-

proach when you want to manually find a star in the sky, because you then need to work your way around the obvious signpost stars and constellations to find the object for which you’re looking. This is helped by a huge star catalog, and you can keep adding more objects to the view until there’s no room left on the screen. You can also change the appearance of the charts, removing the color or the way the stars are rendered, and this is ideal if you want to print out your plans. You can also create a list of the stars you wish to observe directly within the application, and there’s a very useful pane that can be opened to list all the things you can currently see in the display. Another window lets you see what’s visible in the solar system, including what time an object rises and when it sets. This window also generates a real-time view of your selected solar body, which helps you see the red spot on Jupiter, for example, or the current phase of the moon – even on a cloudy night!

One of Cartes du Ciel unique features is that you can open multiple windows on the sky at the same time. This helps you see two different zoom levels of the same area, for example. By adding the configurable telescope reticle to predict the field of view, you’ll be able to see through your hardware. You can also control this hardware directly, thanks to the integrated support for LX200, INDI, and ASCOM telescope controllers. With comet and asteroid data, you can target even transitory objects as they cross the sky. It’s an application, perhaps, without the shine of Stellarium or the OpenGL acceleration of KStars, but it does offer a sensible clear and predictable way to learn about the night sky and plan an evening out beneath the stars.



1. Mirror or flip views: Useful for previewing what you’ll see through a real telescope. **2. Sky chart:** Generate clear and easily navigable sky maps. **3. Split views:** See multiple parts of the sky at once, and sync views together. **4. FOV:** Quickly change the field of view and viewing direction. **5. Planet previews:** Open a render of how each planet currently appears. **6. Visibility:** An easy-to-view chart to see which objects are in the sky tonight. **7. Observations:** Make your own notes on what you want to see or have seen. **8. Object details:** Access all the important info on the selected celestial body.

Project Website

<https://www.ap-i.net/skychart/en/start>

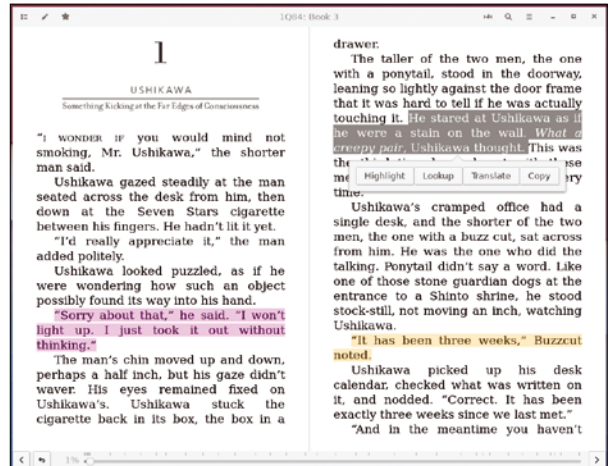
E-reader

Foliate

It goes without saying that we love print, and it's great to see real physical and tangible books once again increasing in popularity after suffering somewhat at the hands of e-readers like Amazon's Kindle. But it's also true that digital content is very useful, especially for reference and archival material when travelling. In those cases, you need a beautiful, distraction-free reader application on which you can rely, such as Foliate. The best thing about Foliate is its wonderful GTK-based minimalist user interface (UI). In particular, it feels like an application that always puts the user-experience first, and what the user wants most of all is a zero-clutter portal to the words in the book. Foliate does this by defaulting to a sim-

ple, clear, and crisp two-page spread view and uses excellent font rendering with a simple configuration pane for font selection and size (alongside being able to select the publisher's own chosen font).

When reading, the progress bar beneath the pages shows the chapter markers and also lets you easily move from one page to the next or from one section to another, while always keeping your current location in mind. This can be easily turned off, and there are additional options to change the contrast of the view, switch to a dark or a sepia mode, or even adjust the brightness on any display. Selecting some text on a page lets you create an annotation, leave the text highlighted in a color of your choice,



Research has shown that you retain details from an ebook better if you can see your relative position within that book, rather than seeing an endless wall of text.

or even look up a word in a dictionary. There's a great bookmark system that mimics bookmarks in an application like Firefox. All of this is secondary to simply reading the text, but Foliate excels at these discrete optional extras and design.

Project Website

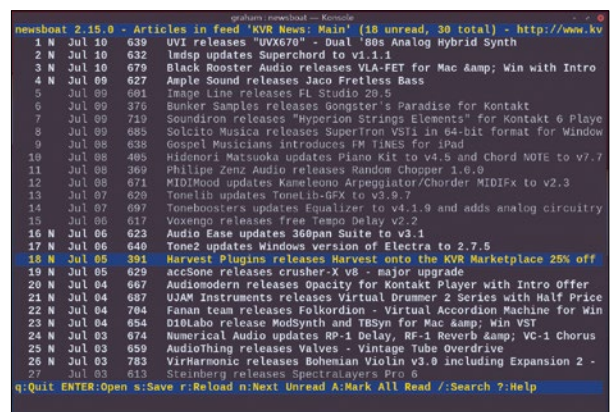
<https://johnfactotum.github.io/foliate/>

RSS reader

Newsboat

RSS, the XML-based synopsis of a website and a podcast's content won't die. It lives on, to fight through the next fashionable wave of social media, and this is definitely a good thing. RSS is a standard that isn't controlled by any one body and can't be easily subverted into a tracking or advertising platform. It remains a useful overview of new updates, without having to revisit a site to see what's changed, and many of us manage dozens of these subscriptions arranged and aggregated into categories. As fans of RSS, we've mentioned much of this before, but what we've not looked at is a command-line RSS reader that can bring all this RSS goodness to your terminal like Newsboat does.

With Newsboat, you first need to source your own list of RSS files, and the easiest way to do this is to import an OPML file. These files are typically generated by online RSS aggregators, and they're common in desktop RSS readers, too. Adding one to the `import` argument when you first launch Newsboat makes configuration almost instantaneous, which is the best thing about reading RSS on the terminal. Everything happens so quickly. With your feeds listed in the main view, pressing `R` will reload the state of every feed in your list, and this happens quickly even with dozens of feeds. You can then easily see which feeds have new posts and select these to see the title synopsis of the new stories. Select-



Newsboat is actually a fork of Newsbeuter, which is no longer actively maintained.

ing a story will then let you read the content, presented in simple distraction-free text. This is another advantage of the command line, as you're unlikely to be tempted to follow links as you might be with a web front end, which means you can quickly scan through your newsfeeds. It's a brilliant way of keeping up to date, and a great vindication for RSS and why it has remained popular, despite the new age of social media sites and walled gardens.

Project Website

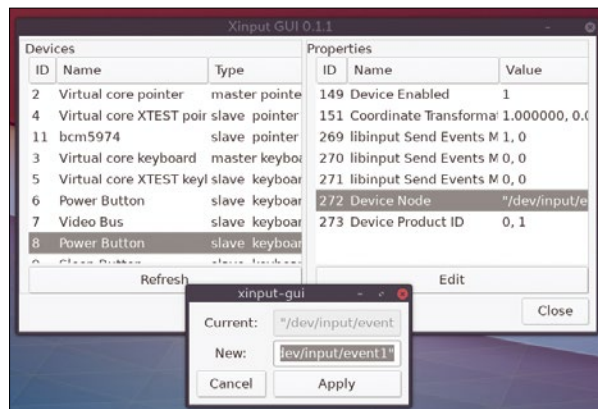
<https://newsboat.org/>

Input configuration

xinput-gui

With even the traditionalist Debian moving to Wayland, these could well be the end of days for the X.Org graphical desktop – either an endgame we’ve been waiting over a decade for or one that could easily stretch a decade more when X makes it into the next Ubuntu LTS that requires 10 years of support. Regardless of how quickly Wayland is able to supplant this ancient protocol, there’s no doubt that X is seriously configurable and can be made to work with almost any kind of input hardware thanks to its Xinput subsystem. But for a subsystem to a graphical environment, it’s perhaps surprising that there’s very little graphical help with its configuration. The Xev tool can help

with feedback and testing, and desktops will often include their own limited configuration panes, but there’s nothing as comprehensive as editing configuration files by hand. `xinput-gui`, however, gets close. `xinput-gui` feels a little like the Windows Registry Editor, only it allows you to edit the many Xinput values being used by your system. When its simple window loads, it lists these editable parameters on the left, while the right side of the UI shows the values for each parameter. Select one of these and press on the *Edit* button to change whatever is assigned to it. What’s most surprising is that you seldom see just how many input devices are connected to your system, and what by default



You never realize how many input devices are connected, or how many parameters they need to work, until you see it all listed in front of you.

they’re all configured to do. A typical list will include the power button, the sleep button, any lid activators, your pointer, and of course your keyboard. You can then change values, such as scroll methods for a touchpad or the touch sensitivity for a click, to change your how desktop operates without having to edit a text file.

Project Website

<https://github.com/IvanFon/xinput-gui>

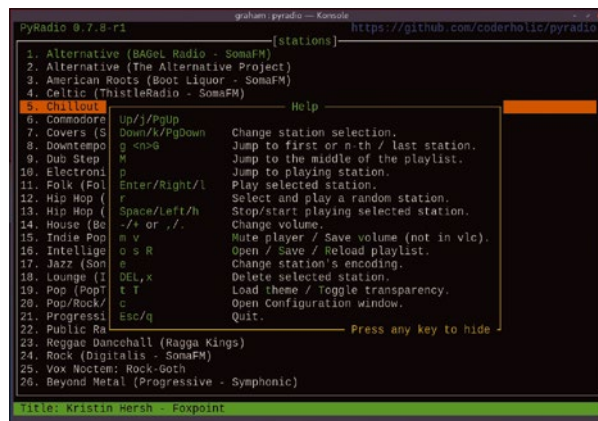
Internet radio

PyRadio

Despite the Internet, radio hasn’t died. Millions of people still listen to live transmissions broadcast across the airwaves every day, and many more have transitioned to radio packets broadcast across the Internet without really thinking about it. There’s something unique about radio that shouldn’t be allowed to die in the age of infinite distraction. Aside from listening to local broadcasts, one of the best things about old school radio was using a cheap shortwave tuner to listen to channels from hundreds, and sometimes thousands, of miles away. There used to be many of these channels, and there still are, compacted into a half twist of the tuning dial. Listening often became a voyage of discovery, and while Internet radio

has removed much of the chance of random listening, it hasn’t replaced the joy of listening to something different. This is where PyRadio can help, because there can’t be many Internet radio players that offer Commodore 64 rock remixes from an Internet radio station as a primary playback source.

PyRadio is a command-line Internet radio player with a great feature if you have MPV, MPlayer, or VLC somewhere in your path for audio playback. When first launched, you’ll see a simple list of curated stations, easily selectable using the cursor keys. Press the space bar to start playback, or if you’re after the true shortwave experience, press *r* to start a random station. Pressing *?* will open the pop-up help shortcuts, so you



Relive the glory days of shortwave radio with the austerity of a command-line Internet radio client.

never need to remember which keys do what. In particular, there are plenty of configuration options. You can load your own playlists, for instance, but there are also several built-in light and dark themes from which to choose. Thanks to this and its low resource usage, PyRadio is perfect for dark nights under the bed sheets, tuning in to pirate radio stations from the other side of the globe.

Project Website

<https://github.com/coderholic/pyradio>

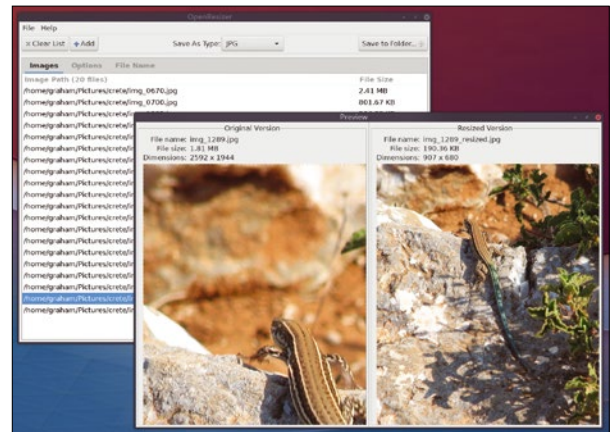
Batch image processor

OpenResizer

Sometimes, you need something simple to do a simple job. You don't want to learn an arcane command-line trick or download a full-featured application to use only one of its features. You want one tool that does exactly what you need and nothing else. OpenResizer loads a selection of images and lets you save resized versions of those images, either with new file names or in a different location. It's the kind of batch image processing conundrum that was solved a long time ago, but there's never been a perfect or memorable way of doing it. OpenResizer couldn't be simpler to use, and the user experience consists of three tabs from left to right. The left tab holds the lists of images you wish to process. These can be loaded individually,

or from a file requester that thankfully supports multiple selections. Only PNG, JPEG, and BMP formats are supported, but that should be enough for the majority of uses. The second tab holds the options for customizing the conversion, while the third tab handles the output file names for your converted images.

For a quick and easy conversion, you simply import the images you want converted and click on *Save to Folder* to choose a destination. Your images will be converted without any further configuration, but if you look a little deeper, especially on the *Options* tab, there are lots of neat additions to help with the conversion. If you double-click an image in the list, for example, you'll see a before and after image of what the conversion will do, and



There's even a Windows version of OpenResizer if you get stuck converting old family photos on your parent's Windows 95-era PC.

the *Options* tab includes the ability to scale by percentage and by pixels. If you're processing a huge number of images, you can also take advantage of however many cores your CPU has or restrict the processing to a single core if you'd rather use the remainder of your CPU power watching YouTube. It is a simple job, but it can't be done much more effectively.

Project Website

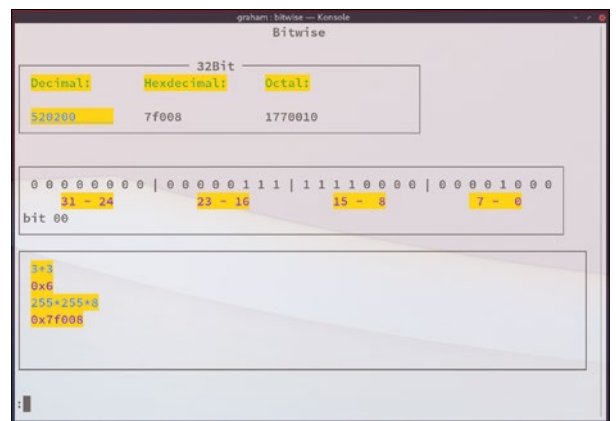
<https://jobinpy.com/>

Number conversion

Bitwise

If there's some task you need to perform regularly, it matters how efficient the process becomes, even if that task is simple. This is why the batch image processor above is so effective, and it's also why this neat little command-line utility also succeeds. In fact, the developer created it after noticing that the desktop calculator was being launched too often and wanting to stay on the command line. That's because Bitwise, as its name suggests, is going to be primarily of use to low-level programmers who often need to convert between binary, octal, and hexadecimal number bases, which is often a requirement when you're getting hardware to work or developing parts to plug into the kernel.

Starting Bitwise without any arguments will launch the small utility in interactive mode, with the cursor flashing in a field that accepts a decimal number. As you type a number in, its value for the other number systems is calculated and shown in real time. You can use either the cursor keys, or the Vim movement keys (*h, j, k, l*) to switch between the fields and enter your chosen number using a different base. With the binary field, you can turn on and off separate bits with the space key. Another similarity with Vim is that you can press `:` to enter the expression calculator mode. This lets you type expressions to be calculated, and the answer alongside the expression is shown in a history



As its name suggests, alongside base conversion and expressions, you can also use Bitwise to perform bitwise operations.

pane. You can also run this mode from the command line without having to enter the interactive mode at all. It's simple, and it's the kind of thing any programmer could build over their lunch hour, but it's also effective.

Project Website

<https://github.com/mellowcandle/bitwise>

Eurorack emulator

VCV Rack

Of all the audio applications we've looked at in these pages, VCV rack remains the most impressive. It's a software emulator of Eurorack hardware, which is a common form factor for many different types of audio processors, synthesizer modules, and voltage manipulators. Many independent creators, alongside long-standing hardware manufacturers, produce their own Eurorack modules. Thanks to the standard, they all fit into the same cases and power supplies, and they all talk to each other with the same voltages and audio signals. This magazine was one of the first media outlets to cover the original release, version 0.4, almost two years ago, and it has since gone on to become hugely popular. This isn't surprising because Eurorack itself is incredibly popular, and Rack allows you to construct a virtual set of your own hardware for free, thanks to its inclusion of so many bread-and-butter modules,

alongside many third-party modules that are perfect recreations of real open source hardware.

All of which makes this release of Rack 1.0, and version 1.1 soon after, a major event. So major, in fact, that it's the second time in the project's history that the API has been broken, and modules will need to be updated to the new version. Fortunately, there's a community effort to help the developers of open source modules to update their modules, and you can still use the previous release if necessary. You also might need to use the previous release if you want to use the bridge for integration with your DAW, because this has been dropped from version 1.0 to make VCV completely stand-alone. This is a contentious move, but it's to make way for a long-promised VST version for the release of version 2.0, and the VST version is going to be a commercial product. This isn't because the project is aban-




There are many native and third-party modules for VCV Rack, but some of the best third-party plugins do need to be ported to the new version of Rack.

doning its open source roots, but because the project doesn't want to use funds given for the open source version to build a plugin bridge to proprietary formats such as VST2/AU and AAX for use in DAWs. It would be nice if the same isn't true of any potential LV2 option for Linux, though.

But the main news for this obviously revolves around the new features that have necessitated such changes. The most important is the ability to create polyphony (playing multiple voices) from monophonic (playing a single voice) sources. Doing the same on a real physical Eurorack system would involve duplicating all your modules and connections for each extra voice, which is why so many classic synthesizers, the original modulars, new modulars, and most Euroracks are monophonic. But software can perform all kinds of magic that isn't possible in the real world, and easily duplicating processes is one of them.

This update makes VCV more powerful than its hardware counterpart. To help with this, the audio engine now supports the multicore architecture of a modern CPU, so you can spread the load more evenly across your system. MIDI has also been overhauled. You can use the shiny new MIDI Polyphonic Expression (MPE) standard to control all that new polyphony. While you could always use MIDI input data to control your modules, you can now generate MIDI from the virtual control voltages sent between modules and from there out into your other software or equipment. If you're at all interested in synthesizers, sound manipulation, open source audio hardware, music creation, and generative audio, VCV is one of the best pieces of software you can waste your time on, and one of the best examples of open source capabilities.



The new release of VCV makes managing the dozens of modules it supports easier, as well as offering parameter input via the number pad and adding the magic of polyphony.

Project Website
<https://vcvrack.com>

Diablo engine

DevilutionX

Diablo was a hugely influential “hack and slash” style, action role-playing game from the late 1990s. The game took an isometric third person view of your protagonist and whoever you happened to be travelling with as you explored the landscape and dungeons of the game’s landscape; levelling up, collecting things, and discovering locations; and ultimately hacking and slashing at hundreds of the land’s denizens. The game’s success wasn’t because any of these elements were particularly unique to Diablo. It was because of their combination and implementation, which felt so finely-tuned and responsive that you just wanted to take your character further. Even with its inevita-

ble sequels, there’s still nothing quite like Diablo 1 in the modern era, and there was never going to be thanks to the original source code being lost forever. Except, there is now, thanks to the totally unofficial and unaffiliated DevilutionX.

DevilutionX represents a remarkable effort by the development team, and hacker GalaXy-HaXz in particular, to reverse engineer the original source code from symbolic debugging code left in a PlayStation port of the original. This effort took over 1,200 hours of work and has resulted in an exacting port that will compile and play on Linux – something unthinkable in the 1990s. The intention is to preserve the code and the playability of the original, which



It’s still early days for the DevilutionX project, but Diablo is already playable to an extent, and playability is improving rapidly.

means there won’t be any new features or new resolution support, but the developers will make it work with OpenGL, and they do hope to modernize the UI. As with many of these modern re-implementations of old games, you will need the original assets if you want to play the original game. For Diablo, that means placing a file called `diabdat.mpq` in the right place. You can get this from the original CD-ROM or from the GOG download if you purchase the game. It also, tantalizingly, opens up the possibility of someone creating a game using free assets on the old classic engine, but that would obviously require considerable work.

Project Website

<https://github.com/diasurgical/devilutionX>

Music game

Performous

Even though the idea behind a game like the classic PlayStation SingStar games sounds horrifying, it’s the perfect example of a game that shouldn’t be judged until played. SingStar is/was a form of karaoke/torture party game. It typically came bundled with a couple of microphones, and its premise was that hapless players would use these to sing along to a track played by the console, with the lyrics and notes shown on screen. The players were then scored according to how close their pitch and timing matched those of the original. Due to the limitations of Playstation processing, this scoring was often crude and approximate, but it all ended up being surprisingly fun – es-

pecially at family get togethers. Normally, everyone playing would sound equally bad, and a play session would often end in fits of laughter and incredulity. This is why, after the sun has set on SingStar’s popularity, it’s great to see a few open source equivalents keeping the beat alive and letting you play a similar game on your Linux box. And Performous is one of the best. Thanks to its support for lots of song formats, including UltraStar, Frets on Fire, and StepMania formats, plus the ability to work with dance mats for additional personal humiliation, you can get started quickly. The game can work with multiple inputs and outputs, provides comprehensive configuration for using one microphone per



Make your parties sing with Performous! You can even impress your friends by using the keyboard as a drum kit.

player, and even allows for a webcam-based background to the play screen, so you can see who is laughing at you from behind. It works brilliantly and plays just the same as the original – complete with cheesy band photos as a background, highlighted lyrics, and dodgy pitch tracking for your singing.

Project Website

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

Decentralized social media Welcome to the Fediverse

If you're looking for social media options where the user has more control, you'll find a range of options to explore in the Fediverse, including the popular Mastodon.

BY PAUL BROWN

Despite what it may seem, despite its promise of unbridled communication possibilities and its supposed gift of giving voice to the traditionally voiceless, current social media is a walled garden at best, although a slimy cesspit with bars over the top would be a more apt description.

The problem with the current social media status quo is that the platform does not have your interests at heart. The companies that run proprietary social media platforms gradually introduce more and more restrictive terms of service, package your personal data and sell it off to other companies and governments, make their algorithms more manipulative, and so on.

In the Beginning

You could renounce social media altogether and go and live as hermit, but why throw the baby out with the bath water? The Fediverse [1] is the FOSS community's response to walled-gardened, personal data-leaching, closed, and proprietary social media. Strictly speaking, the Fediverse would cover all social media services (and some services that are not related to social media at all) that do not rely on a centralized and controlling entity (like, say, Facebook, Twitter, or Google), but instead allow a user to run their own server (often called a "pod") and to connect it to a network of similar servers to share the media between them. This allows user A to post to pod X and user B to read it from pod Y.

The genesis of the Fediverse can be traced back to projects like diaspora* [2] or Identi.ca [3], the earlier attempts to create a non-centralized network of social media. Things evolved and, Identi.ca's creator, Evan Prodromou, developed pump.io [4], a framework upon which developers can build social media services.

It then came to pass that the W3C saw that pump.io was good. It took its essence and breathed new life into it, transforming it into ActivityPub [5] and publishing a recommendation.

The modern Fediverse was born. Let there be posts!

ActivityPub is a protocol and an API that allows the services that follow it to communicate across the Fediverse with nodes of the same kind of service or even with nodes of different services. For example, you can post a picture to your account on Pixelfed [6], a photo-sharing service (Figure 1), and have it pop up on the feeds of people who follow you on Mastodon [7], the Fediverse's most popular microblogging service (Figure 2).

The Fediverse not only shares posts, but it also shares resources among nodes and even, by using P2P technologies, among clients. Say you decide to open your own PeerTube pod on your own server. Sure, you will need a decent amount of bandwidth, but not as much as you think, as the video service has clients watching the same clip share the load.

Apart from Mastodon, PeerTube, and Pixelfed, there is a Fediverse social media platform for audio aimed at musicians and podcasters called Funkwhale [8], and then there is Hubzilla [9], which describes itself as "a powerful platform for creating interconnected websites featuring a decentralized identity, communications, and permissions framework built using common web server technology" – so a federated Facebook? Hubzilla is a bit rough around the edges, and I haven't totally figured it out yet.

To be fair, most of the Fediverse is still rough around the edges. Pixelfed lacks a mobile phone app with camera integration, which, if I understand things right, would make it more attractive to Instagrammers. There does not seem to be a way of seeing who your subscribers are on PeerTube, and Hubzilla's interface looks dated and only implements ActivityPub as a plugin, so I'm guessing that connectivity with other services throughout its network may be spotty.

But this is kind of normal: Many of the Fediverse's services are in a nascent phase, with few users and nearly no posts.

This is not the case of Mastodon.

Mastodon Basics

Mastodon has been the subject of many articles predicting its demise, a nearly certain sign it will survive. Indeed, according to Mastodon's account-counting bot, Mastodon is quickly approaching 3.5 million users and is arguably the most popular, non-proprietary social media platform out there.

One of the reasons for its popularity may be due to its design – Mastodon just looks good (Figure 2 again). It is also simple to set up and use.

Mastodon is laid out in columns. The first column from the left is where you write your toots (up to 512 characters long). You can also add attachments: up to four pictures, or a video, or an audio file. You can also create a poll, decide on the privacy of the post (public, for followers only, direct message, etc.) and set the content as hidden behind a warning, which is ideal for spoilers or for potentially offensive material.

The second column shows your feed. This contains your own toots and those of the people you follow. The third column shows notifications, that is, when people favorite or boost your toots, direct messages they send you, comments, and so on.

The fourth column is a bit of a wildcard. By default, it shows a menu that allows you to pick between showing the local timeline (i.e., all the toots sent in your pod as they happen), the federated timeline, or all the toots posted in your pods network as they happen. You can also turn it into an inbox and show direct messages here, list the toots you have favorited over time, and make lists in which to classify the accounts you follow. Or if you click on a user's profile, it

shows the user's information. Their latest stats and toots will also show up here.

You can also change what is shown in the fourth column using the toolbar across the top of the first column. The *Getting Started* button will take you back to the fourth column's original menu, the *Local Timeline* and *Federated Timeline* buttons will show each of those timelines in the fourth column.

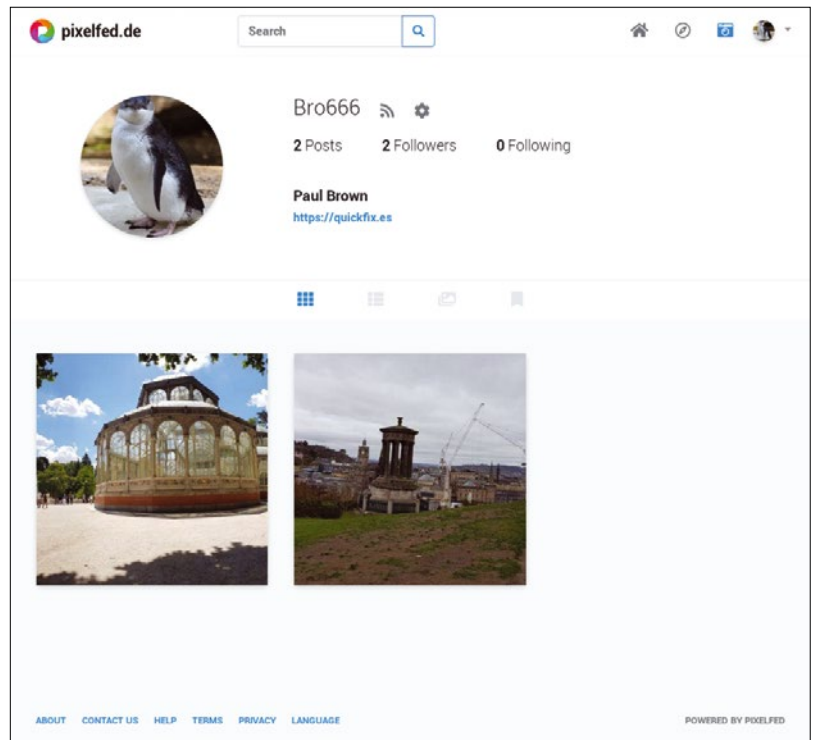


Figure 1: What you post to Pixelfed...

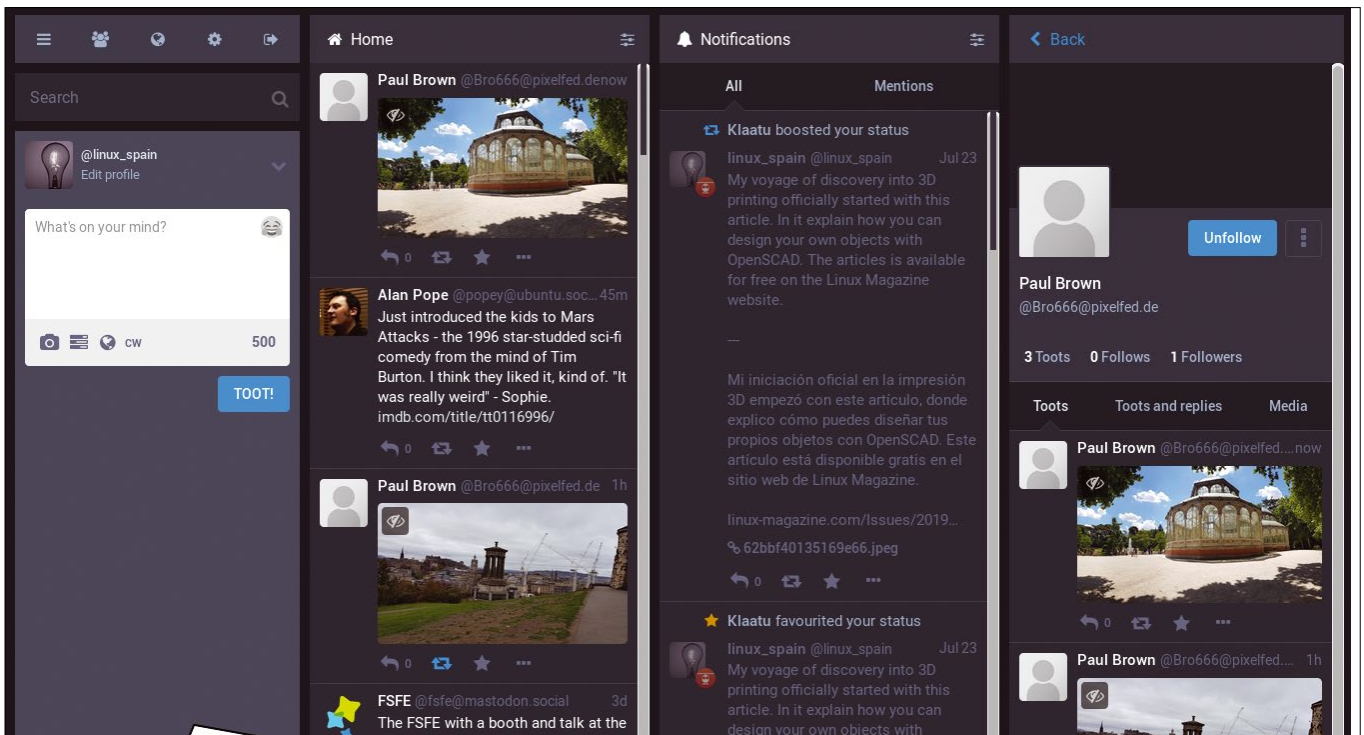


Figure 2: ... will pop up in the feeds of the people who follow you on Mastodon.

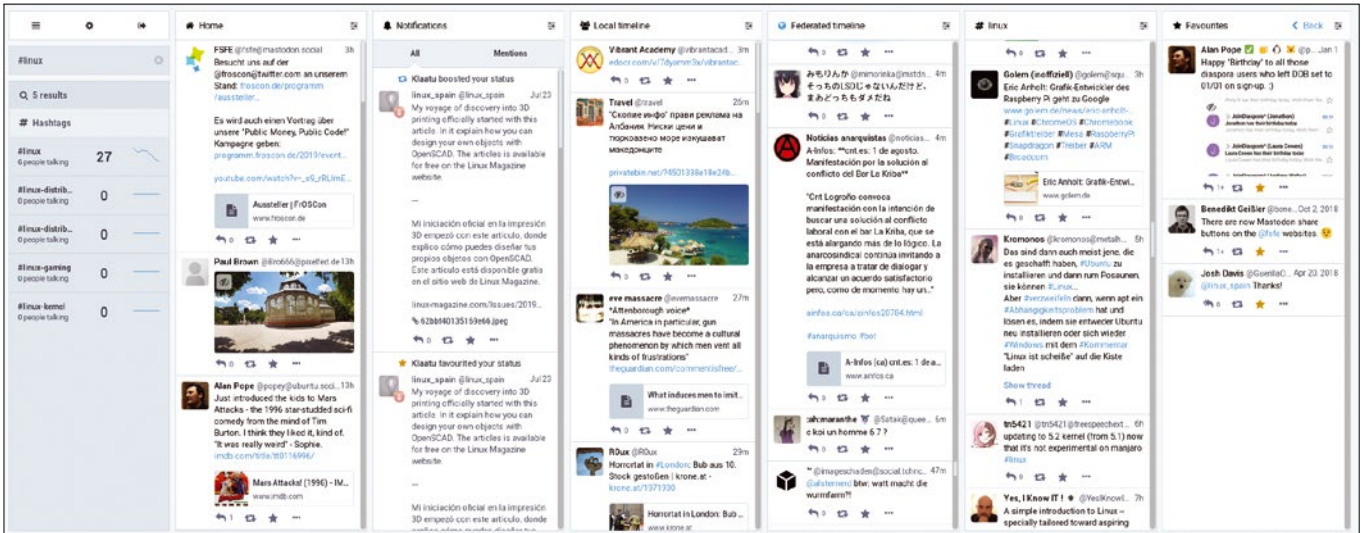


Figure 3: Adding new columns lets you follow several timelines and hashtagged posts at the same time.

The *Settings* button lets you configure everything regarding your account, from your display name and the look and feel of you account, to the custom applications you can use to mine information from Mastodon (more about this later).

In the Settings, by the way, you can also switch on a feature that lets you add more columns to your Mastodon interface. Go to *Settings / Preferences* and check the *Enable Advanced Web Interface* checkbox. The description says:

“If you want to make use of your entire screen width, the advanced web interface allows you to configure many different columns to see as much information at the same time as you want: Home, notifications, federated timeline, any number of lists, and hashtags.”

Annoyingly, it does not say how to do this, so here it is: Say you want to add a new column with the federated timeline. Go to the fourth column, click on *Federated Timeline*, and the federated timeline will appear,

occluding the *Getting Started* panel/menu/column/thingy. At the top right of the timeline, there is a configuration button. It looks like three little sliders. Click it and then click on *+ Pin* in the fold out that appears. The column is now fixed and the *Getting Started* panel/menu/column/thingy will pop up to the right, allowing you to create more columns.

If you are interested in a trending topic, click on its hashtag and a feed with all the toots containing that hashtag will appear over the *Getting Started* column. Again, by clicking the configuration button in the top right of the column, you can pin the column and follow the topic as it develops (Figure 3).

You can search for hashtags (and users and more things) using the search box located in the first column on the left, right beneath the toolbar.

Coding for Mastodon

With the rise of Mastodon's popularity, a few clients have popped up that allow you to access you account from your mobile. My personal preference leans towards Fedilab [10], which lets you manage accounts from Mastodon, PeerTube and soon Pixelfed, all from one single app (Figure 4).

If you can't find a client you like, don't despair: Coding for Mastodon is easy. Its API is well-documented, and there are plenty of hooks in several programming languages.

Let's use Python to do some simple reading of data from accounts.

The module you need is called `Mastodon.py` [11], and the documentation is located at [12].

You can install the modules with

```
sudo pip install Mastodon.py
```

Let's write a short script that will solve a real problem: Mastodon's web stops showing the precise

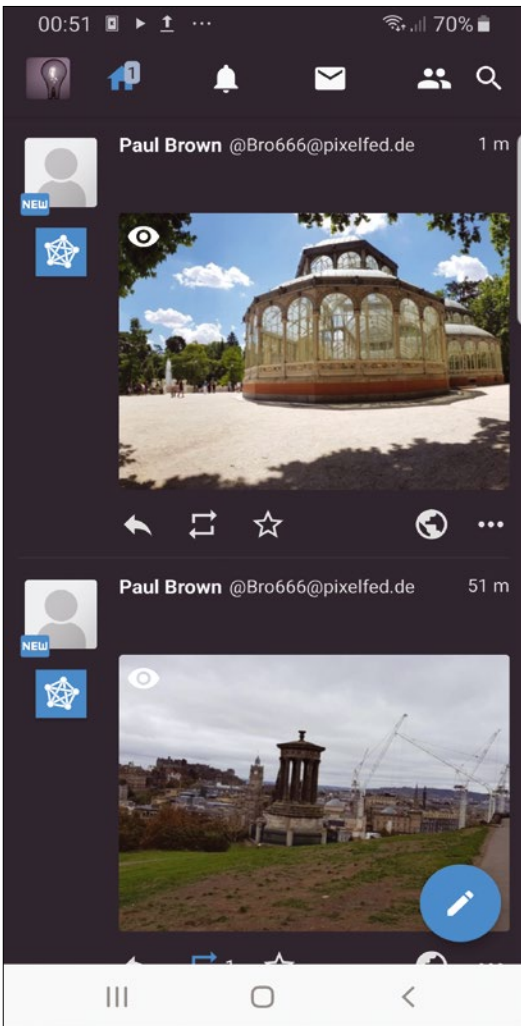


Figure 4: Fedilab lets you manage multiple Mastodon and PeerTube accounts from your phone. It will soon support Pixelfed, too.

number of followers from 1,000 followers up, opting for showing, for example 2.7K, instead of, say, 2,734, 2,755, or whatever. This makes tracking the popularity of an account difficult.

Your first app uses the identification number of a toot to print out the name of the account and the number of its followers.

Before you even start to write your code, you must give your app a name and register it with the federated network of Mastodon pods you are going to query. To do this, log into your Mastodon account. Go to *Settings / Development* and click the *New Application* button. Fill in the *Application Name* text box and choose the scope of your application. In this case, you only have to leave the *Read* checkbox marked, because that is all you are going to do: Read data from the network.

Click *Submit* at the bottom of the page, and this will log your application with the federated network.

Now, if you click on your application's name, at the top of the page it will show a list of three keys. The key labeled *Client Secret* is the one you need to finish registering your application with the network (Figure 5).

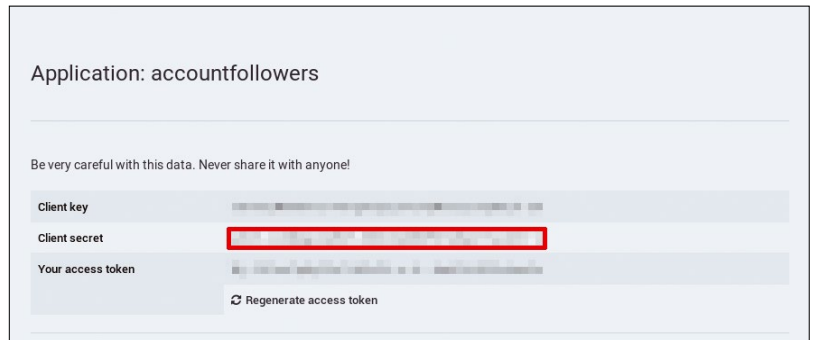


Figure 5: Copy and paste the secret key into your application to register it with the network.

To finalize the registration, create your application with the code shown in Listing 1.

On line 5, you put the URL of the Mastodon instance where you are registering your application. It will look something like:

```
api_base_url = 'https://mastodon.xyz/'
```

On line 6 you put the secret client key you saw above. The line will end up looking like this:

```
to_file = 'HHssuwd_AhTot4LlyMaDEUpK3y_2
aEJHjjIwJj2289usQsjSJqiw-d'
```

You run this version of your app once and then you can comment out or delete the registration code.

The final version of your application reads a list of toot URLs (see the box “Grabbing a Toot”) from the command line and then processes them to figure out the account from whence they came and the number of followers that account has. The application would look like what you can see in Listing 2.

Listing 1: accountfollowers.py (Part One)

```
01 from mastodon import Mastodon
02
03 Mastodon.create_app(
04     'accountfollowers',
05     api_base_url = '[URL OF MASTODON
06         INSTANCE]',
07     to_file = '[CLIENT SECRET]'
08 )
```

Listing 2: accountfollowers.py (Part Two)

```
01 import sys
02 from mastodon import Mastodon
03
04 sys.argv.pop(0)
05 for eachToot in sys.argv:
06     tinstance = eachToot.split("/")[2]
07     taccount = eachToot.split("/")[3]
08     tid = eachToot.split("/")[4]
09
10     mastodon = Mastodon (
11         client_id = '[CLIENT SECRET]',
12         api_base_url = 'https://' + tinstance
13     )
14
15     print ("User:" + taccount)
16     print ("ID: " + str(mastodon.status(tid)['account']['id']))
17     print ("Followers: " + str(mastodon.status(tid)['account']['followers_count']))
18     print ("===")
```

Grabbing a Toot

To grab the URL of an individual toot, click on the time the toot was published. This is in the upper right-hand corner of the toots box. The toot opens in a new window/tab with a URL that will look something like:

```
https://framapiaf.org/@davidrevoy/ 2
102552017830926864
```

Breaking it down, *framapiaf.org* is the Mastodon instance (pod) hosting the toot, *@davidrevoy* is the user, and *102552017830926864* is the unique ID assigned to the toot.

This works is as follows:

- 1 On lines 1 and 2, you import your modules: `sys` to be able to read parameters off of the command line, and `mastodon` for all the Mastodon stuff.
- 2 On line 4, you remove the first element from the list of parameters passed from the command line, as it is the name of the script, and all you want are the links.
- 3 On line 5, you loop through the links.
- 4 On lines 6 through 8, you split the URL for each toot up into the Mastodon pod (instance) it was posted to, the name of the account, and the ID of the toot itself. If you have a URL for toot that looks like this:

```
https://mastodon.technology/@kde/2
102455343978152552
```

`t instance` will contain *mastodon.technology*, `t account` will contain *@kde* and `t id` will contain *102455343978152552*.

- 5 On lines 10 through 13, you create a `mastodon` object that will let your application connect to the instance and identify itself with the secret client code.
- 6 On line 15, you print the user name to the console.
- 7 On line 16 and 17, you use `Mastodon.py`'s `status()` function to read data from the Mastodon instance.

The `status()` [13] function reads a toot (known in ActivityPub jargon as a *status*) given an id for the toot. It returns a `toot` dictionary [14], containing everything you need to know about the toot. The `toot` dictionary includes a subdictionary called `account` [15], which contains information about the account the toot came from, including the accounts id number (accessed in line 16) and the `follower_count` (accessed in line 17).

You can call your script like this:

```
python accountfollowers.py 2
https://mastodon.technology/@kde/2
102455343978152552 2
https://mastodon.social/@fsfe/102563322717544968
```

Both the `toot` and `account` dictionaries contain much more information. You can also use them to pull the content of the toot or maybe make a list of accounts the user follows. The point is, that is it. It is that simple to interact with Mastodon from scripts.

Conclusion

Despite the great ideas behind the Fediverse's services and the simplicity and elegance it provides application programmers, a harsh truth is that a social media service is only as good as the people who use it. Who hasn't said to themselves "This is rubbish, but I can't leave because all these people I need to keep in touch with are here" when browsing through Facebook, Twitter, or LinkedIn? The fact is that the time when a social media platform could attract new users on the merits of its features has long gone.

However, the Fediverse has many things going for it – not least of which is that you can agglomerate all the users as one audience even though each is using a different service – and that may give it a fighting chance. ■■■

Info

- [1] Fediverse: <https://fediverse.party/>
- [2] diaspora*: <https://diasporafoundation.org/>
- [3] Identi.ca: <https://identi.ca/>
- [4] pump.io: <http://pump.io/>
- [5] ActivityPub: <https://www.w3.org/TR/activitypub/>
- [6] Pixelfed: <https://pixelfed.org/>
- [7] Mastodon: <https://joinmastodon.org/>
- [8] Funkwhale: <https://funkwhale.audio>
- [9] Hubzilla: <https://zotlabs.org/page/hubzilla/hubzilla-project>
- [10] Fedilab: <https://fedilab.app/>
- [11] Mastodon.py: <https://github.com/halcy/Mastodon.py>
- [12] Mastodon.py documentation: <https://mastodonpy.readthedocs.io/en/stable/>
- [13] `status()`: <https://mastodonpy.readthedocs.io/en/stable/#reading-data-statuses>
- [14] toot dictionary: <https://mastodonpy.readthedocs.io/en/stable/#toot-dicts>
- [15] account dictionary: <https://mastodonpy.readthedocs.io/en/stable/#user-dicts>

MOBILE USERS

search for us today at your digital newsstand!



Only a swipe away!

Download our convenient digital editions for your iPad, iPhone, or Android device.



Visit our apps page for more information:
shop.linuxnewmedia.com/us/apps

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 <http://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.



ApacheCon Europe

Date: October 22–24, 2019

Location: Berlin, Germany

Website: <https://aceu19.apachecon.com/>

ApacheCon is the official conference series of The Apache Software Foundation. Join us to discuss the latest innovations in containers, cloud, DevOps, IoT, servers, web frameworks, plus many other Apache projects and communities in a collaborative, vendor-neutral environment.

DrupalCon Amsterdam 2019

Date: October 28–31, 2019

Location: Amsterdam, Netherlands

Website: <https://events.drupal.org/amsterdam2019>

DrupalCon brings together thousands of people from across the globe who use, develop, design, and support the Drupal platform. Join us for this opportunity to access and engage with the decision-makers and minds who make Drupal happen.

LISA19

Date: October 28-30, 2019

Location: Portland, Oregon

Website: <https://www.usenix.org/conference/lisa19>

LISA is the premier conference for operations professionals. Sys admins, systems engineers, IT operations professionals, SRE practitioners, developers, IT managers, and academic researchers will share real-world knowledge about designing, building, securing, and maintaining critical systems.

Events

All Systems Go!	September 20-22	Berlin, Germany	https://all-systems-go.io/
Open Networking Summit open-Europe 2019	September 23-25	Antwerp, Belgium	https://events.linuxfoundation.org/events/networking-summit-europe-2019/
Deutsche OpenStack Tage 2019	September 24-26	Berlin, Germany	https://openstack-tage.de/
JAX London 2019	October 7-10	London, United Kingdom	https://jaxlondon.com/
All Things Open	October 13-15	Raleigh, North Carolina	https://allthingsopen.org/
ApacheCon	October 22-24	Berlin, Germany	https://aceu19.apachecon.com/
Open Source Summit + Embedded Linux Conference Europe	October 28-30	Lyon, France	https://events.linuxfoundation.org/events/open-source-summit-europe-2019/
Lisa19	October 28-30	Portland, Oregon	https://www.usenix.org/conference/lisa19
DrupalCon Amsterdam 2019	October 28-November 31	Amsterdam, Netherlands	https://events.drupal.org/amsterdam2019
Linux Security Summit Europe	October 31-November 1	Lyon, France	https://events.linuxfoundation.org/events/linux-security-summit-europe-2019/
Open Source Monitoring Conference (OSMC)	November 4-7	Nuremberg, Germany	https://osmc.de/
OSCamp Foreman	November 7	Nuremberg, Germany	https://opensourcecamp.de/
Linux App Summit	November 12-15	Barcelona, Spain	https://linuxappsummit.org/
Linux Presentation Day 2019	November 16	Cities across Europe	http://l-p-d.org/en:start
SC19	November 17-22	Denver, Colorado	https://sc19.supercomputing.org/
KubeCon + CloudNativeCon North America 2019	November 18-21	San Diego, California	https://events.linuxfoundation.org/events/kubecon-cloudnativecon-north-america-2019/
DevOpsDays Berlin	November 27-28	Berlin, Germany	https://devopsdays.org/events/2019-berlin/welcome/

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.

NOW PRINTED ON recycled paper from 100% post-consumer waste; no chlorine bleach is used in the production process.

Authors

Bernhard Bablok	30
Erik Bärwaldt	26
Swapnil Bhartiya	8
Paul Brown	90
Zack Brown	12
Bruce Byfield	40, 56, 62
Joe Casad	3
Mark Crutch	73
Jon "maddog" Hall	75
Ken Hess	58
Charly Kühnast	60
Christoph Langner	66
Rubén Llorente	50
Vincent Mealing	73
Graham Morrison	84
Scott Nesbitt	80
Dmitri Popov	34, 76
Mike Schilli	44
Mayank Sharma	16
Ferdinand Thommes	20

Contact Info

Editor in Chief

Joe Casad, jcasad@linux-magazine.com

Copy Editor

Amy Pettle

News Editor

Swapnil Bhartiya

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

© Watchara Khamphonsaeng, 123RF.com

Advertising

Brian Osborn, bosborn@linuxnewmedia.com
phone +49 89 3090 5128

Marketing Communications

Gwen Clark, gclark@linuxnewmedia.com
Linux New Media USA, LLC
2721 W 6th St, Ste D
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 © 2019 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 on recycled paper from 100% post-consumer waste; no chlorine bleach is used in the production process.

Distributed by Seymour Distribution Ltd, United Kingdom

LINUX PRO MAGAZINE (ISSN 1752-9050) is published monthly by Linux New Media USA, LLC, 2721 W 6th St, Ste D, 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, 2721 W 6th St, Ste D, Lawrence, KS 66049, USA.

Published monthly in Europe as Linux Magazine (ISSN 1471-5678) by: Sparkhaus Media GmbH, Ziebländstr. 1, 80799 Munich, Germany.

Approximate

UK / Europe Sep 28
USA / Canada Oct 25
Australia Nov 25

On Sale Date

Issue 228 / November 2019

Anonymous File Sharing

The TOR browser is a well known tool for anonymous surfing. But what if you want to do more than surf? OnionShare is a file sharing tool that lets you share files without revealing your identity or location.

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: www.linux-magazine.com/newsletter

Image © Watchara Rojjanasain, 123RF.com

IT Highlights at a Glance

The collage features several newsletters and articles:

- ADMIN HPC**: "The New IT ADMIN - your source for technical solutions to real-world problems." Includes sections for "Further Reading" (GUI of Task Based Interface, Shared Storage with NFS and SDRFS, Intel Linux New Processor Line, Quantum Computing Milestone Achieved, OES Complete Check Master Survey) and "Highlights" (OpenACC - Data Management, HPC Up Close, Further Reading, Highlights).
- ADMIN Update**: "ADMIN Update - Hottest Links" with a "REGISTER FOR YOUR FREETICKET TODAY!" banner for an event on 12-13 March 2019 at ExCel, London. Includes "ADMIN Update - Hottest Links" (Migration in the Cloud, A 19 Year-Old Bug in WinRAR, Hackers Started Exploiting Drupal Bug, Root, Cheat) and "Highlights" (Migration in the Cloud, A 19 Year-Old Bug in WinRAR, Hackers Started Exploiting Drupal Bug, Root, Cheat).
- Linux Update**: "EXPLORING THE WORLD OF LINUX" with "FEATURED ARTICLES" (Remote Repositories in Git, Linux Torvalds Welcomes 2019 with Linux 5.4, GitHub Offers Free Private Repositories, Open Source Tools for Writers) and "FURTHER READING" (Microsoft Gets an Open Source Web Browser, Version Control with Git, Debian, Ubuntu, and Other Distros are Leaving Users Vulnerable, Docker 1.21, Open Source Tools for Writers).
- Linux Update**: "DISCOVER LIBREOFFICE!" with "Explore the free office suite to create your own: Word processing docs • Spreadsheets • Presentations • Databases".
- Linux Update**: "Help your friends and colleagues make the switch to Linux! ORDER NOW!"
- Linux Update**: "2018 Archives Available Now!"

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: <https://bit.ly/HPC-ADMIN-Update>

Linux Update: <https://bit.ly/Linux-Update>

When You Think Cloud Infrastructure Think Supermicro

To Scale Your Compute and Storage
Systems and Optimize Your
Resource Savings Featuring
2nd Generation Intel® Xeon®
Scalable processors



Learn More at www.supermicro.com

© Supermicro and Supermicro logo are trademarks of Super Micro Computer, Inc. in the U.S. and/or other countries.

